

Understanding radio pulsars using modern broad-band instruments



Lucy Sarah Oswald
Hertford College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2021

For the people I love.

Acknowledgements

I'd like to begin by thanking Aris, my supervisor, for years of encouragement, guidance and enthusiasm. A friend once said to me "It's always pulsars with you Lucy" and the only possible answer was "Yes!", because what else could be more exciting? I can't imagine enjoying a DPhil more and so much of that is down to you: thank you for giving me so much to enjoy, learn and think about.

Next, a thank you to Simon, with whom I've had so many interesting discussions and who has given me so much advice and support. Thank you for hosting me in Sydney twice and thank you for all the discussions and Zoom calls. It's been amazing.

So many people in Oxford Astrophysics deserve a mention here. Bettina: thank you for the support, the exciting physics and the pumpkin carving! The pulsar group: Chris, Arun, Isabella and Griffin all with interesting perspectives and fun conversations. Marisa for all the exciting work getting our scattering code publishable. Ben for all the chats, which I usually found a way to make tangentially linked to pulsars: thank you for your friendship and support, and for my pulsar T-shirt! Rebecca for all the fun at Pulsar Christmas parties. All my office mates in the Tower and 650. Garret and Lance have been really helpful as Directors of Graduate Studies, Caroline organised some really useful teaching in first year and Pat was always on hand as my College Tutor. I don't know what any of us would be doing without Ashling: thank you for the fun in South Africa and on Zoom quizzes, and for your endless patience helping me with forms and travel bookings. Also to Martin for the multiple times you helped me with IT issues and for always being ready with an interesting story.

I can't imagine what my experience would have been like without Pulsar Coffee each Tuesday. It has been a highlight of my week throughout the DPhil. Ryo should have a particular mention here as someone always willing to present something interesting at short notice. Thank you to all the regulars for your enthusiasm about pulsar science, in particular to Peter and John for interesting email discussions about new ideas. Special thanks to Jocelyn for your calm guidance and support. You've been so helpful throughout and I have really appreciated it.

A note now for the 12 o'clock lunch club with our captain Eloïse. Regular meals in the canteen may have been ended by the pandemic but the support and fun certainly

didn't! Thanks to all of you in my cohort: Mike, Sam, Nathan, Kit, Andrea, Xing, Norbert, Darsh, Alvaro, David, Dom and Dina.

Thanks now to all of the people who hosted me and helped me have a good time on my travels: Dipanjan and Minhajur at NCRA and the GMRT, Jason and Vlad at ASTRON, Michael Kramer at the MPIFR, Marisa at SARA0, and so many people at CSIRO. Particular thanks to Dick Manchester for the table tennis, Lawrence for the advice with the Parkes observing, and George for the observing advice and the flamingos! Also Charlotte for all your help with the Parkes observations and the interesting chats about pulsar polarization, and to Mark for fascinating conversations and sunshine at Manly.

I should mention now some people without whom I may not have decided to do a DPhil. Mr Wilkinson, whose teaching inspired me to study physics at university. Owen Saxton, whose teaching made me fall in love with physics and made my undergraduate years so fantastic. Prasenjit, who hosted me three times at UZH and taught me so many interesting ways of looking at things.

I wouldn't have enjoyed my DPhil half as much without also having my friends there. Thank you especially to Emily, Lucy, Hannah and Shreya for all the fun and the change of pace you brought with all those brilliant videocalls. Thank you so much to my friends Sheralyn and Eleanor for inspiring me with their stories and showing me that I shouldn't be frightened of writing. Keble College Chapel Choir kept me sane through years of research: particular thanks to Larissa, Matthew, Robert, Beth and Samantha, and hooray for team Alto! My University Challenge team, Matthew, Bridget, Matt and Harry, kept me sane through months of write-up. Thank you also to the Mifsud family for your support, particularly Rachel for the cups of tea and Vincent for all the help and advice. Thank you for all the discussions, for the mock interview and for proof-reading my thesis. It's unbelievably kind of you and so helpful.

Next I want to thank my amazing parents, Caroline and Chris, for their love and enthusiasm. Daddy: thank you for putting so much effort into understanding my work so that you could ask me questions. I've loved our conversations and it's always given me new perspectives to consider. Mummy: thank you for supporting me throughout, encouraging me when I was stuck and being excited with me when I learned something new, it's been amazing. Thank you to you both for your proof-reading efforts, I love you.

My wonderful brother Tim, thank you for all the philosophical discussions that taught me how to ask deeper questions, and for your proof-reading. Thank you also

to my extended family: Clare, June, Crispin, Karen, Jane and David for always being interested and supportive. To Rob and Toby, thank you for my ten weeks staying with you for my undergraduate internship where I really learned how to complete a long-term project (and thanks for the driving lessons too!). Grandad, you have been an inspiration in academia and I've loved our many conversations about all sorts of topics, thank you for being there for me. Granny, you were always so supportive of whatever I chose to do. I'm so pleased to be following in your footsteps by completing a PhD and it makes me happy to think you'd have been proud of me. Papa, you brought me fun, calm support and so much love throughout my life. I would have loved to tell you about pulsars, but I'm glad I got to tell you stories about how much fun I was having at university, and I'm glad you told me when it was time for "serious application": I knuckled down and got through my exams.

I'll end with a big thank you to my wonderful boyfriend Miffy. Thank you for encouraging me through this, along with your team headed up by Little Croc, and for your inexhaustible supplies of patience and cups of tea. Thank you so much for all the effort you put into proof-reading my entire thesis. You've been there for me every day and it's been so much fun. I'm so lucky to have you and I love you.

Statement of Originality

The work presented in this thesis was undertaken at the Department of Astrophysics at the University of Oxford between October 2017 and March 2021 under the supervision of Professor Aris Karastergiou. It was funded by a STFC studentship.

I hereby declare that no part of this thesis has been, or is being, submitted in support of another degree, diploma, certificate or other qualification at the University of Oxford or elsewhere.

This thesis is the result of my own work unless otherwise stated.

Chapters of this work have been published in part as papers in the journal MNRAS.

Chapter 4: *Understanding the radio beam of PSR J1136+1551 through its single pulses*, (Oswald et al., 2019). I am the first author and my co-authors are Aris Karastergiou and Simon Johnston.

Chapter 5: *Pulsar polarimetry with the Parkes ultra-wideband receiver*, (Oswald et al., 2020). I am the first author and my co-authors are Aris Karastergiou and Simon Johnston.

Chapter 6: *The Thousand-Pulsar-Array programme on MeerKAT – V. Scattering analysis of single-component pulsars*, (Oswald et al., 2021). I am the first author and my co-authors are A. Karastergiou, B. Posselt, S. Johnston, M. Bailes, S. Buchner, M. Geyer, M. J. Keith, M. Kramer, A. Parthasarathy, D. J. Reardon, M. Serylak, R. M. Shannon, R. Spiewak, W. van Straten, and V. Venkatraman Krishnan.

This thesis was written in L^AT_EX, using the ociamthesis.cls class¹ and MNRAS bibliography style file. The software I developed was written in Python and made considerable use of numpy (Van Der Walt et al., 2011), matplotlib (Hunter, 2007) and pandas (McKinney, 2010).

¹<https://www.maths.ox.ac.uk/members/it/faqs/latex/thesis-class>

The copyright of this thesis rests with its author. No quotation or figure from it, or information derived from it, may be published without the prior consent and acknowledgement of its author.

Lucy Oswald

March 2021

Abstract

The canonical model of a pulsar is insufficient to describe the variety and variability of its radio emission. Pulsars are neutron stars, with intense gravitational and magnetic fields, which emit bright beams of radio waves that co-rotate with the spinning star. The regularity of pulsar rotation, and hence the arrival times of pulses of radio emission, means that pulsars are used as clocks in space to test fundamental theories of physics and to search for gravitational waves. However, their accuracy as clocks is limited by the fact that we do not fully understand how pulsars produce their radio beams, and so cannot predict their emission behaviour completely.

Throughout the history of pulsar science, new telescopes and updated technology have expanded the complexity observable in the shapes and properties of observed pulse profiles from the radio pulsar population. This growing dataset has raised as many new questions as it has answered about the nature of pulsar radio emission, the properties of the pulsar population and our ability to characterize the behaviour of pulsars and their environments with physical laws. However, modern instruments offer a broad-band view of radio pulsars. This allows us to probe, for the first time, the continuous evolution of pulsar radio emission over a wide frequency range with a single instrument.

This thesis uses the expanded observational capabilities of new broad-band instruments to make progress in answering fundamental questions about pulsar radio emission. I apply a statistical approach to the frequency evolution of single pulses of PSR J1136+1551 observed by the GMRT, in order to constrain the cross-section and frequency-dependent emission heights of its radio beam structure. This work shows that the beam structure of J1136+1551 is best described by a fan beam model and that it is important to include the effects of orthogonal polarization mode interaction to explain the frequency-dependent behaviour. In order to maximize the understanding available with broad-band observations,

it is important to address the interaction between pulsar radio emission and the interstellar medium through which it propagates. I create an algorithm for correcting the effects of the interstellar medium in broadband polarimetric data from the Parkes Ultra-Wideband receiver to reveal the intrinsic polarization behaviour of PSRs J1056–6258 and J1359–6038. Finally, I apply understanding of intrinsic pulsar behaviour to constrain models of pulsar scattering by the interstellar medium, performing a survey of the scattering properties of 84 single-component pulsars observed with the MeerKAT telescope.

This work reveals the capacity of new broad-band observations to expand our understanding of pulsar radio emission. It constrains understanding of both pulsar beam structure, including frequency-dependent emission heights, and the structures that make up the interstellar medium. The results highlight the importance of accounting for the behaviour of both the pulsar and the interstellar medium simultaneously when analysing broadband observations, and future work will focus on applying these modelling approaches to pulsars with complex profile shapes and polarization properties.

*Then, through the vast and gloomy dark,
There moves what seems a fiery spark*
Edward Lear

O oriens, splendor lucis aeternae

and this is the wonder that's keeping the stars apart
ee cummings

*And truly, I reiterate, nothing's small!
No lily-muffled hum of a summer-bee,
But finds some coupling with the spinning stars*
Elizabeth Barrett Browning

Contents

1	Introduction	1
1.1	The motivations of pulsar science	1
1.1.1	Radio pulsars: the canonical model	2
1.1.2	Behaviour of the pulsar population	9
1.1.3	Pulsars as tools	12
1.1.4	The interstellar medium	14
1.1.5	Summary	19
1.2	A history of pulsars and observing bandwidth	19
1.2.1	Early observations and theory	19
1.2.2	The narrowband era	20
1.2.3	Polarimetry	22
1.2.4	Simultaneous multi-frequency observations and wider bands .	26
1.2.5	New broad-band capabilities	28
1.3	Open questions	29
1.4	Structure of the thesis	30
2	Definitions	32
3	Observations and data processing	36
3.1	How pulsars are observed	36
3.1.1	Observational decisions: managing competing factors	36
3.1.2	From telescope to computer: recording and processing of pulsar observations	37
3.2	Observing projects	39
3.2.1	Simultaneous observations with the Parkes and GMRT telescopes	42
3.2.2	Long-term observing programmes with the Parkes telescope .	44
3.2.3	MeerKAT: MeerTime and the Thousand-Pulsar-Array	45
3.3	Observing methods and data post-processing	45

3.3.1	GMRT observations	45
3.3.2	Calibration and RFI excision	48
3.3.3	Data processing for Parkes P574 and MeerTime TPA observations	51
3.3.4	Summary table of data products	54
3.4	Measuring the effects of the ISM	54
3.4.1	Dispersion and Faraday rotation	55
3.4.2	Scattering and scintillation	57
3.5	Summary	58
4	Methodology	59
4.1	Simulating the observables of pulsar physics	60
4.1.1	Assumptions of a pulsar emission model	65
4.1.2	The mathematics of the simulations	68
4.1.3	Derivation of the equations	69
4.1.4	Pulsar beam cross-sections	71
4.1.5	Mathematical derivation of the fan beam cross-section emission peak	72
4.1.6	Orthogonal polarization modes	75
4.2	Quantifying single pulse frequency evolution	76
4.3	Correcting for propagation effects in broadband polarimetric data . .	79
4.4	Visualizing broadband data	82
4.5	Modelling scattered profiles: further application of MCMC	83
4.6	Summary	87
5	Understanding the radio beam of PSR J1136+1551 through its single pulses	88
5.1	Setting up simulations of single pulse evolution	88
5.2	Subpulse independence	92
5.3	Evolution of subpulse separation	95
5.4	Results: comparing observed single pulse properties to different beam models	95
5.4.1	Hollow cone	95
5.4.2	Fan beam	97
5.4.3	Orthogonal mode propagation	99
5.5	Calculation of emission heights	100
5.6	Summary	104

6	Structure in broad-band polarimetric pulsar observations	106
6.1	Testing the modelling process with simulations	108
6.2	Modelled pulsars	114
6.2.1	PSR J1359–6038	117
6.2.2	PSR J1056–6258	119
6.2.3	Broad-band alignment and beam structure	120
6.3	Summary	124
7	Scattering analysis of single-component pulsars with the MeerKAT telescope	126
7.1	Modelling process	127
7.2	Applicability of the model to the dataset	129
7.3	Understanding model break-down	135
7.4	Results of scatter modelling	141
7.5	Table of results and graphs	146
7.6	Summary	149
8	Discussion	150
8.1	Frequency evolution of polarimetric data	150
8.1.1	The information gain from broad-band observations	151
8.1.2	The effect of circular polarization	152
8.2	Modelling philosophy	156
8.3	Pulsar beam structure	160
8.3.1	Emission heights for the hollow cone and fan beam models	160
8.3.2	Emission heights from the delay-radius relation	162
8.3.3	Impact of assumptions about emission region shape	164
8.3.4	Emission heights in the context of understanding pulsars	164
8.4	Understanding inhomogeneous regions in the ISM	165
8.4.1	Scattering parameters in context	165
8.4.2	Scattering at low and high frequencies	172
8.5	Overall discussion and future work	174
8.6	Summary	175
9	Conclusions	176
9.1	Intentions and context	176
9.1.1	Context: ideas and techniques	176
9.1.2	Context: comparison with other work	177

9.2	Approach and goals	178
9.3	Answers and questions	179
9.3.1	Theories ruled out by my results	179
9.3.2	How my work has progressed understanding	179
9.3.3	Progress with the open questions	180
9.4	Going forward	183
A	Scatter modelling results: pulse profiles with best fit models	185
B	Scatter modelling results: frequency dependence of τ and σ	207
	Bibliography	217

List of Figures

1.1	Diagram of canonical pulsar model	3
1.2	Diagram of polarization states on the Poincaré sphere	5
1.3	Vela pulsar, polarized integrated profile	6
1.4	Single pulses of J1136+1551	7
1.5	Diagram of the fiducial plane and fiducial point	8
1.6	PSR J1559–4438: integrated polarized profile and frequency evolution	10
1.7	$P\dot{P}$ diagram for the known pulsar population	13
1.8	Dispersed intensity profile of PSR J1056–6258	15
1.9	Scattered pulse profile for PSR J1818–1422	17
1.10	Diagram of the thin screen model for scattering pulsar radiation . . .	18
1.11	Hollow cone and fan beam models of the pulsar radio beam, side views	23
1.12	PSR J0738–4042: integrated polarized profile and frequency evolution	25
1.13	PSR J0452–1759: integrated polarized profile and frequency evolution	27
3.1	MeerKAT telescope antenna	42
3.2	Single pulse of PSR J1136+1551 observed with GMRT	47
3.3	Continuum bandpass of source 3C 241	50
4.1	Simulated integrated pulse profile	61
4.2	Simulated pulse profile, now broad-band	63
4.3	Figure describing scattering model	64
4.4	Illustration of beam cross-section: hollow cone model	67
4.5	Pulsar emission vectors: rotation, magnetic and line of sight	70
4.6	Parametrization of dipolar field line	71
4.7	Illustrations of beam cross-sections: hollow cone and fan beam models	73
4.8	Parametrization of beam cross-section and line of sight	73
4.9	Relationship between reference field line and line of sight peak for the fan beam model	75
4.10	Illustration of how DM choice affects profile relationships across frequency	78

4.11	PSR J1056–6258: waterfall plots of intensity and polarization	83
4.12	MCMC chains from scatter modelling of three example pulsars	86
5.1	Single pulse histograms for PSR J1136+1551	90
5.2	Relationships between emission region and line of sight for hollow cone and fan beam models	93
5.3	Scatter plots of subpulse separation vs. pulse midpoint: data and randomly combined	94
5.4	PSR J1136+1551: scatter plot of spectral index of subpulse separation vs. separation at a single frequency	96
5.5	Separation spectral index scatter plots: data and hollow cone model simulation	97
5.6	Separation spectral index scatter plots: data and fan beam model sim- ulation	98
5.7	Cumulative distributions of separation spectral index: data and fan beam model simulation	98
5.8	Fan beam model with mode divergence: subpulse histograms	101
5.9	Separation spectral index scatter plots: data and fan beam model sim- ulation with mode divergence	102
5.10	Cumulative distributions of separation spectral index: data and fan beam model simulation with mode divergence	102
5.11	Emission height vs. frequency: fan beam model	103
6.1	Integrated pulse profiles: PSRs J1056–6258 and J1359–6038	109
6.2	Basic simulated profile: alignment results	111
6.3	Simulated profile with flatter PA: alignment results	112
6.4	Simulated profile including emission height variation: alignment results	115
6.5	PSR J1359–6038: alignment results	118
6.6	PSR J1056–6258: alignment results	120
6.7	PSR J1359–6038: broad-band PA before and after alignment	121
6.8	PSR J1359–6038: broad-band intensity before and after alignment . .	121
6.9	PSR J1056–6258: broad-band PA before and after alignment	122
6.10	PSR J1056–6258: broad-band intensity before and after alignment . .	122
7.1	Illustration of scattering strength dependence on frequency and distance	128
7.2	PSR J1818–1422 scatter modelling results: fits to profiles	131
7.3	PSR J1818–1422 scatter modelling results: corner plot	132

7.4	PSR J1818–1422 scatter modelling results: corner plot for α	133
7.5	Scatter modelling results: τ and σ vs. frequency for PSRs J1818–1422, J1743–3153 and J1653–4249	134
7.6	Illustration indicating relative contributions of intrinsic and scattered components of profile under different conditions	137
7.7	Histograms of the ratios of measurement results σ/τ	138
7.8	Simulation: double-component scattered pulse profile at two frequencies	139
7.9	τ and σ vs. frequency for PSR J1653–4249 and for the simulated double-component pulse profile	140
7.10	Histograms of values of α	142
7.11	α vs. τ_0/P for good and bad model fits	144
7.12	Histogram of ΔDM	145
7.13	$\Delta\tau$ vs. frequency for PSR J1630–4733: correction to τ resulting from DM correction	146
8.1	PSR J0452–1759: alignment of discrete profiles and predicted mode decomposition	153
8.2	PSR J0452–1759: broad-band polarization properties	154
8.3	PSR J0908–4913 main pulse: broad-band polarization properties . . .	157
8.4	PSR J0908–4913 interpulse: broad-band polarization properties . . .	158
8.5	Emission heights vs. input footprint parameter ratio for hollow cone and fan beam models	161
8.6	Plot of τ vs. DM	168
8.7	Plot of τ vs. DM including results from NE2001	169
8.8	Histograms of τ residuals for this work and NE2001	170
8.9	Plot of α vs. DM	173
A.1	All scattering results: plots of pulse profiles with successful model fits	186
B.1	All scattering results: plots of τ and σ vs. frequency	208

List of Tables

3.1	Details of pulsar observations	54
5.1	Table of parameters (mean μ , standard deviation σ and their uncertainties) of Gaussian fits to the subpulse distributions for the data and all three simulations.	91
6.1	DM and RM results for PSRs J1056–6258 and J1359–6038	117
7.1	Scatter modelling results for 84 pulsars	147
8.1	Previous and new measurements of scattering spectral index α	166

Chapter 1

Introduction

1.1 The motivations of pulsar science

The enduring appeal of pulsars is that, as astrophysical objects, they represent a combination of simplicity and mystery. From the beginning, the observation of a series of regular pulses revealed pulsars to be rotating neutron stars with the capacity to generate bright beams of radio waves. The regularity of pulsar rotation means that they can be used to do fundamental tests of physics, as precise clocks in some of the most extreme environments in the Universe. Beneath this simplistic description, however, is a world of complexity, variability and unanswered questions. It is not fully understood how pulsars generate their radio beams. Variety and variability is observed that is not yet explained, not just between different pulsars, but from pulse to pulse and at different observing frequencies.

Progress in understanding is led by the development of telescope technology. This thesis explains how new observations made with the latest broad-band instruments can be used to address the open questions about pulsar radio emission. The results reveal information about the beam shape and polarization structure of pulsar radio emission, and how this information can be applied to understanding the structures of the interstellar medium.

In order to explain the questions being addressed by this thesis, I will first describe the basic theoretical framework upon which our understanding of radio pulsars is built. Within this context I will explain how pulsars can be used as tools in fundamental physics experiments and to investigate the interstellar medium, and the ways in which this theoretical framework leaves questions unanswered. In Section 1.2 I will describe the parallel development of pulsar theory and advances in radio technology, up to the broad-band observing capabilities of the present day. This underlies

the context for the open questions investigated in this thesis, which are presented in Section 1.3.

1.1.1 Radio pulsars: the canonical model

The canonical description of a pulsar is a rotating neutron star with a dipolar magnetic field, the magnetic axis of which is misaligned with respect to the rotational axis by an angle α . A commonly assumed model for the generation of the observed radio pulses is that of coherent emission of radio waves by charged particles streaming along the curved dipolar magnetic field lines in the pulsar magnetosphere. This generates a beam of radio waves that co-rotates with the pulsar like a lighthouse beam, so that an observer on Earth will detect the beam only for that fraction of the rotational period when it is directed towards the Earth. Fig. 1.1 depicts the key features of the canonical pulsar. The line of sight from the observer to the pulsar tracks a path across the beam as the pulsar rotates, which can be described fully by the geometrical angle α and the impact parameter β , which is the angle of closest approach of the line of sight to the magnetic axis.

In the assumption of curvature emission, the radio waves are thought to be generated only by particles streaming along the open magnetic field lines, as labelled on Fig. 1.1. Given a rotational period P , there will be a perpendicular distance from the pulsar rotational axis, $cP/2\pi$, at which co-rotating material would have to travel at the speed of light c . Magnetic field lines contained within the light cylinder will form closed loops, however those extending beyond this radius are open. The region of the pulsar surface from which these open field lines originate is centred upon the magnetic axis and is called the polar cap. It is from some or all of these open field lines that pulsar radio emission is expected to originate, so that the last open field line defines an upper limit to the angular size of the radio beam (Lorimer & Kramer, 2005).

If the pulsar beam is sufficiently bright, a radio telescope will detect it as a series of single pulses. Fainter pulsars are detected by folding an observation on the pulsar's rotational period, so that the pulsed emission is added together to increase the signal-to-noise ratio (S/N). Summing thousands of single pulses together produces a stable integrated pulse profile, the shape of which is consistent across time and is inferred to describe the average emission distribution of the radio beam. Polarimetric pulse profiles are recorded as Stokes parameters I , Q , U and V , where Stokes I is the total intensity and the polarization state is given by the relationship between the other three parameters. These parameters are often represented graphically on the

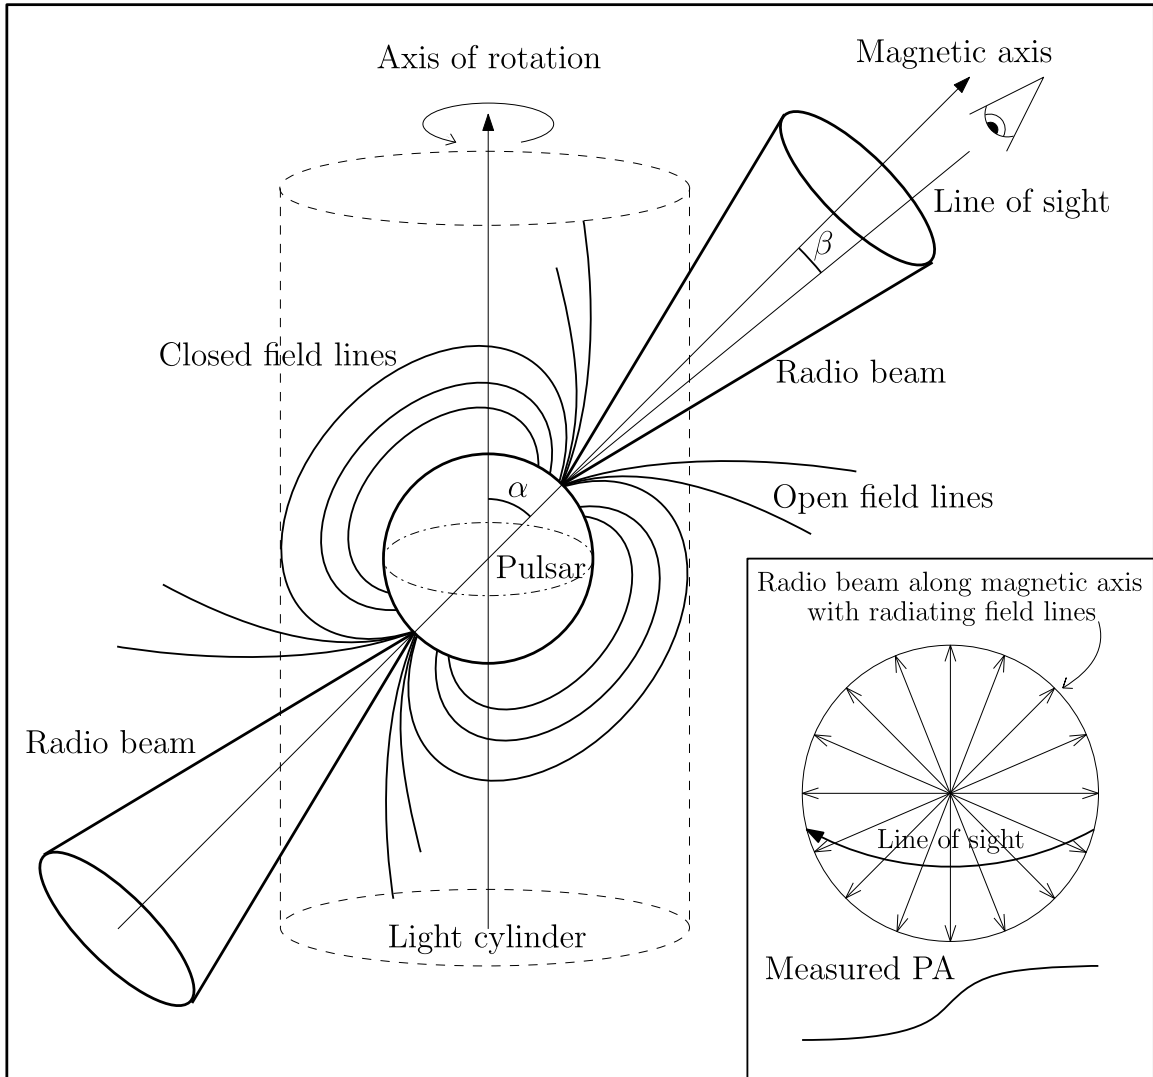


Figure 1.1: Diagram of canonical pulsar model.

Neutron star with axes of rotation and dipolar magnetic field marked with arrows and the radio beam pointing along the magnetic axis. As the star rotates, the magnetic axis revolves about the axis of rotation (which in the diagram is lying in the plane of the page), so that the radio beam pointing along that axis sweeps through space. The angle between these axes is labelled α and the angle of closest approach between the magnetic axis and the line of sight is labelled β . Closed and open dipolar magnetic field lines are marked, where the closed field lines are entirely contained within the light cylinder. Inset: bird's-eye-view of the radio beam looking down the magnetic axis so that the projected magnetic field lines radiate outwards. The line of sight (labelled) cuts across the beam on a curved path and the resultant measured angle of linear polarization or position angle (PA) varies smoothly as the angle of the projected field lines varies along that path. Diagram made using Ipe and Inkscape.

Poincaré sphere with axes Q , U and V (see Fig. 1.2), where Q and U describe the linear polarization $L = \sqrt{Q^2 + U^2}$ and its angle, and V describes the circular polarization. The angle of linear polarization is known as the position angle or PA, and is given by $\text{PA} = -0.5 \arctan(U/Q)$. Fig. 1.3 gives an example of an integrated profile for the Vela pulsar (PSR J0835–4510). It shows total intensity Stokes I , linear polarization L and its angle the PA, and circular polarization V . Fig. 1.4 shows Stokes I only for four of the individual single pulses for PSR J1136+1551.

When analysing pulsar observations the concepts of rotational phase ϕ and the fiducial point ϕ_0 are relevant. The assumption that the time between each observed pulse corresponds to the pulsar period presupposes that the pulsar radio beam is co-rotating with the neutron star, so that there is a point on the pulse profile that can be mapped directly onto a fixed position on the neutron star surface. The pulse profile can then be mapped onto rotational phase, where phase $\phi = 0^\circ$ corresponds to the point where the line of sight crosses the fiducial plane. This is the plane containing the magnetic and rotational axes, which runs through the centre of the beam. In the absence of additional information, the fiducial point of $\phi_0 = 0^\circ$ is commonly assigned to either the peak of the pulse profile or the centre of the profile. Fig. 1.5 depicts the fiducial plane, fiducial point and rotational phase.

The simplest assumption about the pulsar’s magnetic field is that it is dipolar, with the radio beams being emitted out of its North and South poles respectively. This assumption leads to an expectation about the polarization properties of pulse profiles. As indicated on the inset diagram in Fig. 1.1, if we adopt a bird’s-eye-view of the pulsar’s magnetic field looking down the magnetic axis, the projected field lines radiate out from the axis origin. The line of sight to the observer cuts across this projected magnetic field along some path defined by the pulsar geometry, described by the angles α and β . At each point along this path, the angle between the line of sight and the projected field line will vary smoothly, being exactly 90° at the halfway point. This halfway point lies on the fiducial plane of the pulsar. Assuming that the polarization of the radio emission is related to the field line environment of where it is generated, we might expect the measured projected angle of linear polarization, known as the position angle or PA, to be related to the projected direction of the field lines at each point along the line of sight. We would therefore expect the PA profile to evolve smoothly with phase ϕ in the same way as the angle between the line of sight and the projected field line. This generates an S-shaped curve, where the phase of steepest gradient corresponds to the fiducial point, described by the




Figure 1.2: Diagram of polarization states on the Poincaré sphere.

Diagram visualizing the Poincaré sphere construction used to relate the polarization states of the Stokes parameters Q , U and V . The small red curved or double-ended arrows represent the polarization states defined by the red vectors, as linearly, circularly or elliptically polarized respectively. Linearly polarized light is represented by a vector in the QU plane, where the angle of linear polarization Ψ is half the azimuthal angle of the vector. Circularly polarized light points along the V axis, defined such that positive V corresponds to right-handed circular polarization and negative V to left-handed circular polarization. Any vector at another position upon the surface of the sphere, such as the one shown, will have elliptical polarization. Light that is not completely polarized would be represented by a vector that lies inside the Poincaré sphere and does not touch its surface. Figure made using Inkscape.




Figure 1.3: Vela pulsar, polarized integrated profile.

Integrated pulse profile for the Vela pulsar (PSR J0835–4510) at 1087 MHz as observed with the Ultra-Wideband receiver on the Parkes radio telescope and showing only the on-pulse region. Top: PA profile, for which the measured angle is degenerate with a rotation of 180° , making $\pm 90^\circ$ equivalent on the plot. Bottom: total intensity in black, linear polarization ($L = \sqrt{Q^2 + U^2}$) in red and circular polarization (Stokes V) in blue, all normalized so that the peak of the profile has an intensity of 1. Figure made with matplotlib.




Figure 1.4: Single pulses of J1136+1551.

Four double-peaked single pulses for PSR J1136+1551, observed with the Giant Metrewave Radio Telescope. For each pulse the normalized profile shape is shown at seven discrete frequencies between 330 and 470 MHz, as labelled for the top left pulse. The data have been aligned so that the integrated profile widens symmetrically about phase $\phi = 0$, which is marked with a vertical line on the plots. Figure made with matplotlib.




Figure 1.5: Diagram of the fiducial plane and fiducial point.

Three arrows represent three axes with respect to the spherical neutron star: the rotational axis ($\underline{\Omega}$), magnetic axis (\underline{m}) and line of sight (LoS). The fiducial plane is the plane containing axes $\underline{\Omega}$ and \underline{m} , and the fiducial point (red circle) is the rotational phase at which this plane is parallel to the line of sight. α is the angle between $\underline{\Omega}$ and \underline{m} . The impact parameter β is defined as the angle between \underline{m} and the LoS when at the fiducial point. Figure made using Inkscape.

following equation:

$$\tan(\text{PA} - \text{PA}_0) = \frac{\sin(\phi - \phi_0) \sin \alpha}{\sin(\alpha + \beta) \cos \alpha - \cos(\alpha + \beta) \sin \alpha \cos \phi}, \quad (1.1)$$

where PA_0 is the value of the PA at the fiducial point ϕ_0 . This theory, known as the rotating vector model, was first espoused by Radhakrishnan & Cooke (1969) to explain the PA profile of the Vela pulsar (Fig. 1.3), with the mathematical formulation described by Komesaroff (1970). It has since been applied successfully to explain the PA profile shapes of many other pulsars. However, when considering the pulsar population as a whole, the shapes of pulsar PA profiles tend to exhibit considerable additional complexity not explained by this theory alone.

1.1.2 Behaviour of the pulsar population

The integrated profile of each pulsar is unique, and their morphologies are varied, nevertheless there are some key similarities across the population. Profiles are made up of one or more components that are approximately Gaussian in shape and these components can be either clearly separated or overlapping and blended. Considerable effort has been directed into classifying profile shapes into categories, such as doubles and triples (see for example Rankin, 1983), but the majority of pulsars do not fall unambiguously into a single category. Integrated profiles are formed by summing thousands of single pulses (such as those in Fig. 1.4) to form a stable profile. Some pulsars exhibit interpulses: a second pulse at half the rotational period. Those pulsars are inferred to have a value of α close to 90 degrees, so that both beams are observable along the line of sight.

The flux density of pulsar emission exhibits power law behaviour across frequency so that, in general, profiles are fainter at higher frequencies, though some pulsars also show a spectral turnover at low frequencies. The spectral index of this power law is observed to vary for different components of a profile, so that the relative of brightness of different components tends to vary with frequency, dramatically altering profile shape in some cases: see for example PSR J1559–4438 in Fig. 1.6. Further profile evolution with frequency is observed, particularly at frequencies below 1 GHz, where the pulse width, and the spacing between profile components, generally increases with decreasing frequency (e.g. Johnston et al., 2008).

Since the first observations of PSR B1919+21, for which a pulse is detected once every 1.337 seconds (Hewish et al., 1968), radio pulsars have been discovered with rotational periods ranging from 1.4 ms (Hessels et al., 2006) to 23.5 s (Tan et al., 2018),




Figure 1.6: PSR J1559–4438: integrated polarized profile and frequency evolution. Left: integrated profile for PSR J1559–4438 at 1086 MHz, observed with the Parkes Ultra-Wideband receiver. The figure is formatted as for the Vela pulsar in Fig. 1.3, that is, PA profile at the top and intensity profile underneath, where total intensity is shown in black, with linear and circular polarization in red and blue respectively. Right: the total intensity profile at 8 frequencies between 810 and 3780 MHz plotted vertically with increasing frequency, to demonstrate how the profile shape evolves across the observing band. These profiles have been normalized so that the peak is at intensity = 1, which means that the relative shape variation is plotted, rather than the evolution of absolute flux. Figure made with matplotlib.

while X-ray pulsars have been discovered with periods of up to 36200 ± 110 s (Sidoli et al., 2017). Timing the arrivals of pulses reveals that the rotational period increases with time; the spin-down rate is known as \dot{P} . Plotting the periods of pulsars against their spin-down rates on a P - \dot{P} diagram, as in Fig. 1.7, reveals three key populations. Pulsars with periods of order 1 s are commonly known as “young pulsars”, or simply as “pulsars”, with the other categories being defined by the ways in which they differ from this main population. The second population have spin periods of order 1 ms and are known as “millisecond pulsars” (MSPs) or “recycled pulsars”, since these fast spin periods are inferred to result from input of angular momentum during the pulsar’s lifetime, generally through accretion from a binary companion causing the pulsar to spin up, though lone MSPs are also observed (Lorimer & Kramer, 2005). Of relevance also is the small subset of pulsars with particularly long periods and high spin-down rates. The pulsars with these characteristics tend also to exhibit flares, particularly in X-rays, and their behaviour is attributed to very high magnetic fields, resulting in their being termed magnetars. The behaviour of MSPs and magnetars presents challenges to the canonical model described so far. The polarization properties of magnetars are observed to be highly variable over time (Dai et al., 2019) and modelling of hot spots in the X-ray emission from the MSP J0030+0451 gives evidence for a multipolar magnetic field structure, rather than a simple dipole (Kalapotharakos et al., 2020). Understanding the limitations of the canonical model as a descriptor for the young pulsar population can provide a strong basis upon which to build explanations of this complex behaviour seen under more extreme conditions for MSPs and magnetars.

The magnitude of a pulsar’s magnetic field is not directly measurable, but is commonly inferred from the period and period derivative in the following way (Lorimer & Kramer, 2005). The rate of loss of rotational kinetic energy of a spinning body is given by

$$\dot{E}_{\text{rot}} = -I\Omega\dot{\Omega} = \frac{4\pi^2 I \dot{P}}{P^3}, \quad (1.2)$$

where \dot{E} is called the spin-down luminosity, I is the moment of inertia, P is the rotational period and $\Omega = 2\pi/P$. A rotating magnetic dipole with moment $|\mathbf{m}|$ emits radiation with power

$$\dot{E}_{\text{dipole}} = \frac{2}{3c^3} |\mathbf{m}|^2 \Omega^4 \sin^2 \alpha. \quad (1.3)$$

If we assume that the total spin-down luminosity results entirely from magnetic dipole braking, we can set equations 1.2 and 1.3 equal to each other. The magnetic field at

the surface of the neutron star, B_{surf} , is related to the magnetic moment by $B_{\text{surf}} \approx |\mathbf{m}|/R^3$ for neutron star radius R . Rearrangement therefore gives

$$B_{\text{surf}} = \sqrt{\frac{3c^3}{8\pi^2} \frac{I}{R^6 \sin^2 \alpha} P \dot{P}}, \quad (1.4)$$

which means, under this assumption, pulsars in the top right of the $P - \dot{P}$ diagram (Fig. 1.7) are expected to have the largest magnetic fields. Setting $\dot{E}_{\text{rot}} = \dot{E}_{\text{dip}}$ also gives a relationship between P and \dot{P} , usually expressed in terms of frequency $\nu = 1/P$ as

$$\dot{\nu} \propto -\nu^n \quad (1.5)$$

where the moment of inertia I is presumed constant and the braking index $n = 3$. The braking index can be calculated for a given pulsar through pulsar timing, as

$$n = \frac{\nu \ddot{\nu}}{\dot{\nu}^2} \quad (1.6)$$

and measured values for n range between -70 and 2890 (Parthasarathy et al., 2020). This is in part due to the effect of glitching, where a pulsar suddenly and briefly spins up before relaxing back towards its original spin-down rate, a phenomenon thought to be linked to interactions between the neutron star crust and its superfluid interior (Anderson & Itoh, 1975). Given the limitations of the assumptions and the tension with measurements, caution is required when inferring magnetic field from spin-down in the absence of other information.

1.1.3 Pulsars as tools

The regularity and hence predictability of the arrival times of pulses infers a direct analogy between pulsars and clocks. Pulsar timing involves long-term monitoring of pulsars in order to build up accurate predictive models of pulse time arrival. These models need to account for the factors affecting pulse arrival time, including the gradual spin-down of the pulsar, and the relative motions of the pulsar and the Earth. Timing programmes such as the International Pulsar Timing Array (Verbiest et al., 2016) are being used to constrain limits on gravitational wave detection through pulsar timing arrays. Timing of the double pulsar (Kramer et al., 2006) and the discovery of a pulsar in a stellar triple system (Ransom et al., 2014) enabled high precision tests of general relativity in an environment with extreme gravitational and magnetic fields. Diverse alternative applications have been suggested for pulsars, including as timepieces to rival atomic clocks (Hobbs et al., 2020b) and as literal lighthouses for interstellar navigation, such as on the Pioneer plaque (Sagan et al., 1972).




Figure 1.7: $P\dot{P}$ diagram for the known pulsar population.
 Log-log scatter plot of period derivative \dot{P} vs. topocentric spin period P for all measured values recorded in the pulsar catalogue (Manchester et al., 2005, version 1.64, plot made 23rd March 2021 showing 2394 pulsars). Figure made with matplotlib.

1.1.4 The interstellar medium

Proper understanding of pulsar emission must recognise the two limiting factors of astrophysics. First, that telescope observations will always by necessity provide a limited description of the objects they observe. This is due to bandwidth and sensitivity limitations, observing time constraints, and the limitations of the single line of sight to the object in question. Secondly, it is impossible to observe the pulsar alone, separate from the environments in which both it and we exist, and the environment of space that the emission must travel through to arrive at the telescope. The positive result of this is that we can use pulsars as probes of these environments, in particular the structures of gas, dust and plasma that make up the interstellar medium (ISM). We can glean understanding of its structure and composition by analysing its effect on electromagnetic radiation that propagates through it. We can use pulsars to do this by measuring three frequency-dependent ways that the ISM affects the pulses we observe: dispersion, Faraday rotation and scattering.

Dispersion is the slowing down of light through its interactions with charged particles. It has a quadratic frequency dependence, causing lower frequency light to arrive at the observing telescope later than higher frequencies. For the time-limited pulses of a pulsar this effect is observable as a dispersion delay, as the arrival time of the pulse is delayed by an amount that increases quadratically with decreasing frequency (see Fig. 1.8). Knowledge of the total electron column density through which the pulse passes would allow the dispersion delay at a given frequency ν to be corrected through application of the correct dispersion measure (DM) through the equation

$$t = \mathcal{D} \times \frac{\text{DM}}{\nu^2}. \quad (1.7)$$

In this equation the dispersion measure is the integral of the electron column density n_e along the line of sight to the pulsar, given by $\text{DM} = \int_0^d n_e dl$. The dispersion constant \mathcal{D} is given in terms of electron charge and mass, e and m_e , and speed of light, c , as $\mathcal{D} \equiv e^2/(2\pi m_e c)$ (Lorimer & Kramer, 2005).

Another way in which the electron column density of the ISM interacts with pulsar emission is in affecting its polarization properties. Polarized radio waves travelling through a charged medium will undergo Faraday rotation, so that the angle of linear polarization measured at the telescope will be shifted with respect to the angle of the emission originally generated by the pulsar. Like dispersion, Faraday rotation of the PA has a quadratic frequency dependence, with increased rotation at lower frequencies, according to

$$\Delta\Psi_{\text{PA}} = c^2 \times \frac{\text{RM}}{\nu^2}, \quad (1.8)$$




Figure 1.8: Dispersed intensity profile of PSR J1056–6258.

Waterfall plot showing the normalized intensity profile for PSR J1056–6258 across the observing band of the Parkes Ultra-Wideband receiver. The observation has been cleaned of contaminating radio-frequency interference (RFI) (see horizontal lines with zero intensity) but the dispersion delay has not been corrected. This results in the pulse profile arriving at successively later phases at lower frequencies according to a quadratic dependence. The apparent stripes at lower frequencies are the result of the signal wrapping around a complete pulse period. Figure made with matplotlib.

where $\text{RM} = \frac{e^3}{2\pi m_e^2 c^4} \int_0^d n_e B_{\parallel} dl$ is the rotation measure, and is an integration of electron column density multiplied by the component of the magnetic field strength along the line of sight (towards or away from the observer) (Lorimer & Kramer, 2005). Measurements of pulsars' DMs and RMs can be combined to infer

$$\langle B_{\parallel} \rangle = \frac{\int_0^d n_e B_{\parallel} dl}{\int_0^d n_e dl}. \quad (1.9)$$

The pulsar population has been used to map out the Galactic magnetic field by calculating this average magnetic field strength and direction along multiple lines of sight (Sobey et al., 2019), although care must be taken when interpreting the 3-D field based upon average values.

Many pulse profiles are seen to exhibit the effects of scattering. Regions of cold ($< 10,000\text{K}$), dense plasma along the line of sight, which can be approximated as one or multiple thin screens, cause some of the radio flux arriving at the observer to be delayed with respect to the direct line of sight. Observationally, this results in a scattered pulse profile with a characteristic exponential scattering tail (Cronyn, 1970). An example of a scatter-broadened pulse profile is shown for PSR J1818–1422 in Fig. 1.9. The simplest model explaining such an observation is the presence of an isotropic scattering screen, for which the scattering angles have no preferred direction, halfway between the pulsar and the observer along the line of sight (see Fig. 1.10 for a diagram). In this context the key parameter of relevance is the scattering timescale, τ .

It is observed that τ evolves with frequency ν according to the power law relationship

$$\tau \propto \nu^{-\alpha}, \quad (1.10)$$

where α is the scattering spectral index. The simplest thin screen scattering model predicts $\alpha = 4$ (e.g. Cronyn, 1970; Lang, 1971), whilst a medium exhibiting Kolmogorov turbulence would have a value of 4.4 (e.g. Lee & Jokipii, 1976; Rickett, 1977). The simplified spectrum of turbulence is described by a power law for wavenumbers q , for values of q^{-1} that lie well within inner and outer fluctuation scales (k_i^{-1} and k_o^{-1} respectively). This is written as $P_{n_e}(q) = C_{n_e}^2 q^{-\beta}$ (Rickett, 1977), where $C_{n_e}^2$ is the proportionality constant, dependent on the electron density. The fluctuation spectral index β cannot exceed 4 (Romani et al., 1986). It is related to α by $\alpha = 2\beta/(\beta - 2)$, which leads to $\alpha = 4$ being a theoretical lower limit. Observationally however, lower values of α , i.e. flatter spectra, are commonly seen (see for example Krishnakumar




Figure 1.9: Scattered pulse profile for PSR J1818–1422.

Grey: normalized total intensity profile for PSR J1818–1422 at 0.94 GHz, observed with the MeerKAT telescope. The effect of scattering is to smear the profile shape to have an exponential scattering tail. Black line: the best fit model of the scattered profile with scattering timescale $\tau = 73.8 \pm 0.5$ ms, for which the modelling process is described in Chapter 7. Figure made with matplotlib.

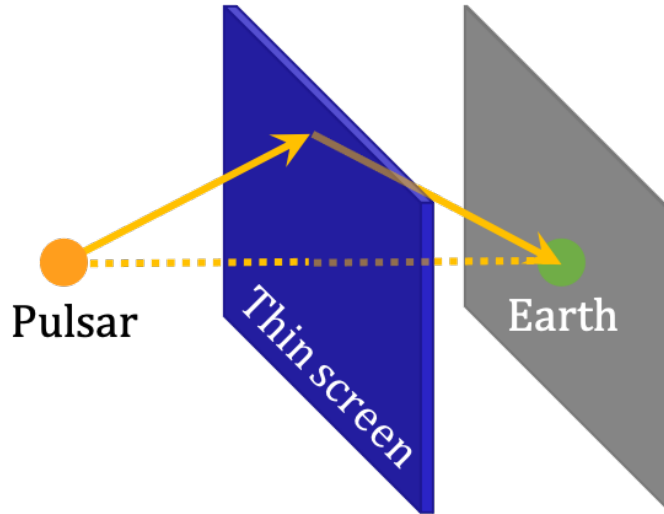


Figure 1.10: Diagram of the thin screen model for scattering pulsar radiation. Diagram indicating the effect of a thin screen scattering pulsar radiation: radio waves are deflected by a thin screen placed half-way between the pulsar and the Earth. This results in some radio flux being scattered into the line of sight and other flux scattered out of the line of sight, distorting the observed pulse profile. Figure made using Microsoft PowerPoint.

et al., 2019). A power law break at low frequencies with the scattering timescale approaching frequency independence, accompanied by a loss of flux, may be attributed to a truncated scattering screen (Cordes & Lazio, 2001). Where q^{-1} drops below the inner scale of the plasma turbulence, k_i^{-1} , the power law description of the turbulence spectrum is no longer a valid simplification. Lewandowski et al. (2013, 2015a) write that, for a single pulsar, this may lead to a flattening of α at lower frequencies. The effect of this inner scale, which corresponds to the shortest scale length in the scattering material, should also become apparent in the scattered pulse shape, particularly at large delays in the profile tail (Rickett et al., 2009).

Furthermore, the exact shape of the scattered profile will depend upon the degree of anisotropy in the scattering material. Evidence of such anisotropy has been seen in measurements of scintillation, which is where scattering causes the pulsar radio intensity to be modulated across frequency and time. The effect of anisotropic scattering is for parabolic arcs to be observable in the secondary spectra of scattered pulsar observations (e.g. Stinebring et al., 2001; Walker et al., 2004). High resolution mapping of the locations of scattering material on the sky by Brisken et al. (2010); Pen et al. (2014); Simard et al. (2019) shows evidence of 1D screen shapes. Geyer & Karastergiou (2016) showed that fitting an isotropic model to simulated anisotropically scattered data could lead to inferring smaller values of α .

The scattering strength is expected to be correlated with distance for a given pulsar, added to which is a further stochastic element due to there generally being a small number of scattering screens along the line of sight. Modelling the scattering material as a power-law electron density spectrum with a Kolmogorov distribution of irregularities, the relationship between scattering timescale and DM is expected to be $\tau \propto C_{n_e}^2 \nu^{-4.4} \text{DM}^{2.2}$ (Romani et al., 1986; Cordes & Rickett, 1998; Krishnakumar et al., 2015).

1.1.5 Summary

As pulsar observations have become more sensitive and comprehensive, it has become clear that the canonical model is insufficient to explain the host of complex phenomena observed in pulse profiles. Our insufficient understanding of the radio emission processes of pulsars affects not only our understanding of these objects in their own right, but also has implications for the accuracy to which we can use pulsars as tools, for testing the laws of physics and searching for gravitational waves through pulsar timing; and for probing the properties of the interstellar medium.

Development in scientific understanding is dependent upon technological progress. In particular, the widening of observational bandwidth that has become available over time has led to more detailed insight into pulsar emission processes, the beam structure and its interaction with the ISM. The next section describes how the increase in observational bandwidth has both answered and raised new questions about pulsar radio emission, providing the context for the questions my research investigates in this thesis.

1.2 A history of pulsars and observing bandwidth

1.2.1 Early observations and theory

After their discovery by Jocelyn Bell Burnell (Hewish et al., 1968), theory quickly developed the early canonical model for pulsars based on the key features of the narrowband observations available. Gold (1968) posited that the speed and regularity with which pulses arrive necessitates a rotating body, and their brightness requires a coherent emission mechanism from a very small emission region with an intense gravitational field, all pointing towards a rotating neutron star being the source of the newly discovered radio pulses. Measurements of the PA profile of the Vela pulsar prompted Radhakrishnan & Cooke (1969) to develop the rotating vector model

(RVM), whilst Goldreich & Julian (1969) developed an axisymmetric solution for the charge density of a rotation-powered pulsar. Observations of pulsars were carried out at multiple different narrow-band frequencies from the start, such as that by Davies et al. (1968), who observed the first pulsar at frequencies 151, 240 and 408 MHz simultaneously, with bandwidths of 4 MHz. The growing catalogue of pulsar discoveries and observations at multiple frequencies prompted Komesaroff (1970) to describe a possible mechanism for the pulsar radio emission process that would account for the key features of pulsar emission. This included the variety of pulse shapes observed, the fact that they vary slowly with observing frequency, and that several pulsars had been measured to be almost 100% linearly polarized. The model proposed that radio emission was produced only by the last open field lines, resulting in a cone of radio emission.

1.2.2 The narrowband era

Observing campaigns over the following years, for example Manchester et al. (1978), resulted in the discovery and analysis of hundreds of pulsars. At higher frequencies, recording sensitive observations with large bandwidths becomes increasingly difficult (Rohlf & Wilson, 1996). However, bandwidths of observations gradually increased with the development of new telescopes and receivers. For example, the original receiver system of the Giant Metrewave Radio Telescope (Swarup, 1991) could perform observations at discrete frequencies between 38 and 1420 MHz, with bandwidths of up to 32 MHz.

The growing number of observations of pulsars at multiple different frequencies made it possible to develop ideas of pulsar beam structure. Ruderman & Sutherland (1975) incorporated the concept of radius-to-frequency mapping (RFM) into their theory of pulsar emission, whereby the frequency of radio emission produced by the pulsar is dependent on the height at which it is generated, such that lower frequencies ν are produced at higher heights r above the pulsar surface according to $r \propto \nu^{-2/3}$. They inferred this by assuming that emission frequency is related to the plasma frequency, itself associated with plasma density and hence emission height. Later, Gangadhara (2004) proposed an emission height model for which the emitting radio frequency is the characteristic frequency of curvature radiation. This depends on the radius of curvature of the emitting field line, ρ_c , and the Lorentz factor γ of the relativistic particle that produces the radiation, so that $\nu = 3c\gamma^3/(4\pi\rho_c)$ (Ruderman & Sutherland, 1975). This would mean that emission produced closer to the magnetic axis at a given frequency must be generated at a lower height.

The concept of RFM has been empirically related to the observation that pulse profiles tend to become wider with decreasing frequency below ~ 1 GHz, and was incorporated into the hollow cone model developed by Rankin and collaborators (see for example Rankin, 1983). This model was developed to explain the variety of pulse profile morphologies observed and the symmetry frequently present in profile shapes. Lyne & Manchester (1988) used a large collection of pulsar observations taken at multiple frequencies between 170 and 8700 MHz to map out the origins of pulse components upon the neutron star surface. The evolution of pulse width with frequency was quantified most comprehensively by Thorsett (1991), who compared integrated profiles observed at multiple different frequencies for seven pulsars and described an empirical width-frequency relationship of a power law plus constant. In addition to affecting profile shape at different frequencies, RFM is thought to affect relative pulse arrival times and relative misalignments of intensity and PA profiles through the differential effects of aberration and retardation (A/R) at different heights. These effects are quantified as timing A/R (Cordes, 1978) and the delay-radius relation (Blaskiewicz et al., 1991; Hibschan & Arons, 2001; Dyks, 2008) respectively. They have been investigated for many pulsar observations: see for example Mitra & Li (2004), Hassall et al. (2012), and the beginning of Chapter 5 for studies of PSR J1136+1551.

Models created to infer the beam structure that leads to the observed pulse profile have focused on two different key features: the apparent symmetry of many, though not all, pulsar profiles, and the presence of multiple components in a pulsar beam. The hollow cone model describes the active emitting regions of the pulsar magnetosphere as having circular symmetry about the magnetic axis, so that one or more cones of emission are intersected by the line of sight. These cones may surround a pencil beam “core” pointing along the magnetic axis. This phenomenology is successful in explaining the symmetrical shapes seen in many pulse profiles, since the line of sight necessarily cuts the circle in two places in all cases other than a grazing line of sight. The central regions of pulse profiles often appear to exhibit greater complexity in polarization characteristics, causing many researchers to designate such regions as “core” emission and seek to characterize their behaviour independent from that of the “conal” components. However, there is no a priori reason to suppose that some circular region of field lines will be active when those originating from a circle of field lines at a different perpendicular radius from the magnetic axis are not. Furthermore, there are many pulsars that do not exhibit any symmetry in their profile structure. Nor is it evident that core components can be unambiguously designated as such, due

to the myriad of pulsar variability and unusual phenomenology witnessed in all parts of pulse profiles.

An alternative model proposed is the fan beam model (Michel, 1987; Dyks et al., 2010; Wang et al., 2015), which describes pulsar radio emission as originating from streams of active field line regions in the magnetosphere, which are placed randomly without symmetry constraints. The distribution of emission wavelengths along these active field line regions will affect how pulse profiles appear at different observing frequencies and may even result in profile width evolution that resembles radius-to-frequency mapping (Dyks & Rudak, 2015). The fan beam model has the advantage of flexibility, but the presence of symmetry in many profiles continues to be a point of relevance. Fig. 1.11 shows illustrative depictions of these models with an example of a one-to-one radius-to-frequency mapping. Another model is that developed by Karastergiou & Johnston (2007), created to describe the population of pulsar profiles as a statistical whole. In this model, the footprints of groups of active field lines form patches randomly distributed within a ring. These active regions emit each frequency at multiple discrete heights along the same field lines. The emission height range is also dependent on \dot{E} in this model, with young, energetic pulsars with high \dot{E} theorized to emit from a narrow range of high heights, and older pulsars with lower \dot{E} emitting from a wider range of heights.

In summary, the rapid increase in pulsars discovered and pulse profiles observed over this time period resulted in a desire to categorize pulse profiles in terms of symmetry and number of components in order to build understanding of pulsar beam structure. However, the increased number of observations also revealed the limitations of this kind of categorization, and many questions about pulsar beam structure and the relative applicability of different models remain unanswered.

1.2.3 Polarimetry

Considerable additional information about pulsar emission can be obtained by studying not only the total intensity profile (Stokes I) but also the polarization properties of pulsars using the full Stokes parameters I , Q , U and V . Further to the initial observations that many pulsar profiles are strongly linearly polarized, evidence began to arise that indicated the presence of orthogonal modes of polarization (OPMs) in pulsar emission. Manchester et al. (1975) showed that orthogonally polarized emission observed in the single pulses of 12 pulsars contributed to the net depolarization of their integrated profiles. The jumps of 90° exhibited by many PA profiles, which coincide with the total measured polarization decreasing to 0 at that pulse phase,




Figure 1.11: Hollow cone and fan beam models of the pulsar radio beam, side views. Illustrative depictions of two proposed models for the structure of the pulsar radio beam. Left: the hollow cone model describes the beam as resulting from a group of curved magnetic field lines whose footprints form a ring on the pulsar surface. The vertical arrow is the magnetic axis \underline{m} and the curved lines represent how the shape of the beam follows the shape of the dipolar field lines (c.f. Fig. 1.1). Assuming radius-to-frequency mapping, higher frequency emission is produced at lower heights, as indicated by the downward direction of the labelled arrow. The three shaded ring-shaped regions show emission from three discrete frequencies at three different heights, where f_1 (darkest shading) is the lowest frequency. The shape of the beam means that lower frequencies are produced from larger rings. Right: the fan beam model describes pulsar radio emission as originating from streams that follow different groups of field lines. The effect of radius-to-frequency mapping is marked on the diagram in a similar way. Figure made using Ipe.

have been taken as an indication of a change in which of the two modes is dominant. Examples of this can be seen in the integrated pulse profile and PA profile of PSR J0738–4042 in Fig. 1.12. Theories about the origins of the orthogonal modes have suggested that, once generated, the radio waves propagate within a birefringent magnetosphere, which results in two modes, labelled the extraordinary (X) and ordinary (O) modes, propagating along different paths. Both coherent and disjoint interactions between these modes have been proposed, and the final polarization state of the observed radio emission has been described as being fixed by the polarization limiting radius (PLR), as described, for example, by Barnard & Arons (1986). They suggest that whereas at low frequencies the orthogonal X and O modes propagate along different paths in the pulsar magnetosphere, at high frequencies the modes do not separate. Such frequency-dependent behaviour of the polarized modes may provide an explanation for the observation that pulse profiles tend to become narrower and depolarize with increasing frequency, and that jumps in the PA are more common at high frequencies (e.g. Johnston et al., 2008). Decomposition of pulse profiles into orthogonal components, assuming incoherent superposition and 100% linear polarization, has been attempted by e.g. McKinnon & Stinebring (2000) and Smits et al. (2006). This decomposition is successful at explaining the polarization properties of some pulsars, but takes no account of the presence of high levels of circular polarization in many pulse profiles.

Efforts to quantify the polarization of pulsar observations, including the behaviour of OPMs, have relied on the increased capability of new instruments to produce high quality polarized pulse profiles. Karastergiou et al. (2005) used the high bandwidth (1024 MHz) of the then new 10-cm receiver on the Parkes radio telescope to generate high S/N polarized pulse profiles of 48 pulsars at the high frequency of 3.1 GHz. They suggested that the variation of mode dominance with frequency could arise from the two modes having different spectral indices, so that their relative intensities vary with wavelength.

New observations have enabled continued development of theories about the extent to which the modes combine together coherently. This includes the single pulse work of Edwards & Stappers (2004), who used the 10 MHz band centred at 328 MHz provided by the Westerbork Synthesis Radio Telescope to observe polarized single pulses for PSR B0329+54. They inferred that their observation of an elliptical spread of polarization states on the Poincaré sphere could only be generated by a coherent superposition of two orthogonal modes with varying phase differences. From a theoretical perspective, such a coherent superposition provides an explanation for the




Figure 1.12: PSR J0738-4042: integrated polarized profile and frequency evolution. Plotted as described for PSR J1559-4438 in Fig. 1.6, the left side shows PSR J0738-4042 at 1087 MHz. Figure made with matplotlib.

presence of circular polarization in pulse profiles, and refraction of wave modes in the magnetosphere has been theorized to be important for the frequency dependence of pulse width and polarization fraction (Lyubarskii & Petrova, 1998; Petrova, 2001).

A common observation relating to circular polarization is that many profiles exhibit at least one change of handedness, suggested often to be in the centre of the pulse profile. A related observation is that orthogonal jumps are usually accompanied by a change of hand in the circular polarization. There also exist, in many pulse profiles, jumps in the PA that are not always 90° in magnitude. They are often characterized by a drop in the linear polarization but a non-zero circular polarization, or even an increase in the circular polarization around the phase of the jump. The profile of PSR J0452–1759 (see Fig. 1.13) shows examples of both two “standard” orthogonal jumps with a change of hand in Stokes V, and one jump where Stokes V is non-zero. A change in hand of circular polarization at a jump suggests that we are seeing a switch in dominance between two elliptically polarized orthogonal modes. A PA jump that is associated with non-zero circular polarization may be related to the overall polarization state transitioning smoothly across the Poincaré sphere from linear to circular to linear again, rather than the disjoint jump seen in the other case, as described by Dyks (2020). Dyks (2019) has proposed that this is further evidence for coherent superposition of the orthogonal modes, which will introduce an element of circular polarization when the modes are superposed with some phase difference. A smooth variation in phase difference of the superposition would then lead to the type of mode transition described.

1.2.4 Simultaneous multi-frequency observations and wider bands

Understanding the behaviour of orthogonal modes in pulsar radio emission requires better characterization of how observed polarization properties evolve with frequency, both for integrated profiles and for single pulses, for a large number of pulsars. When investigating how pulsar emission varies with frequency, the use of integrated pulse profiles observed both with different instruments and at different epochs (dates and times) limits the level to which they can be directly compared. As a result, the European Pulsar Network launched a project in 1995 to perform simultaneous observations of pulsars with multiple instruments, involving various collaborations between the Effelsberg, Jodrell Bank, Bologna, Westerbork, Puchshino, Ooty, Torun and Arecibo telescopes (Karastergiou et al., 2001). These observations collected single pulses, so that the behaviour of every individual pulse could be mapped across frequency. A




Figure 1.13: PSR J0452–1759: integrated polarized profile and frequency evolution. As for Fig. 1.6 and Fig. 1.12, the left side shows PSR J0452–1759 at 1788 MHz. Figure made with matplotlib.

series of five papers (Karastergiou et al., 2001, 2002, 2003; Kramer et al., 2003; Bhat et al., 2007) revealed the multi-frequency single pulse characteristics of pulsars. They focused on polarization including circular polarization and the spectral behaviour of orthogonal polarization modes, flux density spectra and the broad-band nature of nulling, which is where single pulses are observed to switch off randomly. They showed that polarization characteristics of single pulses, including circular polarization and the behaviour of orthogonal modes, are related to each other across frequency, and observed selective frequency-dependent nulling for PSR J1136+1551.

Development of new telescope backends led to large improvements in the pulsar observing bandwidths attainable, such as the GUPPI (Ransom et al., 2009) and PUPPI (Van Rooy et al., 2017) backends on the Green Bank and Arecibo telescopes respectively, which offered a maximum bandwidth of 800 MHz. The multibeam receiver on the Parkes radio telescope afforded a continuous band of 256 MHz between 1240 and 1496 MHz, typically divided into 1024 frequency channels, which enabled high quality polarimetric pulse profiles to be obtained (Johnston & Kerr, 2018). The advantages of increasing bandwidth capability came initially in the form of increased S/N for a given pulse profile, but as bandwidths continued to increase, the opportunity to investigate incremental evolution of pulse profiles with frequency became possible. Hassall et al. (2012) performed wide-band simultaneous observations with the LOFAR, Lovell and Effelsberg telescopes to investigate dispersion and pulse profile evolution. The wide receiver bandwidths of each telescope (48 MHz of bandwidth between 10 and 90 MHz for the LOFAR Low Band Antennas) enabled the observations to be divided into sub-bands. Low frequency pulsar polarization was investigated further by Noutsos et al. (2015) using the LOFAR high-band antennas at 105–240 MHz. Their results showed that the expected effects of magnetospheric birefringence on pulsar polarization, in terms of evolution of pulse width and polarization fraction, were not fully supported by the data. They also estimated emission heights in the magnetosphere for four pulsars, based on the delay-radius relation.

1.2.5 New broad-band capabilities

The progress of pulsar emission theory has relied upon the new information made available through improved observations across wider bandwidths. These observations have revealed the variety of pulse profiles in terms of shape, polarization complexity and the evolution of these parameters with frequency. The most recent developments to telescopes have now moved us into a continuous broad-band era of pulsar astrophysics. Upgrades such as the u-GMRT (Gupta et al., 2017) have updated the

relevance of older instruments, and the development of the new Ultra-Wideband receiver on the Parkes radio telescope has made it possible to observe pulsars with a continuous broad bandwidth from 700–4000 MHz (Hobbs et al., 2020a). The commissioning of the new MeerKAT telescope (Jonas, 2009) is enabling larger and higher sensitivity surveys of pulsar science than previously possible. Making use of these instruments, it is now possible to investigate in detail the ways in which the emission process of pulsar radio waves, and their interaction with the ISM, leave observational signatures on the broad-band frequency evolution of integrated pulse profiles and single pulses. By observing pulsars with these telescopes, I addressed questions about how pulsar radio emission evolves with observing frequency and interacts with the ISM, furthering our understanding of pulsar emission theory and hence our capacity to use pulsars to understand wider questions about physics and the Universe.

1.3 Open questions

The field of pulsar science has seen considerable development in observational technology and, as a result, astrophysical theory explaining pulsar formation, structure and emission properties. There remain, however, many open questions about pulsar science that can be addressed with new data and analysis.

The over-arching open question addressed by this thesis is how pulsars produce the radio beams we observe. This can be broken down into questions about the theoretical mechanisms themselves; how the development of our understanding of theory is dependent upon observations; and how we can therefore use observations to build our understanding of theory. In this thesis I seek to address these three factors by answering the following open questions.

1. How do pulsars produce radio beams?
 - (a) What is the structure of the radio beam?
 - (b) What dictates pulsar beam width, filling fraction and distribution of active emitting sites?
 - (c) How does beam structure vary with frequency?
 - (d) How are single pulses produced and how does this result in the properties of integrated pulse profiles?
2. How do observational constraints affect our understanding of the pulsar population?

- (a) How do the observational properties of pulsars vary across the pulsar population?
 - (b) How does pulsar emission interact with the ISM?
 - (c) How do the polarization properties of pulsar observations map to the emission and propagation processes undergone?
3. How can we use new broad-band radio observations to answer these questions?

I will refer back to these open questions throughout the thesis (e.g. as “Q1a”) in order to place my results in the context of the wider ideas and motivations behind the research.

1.4 Structure of the thesis

The remainder of this thesis is structured into the following chapters. In Chapter 2, I summarize in a list the key definitions introduced in this chapter that are relevant for the remainder of the thesis. Chapter 3 explains the motivations behind the choices made in pulsar observing projects, both from a general perspective and the specific application of this to my own observational work. This encompasses the capabilities of modern instruments and the opportunities they afford observationally. The chapter also describes the practical aspects of observing pulsars and processing the data, including calibration, Radio-Frequency Interference (RFI) excision and standard methods used in pulsar data processing, including correcting for dispersion and Faraday rotation.

Chapter 4 covers the methodology I developed for my research. This includes simulating pulsar radio beams and pulse profiles, in addition to techniques for visualizing and analysing broad-band data in ways that account for the effects of propagation through the ISM.

Chapters 5, 6 and 7 present the results of my research addressing the open questions introduced above. The first describes my work on understanding pulsar beam structure using single pulses from PSR J1136+1551. This is followed by a chapter on the algorithm I developed to account for ISM propagation effects in broad-band polarimetric pulsar observations and the resultant conclusions that can be drawn about pulsar emission heights from the corrected profiles. My work then moves from understanding pulsar emission to applying this understanding to modelling the effects of the ISM. In Chapter 7 I present the results of a large-scale survey of the

frequency-dependent behaviour of interstellar scattering as revealed through modelling broad-band scatter broadened pulsar profiles.

Chapter 8 discusses these results in the context of related work in the field; and addresses its implications for our understanding of both pulsar radio emission and the structure of the ISM going forward. The thesis ends with a summary of conclusions in Chapter 9 and the proposed direction that future research should take in light of these conclusions.

Chapter 2

Definitions

This thesis has so far introduced and defined several terms in the context of their applicability. For clarity, and ease of reference in the remainder of the text, here is a list of some key terms relevant to pulsar science, with their definitions.

1. Pulsar observations

- (a) **Integrated pulse profile**: a stable profile created by summing thousands of single pulses (see for example Fig. 1.3). The shape of the profile and its polarization properties are indicators of the emission processes of the pulsar.
- (b) **Stokes parameters** (I, Q, U, V): an experimentally convenient parametrization of the polarization state of electromagnetic radiation, these are quantities made up of combinations of the time-averaged complex amplitudes of the electric field. Stokes I is the total intensity, Stokes V is the total circular polarization, and Q and U describe the linear polarization and its angle, such that the total linear polarization L is given by $L = \sqrt{Q^2 + U^2}$ and its angle is the position angle (**PA**) calculated as $-0.5 \arctan(U/Q)$. See for example Fig. 1.3 showing I, L, V and the PA for the Vela pulsar. For all the polarimetric data presented in this thesis I apply a bias correction to the linear polarization, L , according to the prescription of Everett & Weisberg (2002). This corrects for the overestimation of L that results from noise on Q and U , due to L being a positive definite quantity. Usage of Stokes parameters is not limited to radio astronomy: they are generally applicable within the field of optics.
- (c) **Poincaré sphere**: a 3-D representation of the polarization state of an electromagnetic wave. Stokes parameters are plotted on the Poincaré sphere

with the axes being Q , U and V respectively. See for example Fig. 1.2. The mathematics of the Poincare sphere and Stokes parameters is more generally applicable outside of descriptions of polarization. In other applications it is known as the Bloch sphere, where the Bloch vector represents a pure or mixed state in a two-level quantum system, for example to describe a qubit (Foot, 2005, Chapter 13).

- (d) **Radio frequency interference (RFI)**: any observed radio emission from a non-astronomical source that contaminates an observation.

2. The canonical model

- (a) α : the angle between the magnetic and rotational axes of the pulsar (see Fig. 1.1).
- (b) β : the impact parameter of the line of sight, that is, the angle of closest approach of the line of sight to the magnetic axis (see Fig. 1.1 and also Fig. 1.5).
- (c) **Rotational phase ϕ** : the angle of rotation that the pulsar makes with respect to a defined axis at a given time. This is used to map the integrated pulse profile onto the rotation of the pulsar.
- (d) **Fiducial point ϕ_0** : a hypothetical fixed point on the pulsar surface that is mapped to rotational phase $\phi = 0$ (see red circle on Fig. 1.5). The fiducial point is identified on a pulse profile, commonly at its peak, in order to link the pulsar rotation to the observed pulse profile.
- (e) **Fiducial plane**: the plane containing the magnetic and rotational axes of the pulsar. The fiducial point is commonly selected to be the point where the line of sight cuts across this plane (see Fig. 1.5).
- (f) **Rotating vector model (RVM)**: a model developed by Radhakrishnan & Cooke (1969) which relates the curvature of the PA profile to the passage of the line of sight across the pulsar beam. The path of the line of sight, and hence the shape of the PA in this model, depend solely on the geometry of the pulsar (its angles α and β).
- (g) **Radius-to-frequency mapping (RFM)**: an idea developed to explain the observed increase in pulse width with decreasing frequency below 1 GHz, RFM relates the frequency of observed radio emission from a pulsar to the height above the pulsar surface at which it is produced. This is defined such that emission height increases with decreasing frequency.

- (h) **Aberration and retardation (A/R)**: factors that affect the arrival time of pulse emission depending on the location of the emission in the pulsar magnetosphere. **Retardation** is the time delay of emission produced at lower heights above the pulsar surface, due to its having to travel a greater distance to the telescope. For pulsar emission generated at heights r_1 and r_2 , the difference in arrival time is, to first order, $\Delta t_{ret} = (r_2 - r_1)/c$. **Aberration** is a beaming effect, where the radio beam is bent towards the direction of the pulsar’s rotation. To first order, this scales pulse arrival time by a factor $\sin \alpha$, where α is the geometrical angle from definition 2a. Time difference of pulse arrival due to aberration is $\Delta t_{ab} = \sin(\alpha)\Delta t_{ret}$. The total time delay due to A/R expected for emission generated at two discrete heights, referred to henceforth as “**timing A/R**”, is then $\Delta t_{A/R} = (1 + \sin(\alpha))(r_2 - r_1)/c$ (Cordes, 1978; Hassall et al., 2012).
- (i) **Delay-radius relation**: the relationship between the radio emission height in the pulsar magnetosphere and the misalignment between the fiducial points of the intensity and PA profiles, given in terms of pulse phase as $\Delta\phi = 4r/R_{LC}$. Originally derived by Blaskiewicz et al. (1991), it is commonly described as resulting from aberration and retardation (A/R) in the pulsar magnetosphere. Dyks (2008) derives it as originating from the contribution of the magnetosphere’s co-rotation to the acceleration of the emitting particles in the observer’s frame of reference, and states that A/R need only be considered when there are different emission heights present, or when describing absolute delays with respect to the fiducial point. A related shift is the vertical shift of the PA in absolute angle by $10r \cos(\alpha)/(3R_{LC})$ derived by Hibschan & Arons (2001) and incorporated into a single description of the altitude-dependent PA by Dyks (2008).
- (j) **Orthogonal polarization modes (OPMs)**: the radio emission from pulsars is thought to consist of two modes with orthogonal polarizations, resulting from the propagation of the radio waves through the birefringent magnetosphere. These are termed the extraordinary (X) and ordinary (O) modes, where the X mode is thought to be independent of frequency and the O mode is thought to follow the curved magnetic field lines up to a polarization limiting radius.

3. Pulsars as tools

- (a) **Pulsar timing**: the practice of using pulsars as clocks, timing the arrivals of pulses in order to build up a model to better predict future pulse times of arrival (TOAs). Pulsar timing compares actual TOAs with those predicted in order to test theories of gravity and search for gravitational waves.
- (b) **Interstellar medium (ISM)**: the structures of gas, dust and plasma which exist between the stars in the Galaxy. Pulsar radio waves interact with the ISM in various ways, including dispersion, Faraday rotation and scattering.
- (c) **Dispersion measure (DM)**: a measure of the electron column density along the line of sight to a given pulsar, the DM quantifies the size of the frequency-dependent time delay that results from pulsar radio emission interacting with this column density (see example in Fig. 1.8). It is measured in units of $\text{cm}^{-3} \text{ pc}$.
- (d) **Rotation measure (RM)**: a measure of the electron column density and magnetic field of the ISM along the line of sight to a given pulsar, the RM quantifies the extent to which the ISM Faraday rotates the PA profile. It is measured in units of rad m^{-2} .
- (e) **Characteristic scattering timescale τ** : strongly scattered pulse profiles are visibly broadened in shape with an exponentially decaying tail (see Fig. 1.9). The scattering timescale quantifies the extent of this scatter broadening process, though the exact functional form of the broadening depends upon the nature of the scattering undergone.
- (f) **Scattering spectral index α** : this quantifies how the scattering timescale varies with the frequency according to a power law as given in equation 1.10. Its theoretical value is determined by the structure of the scattering material present in the ISM.
- (g) **Scintillation**: a phenomenon resulting from pulsar radio waves scattering off structures in the ISM, which results in apparent variation of the brightness of the pulse profile being observed across observing frequency and time. This is because the deflected radio waves combine with varying phase lags, so that the observed intensity is modulated by constructive or destructive interference.

Chapter 3

Observations and data processing

3.1 How pulsars are observed

3.1.1 Observational decisions: managing competing factors

When observing radio pulsars, the key factors to consider that affect the questions a given dataset will be able to answer are as follows: telescope sensitivity, bandwidth, time resolution, frequency resolution, and observing time. These factors must be balanced against the constraints of technological capability, manageable quantities of data, limited time on observing schedules, and the extent to which radio frequency interference (RFI) can affect observation quality. Furthermore, it is important to consider that the properties and environment of a given pulsar can affect how these factors inter-relate, for example, weaker pulsars require longer observing times to build up sufficient signal to detect the source.

Consideration of the above was important when selecting the type of observations and instruments to use to address the open questions presented and numbered in Section 1.3. Questions 1a–d are concerned with probing the structure of the pulsar radio beam in detail. This requires observations with a high S/N, motivating observations of bright pulsars with sensitive radio telescopes. A key focus is to investigate how pulsar beam structure varies with frequency (Q1c), for which broad-band observations are necessary. As described in Chapter 1, the development of new telescopes and receivers with broad-band capabilities means that this is a new observational possibility. To answer these questions requires a combination of both single pulses and integrated pulse profiles (Q1d) in order to investigate both the instantaneous behaviour of the pulsar emission and its average polarimetric properties.

Observational constraints of radio telescopes are not only limited by available telescope capacities, but also by the limitations of there being a single line of sight

along which a source can be observed. The geometry of the pulsar (angles α and β) defines the line-of-sight cut through the pulsar beam. The location of the pulsar in the galaxy will affect its local environment, its distance from Earth and the properties of the ISM material that lies along the line of sight; all of which will have an impact on the observables from a given source. Understanding how the pulsar geometry and ISM impact observations (Q2a–c) requires both detailed analysis of broad-band polarimetric profiles of individual pulsars, and a global approach to studying these factors across the pulsar population.

This chapter addresses how the observational requirements I have described are fulfilled by my choices of radio telescopes and observing specifications. It describes the process of collecting the observations used in this thesis, and the steps for processing and calibrating the data products in order to prepare them for scientific analysis.

3.1.2 From telescope to computer: recording and processing of pulsar observations

The pulsar observations described in this thesis are all recorded in filterbank format. This means that the radio signal has been divided into individual frequency channels that span the observing band of the telescope receiver. Radio telescopes are commonly set up with two orthogonally polarized waveguide feeds. These direct the radio emission through first a low-noise amplifier and then a bandpass filter which retains only those frequencies that lie within the observing band. Where telescopes have multiple antennas, such as the MeerKAT telescope and GMRT, the data streams recorded by each antenna are added in phase by a beamformer. Each of the two polarized signals is passed to a filterbank, which divides the band into multiple frequency channels.

There are three main methods through which this channelization is achieved: an analogue filterbank uses a set of narrow-band filters, each of which retains a small region of the observing band. Alternatively, autocorrelation of the signal over the desired range of lags and Fourier transforming the result gives the power spectrum of the signal, ie. the total power within each frequency channel (Lorimer & Kramer, 2005). The four Stokes parameters (definition 1b) are obtained from the power spectra as combinations of the auto- and cross- correlations of the signals from the two polarized feeds. Finally, it is possible instead to Fourier transform the signal directly, rather than Fourier transforming its autocorrelation. Polyphase filterbank spectrometers are commonly used in current telescope technology: these apply an enhanced filter response in advance of a Fourier transform filterbank (Price, 2016). Doing so retains phase information after channelization.

Since pulsar signals are dispersed by the ISM it is necessary to correct for the frequency-dependent delay of the pulse arrival time. This can be done through incoherent dedispersion once the data have been channelized. This involves introducing lags to the outputs of the frequency channels so that they shift relative to each other in time. This aligns the pulse profiles across the band, correcting the frequency-dependent delay seen in Fig. 1.8. The limitation of incoherent dedispersion is that it takes place after the signal has been converted to powers, so that the phase information of the waves have been lost. The dispersion delay within an individual channel therefore cannot be corrected, and so the pulse profile will be smeared within the channel by the amount of the dispersion delay across that channel’s bandwidth.

The alternative is to perform coherent dedispersion. In order to do this, the signal is sampled using baseband sampling which retains both amplitude and phase information. A set of n data points in the time domain is then Fourier transformed and an inverse transfer function is applied to dedisperse the signal before Fourier transforming back to the time domain. This is then repeated with another set of n data points and so on until the entire time series has been coherently dedispersed. Since dedispersion shifts data points in time, it is necessary for these sets of data points to overlap in time to ensure that the entire sample is dedispersed. Once coherent dedispersion is complete, the power spectra can be obtained through auto- and cross-correlation as described above, and the Stokes parameters calculated. Coherent dedispersion is more data intensive, but solves the problem of dispersion smearing. A similar approach to correcting for Faraday rotation in a coherent way has recently been introduced at the NenuFAR telescope (Bondonneau et al., 2020).

There are two modes of observation used for radio pulsars, known as search mode and fold mode. In search mode, so called because it is used to search for new signals of pulsar emission, incoming data are recorded as a time series as described above, so that single pulses are recorded. Search mode is necessary for single-pulse studies of pulsar emission, which can reveal information about the time-varying properties of the pulses, statistical behaviour and detailed microstructure that makes up the emission observed (e.g. Hankins et al., 2016). Its disadvantages are high data-rates that require large amounts of storage space and are unnecessary for many questions of scientific interest, and the fact that many pulsar signals are insufficiently bright to be observed in single pulses, particularly if high frequency or phase resolution is of interest for the questions being addressed.

The alternative observational set-up is fold mode, where the incoming pulsar signal is folded on the known topocentric pulse period to make subintegrations with an in-

creased signal-to-noise ratio (S/N). These subintegrations are then generally summed over time to generate integrated profiles. The increased S/N of an integrated profile means that resolution is retained if the observation is channelized into a large number of frequency channels, making it an appropriate for studies of pulsar variability across frequency. Another advantage of the integrated profile is that of temporal stability: it provides an average picture of the emission process of a given pulsar. However, folding the single pulses necessarily loses information about the temporal variability of the pulsar emission, such as nulling, mode-changing and subpulse drifting, which may reveal features of the pulsar emission process. I used both search mode and fold mode data for the work in this thesis to address different questions, as described in detail later.

Once the data are formatted as Stokes parameters, further post-processing can be applied. This includes binning the time series into a fixed number of bins per pulse period, and aggregating frequency channels together to build up S/N. The data will need to be calibrated to account for the effects of the bandpass filter and any differences in gain of the two polarized feeds. Flux and polarization calibration may also be applied if these properties are relevant for the science questions being addressed. Further corrections to the data may be applied to account for effects of the ISM: these include additional (incoherent) dispersion corrections, and applying a rotation measure (RM) to correct for the effects of Faraday rotation. Methods that have been employed for doing so are described in Section 3.4.

Another important factor to consider in such post-processing is the effect of radio frequency interference (RFI). This may contaminate entire frequency channels, or affect the whole or part of the observation for some small subset of time. Standard tools to flag and remove RFI have been developed, such as COASTGUARD (Lazarus et al., 2016, see ascl.net/2003.008), but the nature of RFI is that it is not uniformly predictable. It is therefore important to inspect data for evidence of RFI contamination and apply tailored flagging procedures according to requirements. Details of the data post-processing I applied to my observations are given in Section 3.3.

3.2 Observing projects

As described in Section 3.1.1, the core requirements of the types of observation needed to answer my open questions are as follows:

1. High S/N

2. Broad-band observations
3. Single pulse observations
4. Integrated pulse profiles
5. Full Stokes for polarization studies
6. Large-scale survey of the pulsar population

In order to achieve all of these requirements through a combination of datasets, I worked with three sets of observations, each obtained using a different radio telescope. The results of the work using these three datasets are described in three successive chapters in this thesis. I used the GMRT to observe single pulses from bright pulsars at 250–500 MHz (Chapter 5); the Parkes radio telescope for broad-band polarimetric integrated pulse profiles (Chapter 6) and the MeerKAT telescope for integrated pulse profile observations of a large number of pulsars (Chapter 7). Here is a description of the instruments in question, their key features and the reasons why each was particularly suited to the style of observations for which I used it.

The GMRT is a radio telescope array situated 80 km north of Pune, India, consisting of thirty antennas in a Y-shaped formation. Each antenna has a diameter of 45 m and the array consists of 14 antennas in a core region 1 km² in size, with the rest of the antennas spaced along the arms of the formation, so that in total the telescope covers a region of 25 km². Its upgrade (Gupta et al., 2017) has given it the capability to observe in four observing bands covering regions of frequency space between 120 and 1450 MHz, with a fifth (band 1, 50-80 MHz) also planned. The four bands available are at 120–250 MHz (band-2), 250–500 MHz (band-3), 550–850 MHz (band-4) and 1050–1450 MHz (band-5), providing near seamless frequency coverage except for those frequencies already known to be too strongly contaminated by RFI to be observable by the GMRT. In addition to its utility as an interferometric array, the GMRT has long been involved in pulsar science. As the largest array operating at metre wavelengths, it is suitable for both ISM studies (e.g. Krishnakumar et al., 2021) and observations of pulsar behaviour at low frequencies (e.g. Johnston et al., 2008). It is possible to divide the array into multiple sub-arrays which can observe at different frequencies simultaneously (Gupta et al., 2002), a feature that was used by Kramer et al. (2003) and Bhat et al. (2007) to investigate single pulse flux density spectra and broad-band nulling behaviour respectively. The 250–500 MHz band available with the upgraded u-GMRT receiver system is particularly interesting for

studies of the evolution of intrinsic pulse profiles, since it is observed that many pulse profiles evolve most strongly at frequencies below 1 GHz. Using the GMRT gave the opportunity to have dedicated time observing single pulses from a select small number of pulsars. This allowed me to address open question 1d, along with investigating Q1a-c for one pulsar in detail.

The Parkes radio telescope, often referred to as “The Dish”, is a 64 m single-dish telescope located on Wiradjuri country near the town of Parkes in Central West New South Wales, Australia. It was given the name Murriyang by Wiradjuri elders in November 2020. In Wiradjuri tradition Murriyang (literally translated as “Sky-world”) is the home of the creator spirit Biyaami. In addition to famously being used to broadcast the Apollo 11 Moon landing in 1969, Parkes has a long history as a pulsar instrument. It has been very successful at discovering new pulsars, for example, the second Molonglo survey used the Parkes and Molonglo telescopes together to find 154 new pulsars, more than doubling the total known at the time (Manchester et al., 1978). For a summary of the major pulsar searches undertaken by Parkes up until the turn of the millenium, see Manchester (2000). The telescope has also played an important role in observing known pulsars, both for timing with the Parkes Pulsar Timing Array (Kerr et al., 2020), and for studies of pulsar science (see for example Johnston & Kerr, 2018, and further information in Section 3.2.2). The installation of the new Ultra-Wideband receiver (Hobbs et al., 2020a) provides a continuous observing bandwidth from 700 to 4000 MHz. It is therefore now an ideal telescope for detailed study of how pulsar radio emission varies across a wide frequency range, and its continuous broad-band sensitivity is well-suited for polarimetric analysis. This made it an appropriate instrument to use to address my open questions, in particular questions 1c, 2b and 2c.

The MeerKAT telescope (Jonas, 2009) is a new 64-dish radio telescope in the Karoo desert, South Africa, and a precursor instrument for the SKA. It is a highly sensitive telescope with a gain of 2.8 K Jy^{-1} , four times that of Parkes and the largest gain of all radio telescopes in the Southern hemisphere. Fig. 3.1 shows a photograph of one of the antennas, which is set up with an offset Gregorian feed. The L-band receiver, the first to be available for science purposes, records 856 MHz of bandwidth between 856 and 1712 MHz with a system temperature of $\sim 18 \text{ K}$ (Bailes et al., 2020). A key strength of the MeerKAT telescope from a pulsar observational perspective is the possibility of obtaining high-quality pulse profiles using relatively short integration times. It is also possible to split the 64 dishes into sub-arrays in order to observe more



Figure 3.1: MeerKAT telescope antenna.
Photograph of one of the 64 antennas of the MeerKAT radio telescope, taken 6th December 2017 by Lucy Oswald.

than one source simultaneously, all of which makes it suitable for surveys of a large number of known pulsars (open question 2a).

The above explains the ways in which these three telescopes are suited to particular types of pulsar observations. The following sections describe how I applied this information to making observational choices, and the resultant observing projects undertaken using the three telescopes respectively.

3.2.1 Simultaneous observations with the Parkes and GMRT telescopes

In August/September 2018 I was Principal Investigator (PI) of a project to take simultaneous full polarization observations at both the Parkes radio telescope, Australia, and the Giant Metrewave Radio Telescope (GMRT), India, of five bright pulsars. For this I used the Parkes Ultra-Wideband Receiver (UWL) at 700–4000 MHz and

the upgraded GMRT receiver at 250–500 MHz. These were followed up by additional observations with Parkes alone of 6 more pulsars. The purpose of the simultaneous observations was originally intended to result in having a very broad frequency band over which the single pulse behaviour of these sources could be investigated. Ultimately, however, the single pulse work focused on using observations of PSR J1136+1551 from the GMRT telescope only, since it focused on investigating how pulse profile shape evolves with frequency (open questions 1c and 1d), and it is commonly observed that profile evolution is strongest at below 1 GHz. This meant the GMRT observing band of 250–500 MHz was ideally suited to the study.

There were several factors that dictated the choice of pulsars to observe in this project. In addition to being visible to both GMRT and Parkes simultaneously, the pulsars selected for observation had to be bright. In order to be detectable by Parkes, the peak flux density of the single pulses had to be $\gtrsim 10$ mJy over the entire frequency bandwidths of the GMRT and Parkes receivers. This limit was calculated from the radiometer equation as stated below for a single data point with a given time-bin and frequency-channel resolution:

$$S/N = \frac{S_{\text{source}} \sqrt{N_{\text{samp}} \tau \Delta\nu}}{S_{\text{sys}}}, \quad (3.1)$$

where S_{source} and S_{sys} stand for source and system fluxes respectively, τ is the duration of one time bin, $\Delta\nu$ is the bandwidth of a single frequency channel and N_{samp} is the total number of samples at the native time resolution of the instrument which have been integrated to produce the single data point. Parkes samples at a rate of $r = 8 \times 10^8 \text{ s}^{-1}$, and we required a bin resolution τ of 0.5 ms. The bandwidth $\Delta\nu$ of an individual channel is 16 MHz, meaning there are $N_{\text{chan}} = 206$ channels across the entire 700–4000 MHz observing band. N_{samp} is calculated as follows: $N_{\text{samp}} \times N_{\text{chan}} = r \times \tau$. Putting in the numbers, this means that a source flux of at least 10 mJy was required to obtain a S/N of 1.

An additional requirement was that the pulsars observed had to have a high DM (definition 3c) in order to minimize the effects of scintillation (definition 3g). The intensity modulation caused by scintillation is correlated over a scintillation bandwidth $\Delta f \propto f^4$ (Lorimer & Kramer, 2005), which can make it hard to investigate the intrinsic variability of the pulse shape across frequency. More distant pulsars, which tend to have larger DMs, scintillate more quickly in both frequency and time (Wang et al., 2005), whereas with more slowly scintillating pulsars there is a chance of an observation being entirely dominated by destructive interference and associated loss

of flux. When selecting pulsars to be observed, I excluded pulsars with low DMs, which I defined to be those with values below $10 \text{ cm}^{-3} \text{ pc}$. The exception to this rule was PSR J1136+1551, as the wealth of pre-existing literature on this pulsar, focusing particularly on the evolution of its profile shape, means that it is of particular interest for follow-up research. In this thesis therefore, the research based on single pulse observations is focused upon PSR J1136+1551 only, with analysis of the other single pulse data deferred to future work. The context of why this pulsar is most interesting for addressing open questions about pulsar beam structure are given at the start of Chapter 5.

3.2.2 Long-term observing programmes with the Parkes telescope

All of the polarization data presented in this thesis originate from observations made as part of the P574 observing programme on the Parkes telescope. This is a project to monitor a set of young pulsars every month, and has been running continuously since 2007 (Weltevrede et al., 2010). The key information required for a timing solution for a given pulsar, such as its topocentric period P and period derivative \dot{P} , is known as its ephemeris. The P574 project was originally set up to provide ephemerides to be used by NASA's *Fermi* Gamma-ray Space Telescope mission (Smith et al., 2008). In 2014 the list of observed pulsars was altered to widen the science goals investigable with the dataset, removing low flux pulsars in favour of pulsars with higher flux and lower values of \dot{E} . The observing programme runs on a monthly cadence and observes around 180 pulsars for intervals of a few minutes per pulsar. This allows long-term time variability to be studied. High S/N pulse profiles are created by combining together profiles from multiple observations. The polarimetric data presented in this thesis consist of combined integrated pulse profile observations made as part of the P574 project since the installation of the UWL in late 2018 (Hobbs et al., 2020a). The UWL has greatly increased the capacity of this dataset to be useful for understanding the frequency evolution of pulsar emission, including its polarization. This enables high precision measurements of DM and RM, and their variability with time (Johnston et al., 2021), as well as the study of intrinsic polarimetric properties of individual pulsars in great detail. The dataset is therefore very well suited to addressing the open questions laid out in Section 1.3, both in terms of detailed study of a small number of sources and in the possibility of extending this to a large scale survey of the properties of a wide range of pulsars.

3.2.3 MeerKAT: MeerTime and the Thousand-Pulsar-Array

The Large Survey Project devoted to understanding known pulsars on the MeerKAT telescope is the MeerTime collaboration, which seeks to expand understanding of pulsar physics and perform high quality pulsar timing observations through a programme of pulsar observation Bailes et al. (2016). The scientific goals of MeerTime are divided into four themes: the Thousand-Pulsar-Array (Johnston et al., 2020; Song et al., 2020), Relativistic and Binary Pulsars (Kramer et al., 2021), Globular Clusters and the MeerKAT Pulsar Timing Array. Early published work based on the MeerTime dataset has studied giant pulses (Abbate et al., 2020), jitter of millisecond pulsars and single pulse variability (Parthasarathy et al., 2020), a description of the polarization properties of young energetic pulsars (Serylak et al., 2020) and the discovery of eight new millisecond pulsars in the globular cluster census (Ridolfi et al., 2021).

The Thousand-Pulsar-Array (TPA) project seeks to investigate the emission physics of young pulsars and is building up census observations with which to do so. The goals of the TPA are to obtain high quality integrated pulse profiles of 1000 pulsars and 1 million single pulses. Observations have been taking place for the TPA since February 2019. During the initial months of observations, the data from MeerKAT were not yet polarization- or flux-calibrated, however my work on analysing the scattering of pulse profiles did not require these extra calibration steps. The scale of the TPA dataset and the quality of the observations make it ideal for studying population-level questions about pulsars and the ISM (open questions 2a and 2b).

3.3 Observing methods and data post-processing

3.3.1 GMRT observations

In August 2018 I travelled to India to perform observations with the GMRT telescope and coordinate the simultaneous observations of the same pulsars taking place at the Parkes telescope. We recorded pulsar data with GMRT in filterbank format, consisting of auto and cross-polarized powers, sampled with a time resolution of $327.68 \mu\text{s}$ and 2048 frequency channels across a band of 300–500 MHz (Band-3 of the u-GMRT). Observations were intended to be recorded with 20 antennas, consisting of 14 in the core of the array and the first two antennas in each of the arms. The observations were ultimately recorded with 18 antennas in the first session, and 17 antennas in the second session, however, this did not affect data quality.

Observations at GMRT were taken over two sessions on Friday 31st August 2018 and Saturday 1st September 2018. The first session was run simultaneously with observations taken at the Parkes radio telescope. Since the GMRT is a telescope consisting of multiple dishes, its sensitivity can be increased by coherently adding the signals observed by each dish in a phased array. To do this it is necessary to phase up the array by observing a known source, to ensure that the dishes are all pointing towards the same part of the sky. Observations of known continuum sources were therefore made at the beginning of the session and between pulsars to phase up the antenna array, and to check and, if necessary, rephase at intervals if the dishes have drifted out of alignment. The initial power equalizing and phasing was done with two known continuum sources: 3C 48 on Friday 31st August and 3C 286 on Saturday 1st September. Re-phasing observations were done using sources close to each pulsar to be observed. It was also necessary to account for fixed delays between the data arrival times for the two polarized feeds, and a 5 minute observation of PSR J0953+0755 was taken to check that this correction was applied accurately.

From the full set of observations, the two that are of relevance in terms of further science work in this thesis are those of PSR J1136+1551 and the continuum source 3C 241 which was used to check the phasing before this pulsar, and was then used to gain calibrate the data. This latter calibration procedure is described in detail in the next section. We observed PSR J1136+1551 for 94 minutes with GMRT and for 40 minutes with the Parkes radio telescope, the shorter observing time at Parkes being due to necessary stowing of the dish in high winds. We observed a further four pulsars, namely PSRs J0630–2834, J0742–2822, J0738–4042 and J1752–2806. However, I focused on PSR J1136+1551 for the work presented in this thesis, since its double-peaked profile shape and the previous research on this source make it the most relevant for questions about pulsar beam structure and frequency evolution. I therefore did not take analysis of these other observations further.

I processed the filterbank data products to generate individual, single pulse files using the software package DSPSR developed by van Straten & Bailes (2011). These were incoherently dedispersed with a DM of $4.926 \text{ cm}^{-3} \text{ pc}$. Each single pulse file could then be read by the software package PSRCHIVE (Hotan et al., 2004), which I used to read and edit the data to apply gain calibration and RFI cleaning. Fig. 3.2 shows an example single pulse that has been calibrated and cleaned of RFI, with the full channelization and time binning.

It was found in this subsequent processing of data products that the simultaneous observations at the GMRT and Parkes telescopes were misaligned in time by an




Figure 3.2: Single pulse of PSR J1136+1551 observed with GMRT.

The main figure shows a waterfall plot of total intensity across pulse phase (x axis) and observing frequency (y axis) for a single pulse of PSR J1136+1551. The pulse has been gain calibrated and cleaned of RFI, as can be seen for example in the large section of band removed at 370 MHz. Right: the total power as a function of frequency. Top: pulse profile made from summing intensity across the observing band. It shows three distinct components. This figure was made using PSRCHIVE (van Straten et al., 2012).

amount of the order of a few seconds. This meant that it was difficult, and impossible for the pulsars with a small rotational period, to associate the single pulses observed with one telescope with those observed by the other at a different frequency band. Investigation of the data processing used for both pulsars identified that the software package DSPSR automatically truncates clock times to the nearest second, which is compatible with the systems at the Parkes radio telescope but not with the data recording setup of GMRT. However, correcting for this error was insufficient to correct for the total misalignment between the two datasets. Ultimately, I was unable to find the source of this error or to quantify the exact amount of misalignment, and concluded it might be attributable to some further, unidentified, difference in how Parkes and GMRT define their timestamps, particularly since, as GMRT is an antenna array, definition of the phase centre of the interferometer is not straightforward.

3.3.2 Calibration and RFI excision

The raw data products of the GMRT pipeline described above are complex coherence parameters in the left and right circular basis. These are obtained as described in Section 3.1.2 by performing auto- and cross-correlations of the two orthogonally polarized signals E_l and E_r recorded by the polarized feeds. This gives $|E_r|^2$, $|E_l|^2$, $\text{Re}(E_l^* E_r)$ and $\text{Im}(E_l^* E_r)$. These are converted to Stokes parameters as follows:

$$\begin{aligned}
 I &= |E_r|^2 + |E_l|^2, \\
 Q &= 2\text{Re}(E_l^* E_r) \\
 U &= -2\text{Im}(E_l^* E_r) \\
 V &= |E_r|^2 - |E_l|^2.
 \end{aligned}
 \tag{3.2}$$

For notational simplicity, from here onwards I will refer to the coherence parameters as RR, LL, Re(RL) and Im(RL) respectively.

Before converting to Stokes, the coherence parameters must be gain calibrated and cleaned of RFI. Consistent self-calibration of the data is required so that they are comparable across both the bandwidth of the observation and between different points in time. If parameters Q , U and V were required, polarization calibration would also need to be implemented. Furthermore, to be able to investigate the absolute flux and polarization of the data, observations of calibrator sources must be taken, for which these absolute values are known. However, since the goal of the work was to focus on how profile shape evolved with frequency (open question 1c), only Stokes I was needed and absolute flux calibration was not required.

Gain calibration of the recorded signal needs to account for two factors: first that the gain response of the telescope receiver is not uniform across the observing band but has a characteristic bandpass shape. Secondly, the gain responses of the two circularly polarized feeds are slightly different. Any frequency- and time-dependent variation measured in off-pulse region data will be due primarily to the telescope system response, with a Gaussian noise contribution from the sky background. Therefore by bringing the feed inputs into the same frame of reference and then normalizing the whole observation such that the off-pulse region becomes Gaussian noise centred on zero, the on-pulse observations will be correctly calibrated with respect to the telescope response. My approach to gain calibration was therefore to subtract the intrinsic telescope bandpass from the coherence parameters and then scale them by their relative gain responses.

I obtained the intrinsic bandpass by averaging over the offpulse region of the pulsar observation ($RR_{offpulse}$ etc. for the four coherence parameters). I calculated this anew for every single pulse, so that temporal fluctuations in the telescope gain response were accounted for. This gives the calibrated coherence parameters as follows:

$$\begin{aligned}
RR_{cal} &= RR - RR_{offpulse}, \\
LL_{cal} &= G(LL - LL_{offpulse}), \\
\text{Im}(RL)_{cal} &= \sqrt{G}(\text{Im}(RL) - \text{Im}(RL)_{offpulse}), \\
\text{Re}(RL)_{cal} &= \sqrt{G}(\text{Re}(RL) - \text{Re}(RL)_{offpulse}),
\end{aligned}
\tag{3.3}$$

where G is the relative gain response, given as the ratio of the respective gains of the right and left polarized feeds. These I calculated by observing the continuum source 3C 241, and subtracting off the intrinsic bandpass, giving

$$G = \frac{RR_{continuum} - RR_{offpulse}}{LL_{continuum} - LL_{offpulse}}.
\tag{3.4}$$

Fig. 3.3 is a plot of $RR_{continuum}$ and $LL_{continuum}$. This shows the overall bandpass shape, the differences for the two polarized feeds, and the presence of strong RFI signals in some channels, particularly around 370 MHz. I needed to ensure that the channels consistently affected by RFI would not affect the calibrated data products. I therefore smoothed the intrinsic bandpass and the relative gain response G using a median filter, with the size of the region over which the median is calculated being 10 and 100 frequency channels respectively out of the total channel resolution of 2048 channels. Given the shape of the continuum bandpass and the distribution of RFI structures within it as seen in Fig. 3.3, these filter regions were found to be most effective for smoothing. I then performed RFI filtering on the calibrated parameters.

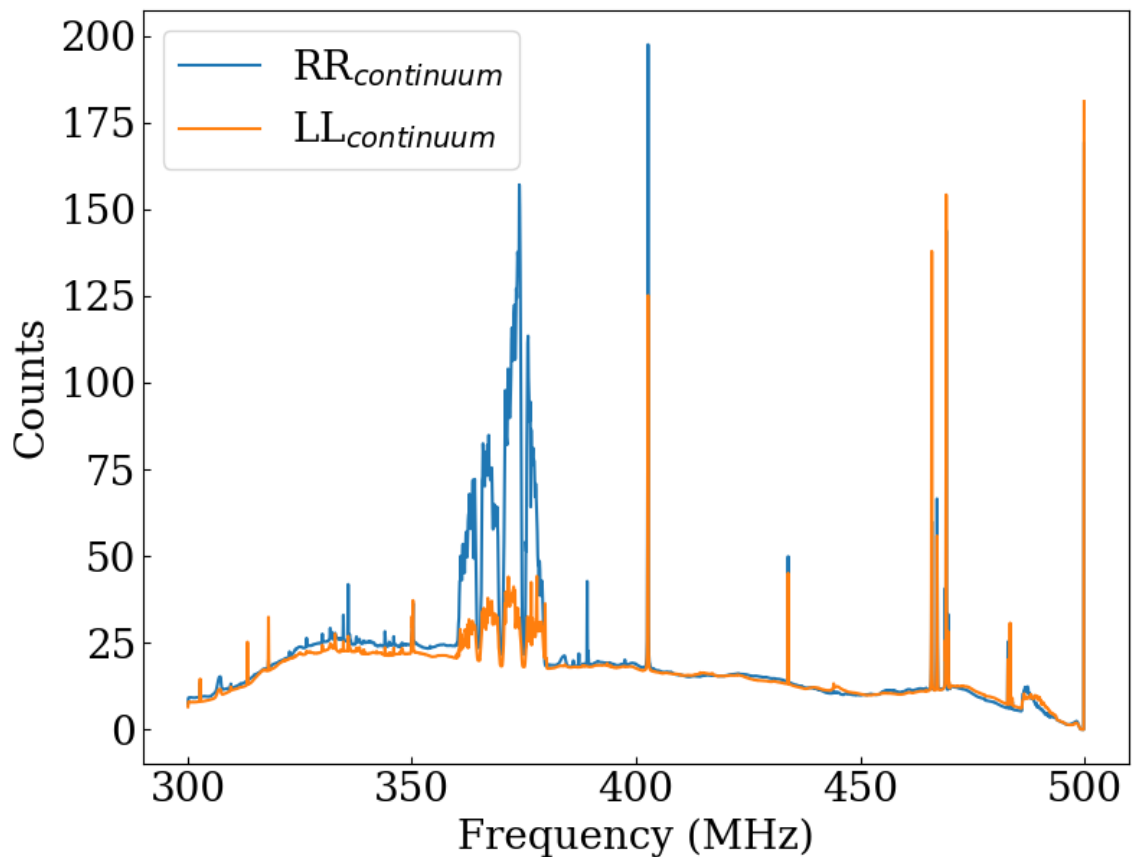


Figure 3.3: Continuum bandpass of source 3C 241. This shows how flux counts vary with frequency for the continuum source 3C 241 observed with the GMRT band-3 receiver. The two lines mark the bandpasses for the two polarized feeds, demonstrating the difference in response of the feeds to the input signal. RFI is visible as individual spikes at various frequencies and a sustained region of high RFI around 370 MHz. Figure made with matplotlib.

RFI filtering has two stages: those frequency channels for which there was persistent RFI were set to 0, then narrow instantaneous spikes of RFI found in individual channels were identified and smoothed. For stage 1, I calculated the bandpass of LL_{cal} as the mean across phase and create a smoothed copy by median-filtering with filter region size equal to half the total number of channels. Subtracting one from the other gives a residual that is close to 0 everywhere other than for channels contaminated by RFI. I defined a cutoff of 5 times the standard deviation of the bandpass, and set to 0 all channels of each coherence parameter for which the absolute residual was larger than this cutoff. For stage 2, the pulse profile in each channel was smoothed with a median filter with region size of 5 phase bins out of a total 3600. The profile residual is then the single pulse with the smooth profile subtracted. Again, spikes in RFI, now instantaneous, were identified by comparing the residual profile to a cutoff, now of 5 times the standard deviation of the profile for the given frequency channel being treated. The RFI spike is replaced with the value of the smoothed profile at the corresponding phase bin.

Having cleaned the data of RFI, they were then aggregated across frequency into 10 frequency channels, each of bandwidth 20 MHz spanning the observing band. The choice of 10 channels was to separate out and discard the three regions of high RFI, at 310, 370 and 490 MHz respectively, to leave seven profiles with high S/N that span what remains of the observing band.

One further contaminating source of RFI was identified in the GMRT data in the form of an oscillating signal. Since the the frequency of oscillation is 50 Hz, this suggests the presence of a source of mains electricity. I attempted to remove this signal from the data by cutting out the 50 Hz peak in Fourier space with the IIR notch digital filter from SCIPY (Virtanen et al., 2020). However, since the widths of profile components in the single pulses also correspond to a peak in the same region in Fourier space, it was not possible to remove the 50 Hz oscillation in this way without having a strong impact on the single pulse shape. Since the magnitude of the 50 Hz signal fluctuated with time it did not uniformly contaminate the entire observation. I therefore elected ultimately to remove those single pulses strongly contaminated by the signal from the dataset.

3.3.3 Data processing for Parkes P574 and MeerTime TPA observations

Since the P574 and TPA programmes involve multiple collaborators, the data processing was done by others, using pipelines written especially for that purpose. The

key concepts of necessary processing of pulsar data are the same for all pulsar observations, and so the pipelines follow similar physics principles to those laid out above in the context of processing the GMRT data. The specific details are as follows, and a summary of all the key data products presented in this thesis, along with the observation type and time and frequency resolutions, is given in Table 3.1 (see Section 3.3.4).

The P574 data are coherently dedispersed and folded on the topocentric pulsar period to generate subintegrations of length 15s. This is digitized to 1024 phase bins and 3328 frequency channels, each of 1 MHz bandwidth. Polarization calibration is necessary to allow us to investigate the polarimetric properties of the pulsars studied. This involves calibration of the relative gain responses of the two orthogonal feeds, along with any relative phase difference of the feed signals, which may be introduced due to differences in cable lengths in the hardware. Throughout the observing programme we make observations every 60 minutes of an injected pulsed square-wave signal, aligned at 45° to both orthogonal feeds. Since the true polarimetric behaviour of this injected signal is known, calibration matrices can be calculated from these observations according to the formalism of Hamaker et al. (1996). These may then be used to correct the pulsar observations. After cleaning the calibration observations of RFI, this technique is applied using PSRCHIVE (van Straten et al., 2010). Absolute flux calibration is done by making observations of known source Hydra A and so converting the units of the digitizer to mJy.

At the time that the research presented in Chapter 6 was completed, polarization calibration efforts for the P574 data were restricted to the technique described above. Subsequently, further calibration was applied to account for instrumental leakage between linear and circular polarization, using the *pcm* technique described by van Straten et al. (2013). This technique uses observations of PSR J0427-4715 at multiple parallactic angles to map out the polarimetric response of the observing system. In the case of the P574 data, *pcm* corrects leakage of linear polarization into circular polarization at the high frequency end of the observing band for highly polarized pulsars. For the two pulsars presented in Chapter 6 (PSRs J1056-6258 and J1359-6038, see also the same data in Fig. 4.11), we find that the *pcm* correction has a negligible effect on the linear polarization, on which the research in that chapter is focused. As a result, application of this correction would not affect the results presented. Apart from PSRs J1056-6258 and J1359-6038, all other UWL data presented in this thesis have been *pcm*-corrected.

The edge channels and any data affected by RFI are flagged and the calibrated subintegrations are summed over time to produce a single integrated profile with all four Stokes parameters. This is aggregated in frequency to produce either 8 profiles spanning the observing band, or 52. These choices were made for two purposes: the 8-channel data are a set of high S/N pulse profiles for detailed study of profile shape and polarization properties. The 52-channel data sacrifice S/N for resolution across the broad observing band, to show the frequency structure of the data in detail. The choice of 52 channels does however have increased S/N relative to the native 3328 channel resolution, since it has been aggregated by a factor of 64. The pulse profiles already shown in Chapter 1 are examples of this dataset. Examples of the 52-channel data products are shown in Chapters 4 and 6. This thesis presents figures from many pulsars observed as part of this dataset, but analysis work in Chapter 6 focuses in particular on two pulsars: PSRs J1056–6258 and J1359–6038, as noted in Table 3.1.

Out of 1164 pulsars observed between commence of observations in February 2019 and 1st June 2020 as part of the TPA programme, I selected 205 by eye for a survey of the scattering properties of single-component pulsars. The pulsars chosen had a profile shape resembling a single Gaussian component convolved with an exponential decay function, so that a scattering model with these properties would be likely to be appropriate to describe their behaviour. The initial sample was deliberately selected to be as broad as possible within these constraints, so that the sample was not biased by visual selection. Specifically, I selected all pulsars that appeared to have only one single component, excepting those very few for which that single component was so idiosyncratically shaped that it could not reasonably be modelled by a Gaussian. My approach to the exponential decay criterion was similarly broad: I included profiles with little to no apparent scattering tail on the basis that non-scattered pulsars would be more accurately identified after applying the modelling process. This allowed me to investigate the limitations of the modelling choices, but required pulsars to be discarded at later stages if their profiles were not well modelled with single Gaussians modified by scattering tails. Observations were made with the L-band receiver on the MeerKAT telescope, generally using a tied-array beam formed out of at least 58 antennas of the total 64. Data were coherently dedispersed and recorded in fold-mode as 8 s subintegrations. They were cleaned of RFI using COASTGUARD (Lazarus et al., 2016, see ascl.net/2003.008) and time-integrated to a single Stokes I profile per observation with a resolution of 1024 bins per pulse period. Observations from the early commissioning phase of the MeerKAT telescope had only the band between 895 and 1670 MHz available out of the ultimate full band of 856–1712 MHz. In order

to make all the observations directly comparable with each other, I therefore kept the 928 frequency channels between 895 and 1670 MHz, divided into 8 sub-bands, in the final data products. Where a pulsar has multiple good quality observations taken at different epochs, these were aligned and added to produce a single profile with increased S/N. This was done through the following method. Within the TPA project, templates of the pulse profile shape were obtained for each individual pulsar and the observational parameters that constitute the ephemeris for a given pulsar, such as the topocentric period, were updated where necessary. These ephemerides were then used to measure the phase shifts between individual observations of a pulsar, using the software TEMPO2 (Hobbs et al., 2006; Edwards et al., 2006) and PSRCHIVE (van Straten et al., 2012). This enabled the profiles to be correctly aligned in phase prior to adding. Full details of the MeerKAT backend are given in Bailes et al. (2020) and a summary of the TPA observing strategy is described in Johnston et al. (2020) and Serylak et al. (2020).

3.3.4 Summary table of data products

Table 3.1: Details of the pulsar observations analysed in this thesis. Note that PSR J1136+1551 had ten channels spanning the 200 MHz band, but three of those channels were discarded due to persistent RFI, leaving seven channels for subsequent analysis. This thesis also presents observations of multiple pulsars observed with the Parkes telescope, for which the observational details are the same as for the two pulsars upon which work was focused.

PSRJ	Telescope	Band (MHz)	N_{chan}	N_{bin}	Data products
J1136+1551	GMRT	300–500	7	3600	Single pulses, Stokes I
Focused work: J1056–6258, J1359–6038	Parkes	700–4000	8 or 52	1024	Integrated profile, full Stokes
205 pulsars	MeerKAT	895–1670	8	1024	Integrated profile, Stokes I

3.4 Measuring the effects of the ISM

An important part of the processing and analysis of pulsar data is to understand and correct for the effects of the ISM. This section summarizes the ways in which correction of dispersion and Faraday rotation has been done in the past, and explains

the methods used to identify the effects of time domain scattering and scintillation in pulsar observations. This summary provides context for the work presented in Chapters 6 and 7, in which I describe my own approaches to these measurements.

3.4.1 Dispersion and Faraday rotation

Section 1.1.4 introduced the concepts of how the ISM disperses and Faraday rotates pulsar radio waves, and how these effects can be encompassed theoretically by the dispersion and rotation measures (DM and RM, see definitions 3c and 3d in Chapter 2). It is important here to stress the difference between the theoretical descriptions of the DM and RM, and the values that are measured for a given pulsar. Taking the DM as an example, according to equation 1.7, it is a constant of proportionality relating the time delay of a pulse to its radio frequency. It is therefore identified observationally as the constant required to correct for this time delay and so align pulse profiles across frequency. It is taken as read that this alignment constant is equal to the measurement of integrated electron column density along the line of sight to the pulsar, the theoretical definition of the DM. However, this connection requires the underlying assumption that we, the observers, know how the pulse profile should be aligned across frequency. In reality, there is not a single unambiguous way to align pulse profiles, since they tend to evolve in shape with frequency (see for example Fig. 1.6). In addition to this, other factors may affect observed pulse shape and alignment, including scattering, which alters profile shape in a frequency-dependent way, and the delay-radius relation (definition 2i), wherein pulse profiles are observably shifted depending on the height above the pulsar surface at which the emission is produced. Following on from this, the RM that is measured will depend upon the choice of DM, since the way in which the PA varies across frequency will depend upon the alignment of pulse profiles relative to each other in time. With the above in mind, here is a summary of the ways in which the DM and RM have been identified from pulsar observations.

When pulsars are observed for the express purpose of pulsar timing, the goal is to create a high S/N profile so that a time of arrival can be calculated. The logical choice of DM then, and the simplest to calculate, is that which maximises this S/N of the pulse profile when it is aggregated across frequency. Pulsar timing software such as TEMPO2 (Hobbs et al., 2006) chooses the DM which minimizes the spread of measured times of arrival (TOAs). The TOA value calculated at a given frequency will depend on the template profile applied, which may or may not incorporate some level of profile evolution across the observing band (e.g. Pennucci, 2019).

By contrast, when studying intrinsic pulsar emission behaviour, care has been taken to ensure that the DM chosen results in alignment of pulse profiles across frequency that is in accordance with some set of theoretical expectations. For example, Hassall et al. (2012) chose to align pulse profiles such that either the brightest peak or the mid-point between the two brightest components were aligned, and inferred the DM from this. Such a method ensures that the DM is not being skewed by the effect of profile evolution, but simultaneously makes an assumption about the location of the fiducial point of a pulse profile based upon symmetry. Through their modelling of scattered pulsars, Geyer et al. (2017) identified corrections to their DMs by modelling the intrinsic pulse profile as a single Gaussian which has been distorted by the scattering transfer function. By locating the position of the intrinsic Gaussian, they inferred the DM required to align the Gaussian peaks across observing frequency. This corrects for the effect of scattering shifting signal to later phases in a frequency-dependent way, but also requires the assumptions that the pulse profile can be modelled by a single Gaussian component, and that the fiducial point of the pulse profile is such that the Gaussian peaks can be aligned across frequency.

To correct for Faraday rotation, the Stokes parameters Q and U are rotated by an angle that depends on both the observing frequency and the RM of the pulsar (definition 3d), according to equation 1.8. Having rotated Q and U , the de-rotated PA is then given by $PA = 0.5 \arctan(U/Q)$. The relevant RM to apply is usually identified by calculating a single average value of the PA for a pulse profile, and then fitting for the quadratic frequency dependence of the PA shift due to Faraday rotation. Traditional techniques for measuring the RM are as follows.

- Finding the value of RM that maximizes the S/N of linearly polarized intensity L averaged across frequency (this is implemented in PSRCRIVE, van Straten et al., 2012).
- Directly fitting either a quadratic relationship between PA and frequency (Noutsos et al., 2008) or a straight line to the PA in frequency-squared space.
- Performing a Fourier Transform on the PA in frequency-squared space and finding its maximum; this is known as RM-synthesis. It was developed by Brentjens & De Bruyn (2005) and has been widely used, see for example Sobey et al. (2019).

Work has also been done calculating phase-resolved RMs, such as that by Noutsos et al. (2009) and Ilie et al. (2019). This has revealed variation across pulsar profiles

which has been attributed both to interstellar scattering and to effects intrinsic to the pulsar magnetosphere. The best fit RM for a given pulsar will depend upon the choice of DM used to dedisperse the profiles. In their measurements, Ilie et al. (2019) chose the value of DM that minimized the resulting variations in measured RM.

3.4.2 Scattering and scintillation

As described in Section 1.1.4 and defined in Chapter 2, the effect of scattering by the ISM on pulse profiles is described by the characteristic scattering timescale τ and the spectral index describing its evolution with frequency, α . As with DM and RM, it is important to stress the difference between the theoretical descriptions of what these values represent, and the quantities that are measured in observations through a variety of techniques. The best method for measuring τ and hence inferring α is dependent on various features of the pulse profile, including its intrinsic complexity, radio frequency and scattering strength. In the literature, values of α are usually calculated by performing a power law fit to τ against frequency, where τ is measured through one of two ways. In the cases where the scatter broadening is evident in the pulse profile, we measure τ in the time domain. This is done either through forward modelling (e.g. Geyer et al., 2017) or through deconvolution analysis, such as the CLEAN algorithm (Högbom, 1974; Bhat et al., 2003). Scattering surveys that have employed forward modelling, and are thus directly comparable to the work presented here, are as follows: Cordes et al. (1985); Löhmer et al. (2001, 2004); Kuzmin & Losovsky (2007); Lewandowski et al. (2011, 2013, 2015a,b); Krishnakumar et al. (2017); Geyer et al. (2017) and Krishnakumar et al. (2019). The CLEAN algorithm has been employed by, for example, Bhat et al. (2004) and Kirsten et al. (2019). A complementary technique is to infer τ from the scintillation bandwidth $\delta\nu$ (Cordes et al., 1985), where scintillation is a variation in pulse intensity over both frequency and time, caused by the constructive and destructive interference of radio waves scattered along different paths by the inhomogeneous ISM. The scintillation bandwidth is identified from the dynamic spectrum, which is formed by plotting pulse intensity across both time and frequency as a waterfall plot. It is related to the scattering timescale by the equation $2\pi\tau\delta\nu = C_1$, where C_1 is a factor with a value that depends on the electron density wavenumber spectrum and distribution of scattering material. When transforming observational values it is usually set to $C_1 = 1$ (e.g. Bhat et al., 2004), the solution for a thin screen, and other values are presented in Cordes & Rickett (1998). This technique is relevant for observations where τ is too

small to be measured directly, generally at higher frequencies. The inverse is also true: where τ is large, $\delta\nu$ will be too small to be measurable.

Measurements of other scintillation parameters give additional insight into the scattering properties of the ISM. The dynamic spectrum can be Fourier transformed to generate the secondary spectrum, within which parabolic scintillation arcs are often visible. These arcs are evidence for scattering being dominated by a thin, anisotropically scattering screen. An arc's curvature depends upon the location of the screen relative to both the pulsar and the observer, which can then be measured (e.g. Main et al., 2020). Multiple arclets, themselves parabolic, located along the main arc parabola, result from further interference of the diffracted images with each other. Brisken et al. (2010) converted the points on the secondary spectra, for PSR B0834+06, into positions mapped out on the sky. Doing so revealed an extremely anisotropic screen shape. The presence of multiple screens along the line of sight will result in there being more than one parabola in the secondary spectrum. Recent work by Simard et al. (2019) has developed the mapping technique of Brisken et al. (2010) to apply to pulsars with multiple screens.

3.5 Summary

This chapter has described the process for observing a pulsar using a radio telescope and for processing the data to obtain broad-band integrated pulse profiles and single pulses with high S/N. It has summarized the factors that must be taken into account when choosing a particular instrument and set of data products for a given science goal, and explained the reasons behind my choices for the data presented in this thesis. The details of the data presented here, along with the discussion of methods for measuring the effects of the ISM, are an important basis upon which I built my methodology for investigating the open questions. Details of the techniques I developed for analysing broad-band data and simulating pulsar emission are given in the next chapter.

Chapter 4

Methodology

In the introduction I described the motivations for addressing the key questions about the beam structure and energetics behind pulsar radio emission, and how these factors vary across the pulsar population (see the open questions in Section 1.3). When using pulsar observations to investigate these questions it is important to be aware of the observational constraints on the information available about the pulsar population (open question 2). Pulsar radio emission is produced by a three-dimensional region in the pulsar magnetosphere, yet the perspective of the observer is a single line of sight cut through the pulsar beam. It is from this view alone that we must attempt to infer the radio beam structure as a whole. In addition, pulsar emission is altered by its passage through the ISM. Conclusions drawn from inspection of pulse profiles must therefore take into account the entire emission and propagation process. Working with these constraints can in fact provide additional relevant information. For example, it is the very signatures of propagation that occlude understanding of the intrinsic behaviour of the pulsar which reveal information about the structure of the ISM. In order to extract answers from my observations it was necessary to reduce the data to the key parameters and metrics that describe the features that are relevant for a given question. This meant quantifying features such as single pulse component position, the interaction of intrinsic profile shape with a scattering function, and the ways in which Faraday rotation and dispersion affect alignment and polarization of pulse profiles.

The methods used for extracting the relevant information from observations must be grounded in the science they intend to illuminate. For example, fitting for a quadratic dispersion delay in pulse arrival time should take account of the fact that other sources of delay may also be present, due to profile evolution and the effect of the rotating pulsar reference frame. Prior knowledge of such factors is hard to quantify, particularly when the purpose of the data analysis is to obtain answers about

these very questions of intrinsic pulsar emission behaviour. I addressed this problem by simulating how pulsar observations behave under certain theoretical conditions. Through making comparisons with the real phenomena observed in pulsar data, I could draw conclusions about pulsar emission and propagation.

4.1 Simulating the observables of pulsar physics

When creating simulations, it is necessary to define the factors which the simulation is being designed to investigate. From an observational perspective the key attributes of an integrated pulse profile are its profile shape, its polarization properties, and how these evolve with frequency. The origins of these features are based in the emission processes of the pulsar itself, so it is necessary to consider the observational basis and assumptions that underlie simulation choices. My simplest simulations of pulse profiles employed a single normalized Gaussian component, 50% linearly polarized, with the PA described by the rotating vector model. I selected the width of the Gaussian so that the pulse profile took up roughly 10% of the pulse period, which is comparable to the majority of young pulsar profiles, though not, for example, to millisecond pulsars. Pulse polarization fraction is observed to be anything between 0% and 100%, so a choice of 50% is reasonable but should also depend on the comparisons being made. For the geometry of the pulsar, which underlies the PA shape, I considered that constraints on beam opening angle, inferred from the last open field lines at a given rotational period, provided an upper bound on possible measurements of line-of-sight impact parameter β . I therefore constrained this parameter to small opening angles (the choices of relevance being -2° , 3.7° and -10.5°). Again such a constraint would not accurately reflect the properties of millisecond pulsars, for which the small rotational period leads to a narrow light cylinder and hence wide beam. Based on my assumptions I obtained the Stokes parameters (I , Q , U , V) for a simulated profile through the following steps, adding Gaussian noise in order to make the profile more directly comparable to data.

- I is defined as the shape of the Gaussian profile chosen,
- $V = 0$,
- Q and U are obtained by transforming $L = \sqrt{Q^2 + U^2}$ and $\text{PA} = -0.5 \arctan(U/Q)$.

Fig. 4.1 shows an example simulated pulse profile with $\alpha = 100^\circ$ and $\beta = -2^\circ$. The assumption of $V = 0$ means that the described simulation is only relevant for

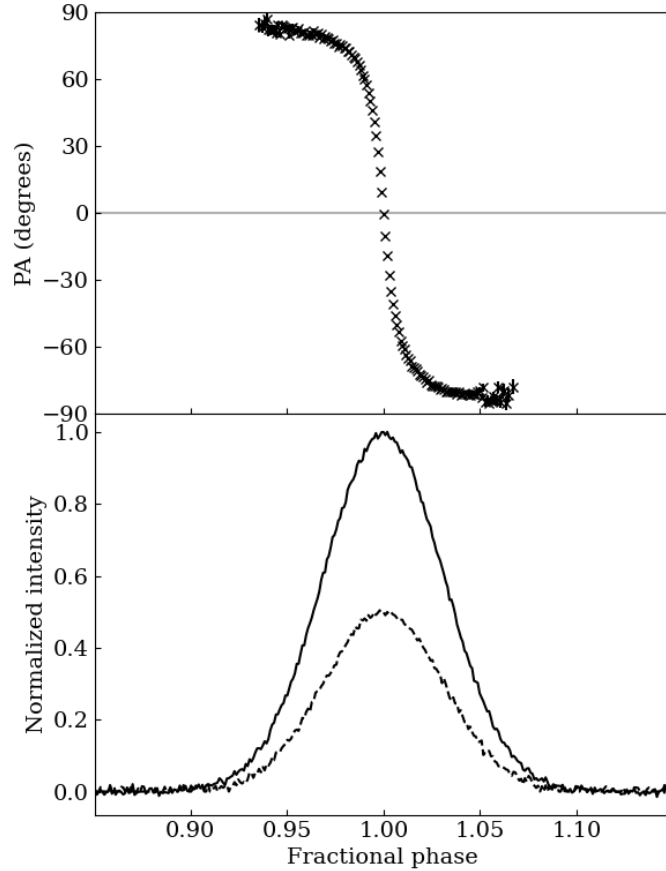


Figure 4.1: Simulated integrated pulse profile.

Top: PA profile showing only those points where the signal-to-noise ratio of the linear polarization $L_{\text{SNR}} > 8$. The PA was generated from the rotating vector model with $\alpha = 100^\circ$ and $\beta = -2^\circ$. Bottom: the total intensity profile (solid line) is a Gaussian with noise level equivalent to that for PSR J1056–6258 at 827 MHz. This simulated profile has 50% linear polarization, L , plotted as a dashed line. Figure made with matplotlib.

comparison with pulsars with little to no circular polarization in their pulse profiles. The ways in which circular polarization affects observed polarimetric characteristics are discussed in Chapter 8, particularly Section 8.1.2.

Certain additional factors need to be added into the simulation to account for profile evolution with frequency. Although the above simulation is created without making assumptions about absolute flux, so that I is normalized to have a maximum value of 1, the total flux of observed pulse profiles generally decreases with increasing frequency. This means the comparable signal-to-noise ratio tends to decrease. In order to account for this in simulated profiles, I set the standard deviation of the Gaussian noise added to be equal to the corresponding value I measured for a

given frequency from the pulse profile of PSR J1056–6258. This was an integrated profile observed using the UWL receiver on the Parkes radio telescope and frequency-aggregated into 8 profiles spanning the band, as described in Section 3.3.3. Since these pulsar simulations were developed for direct comparison with this style of observation, taking noise measurements directly from the observation improved the comparability of simulations and data.

Observing frequency also affects pulse profiles through the effects of both ISM propagation (dispersion, Faraday rotation and scattering) and intrinsic profile evolution due to pulse beam shape. Dispersion can be simplistically modelled as misalignment of pulse profiles by the appropriate quadratic time delay (see Section 1.1.4). A relevant factor is that a pulse profile, if formed from adding signal over a wide, dispersed, band, will necessarily have a smeared shape due to a lack of dispersion correction before adding. It is therefore important to bear in mind the bandwidth that a simulated profile represents, and the extent of the dispersion delay being introduced, when creating a dispersed simulation that does not include such smearing. The same effect is also relevant for the misalignment of PA profiles due to Faraday rotation. Fig. 4.2 shows the same simulated profile as in Fig. 4.1, now broad-band and with dispersion delay and Faraday rotation applied such that $DM = +1 \text{ cm}^{-3} \text{ pc}$ and $RM = +2 \text{ rad m}^{-2}$. The effect of changing S/N is visible: higher frequency profiles have lower S/N and therefore appear more noisy.

Pulse profile shapes are distorted by effect of scattering by the inhomogeneous ISM. In order to simulate a scattered pulse profile I convolved the intrinsic profile shape with an isotropic scattering transfer function, which is described in terms of the scattering timescale τ as

$$\frac{e^{-t/\tau}}{\tau}U(t), \quad (4.1)$$

where the unit step function $U(t)$ constrains the equation to time $t > 0$ (Cordes & Lazio, 2001). Fig. 4.3 demonstrates how this alters the profile shape. Details of how I used simulations to investigate the effect of modelling scattered multi-component profiles with a single-component model are given in Section 7.3.

The choices required to describe the evolution of intrinsic pulse shape with frequency are strongly dependent on the assumptions made about the pulsar emission mechanism and beam structure. This is relevant both for the style of integrated pulse profile simulation described so far, and for simulations of the frequency evolution of the phase positions of single pulses. The next sections address these factors in detail.

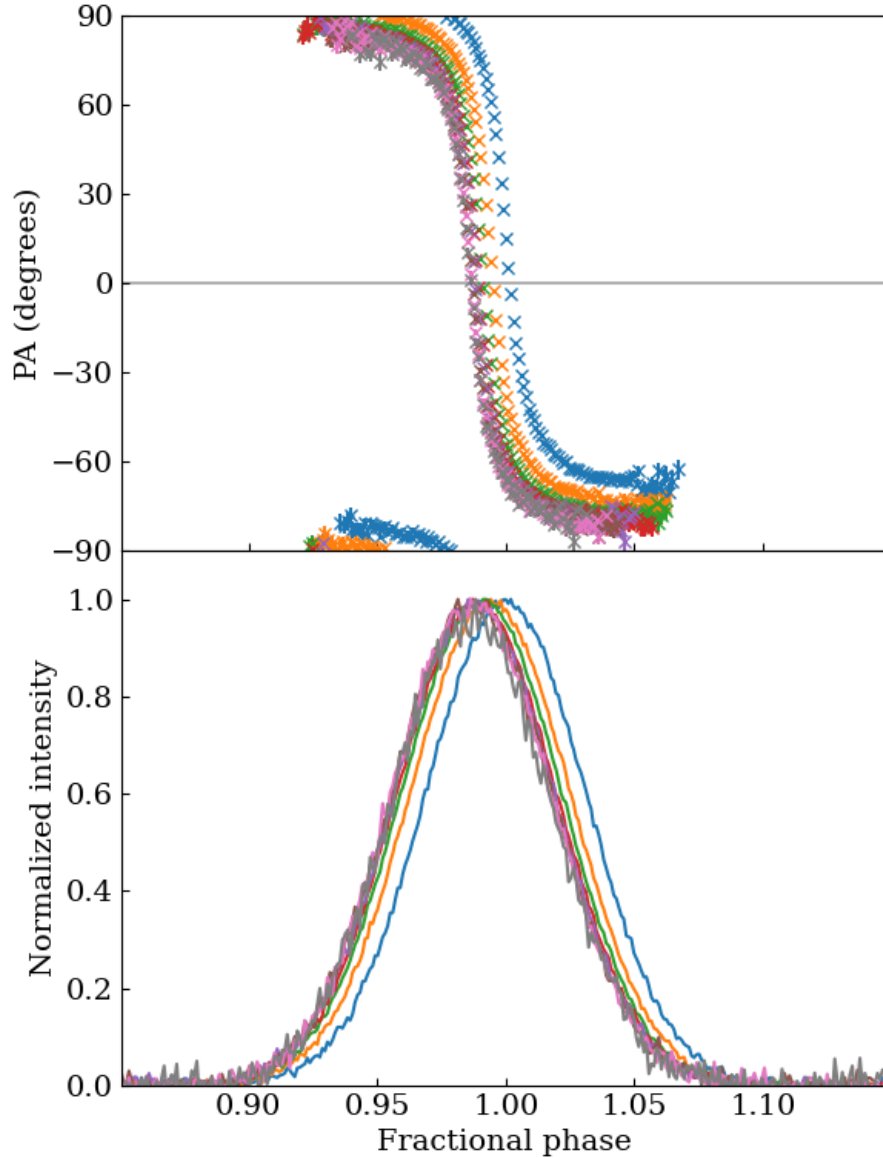


Figure 4.2: Simulated pulse profile, now broad-band.

As for Fig. 4.1, but now showing 8 pulse profiles simulated to be comparable to the 8 frequency channels of the Parkes UWL data. The simulated profiles are dispersed and Faraday rotated, shifting the profiles relative to each other in phase. A DM of $+1 \text{ cm}^{-3} \text{ pc}$ and a RM of $+2 \text{ rad m}^{-2}$ are required to align the profiles. Each of the 8 profiles has a different colour to represent its central frequency, with the same colour scheme repeated for all 8-channel simulations and data in subsequent figures. The colour order from lowest to highest frequency is as follows: blue, orange, green, red, purple, brown, pink, grey. Figure made with matplotlib.

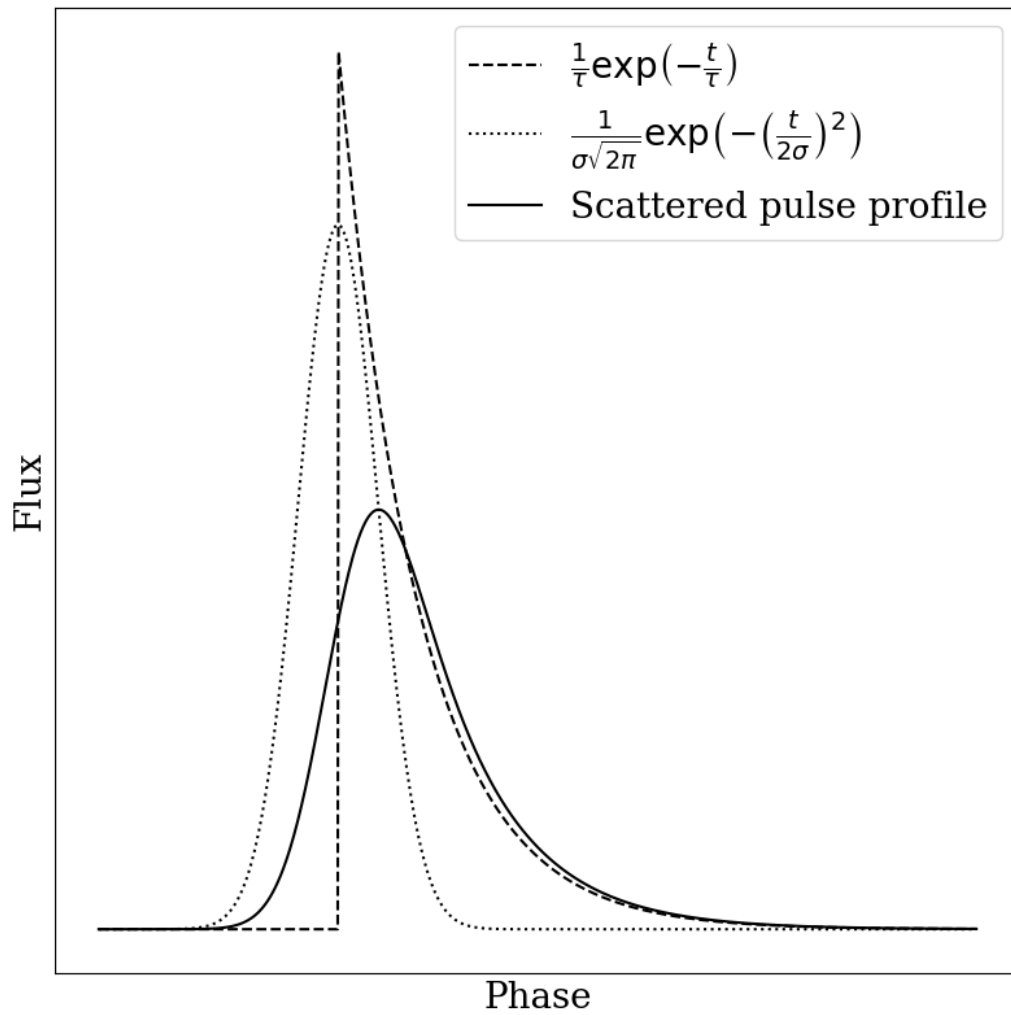


Figure 4.3: Figure describing scattering model.

This figure shows how a Gaussian pulse profile shape (dotted line), when convolved with an isotropic exponential scattering transfer function (dashed line) is distorted to produce a scattered pulse profile (solid line). For this figure, the standard deviation of the Gaussian, σ , and the scattering timescale, τ , are in the ratio $\sigma/\tau = 0.5$. Figure made with matplotlib.

4.1.1 Assumptions of a pulsar emission model

When simulating how both the distribution of single pulse positions and the integrated profile shape evolves with frequency, I required a comprehensive description of the pulsar beam. The key assumptions I made are as follows:

1. The pulsar magnetosphere has a **dipolar** magnetic field.
2. Emission is produced **tangential** to open magnetic field lines.
3. Each frequency of emission is produced at a **single height**.
4. Lower frequencies originate from higher heights: this is **radius-to-frequency mapping** (RFM).
5. The emission at a given height is generated by some **region of active field lines**.
6. The **same** active field lines at different heights are responsible for a subpulse observed across a broad band.

The reasoning behind these assumptions is as follows. Assumptions 1 and 2 are commonly assumed in pulsar beam theory (see for example Gangadhara & Gupta (2001)), particularly when considering the coherent curvature radiation model of pulsar emission. However, it should also be remembered that this is not the only theory of the generation of pulsar radio emission; see for example the review of coherent radio emission mechanisms by Melrose (2017) and the proposal of a rotation-driven pulsar radio emission mechanism by Melrose et al. (2021). The third assumption, of a one-to-one relationship between frequency and emission height, is unlikely to be entirely accurate. However, it is the simplest implementation of radius-to-frequency mapping (assumption 4), which is itself an empirical theory inferred from observations. In the absence of information about the real frequency-height behaviour, it is logical to begin from the simplest assumption and test to what extent it can replicate the observations.

Laying out the first four assumptions and adding in our knowledge that the variety of pulse profile morphologies observed means that the beam is unlikely to be uniformly filled, it is logical to infer assumption 5. From this, we are left with the question of how to relate single pulse emission at different frequencies. Strong correlation has been shown between single pulse profile shapes at different frequencies, as demonstrated by Karastergiou et al. (2001) for PSR B0329+54 at 1.41 and 2.69 GHz. Similarly,

the single pulse shapes appear broad-band for my observations of J1136+1551. It is therefore reasonable to assume a relationship between the emission at different frequencies. Since the frequency-height relationship is constrained, the relationship between subpulses at different frequencies must be due to their being produced by the same field lines.

Building from these assumptions, I converged upon simulations of the 3D emission region. These are a combination of the 2D beam cross-section at a single emission height and the frequency-height relationship, where the two are linked by the dipolar curvature of the field lines. In Chapter 1, Fig. 1.11 shows a visualization of such emission regions in the hollow cone and fan beam cases respectively. The difference between these two models, in simulation terms, is the 2D shape of a slice across the beam at a single frequency and height. I will describe here how beam cross-section affects observed profile evolution in the context of the hollow cone model, and contrast the fan beam model in Section 4.1.4, after explaining the mathematical set-up of the simulation.

Fig. 4.4 shows a visualization of the hollow cone pulsar beam from a bird's eye view, so that the magnetic axis comes out of the page. We observe a cut through the emission region along the line of sight, marked on the figure with the curved solid black line. For the hollow cone model, it can be seen that the assumption of RFM leads to lower frequency emission being produced by a ring of larger diameter. This means the two places where the line of sight cuts the ring are positioned further apart at lower frequencies, leading to the wider pulse profiles observed. However, as the visualization in Fig. 4.4 demonstrates, the emission produced by a given active field line can only be observed at one point along the line of sight, where that field line crosses the line of sight. At other frequencies, we observe emission at different points along the line of sight, which must correspond to being produced by different field lines (radial dashed lines). In Fig. 4.4 the points where we observe a subpulse peak along the line of sight are marked with dots. The points where emission is produced by the same field line at the three indicative frequencies are marked with crosses. Only at one frequency do the cross and dot overlap.

Does this contradict assumption 6, that an observed subpulse is produced by the same active field lines across a broad band? No, because the same region of field lines is producing the emission. It is simply that the field line that generates the emission that corresponds to the peak of the subpulse at one frequency is not the same field line that generates the subpulse peak at another frequency. These two field lines are, however, related by the beam shape properties: in the hollow cone case, they are part




Figure 4.4: Illustration of beam cross-section: hollow cone model.

Bird's-eye-view of the pulsar beam looking down the magnetic axis (central cross) so that the magnetic field lines extend radially (dotted lines). The active emission region for the hollow cone model is marked at three frequencies as three concentric rings, where the lowest frequency ring is darkest grey and has the largest radius. The line of sight (solid line) cuts each ring in two places, marked by dots. The points where the same field line intersects the three rings are marked with crosses. Figure made using Ipe and Inkscape.

of a group of field lines whose footprints on the pulsar surface trace out a ring. This means that, given the phase position of a subpulse at one observing frequency, the simulation of the pulsar beam can be used to infer at what phase the same subpulse will appear at every other observing frequency. It is therefore important to consider the 2D beam cross-section when relating observed subpulse peak phases at different frequencies.

4.1.2 The mathematics of the simulations

The mathematical relationship between the observed pulse phase and the corresponding emission point within the pulsar magnetosphere relies on the following equations, the derivations of which are given in Section 4.1.3. We can describe an emitting dipolar field line in terms of spherical polar coordinates relative to the magnetic axis. This coordinate system must be distinguished from the coordinate system where angles are described with respect to the rotation axis. The rotational phase ϕ was defined in Chapter 2 and is by definition measured with respect to the rotation axis. I therefore describe the magnetic axis coordinates as (r, θ', ϕ') , so that the azimuthal angle ϕ' is distinguishable from ϕ . Since the two coordinate systems have the same origin (the centre of the neutron star) there is no need to distinguish r from r' . The dipolar field line equation then becomes

$$r = K \sin^2 \theta' \quad (4.2)$$

where K is a constant specific to a given field line (see, for example, Gangadhara & Gupta, 2001). This constant is defined for the i^{th} field line by the footprint parameter ratio s_L^i as described by Gangadhara & Gupta (2001), which is the distance from the footprint of the magnetic axis to the footprint of the field line across the surface of the pulsar, divided by the equivalent distance to the last open field line. The relationship is given by

$$K^i = \frac{r_s}{\sin^2 \left(s_L^i \arcsin \sqrt{r_s/R_{\text{LC}}} \right)}, \quad (4.3)$$

where r_s is the radius of the neutron star (conventionally 10 km; this can be estimated roughly by balancing gravitational attraction with the degeneracy pressure of a Fermi gas of non-relativistic neutrons) and R_{LC} is the radius of the light cylinder.

If the emission generated at some point (r, θ', ϕ') is directed along the line of sight, there will be some corresponding rotational phase ϕ at which it is observable. We can relate observing phase ϕ to polar angle θ' via the pulsar geometry (α and β as defined

in the introduction) and the beam half-opening angle ρ using the following equations derived by Gangadhara & Gupta (2001) and Gangadhara (2004) respectively:

$$\cos \rho = \cos \alpha \cos (\alpha + \beta) + \sin \alpha \sin (\alpha + \beta) \cos (\phi), \quad (4.4)$$

and

$$\cos(2\theta') = \frac{1}{3} \left(\cos \rho \sqrt{8 + \cos^2 \rho} - \sin^2 \rho \right), \quad -\pi \leq \rho \leq \pi. \quad (4.5)$$

4.1.3 Derivation of the equations

The derivation of equation 4.3 is as follows. Following from equation 4.2, the footprint parameters associated with the emitting field line (s^i) and the last open field line (s^L) are related to the height of the neutron star surface r_s as

$$r_s = K^i \sin^2 (s^i/r_s) = R_{\text{LC}} \sin^2 (s^L/r_s). \quad (4.6)$$

Rearranging these equations to obtain the footprint parameter ratio s_L^i gives

$$s_L^i = \frac{s^i}{s^L} = \frac{\arcsin \left(\sqrt{r_s/K^i} \right)}{\arcsin \left(\sqrt{r_s/R_{\text{LC}}} \right)}. \quad (4.7)$$

The field line constant K^i of the emitting field line is then obtained by further rearrangement:

$$K^i = \frac{r_s}{\sin^2 \left(s_L^i \arcsin \sqrt{r_s/R_{\text{LC}}} \right)}. \quad (4.8)$$

Equation 4.4 comes directly from the spherical law of cosines. To see this, it is necessary to visualize the relevant vectors and angles in the emission region. Fig. 4.5 sets up the pulsar in a Cartesian basis so that the rotational axis $\hat{\Omega}$ points along the z axis and the unit vector pointing along the line of sight to the observer, labelled \hat{n} , lies in the x-z plane. As the pulsar rotates, so the magnetic axis \hat{m} also rotates about $\hat{\Omega}$. The rotational pulse phase ϕ is the azimuthal angle about the rotational axis, that is, the angle between the x axis and the projection of \hat{m} into the x-y plane. The angles that \hat{m} and \hat{n} make with the rotational axis $\hat{\Omega}$ are defined by the pulsar geometry as α and $\zeta = \alpha + \beta$ respectively. Now, since I assume that pulsar emission is produced tangent to the magnetic field line, the beam half-opening angle ρ must be the angle between vectors \hat{m} and \hat{n} , which is obtained immediately from the spherical law of cosines.

Next, Fig. 4.6 shows the relationship between ρ and θ' , where I have now changed basis so that the z axis is aligned with the magnetic axis \hat{m} . In the case that a field




Figure 4.5: Pulsar emission vectors: rotation, magnetic and line of sight. Geometrical construction of how the major vectors of the pulsar interact and what angles relate them. Magnetic axis \hat{m} rotates about rotational axis $\hat{\Omega}$ and \hat{n} is the vector pointing towards the line of sight. The angles ρ , γ , ϕ , α and $\zeta = \alpha + \beta$ are defined in relation to these vectors. A minor subtlety not shown in this diagram is that the three vectors do not in fact intersect at a common origin. $\hat{\Omega}$ and \hat{m} intersect at the centre of the neutron star and \hat{m} and \hat{n} intersect at the pulsar surface. The $\hat{\Omega}$ vector shown is therefore parallel to the true rotational axis, but shifted. For the purposes of defining the angles there is no mathematical difference, and this parallel shift makes the diagram more legible. Figure made using Inkscape.




Figure 4.6: Parametrization of dipolar field line.

The field lines are defined in polar coordinates (r, θ') with respect to magnetic axis \underline{m} . Beam opening angle ρ is given in relation to the tangent to the field line at the point of emission. Figure made using Inkscape.

line produces emission that is observable along the line of sight, the tangent to that field line will point along the line of sight vector $\hat{\mathbf{n}}$ at the emission point. $\hat{\mathbf{n}}$ can therefore be described as the gradient of the field line,

$$\hat{\mathbf{n}} = \left(\frac{\partial \mathbf{r}}{\partial \theta'} \right) / \left| \frac{\partial \mathbf{r}}{\partial \theta'} \right| \quad (4.9)$$

where $\mathbf{r} = K \sin^2 \theta' (\sin \theta' \cos \phi', \sin \theta' \sin \phi', \cos \theta')$ in spherical polar coordinates. The beam half-opening angle is given by the dot product of the vectors, $\cos \rho = \hat{\mathbf{n}} \cdot \hat{\mathbf{m}}$, where $\hat{\mathbf{m}} = (0, 0, 1)$. Calculating this and rearranging results in equation 4.5.

4.1.4 Pulsar beam cross-sections

In Section 4.1.1 I described how RFM results in the observational line of sight cutting different regions of the pulsar beam at different frequencies, and how the shape of the beam cross-section affects the relationship between the single pulse phase positions at different frequencies. Fig. 4.7 shows the effect of changing the beam cross-section from the hollow cone model to the fan beam model.

As the line of sight cuts across an emitting region the observed intensity increases to a maximum and then decreases again to zero, resulting in a pulse of finite width. For the hollow cone model, I simulated the peak of this pulse as corresponding to the middle radius of the emitting ring at a given frequency. Therefore the beam half-opening angle for the emission peak along the selected field line (marked with crosses) at a given frequency is the same as the beam half-opening angle for the peak of the observed single pulse at that frequency (marked with dots).

With the fan beam, the ring symmetry is lost, and so this relationship is more complicated. For the purposes of the simulation I assumed that the cross-section of a fan beam at a single frequency is completely circular. If the line of sight cuts through the centre of the circle, the observed peak of emission will correspond to the centre of the circle, as is the case for the lowest frequency (darkest colour) on the figure. If instead the line of sight cuts the circular emitting region across some other path, the observed peak of emission will be at that point which lies closest to the centre of the circle.

The assumption of circular symmetry of the fan beam cross-section is unlikely to be replicated in real pulsar emission, just as the initial assumption of one-to-one frequency-to-height mapping is unlikely to be accurate. However, in the absence of other information, this is the simplest assumption to create two models of beam structure that are comparable.

4.1.5 Mathematical derivation of the fan beam cross-section emission peak

Building from equation 4.4 for beam half-opening angle ρ , the spherical law of cosines can again be used to find the angle of intersection of the fiducial plane and the plane containing the magnetic axis and the vector pointing towards the line of sight, that is the planes containing $\hat{\Omega}$ and \hat{m} , and \hat{m} and \hat{n} respectively. This angle is γ , as shown in Fig. 4.5. It is given as a function of the geometry and the beam half-opening angle ρ at each rotational phase as:

$$\cos(\gamma) = \frac{\cos(\alpha + \beta) - \cos(\alpha) \cos(\rho)}{\sin(\alpha) \sin(\rho)}. \quad (4.10)$$

Fig. 4.8 shows the line of sight in the x-y plane projected perpendicular to the magnetic axis. It can be described in terms of ρ and γ as follows:

$$x_L = \rho \sin(\gamma), y_L = \rho \cos(\gamma). \quad (4.11)$$




Figure 4.7: Illustrations of beam cross-sections: hollow cone and fan beam models. As for Fig. 4.4 but now showing both the beam cross-sections of the hollow cone (left) and fan beam models (right). Below: subpulses indicate how the changing beam cross-section with frequency affects the separations of the observed subpulses at different frequencies. Figure made using Ipe and Inkscape.




Figure 4.8: Parametrization of beam cross-section and line of sight. The beam cross-section can be described entirely in terms of angles ρ and γ as defined in Fig. 4.5 and shown here from the bird's-eye-view beam cross-section perspective, with the line of sight (solid black line) cutting across the beam. Figure made using Inkscape.

Equations describing the tangent to the line of sight at point (x_L, y_L) are obtained through differentiation and the chain rule:

$$\begin{aligned} \frac{\partial y_L}{\partial \rho} &= \frac{\cos(\alpha + \beta)}{\sin(\alpha) \sin(\rho)} \left(1 - \frac{\rho}{\tan(\rho)} \right) \\ &+ \frac{\cos(\alpha)}{\sin(\alpha)} \left(\rho \left(1 + \frac{1}{\tan^2(\rho)} \right) - \frac{1}{\tan(\rho)} \right), \end{aligned} \quad (4.12)$$

$$\begin{aligned} \frac{\partial x_L}{\partial \rho} &= \sqrt{1 - \left(\frac{\cos(\alpha + \beta) - \cos(\alpha) \cos(x_L)}{\sin(\alpha) \sin(\rho)} \right)^2} \\ &+ \frac{\rho F}{\sqrt{1 - \left(\frac{\cos(\alpha + \beta) - \cos(\alpha) \cos(\rho)}{\sin(\alpha) \sin(\rho)} \right)^2}}, \end{aligned} \quad (4.13)$$

$$\frac{\partial y_L}{\partial x_L} = \frac{\partial y_L}{\partial \rho} \bigg/ \frac{\partial x_L}{\partial \rho}, \quad (4.14)$$

where F is given by:

$$\begin{aligned} F &= \frac{(\cos(\alpha + \beta) - \cos(\alpha) \cos(\rho))^2}{\sin^2(\alpha) \sin^2(\rho) \tan(\rho)} \\ &- \frac{\cos(\alpha)(\cos(\alpha + \beta) - \cos(\alpha) \cos(\rho))}{\sin^2(\alpha) \sin(\rho)}. \end{aligned} \quad (4.15)$$

Given an observation at phase ϕ^i and a reference observation at ϕ^0 the positions on the line of sight, (x_L^i, y_L^i) and (x_L^0, y_L^0) , are calculated as above. I assume that the peak of the emission region for the reference observation, $(x_{\text{peak}}^0, y_{\text{peak}}^0)$, lies at (x_L^0, y_L^0) . I further assume that the peak of the emission region for our observation at phase ϕ^i lies further along the same emitting field line. Logically then, for a circular emission region, the observation at the line of sight, (x_L^i, y_L^i) , is connected to this emission region peak by a perpendicular bisector to the line of sight at this point. The emission region peak $(x_{\text{peak}}^i, y_{\text{peak}}^i)$ is located at the intersection of this perpendicular bisector and the field line. The relationship between these points is visualized on Fig. 4.9 and described mathematically by

$$x_{\text{peak}}^i = \frac{y_L^0 - g_L x_L^0}{g_f - g_L}, \quad (4.16)$$

$$y_{\text{peak}}^i = g_f x_{\text{peak}}^i, \quad (4.17)$$

where g_f is the gradient of the field line and g_L is the gradient of the perpendicular bisector to the line of sight, with their values given as follows:

$$g_f = \frac{y_L^0}{x_L^0}, \quad (4.18)$$




Figure 4.9: Relationship between reference field line and line of sight peak for the fan beam model.

Set up as in Fig. 4.8, this figure shows how two points along the line of sight are related by a third point off the line of sight but lying on the reference field line, in the case of the fan beam model as simulated in this work. This is a visual representation of the mathematical relationship between these three points (all marked with red crosses) as described in the text. Figure made using Inkscape.

$$g_L = -1 \left/ \frac{\partial y_L^i}{\partial x_L^i} \right. . \quad (4.19)$$

Converting back from the projected x-y plane to the beam half-opening angle of the emission region peak:

$$\rho_{\text{peak}} = \sqrt{x_{\text{peak}}^2 + y_{\text{peak}}^2}, \quad (4.20)$$

which can then be interchanged with θ' and ϕ using equations 4.4 and 4.5.

4.1.6 Orthogonal polarization modes

The above description of beam structure models does not take into account the fact that polarimetric observations suggest that pulsar emission is composed of two orthogonal modes, which are theorized to propagate along different paths due to birefringence in the magnetosphere (definition 2j). As described in Chapter 1, this may either cause, or contribute to, pulse widening with decreasing frequency. Conclusions are yet to be conclusively reached about the extent to which orthogonal modes combine coherently or incoherently to result in the observed polarized pulses. However, since I used the described beam structure models to investigate the statistical behaviour of single pulses collectively, I could also introduce orthogonal mode divergence as a statistical behaviour. I did so by assigning some fraction of the simulated single

pulses to be X mode dominated, and the remainder to be O mode dominated. The O mode pulses follow RFM, whereas X mode emission does not evolve with frequency. The X mode subpulse phases are therefore fixed at all frequencies, and the overall distribution of subpulse phases is made up of a combination of the behaviours of X and O mode pulses.

4.2 Quantifying single pulse frequency evolution

Having developed simulations as described in Section 4.1, I analysed observations and, where relevant, compared them to these simulations, to investigate various features of pulsar emission. I sought to use my observation of the broadband single pulses of PSR J1136+1551 to increase the information the observer has when considering the magnetospheric origins of what is recorded by the radio telescope. Whereas an integrated profile gives a single, average cut across the pulsar beam, using single pulses gives a population of instances of emission in the pulsar beam. I could treat these statistically to understand how the distribution of these single pulses is generated by a three-dimensional emitting region. The key question I addressed is the following: what three-dimensional emitting beam structure is required to produce both the positional distributions of single pulses of PSR J1136+1551, and the frequency evolution of these single pulse shapes? This addresses open questions 1a–d. In order to answer this question, I had to reduce the complexity of single pulse phenomena down to quantifiable features, namely the positions of and separations between the subpulses that make up a single pulse. I could then compare these with those predicted by the simulated pulsar beam described above.

In Chapter 1, Fig. 1.4 shows four example single pulses from the observation, split into seven frequency channels spanning the band, as described in the data processing details in Sections 3.3.1 and 3.3.2. They are plotted as a function of phase ϕ , which is defined such that $\phi = 0^\circ$ corresponds to the fiducial plane of the pulsar. All four of these single pulses have considerable structural complexity. To simplify the dataset to the key attributes required to answer the question, I identified and obtained two metrics that encompassed the required information: the separation between the two components of a double-peaked single pulse in the 470 MHz channel, and a ‘separation spectral index’ describing how this separation varies with frequency.

The concept behind requiring these two metrics is the following: assuming pulse width W evolves with frequency ν , a simple way to capture this evolution is through a power law relationship $W \propto \nu^b$ where the larger the absolute value of b , the more

strongly the width evolves across the band. Since it is observed that pulse width tends to increase with increasing rapidity as frequency decreases, b is expected to be negative. I elected to use PSR J1136+1551 for this study, due to the history of similar research focused on this pulsar, detailed at the start of Chapter 5 and because its integrated pulse profile is double peaked and reasonably symmetrical. The separation between the two peaks is an unambiguous measure of pulse shape, the evolution of which can be tracked across frequency. Having discarded those pulses affected by RFI, including the 50 Hz contaminating signal described earlier, I then selected only those pulses which are double-peaked, leaving 885 pulses for analysis from an original 4759 pulses recorded in the 94 minute observation. For these, I then identified the positions of the two subpulse peaks for each single pulse and took this information forward for further analysis.

The method through which I obtained these subpulse positions is as follows. Having obtained the 7 single pulse profiles across the band, as shown in the examples in Fig. 1.4, I first summed them all together to obtain a single normalized integrated profile and used this to determine the on-pulse phase region. This I did by smoothing the profile with a median filter and then calculating its S/N, where I measured the noise as $\sigma = 1.4826 \times \text{MAD}$, where MAD is the median absolute deviation. I defined the on-pulse region by eye as being bounded by the first and last bins where $\text{S/N} > 4$. Turning now to the single pulses, I first scaled the intensity units for each pulse by dividing by the peak intensity. I then smoothed each scaled pulse profile by correlating its on-pulse region with a normalized Gaussian located centrally within the on-pulse region, with a standard deviation of $\sigma = 1.53^\circ$ equal to that of a Gaussian fit to the widest peak of the integrated profile. I then identified the positions of subpulse peaks as maxima in the smoothed pulse profile. I defined a pulse as double-peaked if its cross-correlation had two maxima with magnitudes greater than 0.1 (in absolute units of the cross-correlation of the scaled pulse profile with the normalized Gaussian). Since pulse profile shapes contain complicated microstructure and subpulses are rarely symmetrical, I sought to capture this in the uncertainty on the subpulse position. I therefore defined the uncertainty as the distance between my calculated position and the point with the greatest flux, treating each half of the profile separately. Having calculated two subpulse positions per double-peaked single pulse, I could then calculate the subpulse separation Σ directly, and the separation spectral index describing how this varies with frequency, η , by fitting the power law $\Sigma \propto \nu^\eta$ for each single pulse.




Figure 4.10: Illustration of how DM choice affects profile relationships across frequency.

Left: illustration indicating how two pulse components move further apart with decreasing frequency (going up in the diagram), aligned such that this evolution is symmetric about the midpoint of the two components (marked with the dashed line). Right: the same pulse components, now aligned as though with a different choice of DM, so that the mid-point is shifted further to the left with decreasing frequency. Figure made using Inkscape.

On plotting histograms of the subpulse positions for the double-peaked single pulses, I found that the mid-point between the two distributions was not positioned at the same phase at all frequencies, but shifted to earlier phases at lower frequencies. By applying a correction of $0.034 \text{ cm}^{-3} \text{ pc}$ to the DM, making the total $4.892 \text{ cm}^{-3} \text{ pc}$, I aligned the subpulse distributions across frequency so that they were symmetric about the $\phi = 0^\circ$ point. Fig. 4.10 shows an illustration of this alignment process, which resulted in symmetrical widening with decreasing frequency. It is usual to choose a DM that results in symmetrical evolution in this way, and an explanation may be that the emission beam of PSR J1136+1551 has two regions positioned equidistant from the magnetic axis. However, it is also possible that my choice of DM here is incorrect, and the beam is in some way asymmetric. For the purposes of the investigation done here, which focuses on intrinsic subpulse separation rather than exact subpulse position, the results for the data are by definition independent of the choice of DM. However, my choice of DM, and hence the position of the fiducial plane, is relevant for making comparisons with my simulations of how emission may be produced within the pulsar magnetosphere. I subsequently explored this concept of DM measurement, and resultant inferences about pulsar beam structure, in work

using broad-band polarimetric pulsar data, the methodology of which I address next.

4.3 Correcting for propagation effects in broad-band polarimetric data

The sources of variability in pulsar polarization profiles can be divided into two categories: those intrinsic to the pulsar’s emission processes, and those caused by the propagation of the radio beam through the ISM. Three key factors in the latter case are dispersion, Faraday rotation and interstellar scattering. All of these are strongly frequency dependent, which is why, when studying broadband data, a robust and scientifically grounded technique for accounting for them correctly is required. As described above in the context of PSR J1136+1551, and in Section 3.4.1, the exact choice of DM to dedisperse an intensity profile will always have a level of subjectivity. This is because, since the shape of the intensity profile evolves with frequency, it is impossible to align differently-shaped profiles at different frequencies unambiguously. However, there is one theoretical feature of polarized pulsar emission that is expected to be consistent regardless of observing frequency, and that is the shape of the PA that results from the rotating vector model. As this is determined by the pulsar’s geometry alone, it is independent of frequency.

Of course, although the pulsar’s geometry is constant, the PA profile of a given pulsar is unlikely ever to be determined by the rotating vector model alone. If there is circular polarization present in the pulse profile, then the polarization vector may track a path across phase such that the linear component will deviate strongly from the track expected from the rotating vector model, and this may also evolve with frequency (Dyks, 2020). The presence and interaction of two modes in pulsar emission can lead to orthogonal jumps in the PA, and many other complex shapes are observed in PA profiles that are not well explained. However, in general, many pulsars are observed to have PA profile shapes that do not evolve with frequency (see for example Karastergiou & Johnston, 2006, who showed that the residual for 17 pulsar observations at 1.4 and 3.1 GHz respectively is distributed tightly around 0). With this in mind, it is logical to use the shape of the PA profile as the basis upon which alignment of broadband polarimetric profiles can be performed.

As I described in Section 3.4.1, the choice of DM used to align profiles will affect the value of RM converged upon. This is because the DM and RM both shift PA profiles relative to each other, but they do so along different axes. Whereas the DM moves PA profiles in phase, the RM shifts the PA profiles by angle. It is therefore important

to identify both of these values simultaneously. I developed a method to do this for those pulsars where the PA does not evolve significantly with frequency. The precept of the method is to search in DM–RM space to find the pair of these values that will successfully align the PA profiles across a broad band for a given pulsar. Unlike previous work based on aligning PA profiles (Karastergiou & Johnston, 2006; Mitra & Li, 2004), this work limits profile shifting to be only according to possible values of DM and RM, and applies it to profiles across a continuous broad frequency band of 700–4000 MHz. By contrast, Karastergiou & Johnston (2006) had only two discrete profiles at 1.4 and 3.1 GHz, for which they aligned the PAs free from constraints and inferred the resultant RM. Mitra & Li (2004) did not consider DM and RM in their alignment process, but sought instead to use their alignment of multi-frequency profiles to investigate lags between intensity and PA profiles. It is important to identify misalignment with the physical properties causing it, here assumed to be dispersion and Faraday rotation, and the continuous broadband polarimetric data from the Parkes UWL gives the resolution in frequency space to do so.

In order to investigate alignment of polarimetric UWL data I first used the data aggregated across frequency to produce 8 high S/N profiles across the observing band. These profiles are generated with the RM and DM values from the pulsar catalogue PSRCAT (Manchester et al., 2005) and I sought to identify the small corrections to these values that would align the PA profiles across frequency. For this work I studied two pulsars, PSRs J1056–6258 and J1359–6038, chosen because previous alignment efforts by Karastergiou & Johnston (2006) revealed anomalies in comparative alignment of the intensity profiles. I calculated a metric to identify how well a given pair of RM and DM align the PA profiles, and then used a Markov-Chain Monte Carlo (EMCEE, Foreman-Mackey et al., 2013) to converge upon the best solution.

My metric for alignment of PA profiles is the normalized χ^2 , given by

$$\chi^2 = \frac{1}{N_{chan}N_{bin}} \sum_i^{N_{chan}} \sum_j^{N_{bin}} \frac{dPA_{ij}^2}{2\sigma_{ij}^2}, \quad (4.21)$$

where dPA is the angular distance of the PA profile from the profile at the lowest frequency. For the i th frequency channel and j th phase bin (where $i = 0$ is the lowest frequency) this is given by

$$dPA_{ij} = \left(\left(|PA_{ij} - PA_{0j}| + \frac{\pi}{2} \right) \% \pi - \frac{\pi}{2} \right) \times \text{sign}(PA_{ij} - PA_{0j}). \quad (4.22)$$

Where $\%$ means modulo and sign encompasses whether the difference between the two PA values is positive or negative. The format of the equation is set up to account

for cyclic boundary conditions, and the resultant uncertainty on this value is obtained by adding the relevant errors in quadrature, as

$$\sigma_{ij} = \sqrt{\sigma_{\text{PA}_{ij}}^2 + \sigma_{\text{PA}_{0j}}^2}, \quad (4.23)$$

where σ_{PA} is calculated as in Everett & Weisberg (2002).

I chose to define $d\text{PA}$ with respect to the lowest frequency because the spectral index of pulsar intensity means that, in the Parkes observing band, S/N tends to be highest at the lowest frequency and so the PA profile is best resolved there. The functional form of χ^2 results from summing over phase and frequency and dividing by the total number of channels and phase bins. This value will be minimized when the PA is aligned, so I sought to maximise the likelihood function $\exp(-\chi^2)$.

A standard tool in model fitting is to use some form of least squares calculation to minimize the residuals difference between the data and the best fit model for a given set of parameters. The limitation of such a method is that it does not capture the extent to which the model represents the data. More information can be obtained by sampling from a probability distribution. This is particularly important in cases where it is known that the model in question does not capture every facet of the features apparent in the data, including additional emission complexities not described by the model physics, along with the errors introduced by the observing process and the environment. For this, Bayesian statistics, and the associated concepts of priors, likelihood and posterior, are important.

For my search for the correct pair of ΔRM and ΔDM , I have selected a χ^2 metric that, in theory, should be minimized by the best alignment of PAs across frequency. However, it is known both that the PA is not described purely by the rotating vector model, and that other effects may be captured in the observations that mean that the PA does not appear identical at every frequency. Furthermore, in the case where the PAs are very strongly misaligned, there will only be a requirement for a small number of data points to overlap, which may well result in a smaller χ^2 than for the more accurate full alignment of the entire profiles. With all this in mind, two things are clear. First, it is unlikely that a single minimization of χ^2 will always be unambiguously accurate, and secondly this extra information should be captured in the priors in order to converge upon the correct answer.

There is not a simple physical formula for quantifying such information into numerical priors. I therefore choose to sample a posterior distribution with uninformative priors in order to understand its near-unconstrained distribution. This meant that I used priors that simply constrain the DM such that the PA profiles must overlap by

at least one data-point, because otherwise it would not be possible to calculate χ^2 . I tested how this distribution behaved for various simulated conditions in order to gain the information required to create appropriate numerical priors for the data. A Markov-chain Monte Carlo method (MCMC), provided it is run for sufficiently long that its samples of the posterior distribution are independent, is a suitable method for sampling a complex and unknown distribution in a finite number of steps (see Chapter 29, MacKay, 2003), so I applied this using EMCEE. Once I had obtained the best fit pair of parameters ΔRM and ΔDM for a given Parkes pulsar observation, I applied these to the data product with the full 3328 channels, in order to correct properly for dispersion and Faraday rotation. I then aggregated the data by a factor of 64 to produce 52 aligned, high S/N profiles spanning the full observing band, which reveal the intrinsic polarization structure of the pulsar radio emission.

4.4 Visualizing broadband data

In order to understand what is happening in the intrinsic pulsar emission, it is important to present data in a way that reveals the evolution of, or consistency of, their polarimetric properties with frequency. The traditional method of plotting pulsar data: as a single profile across phase showing total intensity, linear and circular polarization components and the PA profile, has no capacity for direct comparison of emission across frequency (see for example Fig. 1.6). Since I required a visualization that shows how the profile varies with pulse phase but also reveals evolution across frequency, I sacrificed clarity of profile morphology to obtain a better understanding of these two features. I elected to plot the data for a given pulsar as three waterfall plots, each showing a property of the data as a function of both phase on the x-axis and frequency on the y-axis, with hue indicating the value of the property in question. Fig. 4.11 shows the three plots in question for PSR J1056–6258, along with the standard single profile plot. The first indicates total intensity, normalized by dividing by the summed flux of the on-pulse region. This is to emphasize pulse shape over total intensity evolution, which tends to be steep. The second displays the PA using a cyclic colour scheme modulo 180° . Smooth evolution with phase is expected in the case of pure rotating vector model, and deviations from this would be visible as sharp changes in colour with phase. The third plot shows circular polarization as a fraction of total intensity, or V/I , in a two-toned red and blue colour scheme to emphasize points where the circular polarization changes hands, so that V goes from positive to negative, or vice versa. For the pulsar plotted here, V/I is small and does not vary

strongly. However, pulsars with more complex polarization properties often show considerable complexity in V/I . Other polarimetric variables can be plotted similarly: an example is fractional linear polarization L/I , but I found that variability and features of interest were best captured using the three shown here.




Figure 4.11: PSR J1056–6258: waterfall plots of intensity and polarization. Three waterfall plots showing intensity and polarization information for PSR J1056–6258, as observed with the Parkes UWL receiver, across pulse phase (x axis) and frequency (y axis). Left: total intensity, normalized by integrated energy in the profile so that pulse shape evolution is favoured over flux spectral evolution in the visualization. Middle: PA profile, showing only those points where the S/N of the linear polarization $L_{\text{SNR}} > 4$. Right: Circular polarization Stokes V as a fraction of total intensity I . This is shown only where $L_{\text{SNR}} > 4$ and 20 bins with $|V/I| > 0.2$ (and low total intensity) have been removed in order to increase the colour contrast at the centre of the pulse profile. Note that this figure has been made using the older data products presented in Chapter 6, which means, as described in Section 3.3.3, that the data have not been *pcm*-corrected. This means that a very small excess of Stokes V is seen at the highest frequencies. Figure made with matplotlib.

4.5 Modelling scattered profiles: further application of MCMC

My comments on the challenges facing modelling of pulsar behaviour and the utility of a MCMC (Section 4.3) are applicable in a general sense and not only to my simultaneous measurements of RM and DM. I therefore also used a Markov-chain Monte Carlo for a second piece of model-fitting to pulsar data, this time to investigate the

scattering properties of the ISM through the resultant distortion of pulse profiles. For this work I used pulse profiles from the TPA dataset, as described in Section 3.3.3. As I described in Chapter 3.4.2, forward modelling in the time domain can be used to measure scattering parameters. In this set-up, a scattered pulse profile can be considered to be a convolution of the intrinsic profile with a scattering transfer function. I employed forward modelling to obtain scattering parameters from scattered pulse profiles. In doing this, I had to make certain assumptions about the shape of the intrinsic pulse profile, its evolution with frequency, and the appropriate scattering transfer function to apply. I assumed that the intrinsic profile could be described by a single Gaussian, applying the *train* + *DC* model developed by Geyer et al. (2017). This means the model fits 5 components: the amplitude (A), mean (μ) and standard deviation (σ) of the Gaussian; the scattering timescale (τ); and any DC offset of the profile baseline. The assumption of a Gaussian intrinsic profile is only valid in cases where the pulse profile is either sufficiently similar to a Gaussian in shape, or the profile is dominated by the scattering, in which case the exact shape of the pulse profile is of negligible importance. I therefore selected by eye from the TPA dataset only apparently single-component pulsars, excluding those which could not be reasonably modelled by a single Gaussian convolved with some amount (possibly negligible) of exponential decay. The analysis was therefore applied to 205 single-component pulsars, as described in Section 3.3.3.

The data products for the TPA dataset consist of 8 Stokes I profiles spanning the observing band. I chose to fit my scattering model to each profile independently and imposed flat priors, with the only constraint that the parameters had to be physical, in order to allow the fit to converge freely. This meant I constrained A and τ to be positive, and $0 < \mu < 1024$ and $0 < \sigma < 1024$, where 1024 is the number of bins across the pulse period. I then assessed the modelling success of each pulsar by applying my expectation that scattering evolves smoothly with frequency, something that is not constrained by my independent profile models. Details of how I applied this assessment criterion are given in Chapter 7. The reason behind modelling each profile independently is the similar concept, as applied before, of maximising the freedom available to the modelling process, so as not to impose biased preconceptions of how the data should behave for the cases where we have insufficient information. The relevant example here is the evolution of the intrinsic pulse profile, for which there does not exist a quantifiable physical model. Instead, I allowed each profile to converge and inspected the results afterwards to ensure that they conformed to the broader, non-quantifiable qualities expected.

When applying the MCMC to fit the model to the data, I allowed it to run either until the chains achieved the convergence criterion, or until it reached 20,000 steps. I defined the chains as converged if the chain auto-correlation time estimate was changing by less than 1% and the chain length was more than $100\times$ the auto-correlation time. The necessity of applying a limit of 20,000 steps in order to achieve results for 205×8 pulsar profiles in a manageable period of time, meant that I had to inspect the chains by eye for signs of convergence or non-convergence. Fig. 4.12 shows examples of three sets of chains: one set that meets my convergence criterion (J1840–0559), one set that does not but has nevertheless clearly converged upon a single set of model parameters (J1818–1422), and one set that has in no way approached convergence (J1406–6121). I chose to accept those cases where the chains had visually settled on a single set of parameters and would clearly achieve the convergence criterion if permitted to run for long enough. I rejected all those that either did not converge, or converged upon more than one set of parameters. The clear visual difference between the near-converged and the non-converged chains, as exhibited in the examples in Fig. 4.12, shows that this is a reasonable approach given the computing limitations. My inference about those profiles for which the model does not converge was that the model was in some way incorrect or insufficient for describing the profile. The most likely explanations in such a scenario are either that the profile is too complex to be described by a single Gaussian, or that the pulsar in question is not scattered. I defined a requirement of at least 4 converged profile models to be able to calculate a reliable value of α describing how τ varied with frequency. At least 3 values of τ would be necessary for a power law fit and a minimum of 4 provides a better constraint on the best fit parameters. I therefore discarded all the pulsars with fewer than 4 converged profile models. I calculated the best fit parameters to be the 50th percentiles of the MCMC samples, with the uncertainties taken as the average of the differences between the 16th and 50th, and 50th and 84th percentiles respectively. I made use of GNU Parallel (Tange, 2011) to run the modelling on multiple pulsars simultaneously.

As originally shown by Geyer & Karastergiou (2016), assuming a power-law relationship between scattering timescale and frequency, the appropriate frequency associated with a given measurement of τ is given in terms of the centre frequency, f_c , and bandwidth, δf , of the profile as

$$f_m = 10^{[\log_{10}(f_c+\delta f/2)+\log_{10}(f_c-\delta f/2)]/2}. \quad (4.24)$$




Figure 4.12: MCMC chains from scatter modelling of three example pulsars.
Caption continued on the next page.

Figure 4.12: *Continued caption.*

Three sets of MCMC chains from the scatter model fitting, showing examples of successful convergence (top), visual approach towards convergence (middle) and failure to converge (bottom). There are five model parameters to constrain: the amplitude (A), mean (μ) and standard deviation (σ) of the Gaussian, the scattering timescale (τ) and any DC offset of the profile baseline. Top: PSR J1840–0559 achieves the convergence criterion at a little over 17500 steps. Middle: The chains for PSR J1818–1422 explore parameter space and settle upon the best fit at around 7500 steps. It is clear visually that the convergence criterion would be reached if the MCMC were left for longer, but due to computing limitations the MCMC is halted at 20,000 steps. Bottom: For PSR J1406–6121, the chains spread out across parameter space and fail to converge upon a single set of parameters. Figure made with matplotlib.

Given the values of τ obtained through the MCMC fit, and the associated frequencies f_m , I then calculated α . Initial calculations of α were done with least squares fitting of a power law to τ vs. frequency. However, since it was found that different least squares algorithms tended to produce different answers, I decided instead to use a further MCMC to map out the probability distribution and gain a better understanding of convergence on the solution. I kept only those 142 pulsars which had at least 4 channels for which the τ fit was successful, and then also converged on a successful fit for α . The value of α and its uncertainty I took from the same percentiles of the MCMC samples as described above for τ .

4.6 Summary

In summary, the methodology I applied to analyse broad-band pulsar observations encompassed three approaches. The first was to build simulations of pulsar observables based upon theories about pulsar beam structure and the impact of the ISM upon the radio emission. Secondly, I addressed methods of visualizing broad-band data in order to understand the key features of relevance in individual observations. Finally, analysis of the data involved techniques for quantifying key features of pulsar emission using metrics, such as the spectral index of subpulse separation η , and the use of MCMC models to explore the parameter spaces of model fitting and so understand the factors affecting the modelling results obtained. I tailored these new techniques to the broad-band data from modern instruments, and the results of their application to the data are presented in the following three chapters.

Chapter 5

Understanding the radio beam of PSR J1136+1551 through its single pulses

5.1 Setting up simulations of single pulse evolution

This chapter details the work I completed to investigate the pulsar beam shape for PSR J1136+1551 through a statistical description of the behaviour of its single pulses observed with the GMRT. In Chapter 4 I described the methods I employed to reduce the complexity of single pulse observations down to quantifiable metrics, namely the separation between components for double pulses, and the separation spectral index η encapsulating its evolution with frequency. I also described how, in theory, we might expect the shape of the pulsar beam to affect the observed evolution of pulse profiles, and how I created simulations of the pulsar beam to demonstrate this. By comparing my single pulse observation to these simulations, I could address the open questions about pulsar radio beam structure (Q1a) and single pulse generation (Q1d). In this way I could make progress in explaining how pulsars produce their radio beams, improving our understanding of pulsar physics and our ability to use them as tools, as described in Chapter 1. Details of the observation analysed were given in Sections 3.3.1 and 3.3.2.

Being a bright pulsar with a symmetrical double-peaked profile, PSR J1136+1551 has been studied on multiple occasions from a radius-to-frequency mapping (RFM) perspective. There are three techniques which have been used to calculate emission heights for pulsars based on observable parameters. Cordes (1978), Karuppusamy et al. (2011) and Hassall et al. (2012) investigated timing A/R (definition 2h), and used non-detections to infer upper limits to the emission height ranges in their respec-

tive observing bands. Noutsos et al. (2015) used the delay-radius relation (definition 2i) to relate a misalignment of intensity and PA profiles to associated emission heights (see also similar work for other pulsars by Mitra & Li, 2004). Hassall et al. (2012) also measured heights based on profile widening with frequency, as did Xilouris et al. (1996), Mitra & Rankin (2002), Krzeszowski et al. (2014), Noutsos et al. (2015) and Lu et al. (2016). The results for this method from different authors are all roughly in agreement, and constrain emission heights to be at values only a few percent of the radius of the light cylinder. Converting from frequency widening to emission heights requires a model of the pulsar beam, as described in detail in the previous chapter. All of these authors fixed the measured peaks of the integrated profile to the last open field lines, meaning the footprint parameter of the associated field line is $s_L^{peak} = 1$.

In addition to being important information relevant to the structure of the pulsar beam and hence the origins of the radio emission, measurements of pulsar emission heights can also be used to constrain understanding of the plasma composition of the pulsar magnetosphere (Jones, 2013). Other work on PSR J1136+1551 has revealed the presence of very bright (possibly giant) single pulses (Kramer et al., 2003; Kazantsev & Potapov, 2015) and selective nulling, where nulling does not appear to occur simultaneously at all frequencies of observation (Bhat et al., 2007). PSR J1136+1551 has also been studied in X-rays (e.g. Szary et al., 2017) and recent work by Pétri & Mitra (2020) suggests that the radio and X-ray emission properties can only be reconciled if the magnetic field is an offset dipole.

Looking at the values for μ and their fit uncertainties for the data in Table 5.1, it is clear that there is a net divergence of the subpulse distributions at lower frequencies, with the Gaussians moving further apart. The fit uncertainty on σ for the right subpulse distribution means that it is inconclusive whether this distribution also shows broadening with decreasing frequency. However, frequency-broadening of the left subpulse distribution is also clear, with the Gaussian becoming wider with decreasing frequency.

For the simulations, if we know the relationship between frequency and emission height, and we assume that, for a reference frequency, we know how the phase of an emitted subpulse is related to an associated emitting field line, we can then infer the expected pulse phase of the associated subpulse at all other observing frequencies. For this we need to know how the line of sight cuts across the pulsar beam, which depends on the pulsar geometry. I set $\alpha = 51.3^\circ$ and $\beta = 3.7^\circ$, as derived by Lyne & Manchester (1988) through a fit of the rotating vector model to the PA profile. I chose the reference frequency $f_{ref} = 330$ MHz, the lowest frequency channel of the

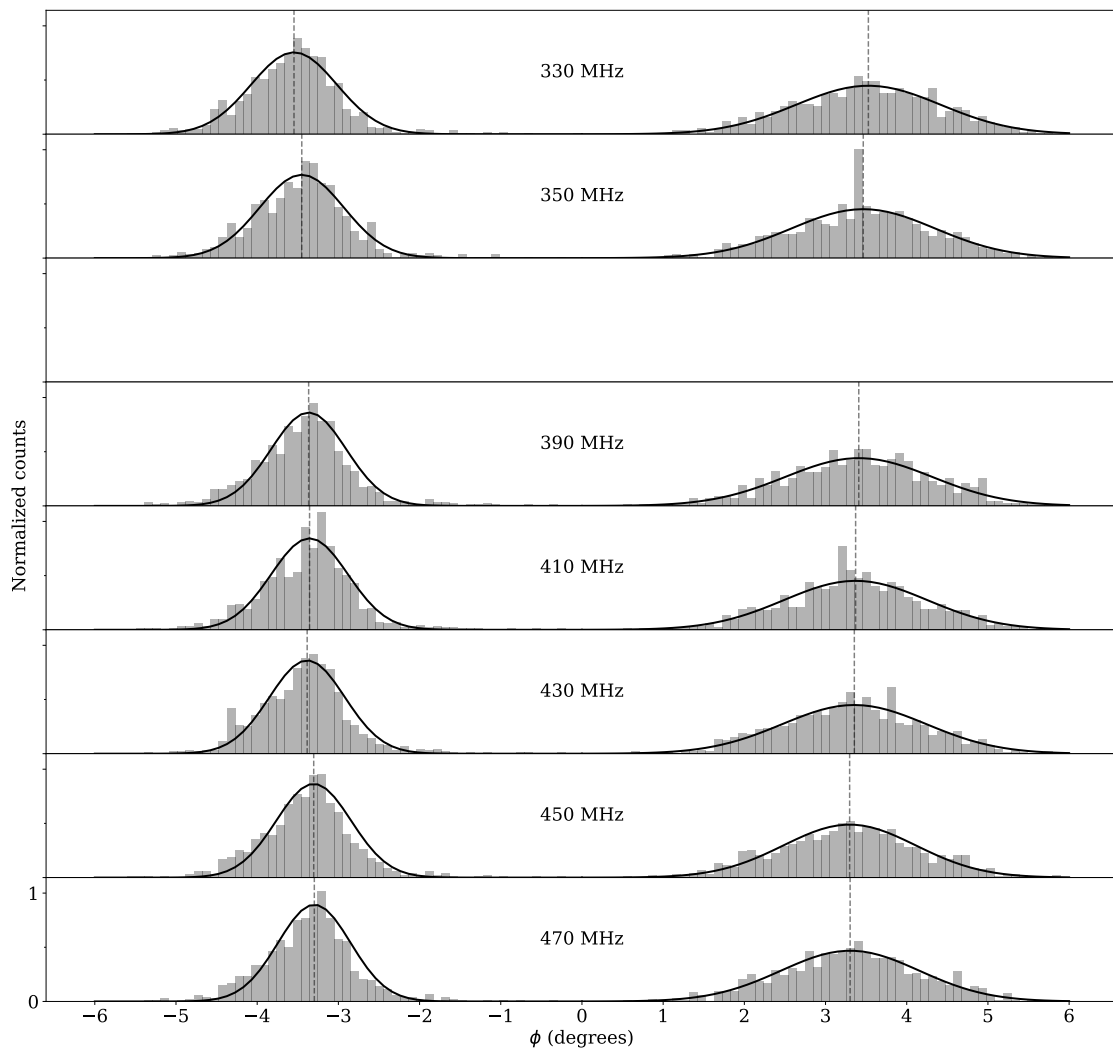


Figure 5.1: Single pulse histograms for PSR J1136+1551.

Histograms of the positions of the subpulses across the seven frequency bands observed. Vertical lines mark the means of the distributions to guide the eye to the distribution broadening. A space is left to mark the frequency band removed due to excess noise, so that the histograms are evenly spaced with frequency along the y axis. Gaussian fits to the subpulse distributions are over-plotted as black lines. The means and standard deviations of these Gaussians, and their uncertainties, are given in Table 5.1. Figure made with matplotlib.

Table 5.1: Table of parameters (mean μ , standard deviation σ and their uncertainties) of Gaussian fits to the subpulse distributions for the data and all three simulations.

Subpulse distribution 1 (left)								
Frequency (MHz)		470	450	430	410	390	350	330
Data	μ ($^\circ$)	-3.30	-3.30	-3.38	-3.35	-3.37	-3.45	-3.55
		± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02
	σ ($^\circ$)	0.45	0.46	0.46	0.47	0.46	0.52	0.53
Hollow cone	μ ($^\circ$)	-3.30	-3.31	-3.36	-3.36	-3.39	-3.46	-3.53
		± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02
	σ ($^\circ$)	0.50	0.50	0.50	0.50	0.50	0.50	0.49
Fan	μ ($^\circ$)	-3.31	-3.30	-3.37	-3.36	-3.39	-3.46	-3.54
		± 0.01	± 0.02	± 0.01	± 0.02	± 0.02	± 0.01	± 0.02
	σ ($^\circ$)	0.47	0.46	0.48	0.48	0.47	0.49	0.50
Fan, mode divergence	μ ($^\circ$)	-3.36	-3.35	-3.40	-3.39	-3.41	-3.46	-3.52
		± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02
	σ ($^\circ$)	0.43	0.43	0.44	0.44	0.45	0.46	0.45
		± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02	± 0.02
Subpulse distribution 2 (right)								
Frequency (MHz)		470	450	430	410	390	350	330
Data	μ ($^\circ$)	3.30	3.30	3.35	3.37	3.41	3.46	3.53
		± 0.07	± 0.06	± 0.05	± 0.07	± 0.05	± 0.05	± 0.07
	σ ($^\circ$)	0.85	0.82	0.89	0.88	0.91	0.88	0.89
Hollow cone	μ ($^\circ$)	3.29	3.30	3.35	3.35	3.39	3.45	3.52
		± 0.03	± 0.04	± 0.03	± 0.03	± 0.03	± 0.03	± 0.03
	σ ($^\circ$)	0.91	0.93	0.91	0.92	0.91	0.92	0.92
Fan	μ ($^\circ$)	3.28	3.29	3.35	3.35	3.36	3.44	3.52
		± 0.03	± 0.03	± 0.03	± 0.03	± 0.03	± 0.03	± 0.03
	σ ($^\circ$)	0.85	0.84	0.86	0.86	0.88	0.88	0.91
Fan, mode divergence	μ ($^\circ$)	3.37	3.38	3.42	3.43	3.44	3.50	3.54
		± 0.05	± 0.05	± 0.06	± 0.06	± 0.05	± 0.04	± 0.04
	σ ($^\circ$)	0.84	0.83	0.85	0.84	0.84	0.85	0.87
		± 0.05	± 0.05	± 0.06	± 0.06	± 0.05	± 0.04	± 0.04

observing band. I calculated the emission height at that frequency by assuming that peak of the emission region at f_{ref} lay on the line of sight. This means I assumed that the mean phase μ of the distribution of observed subpulse peaks at f_{ref} lies on the intersection of the line of sight and the field line responsible for the peak of

the overall emission region. This corresponds to the point where the cross and dot lie on top of each other in the illustrations of beam cross-section in Fig. 4.7. This mean phase μ is marked on Fig. 5.1 by a vertical dotted line, for each of the two distributions of subpulses at the reference frequency of 330 MHz. I further assumed that this field line, FL_{peak} , had a footprint parameter ratio of $s_L^{\text{peak}} = 0.5$, an average of the values calculated for seven pulsars in Gangadhara & Gupta (2001) and Gupta & Gangadhara (2003). This is enough information to convert μ into emission height r at f_{ref} using the equations in Section 4.1.2, and the choice of beam cross-section dictates the associated emission heights at the other frequencies. Fig. 5.2 shows how the mean phase position on the line of sight at each frequency for the left subpulse distribution (dots) is related to the emission peak on the reference field line FL_{peak} (crosses), for the hollow cone and fan beam models respectively.

Once I had a relationship between frequency and emission height inferred from the average behaviour of the subpulse distributions, I had enough information to predict how all the subpulse phases at $f_{\text{ref}} = 330$ MHz evolved with frequency. From each phase I identified the footprint parameter of the corresponding emitting field line, and then used the beam cross-section relationship to find the corresponding observed phase at every other height/frequency. I added a small Gaussian noise contribution to each predicted subpulse position, with $\sigma = 0.1^\circ$. Having predicted the subpulse phases separately for the two subpulse distributions, I calculated the phase separations and hence η .

5.2 Subpulse independence

In the above description of single pulse simulation, I generated the two subpulses that make up a double-component pulse separately. If there were a correlation between their positions this would have ramifications both for understanding of the pulsar emission mechanism and for how the simulations would need to be constructed. Fig. 5.3 shows the phase separation between the two subpulses at 470 MHz, plotted against the pulse mid-point, for each of the single pulses. The mid-point is the phase value half-way between the two subpulses, where a mid-point of 0 means that the subpulses are symmetrically positioned about the defined fiducial plane. The shading in this figure, and in all of the subsequent scatter plots in this chapter, indicates the densities of the scatter points, calculated using Gaussian kernel density estimation. This aids in comparison between the distributions of variables for the data and the respective simulations.

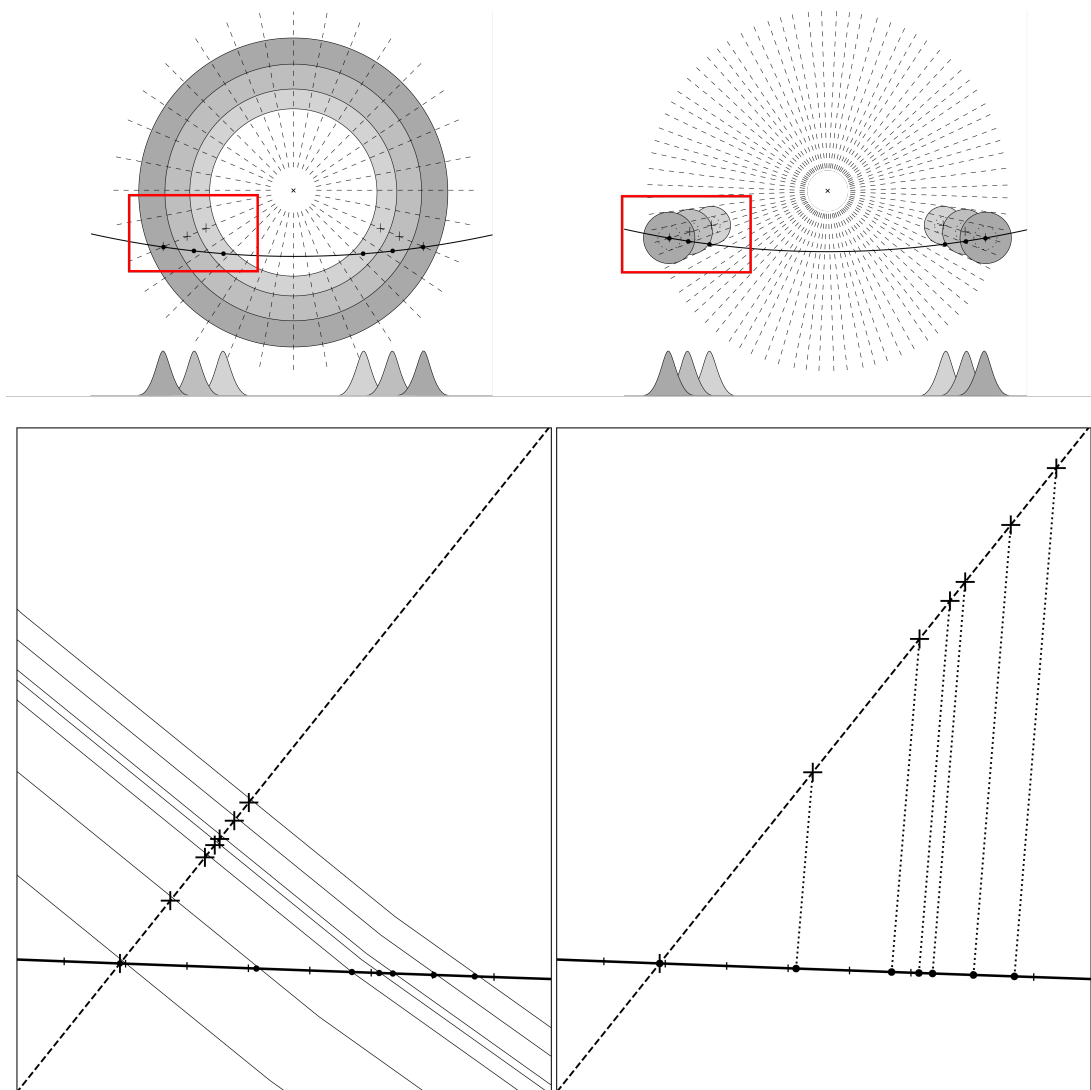


Figure 5.2: Relationships between emission region and line of sight for hollow cone and fan beam models.

Above: illustrative representations of beam cross-sections for the hollow cone and fan beam models, repeated from Fig. 4.4 and Fig. 4.7. Below: accurately modelled beam cross-sections, zoomed to a small region around the line of sight that is indicated on the illustrations above. As in the illustrations, the solid, near-horizontal, black line is the line of sight cutting across the beam cross section, whilst the dashed diagonal line is a magnetic field line. The dots mark the phase locations along the line of sight for the means of the left subpulse distribution at the seven observing frequencies. The crosses mark the locations on the designated reference field line that are related to these dot locations for the hollow cone model (left) and fan beam model (right). For the hollow cone model, these points are linked by circles of constant radius (these appear as diagonal lines on this zoomed region) and for the fan beam model they are connected by a line that is the perpendicular bisector of the line of sight at the dot location (see Fig. 4.9). Tick marks along the line of sight indicate intervals of $\pi/4000$ radians, going from $-73\pi/4000$ on the right to $-80\pi/4000$ on the left in both cases. Figure made using matplotlib, Ipe and Inkscape.




Figure 5.3: Scatter plots of subpulse separation vs. pulse midpoint: data and randomly combined.

Left: scatter plot of subpulse separation vs. the mid-point between the subpulses for all single pulse profiles in the highest frequency channel 470 MHz. Right: equivalent scatter plot generated by randomly pairing subpulses from the two distributions. Both scatter plots are shaded to indicate the density distribution of points more clearly, with the distribution calculated using Gaussian kernel density estimation. Figure made with matplotlib.

Randomly combining the two sets of subpulses from the data and plotting the resulting values for pulse mid-point and separation generates the same distribution as that from the unaltered data. The positions of the two subpulses making up a single pulse are therefore independent of each other, as expected.

5.3 Evolution of subpulse separation

Fig. 5.4 is a scatter plot of spectral index of subpulse separation, η , against the separation of the subpulse components at 470 MHz, the highest frequency in the observing band, for the double-component single pulses of PSR J1136+1551. The average profile evolution is overlaid, with a vertical black line marking the separation between the means of the Gaussian fits to the data histograms at 470 MHz, and a horizontal line marking the separation spectral index for the evolution of those Gaussian means across frequency. Both the subpulse separations and the spectral indices of their frequency evolution are spread over a range of values. The majority of single pulses have negative spectral indices, meaning that the single pulses get wider with decreasing frequency, as is expected from RFM. However, there is a proportion which have positive spectral indices. The data presented here only describe the double-component single pulses produced by PSR J1136+1551. Pulses where only one subpulse or the other is present are excluded from the analysis, as are other more complex pulses, including those with more than two components and those contributing to the “bridge” emission modelled in the integrated profile by Hassall et al. (2012). It is logical to assume that all such pulses evolve with frequency in similar ways, but this evolution cannot be quantified so straightforwardly.

5.4 Results: comparing observed single pulse properties to different beam models

5.4.1 Hollow cone

Simulating the single pulse frequency evolution expected for the hollow cone model fixes the emission to a ring of field lines with footprint parameters s_L^i distributed about $s_L^{peak} = 0.50$ with standard deviations of 0.02 and 0.04 for the two subpulse distributions. Fig. 5.5 shows the relationship between spectral index of subpulse separation, η , and pulse peak separation at the highest frequency of the observations (470 MHz) for both the data and the simulation generated with the hollow cone model emission region. Whilst the distributions occupy roughly similar regions of parameter

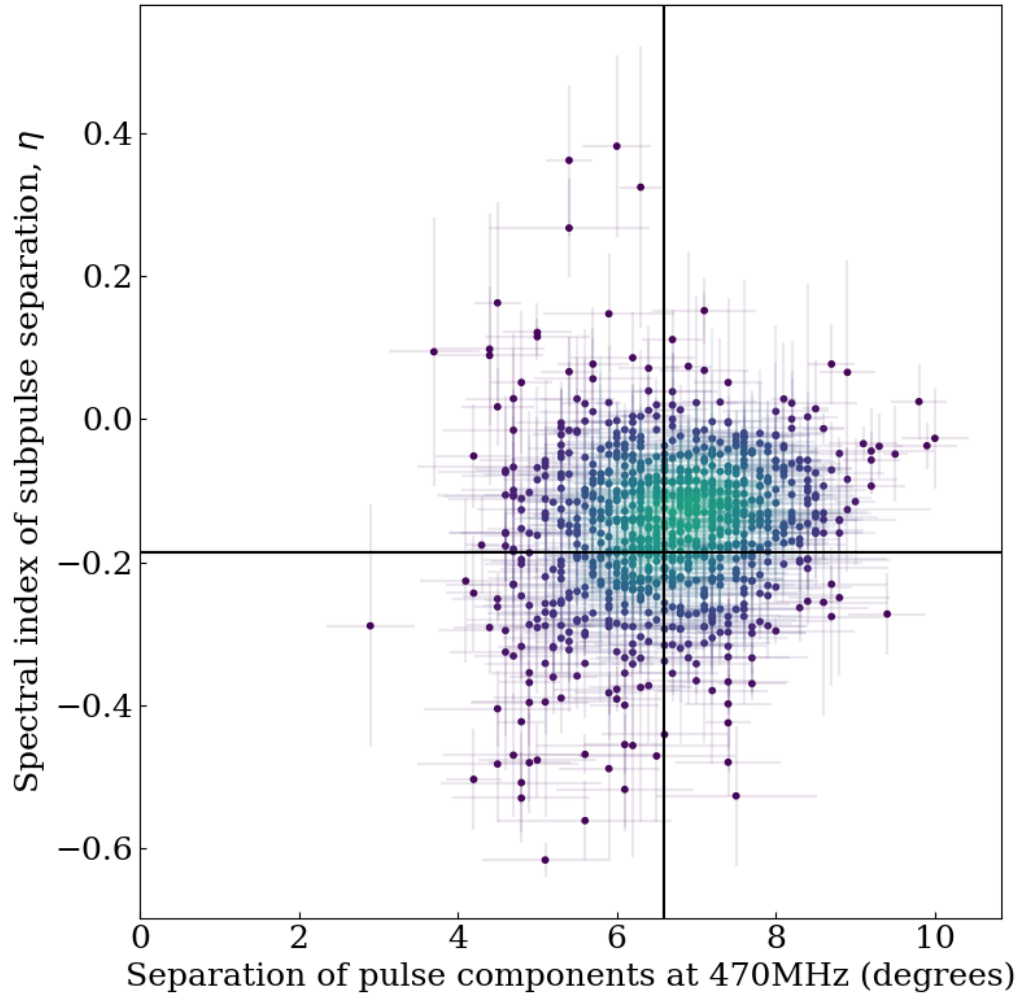


Figure 5.4: PSR J1136+1551: scatter plot of spectral index of component separation vs. separation at a single frequency.

Scatter plot showing the distribution of spectral index of subpulse separation, η , against pulse separation at the highest frequency channel (470 MHz). As for all the scatter plots in this chapter, the plot is shaded to show density distribution. The separation between the component distribution means at 470 MHz is overlaid as a black vertical line. The spectral index for the separation evolution of the component distribution means is overlaid as a black horizontal line. Figure made with matplotlib.

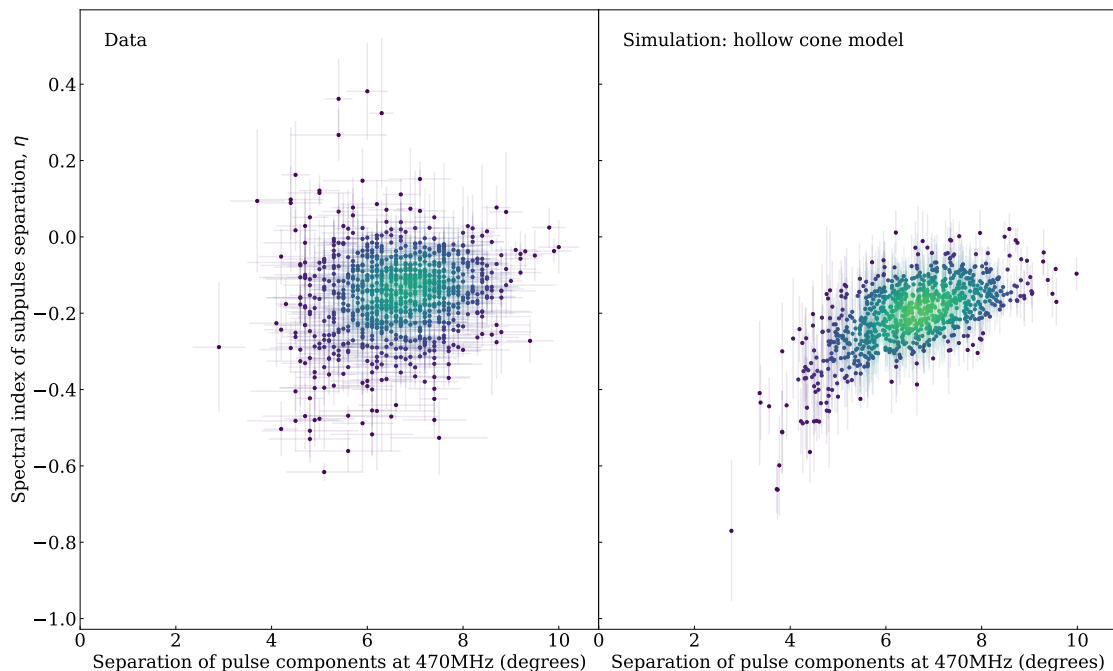


Figure 5.5: Separation spectral index scatter plots: data and hollow cone model simulation.

Scatter plots as in Fig. 5.4, for the data (left) and the RFM simulation using the hollow cone emission region (right). The simulation has a clear slope towards more negative spectral indices at lower separations that is not seen in the data. Figure made with matplotlib.

space, there is a clear curved shape in the simulated distribution, such that smaller pulse separations result in steeper spectral indices, in contradiction both with these observations and with those done on multi-component profiles by Mitra & Rankin (2002). This is an effect arising from how the line of sight cuts ring-shaped regions of different sizes for this particular geometry.

As can be seen in Table 5.1, although the mean positions of the subpulse distributions for the hollow cone model are similar to those of the data, the standard deviations of the subpulse distributions do not increase with decreasing frequency as expected, rather they stay roughly constant. This too is a geometrical effect of a ring-shaped emission region for this pulsar and is not seen in the data.

5.4.2 Fan beam

For the two fans in the fan beam simulation, the distributions of s_L^i have mean 0.50 and standard deviations 0.03 and 0.05. It can clearly be seen in Fig. 5.6 that the relationship between η and subpulse separation at 470 MHz for the fan beam model




Figure 5.6: Separation spectral index scatter plots: data and fan beam model simulation.

Scatter plots showing the distribution of η against pulse separation as for Fig. 5.5, now with the fan beam model for the emission region.




Figure 5.7: Cumulative distributions of spectral index of subpulse separation, η , for the data (solid) and the RFM model with the fan beam emission region (dotted). Figure made with matplotlib.

is more similar to the results from the data than those of the hollow cone model. The distributions of subpulse positions also replicate those of the data more closely (see Table 5.1).

The geometrical construction of the fan beam model produces results that are more compatible with the single pulse frequency evolution. However, collapsing Fig. 5.6 along the x-axis to look at the distribution of η indicates that the model predicts overly steep frequency widening, as shown in the cumulative plot in Fig. 5.7. In particular, the model makes no allowances for η at or close to 0, so that a two-sample Kolmogorov-Smirnov test on the two distributions yields a negative result: KS value of 0.27 and corresponding p-value of 3.33×10^{-16} .

Since at different frequencies the peaks of the intensity contours are at different distances from the line of sight, the fan beam model predicts that the intensity of the emission will also vary with frequency as a purely geometrical effect. The further the peak of the emission contours from the line of sight, the lower the intensity observed. The assumption made in the modelling—that the emission peaks lie on the line of sight at the lowest observed frequency—means that for this simulation the higher frequency emission lies further from the line of sight than the lower frequency emission, so that intensity decreases as frequency increases. Since the data were not flux calibrated, I could not test this, but in terms of a qualitative trend this is as expected. However, this particular set-up also means that the intensity would decrease for frequencies below this observing band. Moving all contour peaks off the line of sight and closer to the magnetic axis would be more correct, but there is no information as to where they should be placed, hence the choice made to fix them on the line of sight at the lowest frequency. This choice also means that the resultant emission heights estimated must be an upper bound.

5.4.3 Orthogonal mode propagation

The distribution of η for the data is skewed towards flatter spectral indices than those predicted by the fan beam model. The theory of OPMs provides a justification for this asymmetry, to account for the frequency-independence of the X mode propagation (see Section 1.2.3 and definition 2j in Chapter 2). Manchester et al. (1975) studied the polarization properties of the single pulses of twelve pulsars, including PSR J1136+1551, and found that for this pulsar, “about 25%” of the single pulses had a polarization angle orthogonal to that of the other 75% when averaging over long time intervals, suggesting an approximate 3:1 split in mode dominance for the single pulses of PSR J1136+1551.

I split the simulated subpulse distributions into X and O modes by assuming that both modes had the same distribution at the highest observing frequency of 470 MHz. For the X mode I therefore randomly drew a fraction of the subpulses from the 470 MHz phase distribution and used that same distribution for those subpulses at all the other frequencies, since the X mode does not evolve with frequency. For the remaining subpulses I followed the fan beam modelling process as already described.

The simulation requires as an input parameter the value of the height of emission at every frequency. I found that if I used those heights calculated from the average frequency evolution of the O mode distribution, then the height range was too broad and resulted in an overly steep distribution of η . However, the assumption that the X and O mode subpulses have identical distributions at the highest frequency is likely to result in an overestimation of the emission height range for the O mode pulses alone, so this is unsurprising. Since it was not possible to misalign the distributions of X and O modes at the highest frequency, I instead used the heights calculated earlier for the fan beam model without mode divergence. Using these heights and an X mode fraction of 30% gave the closest replication of the data.

Fig. 5.8 shows the distributions of simulated subpulses from the fan beam model with mode divergence included, with the Gaussian fits to the data distributions overlaid. The subpulse distributions closely follow the shapes of the data distributions, with the left component distribution in particular showing the same asymmetry. The values of the fit parameters, the mean μ and standard deviation σ , along with the uncertainties on these fits, are given in Table 5.1. As before, the scatter plot of η against separation in Fig. 5.9 is morphologically very similar for the model and the data. It is clear from Fig. 5.10 that the distribution of η for this model replicates the data much more closely, with a two-sample Kolmogorov-Smirnov statistic of 0.07 and corresponding p-value 0.02.

5.5 Calculation of emission heights

Fig. 5.11 shows the emission heights generated assuming a fan beam model and $s_L^{peak} = 0.5$. Neglecting uncertainties on the emission height values, the best power law fit to the heights, r , measured in km, against frequency, f (GHz), is given by:

$$r(f) = (424 \pm 105) f^{-0.39 \pm 0.04}, \quad (5.1)$$

where the constant of proportionality depends on s_L^{peak} , hereto set to 0.5. If the uncertainties on the emission heights are used to weight the power law fit, this changes

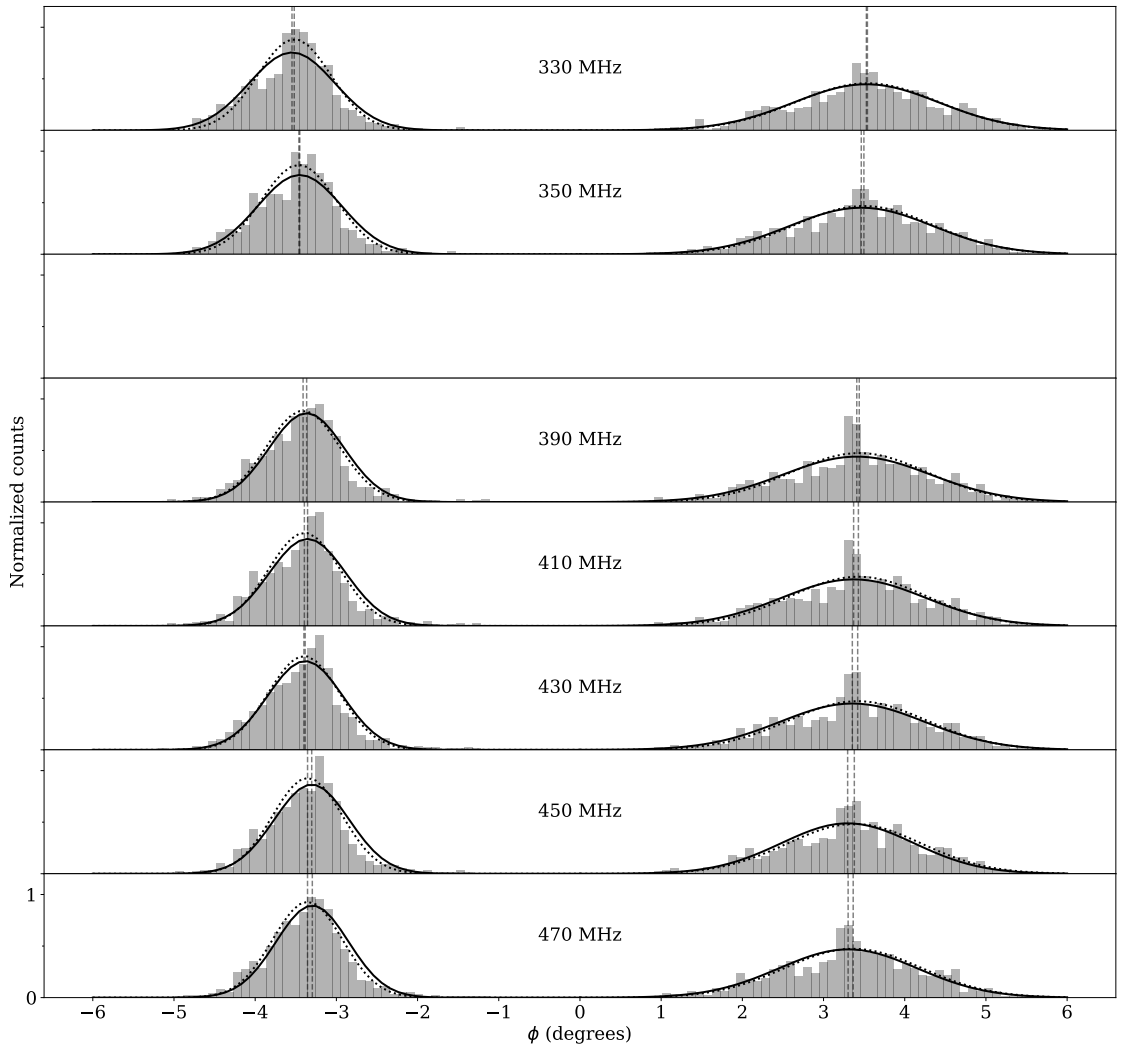


Figure 5.8: Fan beam model with mode divergence: subpulse histograms. Histograms of subpulse positions generated by the mode divergence simulation with fan beam emission region at all seven frequency bands. Gaussian fits to histograms are over-plotted as lines: the fit to the data as a solid line and the fit to the simulation as a dotted line. See Table 5.1 for the means and standard deviations of both sets of Gaussians. Figure made with matplotlib.




Figure 5.9: Separation spectral index scatter plots: data and fan beam model simulation with mode divergence.

Scatter plots as in Fig. 5.5 and Fig. 5.6, showing the distribution of spectral index of subpulse separation, η , against pulse separation at the highest frequency channel (470 MHz), for the data (left) and the mode divergence simulation with fan beam emission region (right). Figure made with matplotlib.




Figure 5.10: Cumulative distributions of the separation spectral index η for the data (solid) and the mode divergence model with fan beam emission region (dotted). Figure made with matplotlib.




Figure 5.11: Emission height vs. frequency: fan beam model. Emission heights calculated using the RFM model with fan beam emission region and $s_L^{peak} = 0.5$, as described in Section 5.1. The axes are log-scaled, with the power law best fits shown as straight black lines. The solid line shows an unweighted power law fit, whilst the dashed line shows the power law fit weighted by the uncertainties on the emission heights. Figure made with matplotlib.

to:

$$r(f) = (447 \pm 125) f^{-0.33 \pm 0.05}. \quad (5.2)$$

These heights are actually an upper bound, caused by fixing the peak of the emission region for the lowest observed frequency to the line of sight as discussed in Section 5.1. If I shifted the emission region to be closer to the magnetic axis, this would reduce the emission heights. However, it would not substantially alter the range over which these heights are spread. When drawing comparisons with this work in subsequent work presented in this thesis, I use the first of these two equations. Since all of the comparisons are qualitative, and this height-frequency relationship is meant to be indicative rather than an absolute description for this pulsar, the conclusions drawn for such comparisons would be unchanged were I to use the second equation instead.

The fan beam model results in a larger spread of heights across the observing band than the hollow cone model. This can be readily understood by considering the geometry on the 2D projection of the beam as described in Section 5.1: see Fig. 4.7 and Fig. 5.2. The observed peaks of emission from the fan beam model are related to points further along the reference field line, and therefore at a greater height difference, than the equivalents for the hollow cone model.

5.6 Summary

I used the single pulse frequency evolution of PSR J1136+1551 to compare two different models of the emission region within the magnetosphere. The metrics describing the single pulse evolution cannot be explained by the hollow cone model, which incorrectly predicts a shrinking of the subpulse distributions with decreasing frequency and a geometry-based relationship between spectral index of separation, η , and separation of subpulse components at a given frequency (Fig. 5.5).

By contrast, the fan beam model does not have symmetry requirements and the subpulse evolution it predicts is much closer to that seen in the data. This model is however limited by the constraint of a one-to-one mapping of frequency and height. It also only generates an upper bound on emission heights, since lack of other knowledge necessitated my positioning the lowest frequency contours so that they lie exactly on the line of sight.

The distribution of separation spectral indices generated by the fan beam model does not replicate the subset of single pulses in the data that have η close to zero. This can however be explained by the differing frequency behaviour of OPMs. My statistical implementation of OPM behaviour is subject to the same limitations as

those of the fan beam model, plus the added limitation of having to use the subpulse distribution at the highest observed frequency as the fixed point where the two modes converge in emission height. Despite this, the fan beam model with mode divergence results in the best replication of the single pulse evolution observed in the data.

Chapter 6

Structure in broad-band polarimetric pulsar observations

The new broad-band observations made with the UWL on the Parkes radio telescope (see Section 3.2.2), open up our capability to understand in detail how pulsar emission evolves with frequency. This is important for questions of beam structure, as addressed in the previous chapter, and also for understanding the polarization properties of pulsar emission through its evolution with observing frequency (open question 2c). In order to be able to address these questions on the intrinsic emission properties of the pulsar radio beam however, it is necessary to account for the effects of propagation of the radio waves through the ISM (Q2b). My initial work with UWL data therefore focused on disentangling the frequency-dependent effects of propagation through the ISM from the frequency-dependent signatures of the pulsar beam structure. For the purposes of this work, I focused on the linear polarization component of pulsar emission, in order to provide the initial base work from which future efforts involving circular polarization can be built.

Accounting for the effects of the ISM requires precise alignment of pulse profiles across frequency, to correct for both dispersion and Faraday rotation. The choice of DM and RM applied to correct for these effects depends upon the purpose for which it is being applied. In cases where the goal of dedispersion is to obtain a profile with high S/N, the DM which maximises this property for the total intensity is the correct choice. However, since pulse profiles tend to evolve with frequency, such a choice of DM does not necessarily correspond to correctly aligned profiles. In order to be able to link broad-band observations to the emission processes of the magnetosphere we implicitly require some concept of a fiducial point upon the observed pulse profile, an idea introduced in Section 1.1.1 and defined in Chapter 2. The correct dispersion measure should hypothetically allow us to map rotational

phase onto the pulse profile in a consistent way across the observing band, which is then relevant for the conclusions to be drawn about the positions and frequency structures of emitting components of the pulsar beam.

In reality, there is no way to obtain the unambiguously correct DM. Any dedispersion applied to a pulsar observation requires underlying assumptions, as discussed in Section 3.4.1. The appropriate path to take, therefore, is to base those assumptions in the theory of the pulsar beam. My approach was to use the fact that the geometry of the pulsar is constant (certainly on the timescales of individual observations at least) and independent of observing frequency. My key assumption was that the pulsar geometry leaves a signature on pulsar observations through the rotating vector model. This means I assumed that the underlying shape of the PA profile, though likely to be affected by other factors, was linked directly to the pulsar geometry and was hence unaltered by observing frequency. As a result, I developed a method of alignment of broad-band pulse profiles, described in Section 4.3, which obtains measurements for the DM and RM that best align the PA profiles across the band.

The initial limitations of this assumption are immediately apparent with further consideration of pulsar theory. It has been found to be very difficult to perform fits of the RVM to PA profiles, and such fits rarely constrain the geometry of the pulsar (see e.g. Rookyard et al., 2015). Furthermore, it is observed that pulse profiles commonly show complexities not encompassed by the RVM, including orthogonal and non-orthogonal jumps, and unexplained deviations from RVM tracks. These are commonly associated with the presence of a high circular polarization pulsar fraction and variability in that circular polarization across the pulse profile. The RVM does not take account of the presence of circular polarization and its potential origins in the production and interaction of polarized modes in the pulsar magnetosphere. In addition, many pulsars are observed to have low polarization fractions, or to be strongly scattered. The effect of scattering on pulsar polarization is to smear out variation across the pulse profile so that some average PA value is measured (Karastergiou, 2009). Nevertheless, despite these limitations, the fact remains that the majority of pulse profiles bear at least some resemblance to a shape inferred from the rotating vector model. Past work, too, has shown that this shape tends to be consistent across frequency: Karastergiou & Johnston (2006) (hereafter KJ06) showed the difference between the PA profiles of 17 pulsars at 1.4 and 3.1 GHz is tightly distributed around zero. Similarly, fig. 1 in Mitra & Li (2004) shows very consistent PA profiles at frequencies between 0.4 and 1.6 GHz for the pulsars studied in that paper.

My goal with this research, therefore, was to develop a method to align broadband pulsar observations in order to reveal the intrinsic behaviour of the pulsar emission, and to test the extent to which this method was applicable to pulsar observations, given the breadth of observational features recorded for pulse shapes and the behaviour of the PA profile. For this work I focused on aligning two pulsars, PSRs J1359–6038 and J1056–6258, for which previous work by KJ06 had indicated that alignment based on either the total intensity or the PA profile yielded different results. Profiles for these pulsars are shown in Fig. 6.1. My method seeks to find the best fit values of DM and RM simultaneously. This is important because there is a covariance in the choices for these two values: the best fit RM will always depend on the choice of DM used to align the profiles in time before comparing the PA values across frequency. Previous work that sought to align pulse profiles across frequency, and so investigate the intrinsic beam structure, have mostly not done so through dedispersion and correcting for Faraday rotation. Mitra & Li (2004) selected the best alignment for their PA profiles without reference to the DM and RM. KJ06 inferred RM values based on the best fit alignments of two PA profiles at discrete frequencies, but the use of only two discrete frequencies meant that the frequency-squared relationship of those RMs could not be tested. More recently, Ilie et al. (2019) selected the DM values for which variations of the phase-resolved RM values measured across the profile were minimized. This is an alternative approach which takes as its precepts similar assumptions to those I make: that we would expect the PA shape to remain largely consistent across frequency and hence alignment should be possible through a combination of DM and RM. Ilie et al. (2019) attribute the phase-resolved RM variations they measure to either scattering or to magnetospheric effects.

6.1 Testing the modelling process with simulations

I tested my alignment algorithm on simulated data in order to understand the ways in which PA geometry and S/N affected the algorithm fitting results. I generated these simulations using the methods described in Section 4.1. The simplest simulation involved a Gaussian intensity profile centred upon the fiducial plane with a PA profile with a steep gradient due to its geometry of $\alpha = 100^\circ$ and $\beta = -2^\circ$. I gave my simulated pulsar a DM of $+1 \text{ cm}^{-3} \text{ pc}$ and RM of $+2 \text{ rad m}^{-2}$. These values are deliberately small to replicate the situation for the real data: that the DM and RM have been roughly corrected prior to further analysis and the modelling process is optimized on applying the final accurate corrections to the alignment. The steepness




Figure 6.1: Integrated pulse profiles: PSRs J1056–6258 and J1359–6038. Integrated pulse profiles for PSR J1056–6258 at 1410 MHz (left) and for PSR J1359–6038 at 1410 MHz and 830 MHz (right). The top half of each graph shows the absolute PA profile, showing only those values for which $L_{\text{SNR}} > 20$ (left) and $L_{\text{SNR}} > 5$ (right), with points marked by black crosses with error bars. The bottom half shows the total intensity (black), linear polarization (red) and circular polarization (blue). For PSR J1359–6038 the profile at 830 MHz is overlaid, with faint dashed lines for the profile and faint dots with error bars for the PA, to show the effect of temporal scattering on the profile shape at lower frequencies. The RM (33 rad m^{-2}) and DM ($293.736 \text{ cm}^{-3} \text{ pc}$) used to align these two frequencies for the figures were taken from the pulsar catalogue (Manchester et al., 2005). It is immediately clear that the catalogue RM is incorrect, since the two PA profiles do not overlap on the unscattered left-hand side of the profile. Figure made with matplotlib.

of the PA gradient meant that the possibilities for profile alignment were tightly constrained. As a result, the modelling process returned the correct input results for DM and RM, with tight constraints upon those values.

Fig. 6.2 shows the two-dimensional histogram for the alignment probability posterior distribution for the simulation, along with the results of aligning the simulated profiles across frequency with the correct values. Even in this example, where the results are tightly constrained, it is clear that there is a correlation between smaller ΔDM values and larger ΔRM values, and vice versa. This reflects the slope direction of the PA profile. Simplistically, a smaller value of ΔDM shifts profiles to earlier phases, or to the left on the figure, and a larger ΔRM shifts the PA profiles to larger angles, or up on the figure. Since the PA profile has a negative slope, the best chance of alignment comes from shifting profiles diagonally along that direction. This factor is magnified for PA profiles that are less curved in shape, and is taken logically to the extreme in the hypothetical scenario of the PA shape being a straight diagonal line.

I investigated this aspect of model fitting in more detail by altering the PA geometry and profile shape in the simulation. I used $\alpha = 156^\circ$ and $\beta = -10.5^\circ$ to generate a PA with much less curvature, and shifted the Gaussian intensity profile off axis relative to the fiducial plane, so that the part of the PA with sufficient S/N to be measurable was asymmetric. This simulation is shown in Fig. 6.3. I made these choices deliberately in order to create a simulation that, within the parameters of model simplicity, was as similar to the pulse profile for PSR J1056–6258 as possible (see Fig. 6.1). Performing the alignment correctly recovers the values of ΔRM and ΔDM again, but this time the 2D alignment histogram is considerably wider. This reflects the increased flexibility in alignment probably introduced by the change in PA geometry. There are also non-negligible regions of alignment probability at large shifts away from the recovered values, respectively where ΔDM is large and positive whilst ΔRM is large and negative; and where ΔDM is large and negative whilst ΔRM is large and positive. These alignments represent the extremes, where the profiles have been misaligned so strongly that only a few data points from the leading edge of the PA profiles at one end of the band are overlapping with a few data points from the trailing edge of the PA profiles at the other end of the band. These alignments have a non-zero probability because the alignment algorithm does not penalize for low numbers of overlapping points. It is clear however that such alignments do not in fact correspond to the theoretical precept of the algorithm, seeing as they are instead extreme misalignments. It is important to be aware therefore that a similar triple alignment probability distribution is likely to appear when aligning real data.

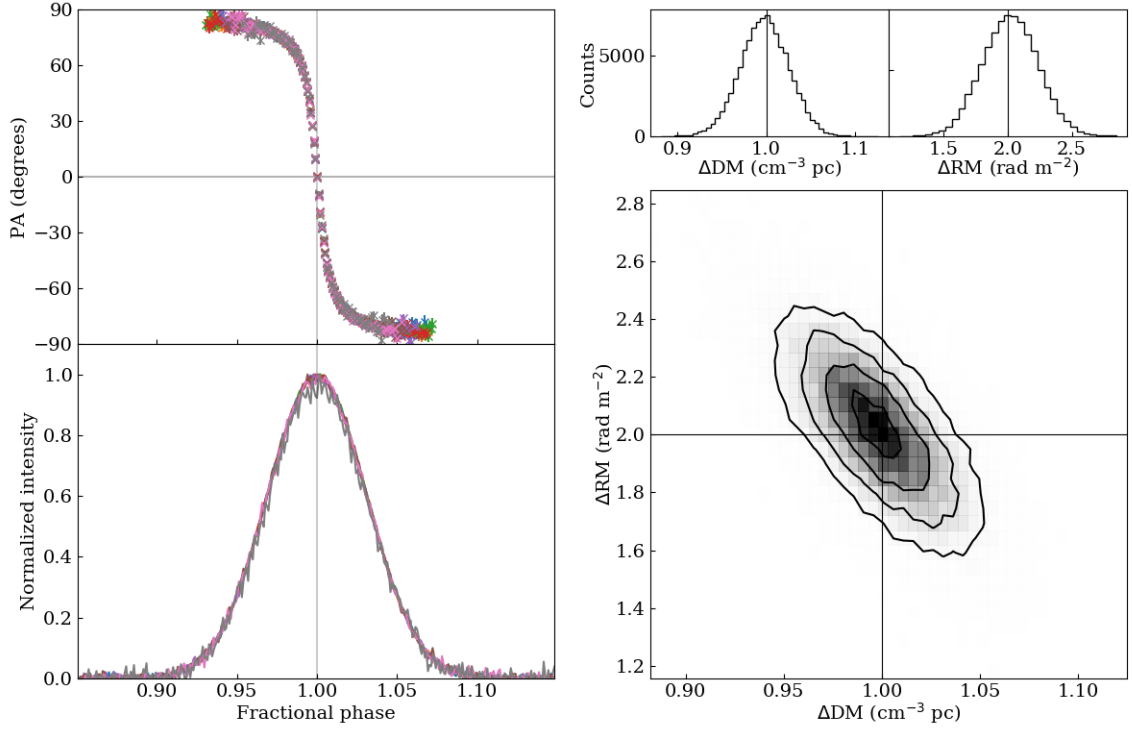


Figure 6.2: Basic simulated profile: alignment results.

Plots showing the alignment of a simulation of pulsar emission with a PA geometry of $\alpha = 100^\circ$, $\beta = -2^\circ$. Right: the main plot shows a 2D histogram, showing the alignment probability distribution for pairs of parameters ΔRM and ΔDM . The vertical and horizontal black lines overlaid indicate the true values of these parameters for the simulation: $\Delta\text{RM} = 2 \text{ rad m}^{-2}$ and $\Delta\text{DM} = 1 \text{ cm}^{-3} \text{ pc}$. Above are the marginalized distributions of ΔDM and ΔRM with the true values marked with vertical black lines. Left: simulated pulsar profiles aligned using the ΔRM and ΔDM values marked on the 2D histogram, with all 8 frequencies shown. The frequencies are overlaid from lowest to highest just as in Fig. 4.2, so that the lowest frequency of 827 MHz (coloured blue but in this case hidden behind other overlaid data) is plotted first and the highest frequency of 3775 MHz (coloured grey) is plotted last. The upper panel shows the position angle and the lower panel shows the normalized total intensity. Figure made with matplotlib.




Figure 6.3: Simulated profile with flatter PA: alignment results. Simulation of pulsar emission designed to be similar to the observation of PSR J1056–6258. See Fig. 6.2 for details of the figure contents. Figure made with matplotlib.

Another factor to consider is the effect of S/N upon pulsar alignment. Pulsar intensity tends to decrease with increasing frequency, so that S/N is often lower at higher frequencies. Since accurate calculation of the PA is only possible where the linear polarization L has high S/N, this reduces the number of data points available for alignment at higher frequencies. However, since the largest misalignments due to dispersion and Faraday rotation take place at lower frequencies, the alignment accuracy is most dependent upon the lower frequency end of the observing band. In a wider survey of pulsar polarization therefore, the low frequency channels are the most important for performing alignment of pulse profiles, and alignment accuracy is likely to be poorest for pulsar profiles that have RFI at lower frequencies, along with PA profiles with minimal curvature.

For the third simulation I added in additional misalignments of intensity and PA profiles, to investigate the extent to which the algorithm would correct for dispersion and Faraday rotation when there were competing effects also at play. The rationale behind testing this is that we expect the frequency structure of the beam to be apparent observationally in terms of additional misalignments independent of those caused by dispersion and Faraday rotation. The origin of this is the delay-radius relation originally proposed by Blaskiewicz et al. (1991) and re-derived mathematically by Dyks (2008) (see definition 2i in Chapter 2). This relation predicts a lag of the PA with respect to the intensity profile of $4r/R_{LC}$ in pulse phase, caused by the rotation of the pulsar reference frame in comparison to that of the observer. Dyks (2008) also describes a height-dependent PA rotation of $10r \cos(\alpha)/(3R_{LC})$ as originally calculated by Hibschan & Arons (2001).

The amount by which the intensity and PA profiles shift depends upon the height at which the pulsar emission is produced. I used the height-frequency relationship calculated in Section 5.5 for PSR J1136+1551 of $r(f) \propto f^{-0.39}$, with height r in km and frequency f in GHz. This relation comes from applying the fan beam model to the single-pulse frequency evolution of PSR J1136+1551 across the frequency range 330–470 MHz, assuming a 1-to-1 relationship between frequency and emission height and the field lines of the emitting region having a footprint parameter ratio, as defined in Gangadhara & Gupta (2001), of 0.5. The particular height-frequency relationship and footprint parameter ratio are only taken as indicative examples. Were I instead to show a simulation that uses the error-weighted height-frequency relationship also described in Section 5.5, $r(f) \propto f^{-0.33}$, the frequency evolution of the simulation would be qualitatively very similar and the conclusions drawn would not be changed at all. I also used the full simulation of the fan beam model described in Chapter 5

to further calculate how the shape of the intensity profile evolves with frequency as emitting field lines diverge.

I followed the prescription of Dyks (2008) to implement the phase shifts of the intensity and PA profiles, namely that they are shifted in phase by equal amounts of $2r/R_{LC}$ earlier and later respectively. Here R_{LC} is the radius of the light cylinder of the pulsar, given in terms of the pulsar period as $cP/2\pi$. This is found to be valid for a range of theoretical models provided emission is produced at heights below $0.1 R_{LC}$ (Craig & Romani, 2012). I also included the PA rotation of $10r \cos(\alpha)/(3R_{LC})$. The PA rotation between the top and bottom of the band for the heights mentioned above (assuming a pulsar rotational period of 0.42 s as for PSR J1056–6258) is 1.8° , which is small compared to the effect of Faraday rotation.

Fig. 6.4 shows the alignment for this height-varying simulation based upon its PA profiles. I fit Gaussians to the marginalized ΔDM and ΔRM distributions (marked in blue) to identify the best fit and its uncertainty. As for the previous simulation, I neglected the probability region corresponding to a large ΔDM and clear misalignment. It is a point of note that only the positive ΔDM side-lobe is seen in the marginalized posterior distribution, as was almost the case in Fig. 6.3. This is likely a result of the asymmetry of the PA profile due to the intensity profile being offset with respect to the fiducial plane. In this simulation, the peaks of the ΔDM and ΔRM marginalized posterior distributions do not correspond to the true values for these parameters. ΔDM stems from the delay-radius relation shifting the PA with respect to the total power profile in pulse phase, whereas ΔRM additionally stems from the simulated PAs including both Faraday rotation and rotation due to different emission heights. Shifting the profiles based on the best recovered ΔDM and ΔRM results in aligned PA profiles but a frequency-dependent offset in the total power. If the conditions describing the beam emission height structure are representative of reality, a similar effect should be observable in real data.

6.2 Modelled pulsars

Fig. 6.1 shows the integrated profiles for the two pulsars I selected for the alignment modelling to correct for the effects of the ISM. The profiles show total intensity, linear and circular polarization, and the PA profile, at 1410 MHz for both pulsars. For PSR J1359–6038 I also show the profile at 830 MHz overlaid, to demonstrate the effects of interstellar scattering on profile and PA shape. In addition to smearing out the profile to have a characteristic exponential scattering “tail”, the curvature of the




Figure 6.4: Simulated profile including emission height variation: alignment results. Simulation as in Fig. 6.3 but now including the effects of varying emission height with frequency. The alignment used to plot the profiles on the left is the best fit indicated by dotted horizontal and vertical lines on the 2D histogram, found by fitting Gaussians to the marginalized distributions (blue dashed lines). The true values of $\Delta RM = 2 \text{ rad m}^{-2}$ and $\Delta DM = 1 \text{ cm}^{-3} \text{ pc}$ are indicated on the histograms with solid lines. Figure made with matplotlib.

PA profile is flattened, as expected from the effects of scattering to later timescales as described by Karastergiou (2009). The alignment of these two profiles for this pulsar was done using the DM and RM values given in the pulsar catalogue (Manchester et al., 2005), which I give in Table 6.1, and it is immediately clear that there are limitations to these values for PSR J1359–6038. First, the peaks of the intensity profiles are aligned, even though it is known that scattering alters profile shape, shifting its peak to later times. The DM used here was clearly chosen to maximise S/N across the observing band: a valid choice for applications where this is the key factor, but scientifically inaccurate in terms of using the DM to characterize the electron column density along the line of sight. We would expect the RM to align the PA profiles in angle, particularly at the leading edge of the profile where the smearing of the profile due to scattering has a less pronounced effect. It is clear therefore that the catalogue RM value for this pulsar is incorrect.

The process I have developed for aligning pulsars using the DM and RM simultaneously is relevant even for pulsars where the alignments due to the catalogue DM and RM are not visibly incorrect in the same way as for PSR J1359–6038. This is because DM values for a pulsar are commonly calculated with the purpose of maximizing the total S/N of the profile, rather than accounting for the effects of the ISM and revealing the intrinsic frequency evolution of the pulse profile. The choice of RM that best aligns polarization across frequency will depend on the DM that has been used to align profiles, so the choice of DM will always influence the outcome RM measured. Therefore, simultaneous fitting for DM and RM is useful for better understanding of both the pulsar emission itself, and what these values actually mean in terms of the electron column density in the ISM.

The resultant values of DM and RM that I obtained by applying my modelling technique to these two pulsars are given in Table 6.1. The table presents these results in terms of the catalogue RM and DM values and the corrections to them required to align the PA. For comparison, I also present the ΔRM and inferred ΔDM values resulting from the methods of KJ06 and Ilie et al. (2019). I calculated the values of ΔDM for KJ06 by estimating the DM required to align the intensity profiles at the frequencies they used and then converting the time delay given in the paper to its associated ΔDM , given the observational frequencies. I converted the ΔDM corrections from Ilie et al. (2019) to be defined with respect to the catalogue values, to make them directly comparable to my work, and used their average of their phase-resolved RM calculations to obtain ΔRM . These values are also plotted on Fig. 6.5

Table 6.1: The DM and RM for the pulsars studied in this paper, given as the catalogue values and the difference to these from this work and from that of both KJ06 and Ilie et al. (2019). The values of \dot{E} are from Manchester et al. (1978). The catalogue values for RM and DM have the following references. PSR J1359–6038 DM: Petroff et al. (2013); RM: Han et al. (1999). PSR J1056–6258 DM: D’Alessandro et al. (1993); RM: Costa et al. (1991).

Name	DM_{cat}	ΔDM (cm⁻³ pc)		
	(cm ⁻³ pc)	KJ06 (est.)	Ilie et al. (2019)	This work
J1359–6038	293.736 ± 0.003	+0.4 or 0	0.498 ± 0.004	+0.2 ± 0.2
J1056–6258	320.3 ± 0.6	+6.6 or +0.5	0.79 ± 0.01	+1.5 ± 0.8
\dot{E}	RM_{cat}	ΔRM (rad m⁻²)		
	(ergs s ⁻¹)	KJ06	Ilie et al. (2019)	This work
1.2 × 10 ³⁵	33 ± 5	+6 or +4 ± 0.5	5.5 ± 0.1	+6 ₋₂ ⁺⁰
1.9 × 10 ³³	4 ± 2	-5 or +2 ± 0.5	2.5 ± 0.1	+2 ± 1

and Fig. 6.6 respectively for the two pulsars, to show the effect of different data products and different methodology on the measurements of RM and DM.

6.2.1 PSR J1359–6038

Since scattering affects the PA profile shape in a frequency-dependent way, this presents a limitation on the underlying assumption of the modelling that the PA profile is consistent across the observing band. However, as scattering is strongly frequency-dependent, the deformation of the PA is only important at the lowest observed frequencies. When fitting for alignment of data for PSR J1359–6038, I therefore excluded the three scattered profiles at the low frequency end of the observing band from the fitting process. Plotting them after the alignment then gave additional information constraining the RM correction. I performed the fit for this pulsar using the PA values where the S/N of the linear polarization was greater than 5.

Fig. 6.5 shows the alignment results for PSR J1359–6038. Similar to the results from simulation fitting, there is some probability associated with large DM shifts which can be neglected. As described for the height-varying simulation, I fit Gaussians to the marginal Δ DM and Δ RM distributions (marked with blue dashed lines) to obtain the best fit Δ DM = +0.2 ± 0.2 cm⁻³ pc and Δ RM = +6 ± 2 rad m⁻². The best fit is marked on the 2D histogram with vertical and horizontal dashed lines. The catalogue values and those calculated by KJ06 and Ilie et al. (2019) are indicated with markers. The catalogue values lie outside the contoured region whilst the two fits by KJ06 (with respect to the intensity and PA profiles) both lie close to the contour peak, as does the result of Ilie et al. (2019). The plot of the PA profiles on the left-hand




Figure 6.5: PSR J1359–6038: alignment results.

Plots as described in Fig. 6.2, now showing PSR J1359–6038. The alignment on the left is $\Delta\text{RM} = +6 \text{ rad m}^{-2}$ and $\Delta\text{DM} = +0.2 \text{ cm}^{-3} \text{ pc}$, identified by Gaussian fits to the marginal ΔDM and ΔRM histograms (blue dashed line), and marked with dashed lines on the 2D histogram. Overlaid on the 2D histogram are the ΔRM and ΔDM values from the catalogue (black circle) and those computed by KJ06 aligning by intensity (white inverted triangle) and PA (grey triangle). The corrections calculated by Ilie et al. (2019) are marked with a white star. An additional marker shows the value of $\Delta\text{DM} = -0.075 \text{ cm}^{-3} \text{ pc}$ calculated from a scattering fit to the intensity profiles (white square). Figure made with matplotlib.

side of the figure shows that alignment has been successful even for the three lowest frequencies, which were not included in the fitting. Since the lower frequencies are the most sensitive to changes in RM, adding them in provides a tighter constraint. Consequently I can quote $\Delta\text{RM} = +6_{-2}^{+0}$ rad m⁻² in Table 6.1.

Ilie et al. (2019) proposed that the variation in phase-resolved RM they detected could be explained partly by scattering, but that the lack of corresponding variations in circular polarization may suggest a further magnetospheric origin for the RM variations across the profile. Although I do not investigate circular polarization in this work, the results are in agreement with Ilie et al. (2019) on the effect of scattering on the phase-resolved RM. My technique therefore provides a complementary method to that of Ilie et al. (2019) to investigate frequency-dependent effects on polarization beyond dispersion and Faraday rotation.

6.2.2 PSR J1056–6258

Although unaffected by scattering, at the highest 2 frequency channels the PA profile of PSR J1056–6258 deviates from the rest with the presence of a pair of orthogonal jumps. I therefore excluded these two profiles when performing the alignment fit for PSR J1056–6258. For this pulsar I used only those PA values where the linear polarization S/N was greater than 20. The same edge effects are present as for the second and third simulations, though the lack of curvature of the PA means that they are more dominant. I therefore ran the alignment modelling twice for this pulsar, first with uninformative priors as standard, and then used these results to infer priors that I applied for a second run. To do this, I fit three Gaussians to the marginal ΔDM distribution from the first run’s output. I then used the central Gaussian of these three as the prior on ΔDM . Running the fit a second time with the ΔDM prior gave outputs for ΔDM and ΔRM that follow Gaussian distributions. These are marked on the marginalized distributions in Fig. 6.6 with dashed blue lines. The best ΔDM and ΔRM values and their uncertainties are given directly from the Gaussian posteriors, with $\Delta\text{DM} = +1.5 \pm 0.8$ cm⁻³ pc and $\Delta\text{RM} = +2 \pm 1$ rad m⁻².

As shown in Fig. 6.6, the catalogue values for RM and DM each individually could lie within the fit region for this pulsar, however, together they lie far outside the contours. This suggests that a different DM may have been used in the calculation of the catalogue RM. The best PA fit of KJ06 lies along the direction of the contoured region but at too large a shift. It therefore corresponds to the edge effect situation, something that is revealed due to the increased frequency coverage and resolution of these new data. Their other best fit, based on aligning the intensity profiles and




Figure 6.6: PSR J1056–6258: alignment results.

As for Fig. 6.2, now showing PSR J1056–6258. The best fit alignment (shown on the left) is that marked with dashed horizontal and vertical lines on the 2D histogram: $\Delta\text{RM} = +2 \text{ rad m}^{-2}$ and $\Delta\text{DM} = +1.5 \text{ cm}^{-3} \text{ pc}$. These numbers were obtained by fitting three Gaussians to the marginalized ΔDM distribution (solid blue line) and using the central Gaussian as a prior for a second run of the fitting algorithm. This produced a Gaussian distributed output in ΔRM and ΔDM (dashed blue lines), from which the best fit parameters are obtained. The symbols overlaid on the 2D histogram are as described in Fig. 6.5. Figure made with matplotlib.

calculating the corresponding RM, sits close to the contoured region. The extension to lower frequencies in these new data provides higher fidelity for the measured RM given the DM, which is likely to explain this discrepancy. The result from Ilie et al. (2019) lies in the contoured region close to the best fit parameters from my results.

6.2.3 Broad-band alignment and beam structure

Figures 6.7–6.10 show broad-band waterfall plots for the two pulsars, showing alignments across frequency before and after applying the DM and RM corrections given in Table 6.1. The methodology behind this style of visualization is given in Section 4.4.

At a single epoch there can be only one true value for the interstellar DM, defined




Figure 6.7: PSR J1359–6038: broad-band PA before and after alignment. Waterfall plot showing how the PA profile varies with frequency for PSR J1359–6258. Left: aligned using the catalogue RM and DM. Right: aligned using the DM and RM correction picked out in Fig. 6.5. The PA is only shown where the SNR of the linear polarization exceeds 4. Figure made with matplotlib.




Figure 6.8: PSR J1359–6038: broad-band intensity before and after alignment. Waterfall plot showing how the intensity profile varies with frequency for PSR J1359–6038, with DM and RM applied as in Fig. 6.7. The intensity is normalized by dividing by the summed flux of the on-pulse region. The white line indicates the profile peak and the dots indicate the peak of the unscattered profile for the 10 lowest frequency channels, estimated by fitting to the profile a Gaussian convolved with an exponential scattering function. Figure made with matplotlib.




Figure 6.9: PSR J1056–6258: broad-band PA before and after alignment. The PA profile for PSR J1056–6258, with the RM and DM values from the catalogue (left) and this work (right) respectively. Figure made with matplotlib.




Figure 6.10: PSR J1056–6258: broad-band intensity before and after alignment. The intensity of PSR J1056–6258, normalized by dividing by the on-pulse flux. DM and RM values applied are as in Fig. 6.9. Figure made with matplotlib.

as the line integral of the electron column density along the line of sight. However, as can be seen by comparing Figures 6.7–6.10, the DM that best aligns the PA across frequency is different from the DM that best aligns total intensity for both of the pulsars investigated. This measurement difference between the two obtained DMs must therefore result from an effect unrelated to interstellar dispersion. The results for these pulsars are qualitatively similar to the results for the height-varying simulation, in which the frequency-dependent PA lag, caused by the delay-radius relation, was absorbed into the best fit DM. This resulted in the intensity profile being visibly misaligned across frequency. Guided by this simulation result, it is possible to assume that the difference in the two dispersion measures originates from the effect of the delay-radius relation when pulsar emission radius varies with frequency. Doing so, I can associate the measured ΔDM with a difference in emission height between the top and bottom of the observing band.

The value of ΔDM in question is the difference between the DM that best aligns the PA, and the DM that best aligns total intensity, where I defined this latter neglecting any intensity profile evolution due to field line curvature. I then converted ΔDM into a frequency-dependent time lag. The difference between the time lags at the top and bottom of the band can then be used to calculate an emission height range across the band using the delay-radius relationship. This, converted into units of time, is $\Delta t = 4r/c$, where r is the emission height and c is the speed of light. Comparing lags at different frequencies therefore gives $\Delta t_1 - \Delta t_0 = 4\Delta r/c$.

The purpose of performing this calculation is not to measure the absolute emission heights for this pulsar, because the DM difference does not give sufficient information about how observable profile evolution might be related to emission heights. Comparing differential shifts across frequency, however, meant that I could investigate whether the misalignment is of a magnitude that, through the delay-radius relation, leads to an emission height range that is plausibly comparable to the expected range of emission heights in a radius-to-frequency mapping scenario.

The calculation for PSR J1359–6038 gives $\Delta r = 90 \pm 90$ km over the observing band. At the lowest frequencies, scattering delays the observed profiles by some further Δt . Correcting the total delay, including the effects of scattering, then results in an estimated indicative height range of $\Delta r = 120 \pm 120$ km. I can compare this result with the corresponding value from the height-frequency relation of $r(f) \propto f^{-0.39}$, which I calculated in Section 5.5. That equation is meant to be indicative of emission heights for the fan beam model under a particular set of criteria, so again the height ranges are not absolute measurements, but are nevertheless a useful

comparison for results from different approaches. For the UWL observing band this gives $\Delta r = 200$ km, which is comparable in magnitude.

I performed the same height range calculation for PSR J1056–6258, although this time without the need to correct for the effects of scattering. The catalogue DM does not appear to align the total intensity well for this pulsar, so instead I estimated that correcting the catalogue value by approximately $+0.5 \text{ cm}^{-3} \text{ pc}$ visually aligns the intensity profile across frequency. The difference between the two DMs aligning intensity and PA respectively is then $\Delta \text{DM} = +1 \text{ cm}^{-3} \text{ pc}$. This corresponds to an emission height range of $\Delta r = 430 \pm 350$ km, which is larger but still comparable in magnitude to the simulated value of $\Delta r = 200$ km. These results are not providing a strong constraint upon emission heights or the science that can be inferred from them. However, the fact that these results are consistent with theoretical expectations merits further investigation for a larger number of pulsars.

6.3 Summary

Having developed an algorithm to correct for the effects of the interstellar medium in broad-band pulsar observations, I tested its efficacy on three simulations of increasing complexity and then applied it to observations of two pulsars, chosen because previous work had raised questions about their profile alignment across frequency. The underlying assumption of my modelling, that the PA profile is defined primarily by the geometry, which is consistent across frequency, is not relevant in cases of interstellar scattering, or where the PA profile exhibits orthogonal (or non-orthogonal) jumps that evolve across the observing band. However, for pulsars and frequency bands where these effects are localized to a part of the band (for example scattering is only a strong effect at lower frequencies) I showed that observations from the remainder of the band may still be used effectively. Simulation testing revealed that S/N and PA profile shape both affect the accuracy to which the modelling can return the input values for RM and DM. Furthermore, simulation shows that the effects of the delay-radius relation under a RFM scenario can be absorbed into the estimates of DM and RM, at least within the constraints of what we expect the magnitude of these effects to be theoretically. However, doing so leads to visible misalignment of the total power which may be a recognisable signal of this effect.

Applying the modelling to the two pulsars gave results for the DM and RM that could be compared with previous efforts, where the posterior distributions of the modelling process provide insight into many of the causes of discrepancy between

these values for the different modelling approaches. Aligning the PA was found to result in misalignment of the total power for both pulsars, similar to the effect seen in the height-varying simulation. I converted the difference in alignment between PA and intensity profiles into an estimate of the emission height range for the observing band for each pulsar, and found that the results were compatible with theoretical expectations.

Correcting for the effect of the interstellar medium makes it possible to understand in detail the broad-band behaviour of pulsar emission, and so relate that to the structure of the pulsar radio beam and its emission processes. This work also highlighted the importance of understanding the impact of the idiosyncrasies of individual pulsars when using them as tools to understand the interstellar medium, particularly in the way that modelling accuracy is so dependent upon key attributes of the pulsar in question, such as the geometry underlying its PA profile shape. The theoretical definitions of the DM and RM are straightforward, but in practice measuring these values requires a set of assumptions about pulsar emission behaviour and how it varies with frequency. It is important to recognise and understand these assumptions for what they are, so that understanding about both the ISM and the pulsars themselves can progress.

Chapter 7

Scattering analysis of single-component pulsars with the MeerKAT telescope

In the last chapter I emphasized the importance of understanding how both the pulsar itself and the effect of the ISM have an impact upon the radio pulse profiles we observe. In that chapter I focused on the effects of RM and DM, which must be calculated and corrected for every pulsar observation. Many pulsar observations are also affected strongly by interstellar scattering, and so it is important to consider its impact. The effect of plasma structures in the ISM lying along the path of the radio beam is to scatter the radio waves, so that they travel along deflected paths, and the resultant pulse profile is smeared out over time. The effect of this scattering, provided it is the right strength to be observable in this way, is that the profile has a characteristic exponential tail. By studying the properties of a scattered pulse profile, it is possible to relate its shape to the structures in the ISM that caused the scattering.

To do so, however, requires accurate understanding of the pulsar emission itself, so as to be able to account for how the intrinsic profile shape, combined with the effects of scattering, results in the profile observed. Studying scattered pulsars in the time domain therefore gives new information, both about pulsars themselves (open questions 1a-c) and about the ISM (Q2a,b).

In order to make the most progress in understanding the ISM through observations of scattered pulsars, we require surveys with high quality observations of a large number of pulsars. The Thousand-Pulsar-Array (TPA) project on the MeerKAT telescope presents just such an opportunity. The TPA (Johnston et al., 2020) is an observing project that is being carried out as part of MeerTime, a large scale project to observe known pulsars using the MeerKAT telescope (Bailes et al., 2020). MeerKAT

provides high sensitivity in the observing band 856–1712 MHz, which results being able to measure the scattering properties from high quality broad-band observations that have been obtained using a single instrumental set-up.

Since scattering scales strongly with frequency and has the greatest effect at low frequencies, the MeerKAT observing band is, from a scattering perspective, relatively high frequency. This is relevant for considering which pulsars will have observations that are measurably scattered at this observing band. Broadly, there are two factors at play that affect scattering strength of a given observation. The first is frequency: the lower the frequency, the more strongly scattered the pulse profile. The second is distance. As a general rule, if a pulsar is further away from Earth, the radio waves must travel through more of the ISM before arriving at the telescope, increasing the possibility for the radio waves to be more strongly scattered. The ISM is not isotropic – as evidenced immediately by the fact that not all pulsars are observed to be scattered – however it is observed that scattering timescale tends to increase for larger DMs, where DM can be taken as a rough proxy for distance.

These two factors (the effects of frequency and distance/DM) combined mean that only certain subsets of scattered pulsars will be observable as such for a given frequency band, as depicted in Fig. 7.1. Low frequency observations will favour nearby pulsars with smaller DMs, since more distant pulsars may be so strongly scattered that there is a loss of signal, making the pulsar impossible to observe. Higher frequency observations will favour distant pulsars, since nearby pulsars may be scattered so weakly as to appear not scattered at all at the observing band. The MeerKAT scattered pulsar sample will therefore correspond to pulsars with larger DMs.

What this means is that there is a high possibility of there being pulsar observations where the pulsar is scattered, but weakly enough that the scattering timescale cannot be measured accurately. Such edge cases have the potential to skew the conclusions drawn about the ISM. It is therefore very important to consider the ways in which modelling might break down in order to ensure that results are not biased.

7.1 Modelling process

As described in Chapter 3.3.3, the dataset to which I applied the scatter modelling was a set of 205 single-component pulsars observed with the TPA programme, selected by eye. The observations consist of 8 pulse profiles across the band 895–1670 MHz, and I modelled the scattering parameters for each independently. I used the *train + DC*




Figure 7.1: Illustration of scattering strength dependence on frequency and distance. This illustration indicates qualitatively how scattering strength, and hence profile distortion, tends to vary with frequency (x axis) and distance to the pulsar (y axis). Since scattering is stronger at lower frequencies and at larger distances, observations with one or other of these attributes tend to have observably distorted profiles. Distant scattered pulsars observed at low frequencies are likely to suffer from a loss of S/N, and nearby pulsars at higher frequencies may be too weakly scattered for time domain scatter modelling to accurately constrain the scattering timescale. Figure made using Inkscape.

method as described in Geyer & Karastergiou (2016), updating the fitting process to sample the log-likelihood function over parameter space using the MCMC algorithm EMCEE (Foreman-Mackey et al., 2013). This models the scattered pulse profile as a single Gaussian component representing the intrinsic emission, convolved with an exponential scattering function. I assumed that the scattering material took the form of a thin screen with isotropic scattering properties, so that the scattering transfer function takes the form given in equation 4.1. Early modelling also investigated the applicability of an extreme anisotropic model, as given by equation 3 in Geyer et al. (2017). However, since this was found to be generally less successful at fitting the profiles than the isotropic model, I did not take this further.

Other models used in time domain scatter modelling have included similar forward modelling to *train* + *DC*, but describe the intrinsic profile shape using a high frequency non-scattered profile for that pulsar (Krishnakumar et al., 2019), or using deconvolution analysis, such as the CLEAN algorithm (Högbom, 1974; Bhat et al., 2003). These methods have the advantage of being applicable to pulse profiles with multiple components, which cannot be successfully modelled with a single Gaussian. However, using a high-frequency profile does not allow measurement of profile evolution, whilst applying the CLEAN algorithm requires a more complex set of assumptions in rather than free convergence onto the best solution. It is important with any modelling to understand both the possibilities and the limitations afforded by the choice of model being applied, hence my focus here on understanding the circumstances under which my scatter modelling might break down. The advantage of *train* + *DC* is that it does not constrain the intrinsic pulse profile, which is known to evolve with frequency.

7.2 Applicability of the model to the dataset

The MCMC modelling process resulted in 142 pulsars for which there were at least 4 profiles across the band with a successful fit. Comparing the results across frequency for each pulsar narrowed this down to 84 pulsars for which I could be confident in the fit. The first test is to check the diagnostic output of the methodology showing τ as a function of frequency, which resulted in my rejecting all 21 sources which showed no significant evolution of τ across the band. Even though some pulsars matched the selection criterion of a Gaussian profile convolved with an exponential, they showed no frequency evolution consistent with scattering. There is an expectation that time domain fits will break down when τ is too small to measure. A further two pulsars

(J1320–3512 and J1738–3211) were discarded due to the model being visibly inaccurate at the leading edge of the profile. The rising edges of the model and profile respectively did not overlap for these profiles. PSR J1738–3211 has an asymmetric profile with the trailing edge steeper than the rising edge, implying that it is not scattered and should be discarded regardless. PSR J1320–3512 has a leading edge that is shallower than the model leading edge. This suggests either that the intrinsic profile shape is not well modelled by a Gaussian, or that the scattering function for this pulsar might be better described by a thick screen model (Williamson, 1972; Kirsten et al., 2019), which has a more gradual leading edge shape than the thin screen model applied in this work. I also rejected PSR J1655–3844 after closer inspection at high frequencies revealed it not to fit the single-component criterion.

The reasons for rejecting the remainder of the sample are best explained through comparison of some example results, both for a successful fit, and for cases where the results of the fit are unreliable. Fig. 7.2 shows the eight observed pulse profiles of PSR J1818–1422 with the best fit model at each frequency overlaid. An example of how the parameters for the model were obtained is given by Fig. 7.3, which represents the MCMC chain for the profile at 0.94 GHz as a corner plot. Fig. 7.4 shows the corner plot for the power law fit of $\tau = A\nu^{-\alpha}$, which gives $\alpha = 3.787 \pm 0.008$. The figure shows that A and α are covariant, but that α is nevertheless tightly constrained.

The top subplot of Fig. 7.5 shows how both τ and σ evolve with frequency for this pulsar, with the power law fit indicated with a straight line. In addition, the modelling indicates intrinsic profile evolution, as σ increases with decreasing frequency below ~ 1 GHz, evolution comparable to that described for many other pulsar observations (see for example Thorsett, 1991).

Whereas PSR J1818–1422 was successfully modelled, the other subplots in Fig. 7.5 show results for unreliable modelling processes. The results for several pulsars, particularly those that are less strongly scattered, exhibited an inter-relation between the measured values of σ and τ that is not inherent to the pulsar and must result from the fitting process. An example is PSR J1743–3153, shown in the middle subplot of Fig. 7.5. The decrease in τ with frequency seen for the first three measurements is mirrored by an increase in σ . The converse is then true for the next measurement: a drop in σ is counteracted by an increase in τ . Similar mirroring behaviour can also be seen in the measurements for the highest frequencies. This suggests that the modelling process has a covariance in σ and τ : multiple different pairings of these parameters all generate similarly shaped profiles and are therefore indistinguishable. Since each sub-band was modelled individually, the inaccuracy of model parameters

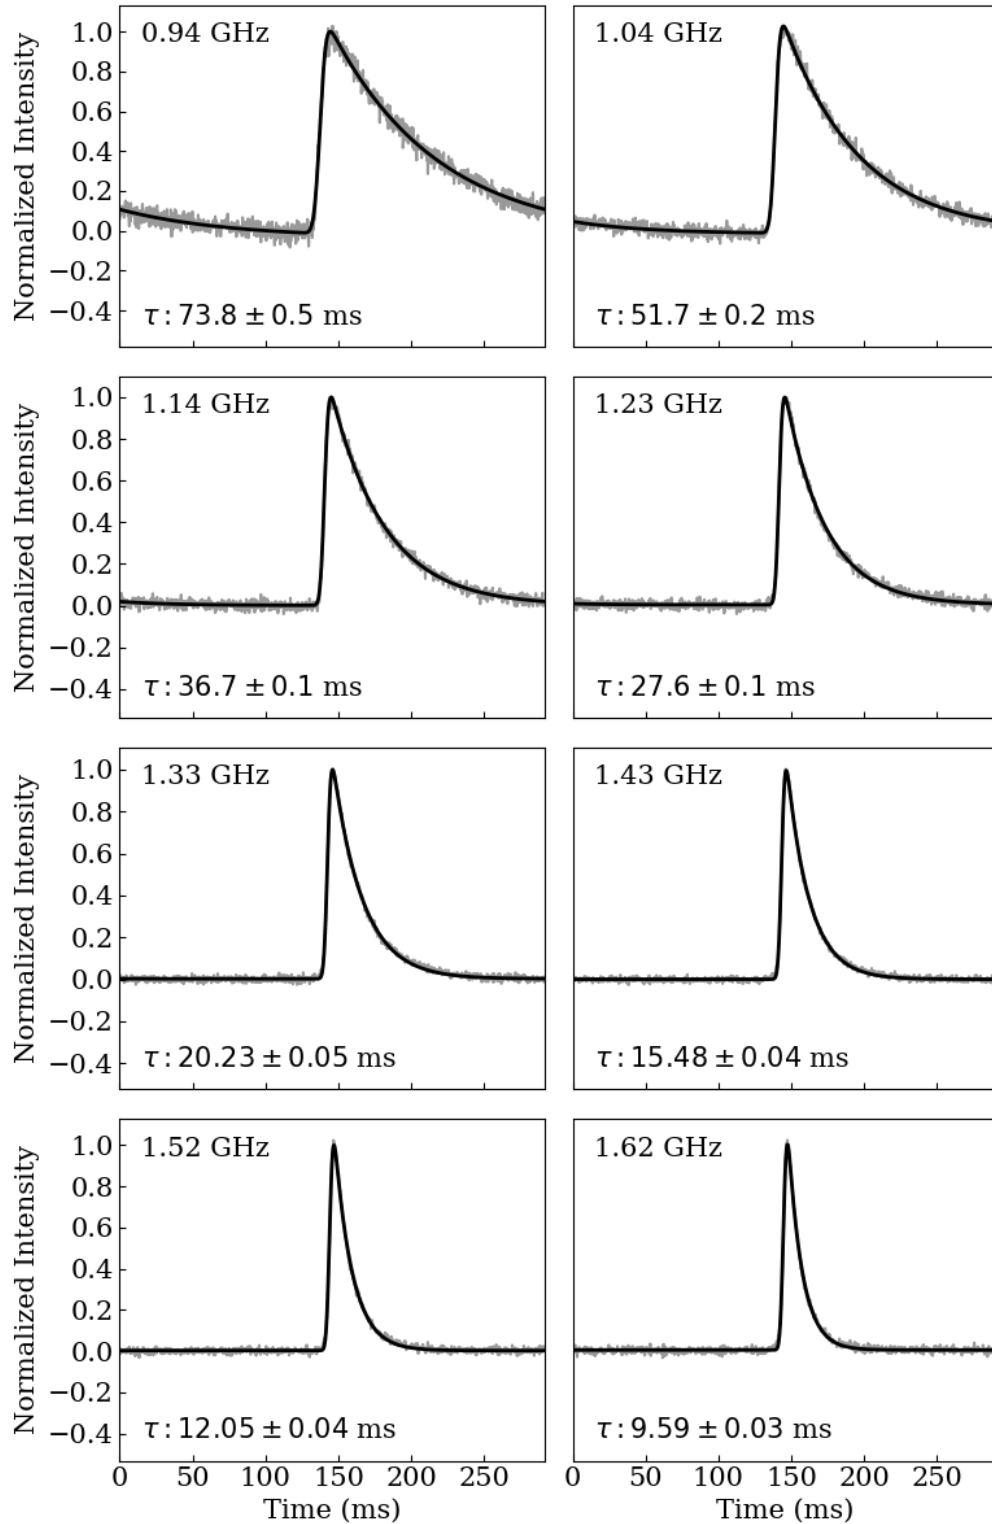


Figure 7.2: PSR J1818–1422 scatter modelling results: fits to profiles. Profiles of PSR J1818–1422 at 8 frequency channels across the observing band, plotted in grey, with the best fit scattering models over-plotted in black. The corrected frequency f_m and scattering timescale, τ , are shown on each profile. The profiles are ordered by increasing frequency, reading from left to right and top to bottom. Figure made with matplotlib.

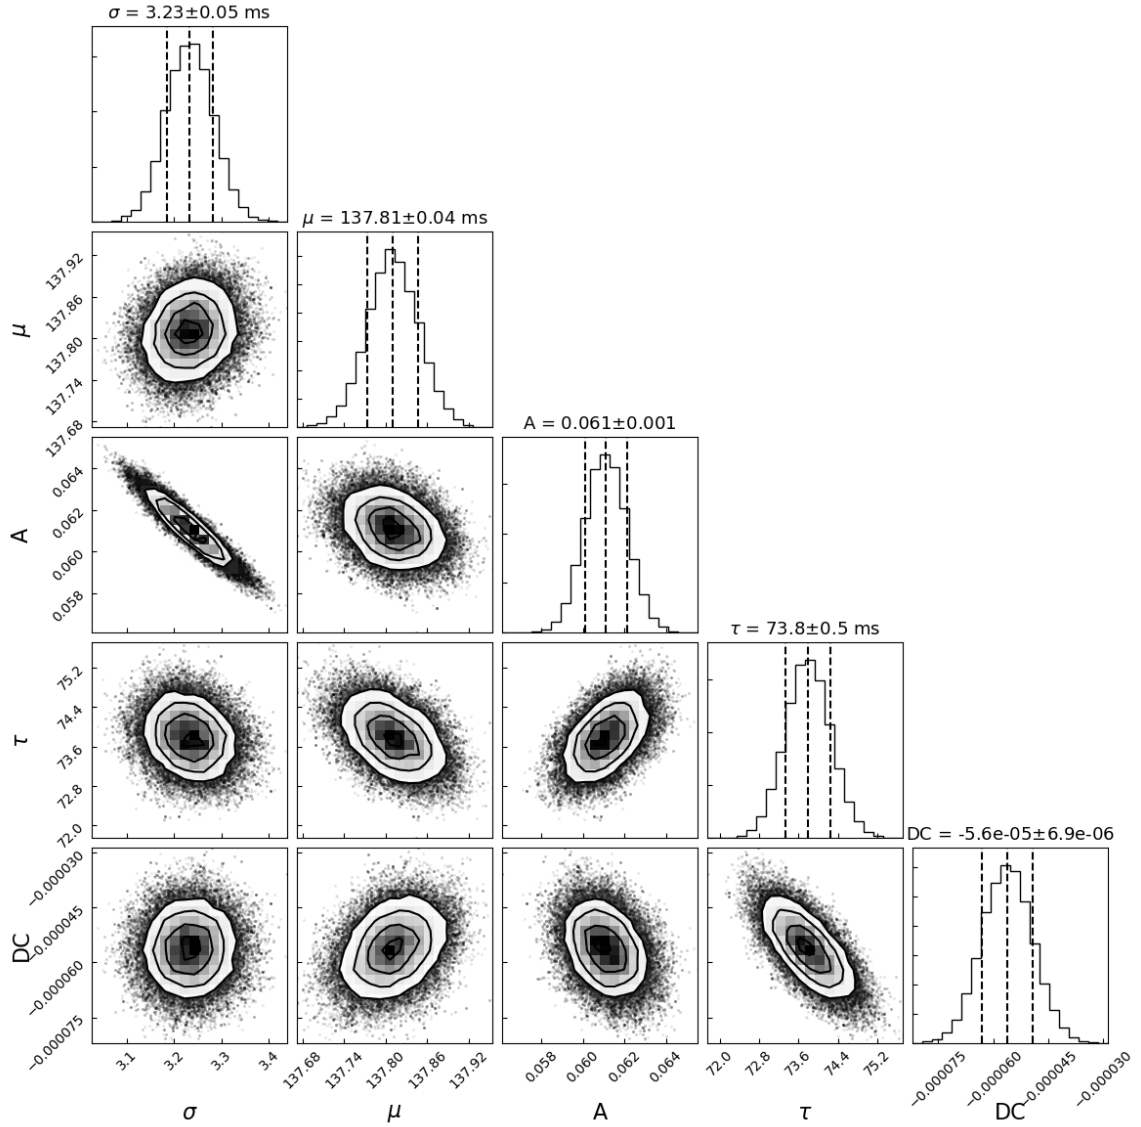


Figure 7.3: PSR J1818–1422 scatter modelling results: corner plot.

Corner plot for the scatter modelling parameters of the lowest frequency (950 MHz) profile of PSR J1818–1422. In each 2D subplot the black points indicate all of the sample values explored by the chains. On top of these points is plotted a 2D histogram and a contour plot to indicate the overall density distribution. The five model parameters are the standard deviation σ , mean μ and amplitude A of the Gaussian, the scattering timescale τ and any DC offset of the pulse baseline from 0. I mark the 16th, 50th and 84th percentiles on the marginalized histograms. Figure made with corner (Foreman-Mackey, 2016).

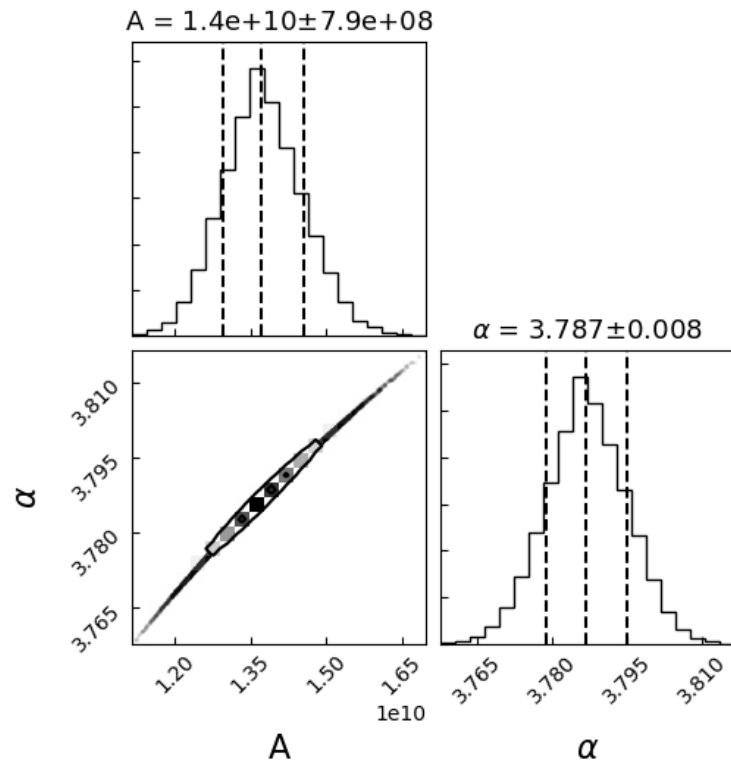


Figure 7.4: PSR J1818–1422 scatter modelling results: corner plot for α . Corner plot for the power law fit parameters for scattering timescale τ against frequency ν , where $\tau = A\nu^{-\alpha}$, α is the spectral index we wish to measure and A is the constant of proportionality. The corner plot is formatted in the same way as that in Fig. 7.3. Figure made with corner (Foreman-Mackey, 2016).




Figure 7.5: Scatter modelling results: τ and σ vs. frequency for PSRs J1818–1422, J1743–3153 and J1653–4249.

Log-log plots of scattering timescale τ (crosses) and intrinsic Gaussian standard deviation σ (dots) against frequency for three pulsars, with the power law best fit to the scattering timescale plotted as a black line. Top: PSR J1818–1422, an example of successful modelling. Middle: PSR J1743–3153, an example of $\sigma - \tau$ mirroring; this pulsar and 32 similar sources are rejected from the final sample. Bottom: PSR J1653–4249, an example of inaccurate modelling at high frequencies. For this pulsar I use only the lowest 5 frequencies to calculate the power law relationship; similar rejection of high frequency model results is done for a further 14 pulsars. Figure made with matplotlib.

converged upon for a given sub-band is only apparent when compared to the results for the other sub-bands, and this $\tau - \sigma$ covariance is revealed. It seems likely that PSR J1743–3153 is scattered, but the mirroring behaviour of σ and τ means that the accuracy of the measurements cannot be trusted. I therefore excluded 33 pulsars showing this behaviour, as identified by visual inspection of similar figures, from further analysis.

PSR J1653–4249, for which τ and σ are plotted against frequency in the bottom subplot of Fig. 7.5, is one of 15 pulsars which show a visible deviation from a single power law at the higher frequencies, with τ larger than expected, something I will refer to as spectral flattening. I hypothesized that this may be caused by the intrinsic profile, and not the scatter broadening, being the dominant factor in the overall pulse shape at high frequencies for these pulsars. If so, the assumption of a single Gaussian component would be insufficient to replicate the profile shape, resulting in over-estimates of τ . I tested this hypothesis by investigating the theoretical model applied and by testing a simulation, details of which are in the next section, and found that the results from these tests supported the hypothesis. For these 15 pulsars therefore, I kept only that frequency range where τ follows a power law, selecting the values through visual inspection, and inferred α from those. This resulted in 3 more pulsars being removed from the sample, due to having fewer than 4 successful results for τ . Hence the final successfully modelled sample consisted of 84 pulsars.

7.3 Understanding model break-down

I looked to quantify theoretically the circumstances under which we would expect τ – σ covariance or spectral flattening to appear in when modelling scattered pulse profiles. I considered the functional form of the model being applied—the convolution of a Gaussian with isotropic exponential decay—for which intensity I as a function of time t is given as follows:

$$I(t) = \frac{1}{2\tau} \exp\left(-\frac{t}{\tau} + \frac{\sigma^2}{2\tau^2}\right) \left(1 + \operatorname{erf}\left(\frac{t}{\sqrt{2}\sigma} - \frac{\sigma}{\sqrt{2}\tau}\right)\right). \quad (7.1)$$

As σ tends to 0, this equation tends to pure exponential decay, as expected. The extent to which σ is relevant to the shape of the profile depends on the error function term $x = \left(\frac{t}{\sqrt{2}\sigma} - \frac{\sigma}{\sqrt{2}\tau}\right)$. For a given ratio of σ/τ , the Gaussian part of the model will make a large contribution to the overall model shape up to some time t_0 , after which the exponential decay will dominate the pulse profile shape. I define t_0 as the time that, for a given ratio of σ and τ , gives $x = 2$. This is based on the shape of

the error function, since at $x = 2$ the functional form starts deviating strongly from $\text{erf}(x) = 1$.

In Fig. 7.6 I have plotted equation 7.1, along with its constituent Gaussian and exponential components, for 4 different values of σ/τ . For each, I have shaded that part of the profile where $t < t_0$. As can be seen in the figure, a small σ/τ results in small t_0 , so that the majority of the pulse profile shape ($t > t_0$) is dominated by exponential decay and τ can be determined easily. Conversely, if σ/τ is large, then one may expect to see covariance between the measured values of σ and τ .

For the purposes of understanding the comparative contributions of the Gaussian and the scattering transfer function, I defined the Gaussian contribution as too large when t_0 coincides with the point at which the total intensity has dropped to 10% of its peak. This 10% intensity point is marked with a vertical line in Fig 7.6. This means that the cut-off ratio of σ/τ (for which the underlying pulse profile shape interferes too much in the scattering measurements to be able to reliably separate σ from τ) is $\sigma/\tau \sim 1$.

Fig. 7.7 shows the histograms of values of σ/τ measured for two groups of pulsars. First, for the pulsars for which I can be confident of the fit at all frequencies, I have calculated σ/τ at the highest observing frequency and plotted the values as a shaded grey histogram. For the pulsars showing the spectral flattening behaviour, I have taken σ/τ at the highest frequency for which the values are still following power law behaviour. I am treating this as the cut-off point: at frequencies above this τ can no longer be measured accurately, so it is expected that the ratio σ/τ should be larger for these values than for cases where there is no problem measuring τ . The histogram of these values is plotted in Fig. 7.7 as a stepped histogram. It can be seen that the majority of pulsars with no spectral flattening have $\sigma/\tau < 0.7$, which is in accordance with the theoretical estimates that the parameters can be measured accurately for smaller values of σ/τ . By contrast, the histogram of cut-off values peaks at $\sigma/\tau \sim 0.8$, implying that the flattening seen at higher frequencies is indeed due to the loss of ability to separate the intrinsic profile shape from the exponential scattering. My theoretical estimate of $\sigma/\tau = 1$ being the cut-off for successful modelling is intended only to be indicative, since it takes no account of the variety of signal-to-noise ratios or intrinsic profile shapes of the observations. It is therefore unsurprising that there are large spreads in both histograms in Fig. 7.7. In particular, the left-hand side of the stepped histogram is consistent with explaining those pulsars which, like PSR J1653–4249, may have extra components that are altering the scattering behaviour of the pulsar in comparison to what might otherwise be expected.




Figure 7.6: Illustration indicating relative contributions of intrinsic and scattered components of profile under different conditions.

Plots indicating how the Gaussian and exponential components of the scattering model contribute to the overall model shape, and how that changes with increasing σ/τ . Dashed line: the isotropic scattering transfer function described by equation 4.1. Dotted line: a Gaussian function. Solid line: the convolution of these two, described by equation 7.1. Grey vertical line: this marks the phase point at which the combined model intensity (solid line) drops to 10% of its maximum. Grey shading: this indicates the region where the Gaussian (dotted line) is contributing strongly to the overall profile shape (solid line). The definition of a strong contribution is given in the text. Each subplot shows the same set of functions: the changed curve shapes result from the changed ratio of σ/τ , which is marked in the top right corner of each subplot. Figure made with matplotlib.




Figure 7.7: Histograms of the ratios of measurement results σ/τ . Grey: values for the highest frequency profile fits for the pulsars with a standard power-law τ -frequency relationship. Transparent, black-edge: values for the highest frequency profile fit that is still following the power-law relationship for those pulsars exhibiting tau-flattening. Figure made with matplotlib.

In addition to considering these factors from a theoretical perspective, I attempted to replicate the output of the scatter modelling process of PSR J1653–4249 with simulated data, to test whether the presence of a hidden second component in the pulse profile could be responsible for the spectral flattening behaviour seen in this pulsar’s modelling results. I simulated a pulse profile made up of two Gaussian components: a large main component and a smaller, narrow secondary component that sits to the right of the main peak. On top of this, I introduced profile evolution: the width of the main component decreases with increasing frequency according to a power law plus constant relationship (Thorsett, 1991), and the flux spectral index of the secondary component is flatter than that of the main one, reflecting what is commonly seen in observations. I defined the scattering timescale at the lowest frequency to be the same as that measured for the data. I then applied a scattering spectral index of $\alpha = 4.4$ to τ to obtain its value at other frequencies (ν), since $\tau \propto \nu^{-\alpha}$. I defined the height, width and position of the second component such that the overall profile shapes of the simulation appear similar to those of the real pulsar. Fig. 7.8 shows how the two components combine to make the simulation profile shapes at the lowest and highest frequencies.

Performing the MCMC fit on this simulation, I obtained results for τ and σ that I




Figure 7.8: Simulation: double-component scattered pulse profile at two frequencies. Profiles for a simulation of a scattered double-component pulsar, showing the lowest (top) and highest (bottom) frequency simulated profiles respectively. Solid line: the simulated profile. Dashed lines: the two noiseless components that comprise the simulation. These are added together and then noise is added to generate the simulated profile. The subplots are zoomed in to show only the on-pulse region. Figure made with matplotlib.




Figure 7.9: τ and σ vs. frequency for PSR J1653–4249 and for the simulated double-component pulse profile.

Log-log plot of scattering timescale τ and intrinsic Gaussian standard deviation against frequency, similar to Fig. 7.5. Black: scatter modelling results for PSR J1653–4249, showing τ (crosses) and σ (points). Blue triangles: modelling results for a the simulated scattered double-component pulsar shown in Fig. 7.8, showing τ (upright) and σ (upside-down). Blue dashed lines: input values of τ and σ used to generate the simulated pulse profiles. Figure made with matplotlib.

have compared to those obtained from modelling the data in Fig. 7.9. The presence of the second component does little to alter the fit parameters at the low frequency end, where the simulation parameters are recovered well, but at higher frequencies there is a flattening off in the spectral behaviour of both τ and σ that mimics the measurement results for PSR J1653–4249. This lends credence to my choice, for pulsars like PSR J1653–4249, to keep only those values of τ where a power law is still being followed.

These results indicate the extent to which individual measurements of τ of single profiles are vulnerable to many sources of bias and error. Scattering properties of pulsars can only be characterized reliably in cases where scattering measurements can be performed across large frequency bands. Further, the scattering results for a single pulsar, and the science that can be inferred from them, are best understood in the wider context of all of the other pulsars observed and analysed with the same method.

7.4 Results of scatter modelling

Fig. 7.10 shows histograms of the distributions of α obtained for the survey. It shows the final sample of 84 pulsars (hatched histogram) and the 33 pulsars which were filtered out due to problems in the modelling (histogram with circle pattern). These have primarily been removed due to covariance between the values of τ and σ , but this subset also includes the two pulsars rejected due to the leading edge not being correctly modelled (PSRs J1320–3512 and J1738–3211). The removal of those pulsars was done without reference to the value of α obtained, so that it is not biased in favour of those pulsars with values of α close to what might be expected from theory. The result of the filtering is that pulsars with small values of α are largely removed from the sample, and those remaining show a near-symmetric distribution. The causes of inaccurate scatter modelling tend to mean both that poor model fits will overestimate τ , and that poor model fits are more likely for smaller values of τ , which are seen more often at higher frequencies. It is therefore expected that poorly modelled pulsars will tend to underestimate α , and so it is unsurprising that the majority of the poorly modelled pulsars returned small values of α .

Fig. 7.10 also shows a histogram of the values of α found in the literature for cases where pulsar profiles observed at ≥ 400 MHz have been fitted with an isotropic time domain model (Cordes et al., 1985; Löhmer et al., 2001, 2004; Lewandowski et al., 2011, 2013, 2015a,b; Krishnakumar et al., 2017, 2019). The histogram of literature values has been weighted so that it represents the same total count of pulsars as those shown for this work. The ≥ 400 MHz cut is relevant because different behaviour has been observed in scattering surveys at lower frequencies. The reasons behind this will be discussed in the next chapter.

Comparing my distribution of α with that of previously published values, it can be seen that, when I include the poorly modelled pulsars, the distributions are similar. The population of pulsars around $0 < \alpha < 3$ shown in the published values has been filtered out from my results distribution. Performing a Kolmogorov-Smirnov (KS) test to compare my sample to that from the literature gives the following values. Using just the final filtered sample in the test gives a KS statistic of 0.247 with a corresponding p-value of 0.001. Adding in the poorly modelled pulsars as well gives KS statistic 0.109 and p-value 0.335. This indicates that, whereas the final filtered sample is clearly not drawn from the same distribution as that of the combined literature values, when the poorly modelled pulsars are included, the null hypothesis cannot be rejected.




Figure 7.10: Histograms of values of α .

Top right (grey): previously published values of α obtained through time domain scatter modelling. Middle right (transparent, circle pattern): values of α for the pulsars rejected from our sample due to evidence of inaccurate modelling. Bottom right (transparent, hatched): histogram of α values for our successful sample of 84 pulsars. Left: the same three histograms, now overlaid (and, for our observations, stacked) for visual comparison. Figure made with matplotlib.

Fitting a Gaussian to the filtered distribution of α gives a mean of 4.0 and a standard deviation of 0.6. The theoretical values of both $\alpha = 4$ and $\alpha = 4.4$ fall within one standard deviation. Fig. 7.11 shows α plotted against τ_0/P , where τ_0 is τ measured for the lowest observing frequency of 950 MHz. It is interesting to note that if I split the population into sources where τ_0 is less or greater than 10% of the period, I find a mean α value for each group of 3.7 ± 0.6 and 4.1 ± 0.4 respectively. For $\alpha = 4$, a pulsar with $\tau_0/P = 0.1$ at the lowest frequency will have $\tau/P \sim 0.01$ at the highest frequency. Small values of τ tend to be overestimated by inaccurate modelling. Fig. 7.11 shows that the sources for which $\tau_0/P < 0.1$ are comprised of pulsars that I have filtered out as being either poorly modelled (orange diagonal crosses), or not scattered (blue vertical crosses), or they are not filtered out (black points) and yet still show systematically lower values of α , as described above. This graph is further evidence that the modelling results in systematically smaller estimates of α for such cases, which is a limitation of the time domain methodology.

I scaled the scattering timescale by the period for the analysis above because all of the observations have the same number of bins across the pulse period (1024 bins). The number of bins containing information about the profile shape will affect the accuracy of the modelling, and a profile that takes up a smaller fraction of the pulse period (smaller τ/P) will have fewer bins spanning the profile. It is therefore the ratio $\tau/P < 0.1$ that is important in terms of likelihood of the model being able to correctly capture the scattering behaviour of a given pulsar. However, the time resolution of the bins is also relevant to consider. The larger a pulsar's period, the lower the time resolution of a single bin, meaning that less information is captured in each bin. I found that the distribution of periods for the $\tau/P < 0.1$ cases is shifted slightly towards longer periods than the distribution for the $\tau/P > 0.1$ cases, which reflects the time resolution limitation.

The scatter modelling process identifies the position of the underlying Gaussian that describes the intrinsic profile in our model (μ). I fit for a correction to the DM that aligned these values across frequency in order to obtain the DM that best aligns the intrinsic pulse profile independent of the effects of scattering. Fig. 7.12 shows a histogram of the measured corrections to the DM. The majority of these are small and negative, as expected. As discussed in the previous chapter in the context of PSR J1359–6038, the value of DM calculated for a given pulsar will depend on the method chosen, and since scattering shifts the intensity to later phases, the DM that maximises the S/N will not be the same as the DM which aligns the unscattered pulse profile across frequency.




Figure 7.11: α vs. τ_0/P for good and bad model fits.

Plot of scattering spectral index α against scattering timescale at the lowest frequency channel of 950 MHz, τ_0 , where τ_0 is scaled by pulse period P . Blue vertical crosses: non-scattered pulsars rejected from the final sample. Orange diagonal crosses: pulsars rejected due to evidence of inaccurate modelling. Black points: final filtered sample of successfully modelled scattered pulsars. Figure made with matplotlib.




Figure 7.12: Histogram of ΔDM .

Histogram of the values of ΔDM calculated for our filtered scattered pulsar sample, where ΔDM is the difference between the original DM used to dedisperse the pulsar, and the best fit DM generated by our scattering model. Figure made with matplotlib.

The tail of large negative ΔDM values in the figure is partially attributable to the small subset of pulsars that are more strongly scattered and therefore require larger corrections to the DM to align intrinsic profiles. However, it also reflects the distribution of periods in the pulsar sample. This can be understood as follows. We can assume very roughly that pulsars have the same duty cycle. For profiles at two discrete frequencies that are misaligned by an amount on the order of the pulse width, the shift required to align the profiles will be the same fraction of the pulse period. For pulsars with very different periods, this shift will comprise different absolute lengths of time; correspondingly the DM correction to perform this shift will be larger for the pulsar with the larger period.

Since the choice of DM affects the profile shape within each sub-band, a question arises as to whether correcting the DM would change the measurements of τ . I tested this on PSR J1630–4733, chosen because its ΔDM of $11.3 \text{ cm}^{-3} \text{ pc}$ results in the largest relative shift of profiles at different frequencies in the sample set. PSR J1850–0006 has a larger absolute ΔDM of $-21.0 \text{ cm}^{-3} \text{ pc}$, however, since its period is also larger, the corresponding time shift between profiles at different frequencies is a smaller fraction of the period and so the relative misalignment of profiles is smaller. I identified the magnitude of the DM correction required through scatter modelling,




Figure 7.13: $\Delta\tau$ vs. frequency for PSR J1630–4733: correction to τ resulting from DM correction.

Plot of $\Delta\tau$ vs. frequency for PSR J1630–4733, where $\Delta\tau$ is the difference between the scattering timescales calculated for the data processed with the initial, visibly incorrect, DM, and when reprocessed with the best fit DM produced by the scattering model. Figure made with matplotlib.

then re-processed the data with the new DM and then repeated the scatter modelling. The results for τ obtained before and after the reprocessing are compared in Fig. 7.13. The difference is largest at low frequencies, as expected. It is also uniformly negative, meaning that the scattering timescales in the uncorrected DM case are larger. However, even for this pulsar, where the change in DM is most extreme (shifting the top of the observing band with respect to the bottom of the band by 6% of the pulse period), the measured values of τ are still equivalent to within 1σ in all but the third channel (counting from lowest to highest frequency), and equivalent to within 3σ for all channels.

7.5 Table of results and graphs

Table 7.1 lists, for each modelled pulsar, the scattering timescale τ at 1 GHz, along with α and the corrected dispersion measure (DM). I have also indicated in the table which of the 8 frequencies have been used to calculate these values, indicating which

channels were excluded due either to failing to reach convergence on a solution, or to deviating from a power law at high frequencies, as previously explained. I used the MCMC power law fit $\tau = A\nu^{-\alpha}$ to compute the best fit value of τ at $\nu = 1$ GHz (where A is just the constant of proportionality). I did this by calculating the value of τ at 1 GHz associated with every pair of values of A and α explored by the MCMC chain. I then took the 50th quantile of these values as τ and the difference between the 16th and 84th quantiles for its uncertainty. The graphs associated with all 84 modelled pulsars are found in Appendices A and B. The first appendix shows the eight pulse profiles spanning the observing band with the best fit models overlaid, as in Fig. 7.2. Appendix B shows the frequency dependencies of τ and σ , formatted as for those in Fig. 7.5.

Table 7.1: Values for τ at 1 GHz, α and DM for our filtered sample of scattered pulsars. The values and their uncertainties are calculated as described in the text. The DM is the value that best aligns the modelled intrinsic profile at each frequency. I indicate which of the 8 sub-bands (ordered in increasing frequency) were used to generate these values with ones, and those not included with zeros. The minimum allowed number of sub-bands is 4. Reasons for exclusion of sub-bands (failure of model convergence or deliberate exclusion) are explained in the text.

PSRJ	Which channels	τ at 1 GHz (ms)	α	DM (cm⁻³ pc)
J0646+0905	11110110	10.03±0.03	3.46±0.05	147.85±0.07
J1055−6028	11110000	2.36±0.08	4.5±0.4	636.9±0.1
J1112−6103	00111111	33±3	3.8±0.2	595.69±0.09
J1114−6100	11101101	28.3±0.1	3.14±0.02	676.7±0.1
J1138−6207	11111111	26±2	3.8±0.2	518.86±0.05
J1305−6203	11110000	10.4±0.2	3.1±0.2	468.78±0.05
J1316−6232	00001111	1099±25	4.435±0.003	966.4±0.3
J1319−6105	11110000	7.05±0.08	3.5±0.1	440.6±0.1
J1341−6220	11111111	23.5±0.2	4.04±0.02	718.26±0.06
J1349−6130	11111100	6.4±0.1	3.9±0.1	283.87±0.04
J1406−6121	00011111	48±1	4.751±0.004	537.29±0.08
J1412−6145	11111111	19.9±0.5	4.0±0.1	512.48±0.08
J1413−6141	11111111	42±2	4.1±0.1	667.60±0.06
J1511−5835	11111100	23.3±0.5	4.8±0.2	329.4±0.2
J1512−5759	11111111	7.14±0.02	3.16±0.01	626.90±0.02
J1514−5925	00111111	16.7±0.6	5.107±0.005	192.51±0.08
J1519−5734	01011101	118±7	3.9±0.2	654.7±0.4
J1538−5551	11111111	22±2	4.0±0.2	602.67±0.06
J1543−5459	11111111	36.2±0.8	3.76±0.08	344.99±0.03
J1551−5310	11111111	131±5	4.4±0.1	485.8±0.6
J1610−5006	11111000	124±1	4.670±0.002	410.0±0.2
J1630−4719	11111110	5.71±0.09	3.3±0.2	487.2±0.1

Continuation of Table 7.1

J1630-4733	11110000	340±2	5.049±0.001	509.3±0.2
J1632-4621	11111110	16.85±0.07	3.87±0.03	559.83±0.07
J1633-4453	11111111	19.59±0.08	3.99±0.02	472.22±0.03
J1638-4608	01111011	17±1	4.6±0.2	422.45±0.10
J1640-4715	11111111	56±1	4.45±0.08	581.0±0.2
J1640-4951	11111111	14.4±0.6	4.5±0.3	407.4±0.2
J1650-4341	00011111	66±7	4.4±0.3	672.8±0.4
J1653-4249	11111000	20.9±0.2	4.17±0.05	415.36±0.06
J1700-4422	11101100	75±4	4.7±0.5	404.9±0.9
J1702-4128	11111111	25.5±0.7	3.97±0.08	365.75±0.08
J1707-4053	11111011	95.2±0.2	3.950±0.007	351.78±0.08
J1715-3859	11111111	228±9	4.1±0.1	806.2±0.6
J1717-3425	11011010	19.82±0.02	3.466±0.007	583.46±0.01
J1717-3737	11111111	35.7±0.5	3.42±0.06	522.70±0.05
J1719-4006	11101000	5.10±0.09	3.4±0.2	386.12±0.04
J1720-3659	11111111	12.6±0.2	3.63±0.07	378.97±0.03
J1721-3532	11111111	113.4±0.7	4.10±0.02	493.0±0.1
J1724-3149	11111000	40±2	4.928±0.007	400.8±0.3
J1725-3546	11111110	50±2	4.1±0.1	738.3±0.4
J1730-3350	11111111	21.0±0.4	4.00±0.05	260.40±0.04
J1731-3123	11111111	15.1±0.5	1.3±0.1	354.5±0.2
J1739-3131	11111111	68.2±0.2	3.836±0.009	596.9±0.1
J1740-3052	11111000	12.1±0.1	3.83±0.10	738.80±0.05
J1801-2304	01111000	547±10	4.529±0.003	1067.85±0.10
J1811-1736	01111111	42±1	3.24±0.09	473.93±0.04
J1812-1718	11111101	31.8±0.2	3.95±0.03	251.4±0.1
J1812-1733	10101111	102±1	3.40±0.04	509.8±0.1
J1816-1729	11111000	14.3±0.1	4.47±0.07	520.7±0.2
J1818-1422	11111111	59.8±0.2	3.787±0.008	619.65±0.07
J1818-1607	11111100	45.0±0.5	4.41±0.08	699.2±0.8
J1819-1114	11111111	51±2	4.4±0.1	309.7±0.2
J1819-1510	11111111	19.7±0.4	3.99±0.08	418.14±0.04
J1820-1346	11111111	110.9±0.5	3.81±0.01	771.0±0.2
J1822-1400	11111111	6.81±0.09	3.31±0.08	649.27±0.04
J1824-1118	11111111	26.9±0.2	3.77±0.02	601.31±0.06
J1824-1159	11111111	18.8±0.3	3.28±0.06	462.96±0.09
J1824-1423	11111111	12.8±0.3	4.2±0.1	427.63±0.08
J1825-1446	11111110	21±1	4.0±0.2	351.6±0.1
J1832-1021	11111111	12.49±0.07	4.05±0.03	474.14±0.03
J1833-0559	11111111	180±6	3.78±0.09	346.70±0.10
J1834-0731	00101111	122±9	3.9±0.2	288.3±0.4
J1835-0643	11111111	110±3	3.77±0.07	464.8±0.1
J1837-0604	00011101	62±6	5.01±0.02	462±10
J1839-0321	00111111	14±2	5.0±0.4	450.5±0.1

Continuation of Table 7.1				
J1839-0643	11111111	47.3±0.7	4.19±0.05	493.5±0.2
J1840-0559	11111111	23.5±0.4	3.85±0.09	319.1±0.1
J1841-0425	11111110	3.22±0.03	3.29±0.07	324.77±0.02
J1842-0153	11111111	32.7±0.3	3.92±0.05	422.9±0.1
J1844-0030	11111000	12.3±0.3	4.1±0.2	603.2±0.1
J1844-0244	11101110	20.0±0.4	3.31±0.09	422.13±0.06
J1844-0538	11111100	11.73±0.07	4.07±0.03	410.51±0.04
J1846-0749	11111100	7.14±0.06	3.62±0.06	389.13±0.02
J1850-0006	11111111	260±7	4.2±0.1	625±2
J1850-0026	11111000	46.5±0.6	4.946±0.002	948.8±0.2
J1852-0127	01111111	57±3	3.6±0.1	427.9±0.2
J1853+0545	11111111	20.0±0.3	3.13±0.04	197.91±0.03
J1857+0143	11111111	50±2	4.0±0.1	247.9±0.1
J1857+0526	11111111	24.0±0.6	3.78±0.07	464.79±0.05
J1859+0601	00011111	125±26	4.4±0.7	272.4±0.5
J1913+1145	01111111	15.2±0.7	3.9±0.2	642.0±0.2
J1916+0844	11111100	12.8±0.1	4.14±0.07	338.01±0.07
J1928+1923	11111110	48.6±0.5	3.94±0.05	476.4±0.2

7.6 Summary

I used the high-quality pulsar dataset provided by the TPA programme to investigate how the ISM scatters pulse profiles. I paid particular attention to understanding the extent to which my model of pulsar scattering was applicable to the diverse population of pulsars. This meant addressing the ways that both intrinsic pulse shape and strength of scattering had the potential to affect the accuracy of the modelled scattering parameters identified. Having identified and accounted for the factors that caused poor profile modelling, I was left with a set of modelled scattering parameters for 84 single-component pulsars. These have an average scattering spectral index α of 4.0 ± 0.6 , which is consistent with theoretical expectations. Comparing my distribution of α values with those from previously published work suggests that previously measured low values of α may have been subject to similar limitations in terms of modelling accuracy as those that I rejected from my final sample.

Chapter 8

Discussion

The work completed in this thesis has increased understanding of pulsars and used that increased knowledge to understand the ISM. This has been possible due to the new broad-band data available from the latest radio instrumentation. In this chapter I will discuss in more detail the approach of the thesis as a whole, in terms of perspectives on both data and modelling processes. I will then discuss the implications of the results presented in Chapters 5, 6 and 7, in terms of what they mean for our increased understanding of both pulsars and the ISM. I will address the applicability of both the methods and the results more widely within pulsar science, and what the results of this work reveal in relation to the directions that future research should take.

8.1 Frequency evolution of polarimetric data

A key aspect of the data-driven approach of this thesis is the question of how best to analyse new types of data, for which older analysis approaches would be incorrectly targeted and hence insufficient. This question is of most relevance when considering the broad-band polarimetric data products of the Parkes UWL. The observations with which I worked are integrated across time in order to build up as much S/N as possible. They are made up of multiple observations taken at different epochs, summed together. This means that, whereas an individual observation may be heavily affected by RFI at some portion of the observing band, combining it with observations where that part of the band is RFI-free means that a full broad-band observation can be put together. What is lost in terms of being able to understand the time-varying aspect of the data is gained in the broad-band frequency coverage that this approach affords using this instrument.

Understanding frequency evolution of pulsar profiles requires visualizations of the data in such a way that the frequency evolution is the factor foremost in the visualization. For the pulsars presented in Chapter 6 the focus was on the PA and the total intensity, and its alignment across frequency. The waterfall plots can however be applied to visualizing other aspects of pulsar polarization. Of particular interest as a follow-on from the work on linear polarization is to consider circular polarization and its relationship to linear polarization.

8.1.1 The information gain from broad-band observations

In the early stages of my research for this thesis, I worked with archival data from the Parkes telescope to investigate the polarimetric structure of pulse profiles and how these appear at different discrete frequencies. These observations were integrated polarimetric pulse profiles from three discrete frequencies: 691, 1370 and 3107 MHz. My goal was to investigate the extent to which I could predict the broad-band polarization structure of the observed pulsars, using only these discrete profiles and interpolating the expected results for the frequencies in between. This work was superseded by the analysis of the broad-band UWL observations in Chapter 6, however it is useful to present it here in comparison, to demonstrate the increased information offered by the new observations.

I used the method of McKinnon & Stinebring (2000) to decompose the pulse profiles into two orthogonally polarized modes. These modes are both 100% linearly polarized, so that their sum and difference respectively combine to give the Stokes I profile and the linearly polarized component, L . The model assumes that no circular polarization is present. By fitting a RVM curve to the PA profile at each frequency I identified the best alignment of the three PA curves and the phase locations of any orthogonal jumps in the PA curves. Under an assumption of disjoint, 100% linearly polarized modes, the orthogonal jumps mark the phase positions where dominance switches from one mode to the other.

Having aligned my three profiles in phase, I predicted the relative intensity evolutions of the two modes across frequency, using a power law fit to predict the spectral indices of the mode evolution. This meant I could infer the frequency at which I would expect mode dominance to switch (if at all) for a given phase. This is following the discussion of Karastergiou et al. (2005) around the concept of orthogonal polarization modes having different spectral indices, and this affecting the resultant observed profile shape at different frequencies. Fig. 8.1 shows the aligned PA and

Stokes I profiles for PSR J0452–1759 and a waterfall plot of predicted mode dominance. This plot shows the pulse phases where either one mode (plotted in blue) or the other (orange) is dominant, and how this varies across frequency. The bright bands indicate the results from my modelling for the three discrete pulse profiles, and the remainder of the band is predicted from the power law fits. The modelling predicts a widening of the central “mode 2” section with increasing frequency. It is also of note that aligning the PA profiles shifts the Stokes I profiles successively earlier in phase for lower frequencies, echoing the alignment results for PSRs J1056–6258 and J1359–6038 in Chapter 6.

Individual discrete observations can only give an indication of the broad-band behaviour of pulsar radio emission. Also marked on Fig. 8.1 is the bandwidth of the UWL receiver, to show the region over which my predictions with these early data could be compared to reality. Fig. 8.2 shows the reality for PSR J0452—1759, indicating the broad-band evolution of total intensity and polarimetric properties for this pulsar. The observations bear out the expectations of my predictive modelling: the PA is split into four regions of colour across phase, reflecting the mode jumps across the profile, and the central region of the PA widens with increasing frequency, as predicted from the early work through its correspondence with the relative mode dominance. From here on, however, the utility of the early predictions are minimal in comparison to the wealth of information that the new UWL receiver provides.

8.1.2 The effect of circular polarization

My early modelling of discrete pulse profiles took no account of the contribution of circular polarization (Stokes V) to the pulse profile. However the broad-band UWL observations reveal considerable complexity to this aspect of pulsar polarization. The plots for J0452–1759 show that the handedness of Stokes V changes between positive and negative across phase, and also evolves with frequency. Whereas in my earlier work I attributed any deviation of the PA from the RVM curve to be associated with an orthogonal jump of 90° , the broad-band data indicate that the more complex distortions of the PA swing correspond to modulations in Stokes V , and that these evolve smoothly over frequency. I have started investigations into the broad-band polarimetric evolution of pulsars by categorizing by eye the 100 brightest pulsars observed with the P574 programme on the Parkes radio telescope. For this sample set (excluding 8 which were adversely affected by RFI), I characterized PA jumps as




Figure 8.1: PSR J0452–1759: alignment of discrete profiles and predicted mode decomposition.

Three pulse profiles of PSR J0452–1759 from archival Parkes data at three discrete frequencies: 691 MHz (blue), 1370 MHz (orange) and 3107 MHz (green). Top: PA profiles, Middle: total intensity profiles. These have been aligned by fitting the rotating vector model to the PA and identifying the closest alignment of each PA profile to that rotating vector model curve. This results in the total intensity profiles being visibly misaligned. Bottom: predicted evolution of orthogonal mode dominance with frequency from these three profiles. The aligned profiles were decomposed into two orthogonal modes and the spectral index of the relative flux evolution of the mode with frequency was calculated at each phase bin. The relative dominance of either mode 1 (blue) or mode 2 (orange) was therefore predicted across phase and frequency. Bright colours: calculated mode dominance for the three discrete profiles. Pastel colours: predicted mode dominance at other frequencies. Hatched grid: continuous bandwidth of the new Parkes UWL receiver which then enabled more detailed study of frequency evolution. Figure made with matplotlib.

J0452-1759

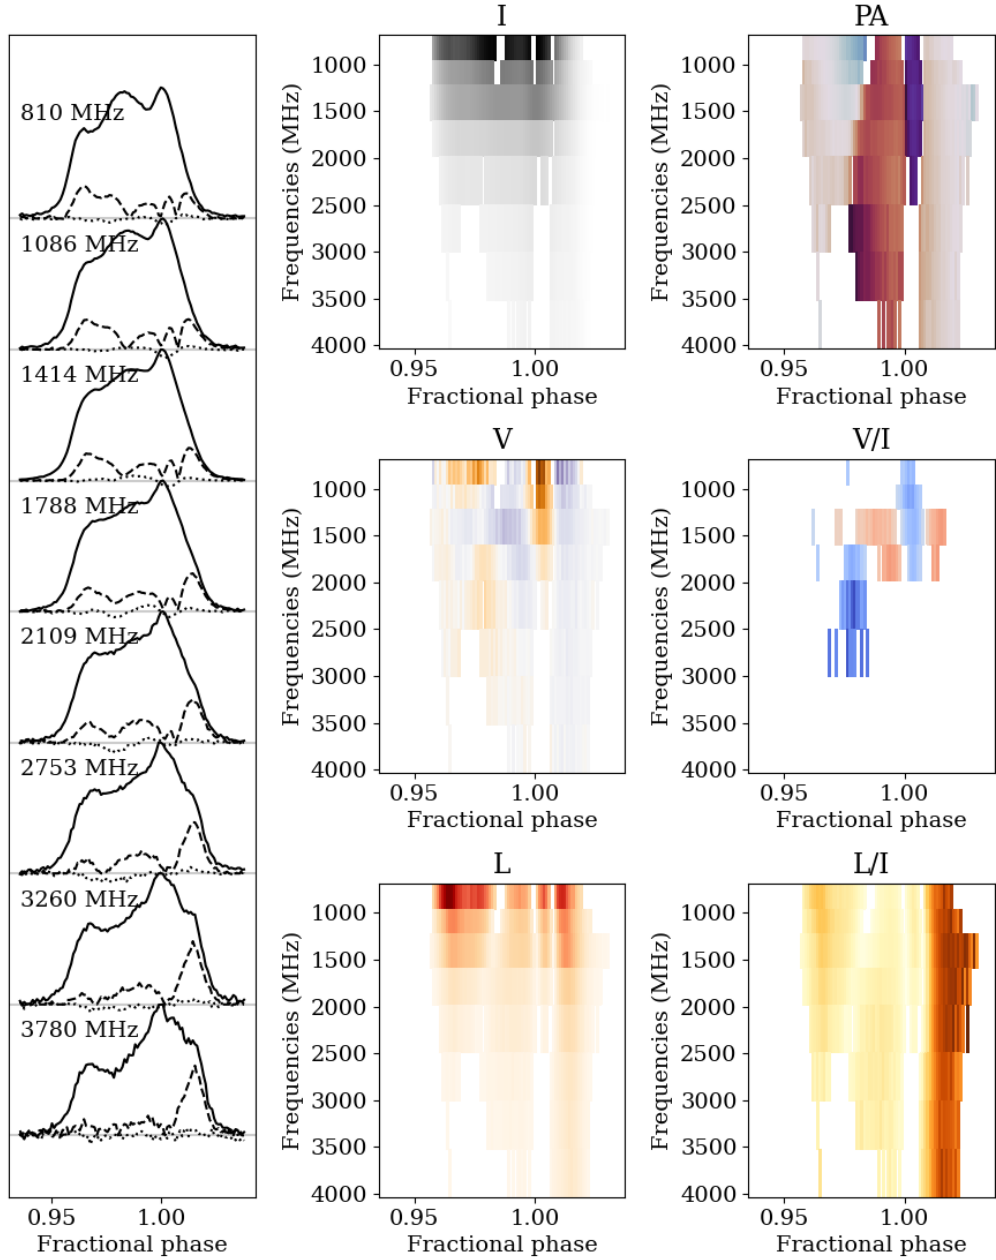


Figure 8.2: PSR J0452–1759: broad-band polarization properties.

The polarization properties of an integrated profile observation of PSR J0452–1759 as observed by the Parkes UWL receiver. The profiles have been dedispersed and Faraday derotated using the catalogue values for this pulsar (Manchester et al., 2005). Each plot shows the on-pulse region for eight pulse profiles spanning the band from 700–4000 MHz. The data are only shown where the S/N of the linear polarization $L_{\text{SNR}} > 5$, except for V/I , which only shows the data where the S/N of the circular polarization $V_{\text{SNR}} > 5$.

Caption continued on the next page.

Figure 8.2: *Continued caption.*

Left: profile shape evolution with frequency, showing total intensity (solid line), linear polarization (dashed line) and circular polarization (dotted line). The remaining six plots are waterfall plots showing, from top left to bottom right, total intensity or Stokes I , PA, circular polarization Stokes V , fractional circular polarization V/I , linear polarization L and fractional linear polarization L/I . The color schemes are chosen to highlight evolution of shape rather than absolute magnitude. For the PA plot, the colour scheme is equivalent to that in Fig. 6.7. For the others, darker colours indicate greater intensity. The plots for V and V/I are two-coloured to represent both positive and negative values. For V , purple is positive and orange is negative, and for V/I , red is positive and blue is negative. Figure made with matplotlib.

either classic orthogonal jumps, where the total polarization drops to 0, and nonzero- V jumps, where a jump in the PA is accompanied by a drop in linear polarization, but the circular polarization does not drop, and may even increase.

Unlike the pulsars presented in Chapter 6, since this is early categorization work, I have not yet performed analysis of the best fit RM and DM required to align the pulse profiles across frequency. The profiles were inspected after aligning them either with the catalogue values for RM and DM, or after applying approximate corrections to these values by eye. Consideration of the relative alignment of polarization features across frequency should therefore take this factor into account, and quantitative alignment is a necessary step to be taken to progress with this work in the future. The results of the qualitative categorization were as follows. 25% of the sample showed only classic orthogonal jumps, 22% showed only nonzero- V jumps, and 12% (including PSR J0452–1759) showed both. In general, pulsars only showed evolution of Stokes V with frequency if the profile included a nonzero- V jump at some frequency, but Stokes V evolution only occurred for roughly half of the cases with a nonzero- V jump. In general, I also found that 57% of the sample showed at least one change in handedness for the Stokes V profile over frequency. Meanwhile, the pulsars with no hand change in Stokes V also exhibited no major evolution of polarization properties with frequency.

PSR J0908–4913 is a particular striking example of the Stokes V handedness changing not merely across phase, but across frequency as well, as can be seen for the main pulse in Fig. 8.3. The plot of V/I in particular exhibits a diagonal transition between negative and positive, where the transition at earlier phases happens at higher frequencies. This pulsar has an interpulse too, shown in Fig. 8.4, and it can be seen that the broadband behaviour of V and V/I is reversed in the interpulse in comparison to that shown for the main pulse. Not only is V negative in the interpulse

where the main pulse has positive V , and vice versa, but the diagonal slope of the transition between positive and negative is angled in the other direction.

These early investigations indicate that the behaviour of circular polarization is a signifier of wider frequency-dependent phenomena observed in pulse profiles. Dyks (2020) showed that apparent orthogonal jumps in the PA swing can result from the passage of a pulsar's polarization vector close to the Stokes V axis of the Poincaré sphere. Such a signature corresponds to the total polarization transferring smoothly from linear to circular polarization, rather than a disjoint jump, and may be an indication of both the magnetospheric origins of the polarization modes, and their coherent and incoherent interactions. Works of note that suggest evidence for coherent and disjoint relationships between orthogonal modes are Edwards & Stappers (2004) and Karastergiou et al. (2011). The former showed deviations from orthogonality in the single pulse distributions on the Poincaré sphere for PSR B0329+54. Karastergiou et al. (2011) demonstrated that the polarization properties of the transient component that appeared in the profile of PSR J0738–4042 could be explained completely by assuming a disjoint superposition of competing orthogonal modes. If two orthogonally polarized modes combine together with some relative phase difference, this will be reflected in the relative amounts of linear and circular polarization observed in the resultant radio emission. If the phase difference between the modes varies with frequency, perhaps due to the modes at different frequencies being produced in different regions of the magnetosphere, then the relative linear and circular polarization observed will also change with frequency. Dyks (2019) investigated modelling pulsar emission as being composed of two entirely circularly polarized orthogonal modes, which interact coherently to result in observed elliptically polarized modes. He found that degeneracy between the mode amplitude ratio and the phase difference between the modes limited the direct comparisons that might be drawn with observed data. However, he also stated that this degeneracy could be broken by including the frequency-dependence of the polarimetric properties of the data, something that should be investigated. Going forward, it is clear that when building understanding the complexities of the pulsar emission process, it is important to take account of the broad-band properties of circular polarization.

8.2 Modelling philosophy

I have taken a data-driven approach to the modelling in this thesis. In practice, this has meant keeping models as simple as possible because there are a limited number

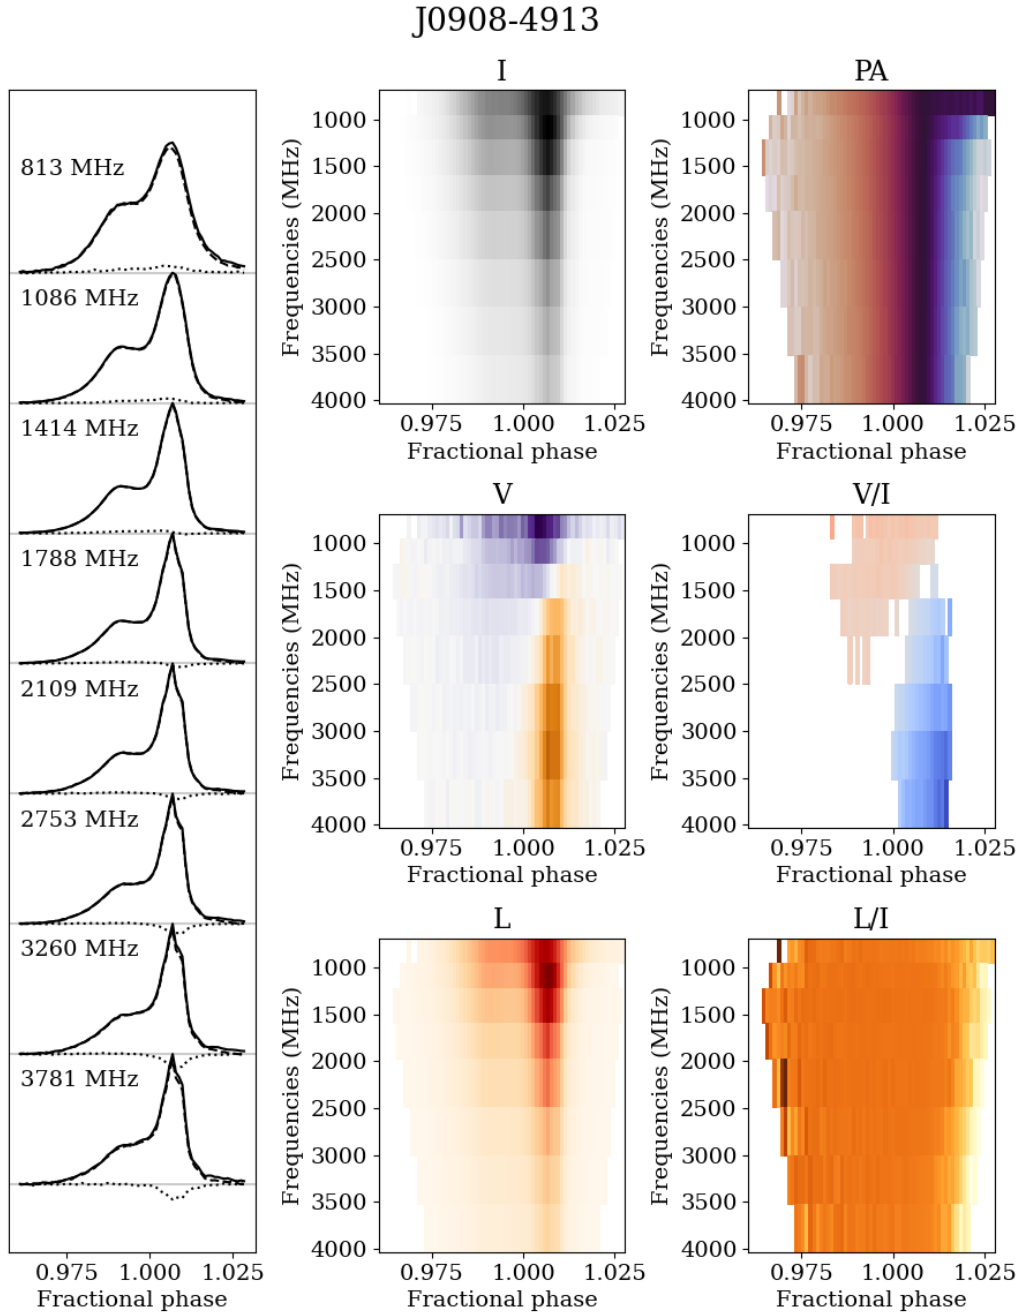


Figure 8.3: PSR J0908–4913 main pulse: broad-band polarization properties. Broad-band observation of PSR J0908–4913 made with the Parkes UWL receiver. This figure shows the main pulse and Fig. 8.4 shows the interpulse. The observation was dedispersed and Faraday de-rotated using the catalogue values (Manchester et al., 2005) and additional corrections of DM $0.18 \text{ cm}^{-3} \text{ pc}$ and RM 4.8 rad m^{-2} were applied to further align the profiles by eye. The figure is formatted as described for Fig. 8.2. Figure made with matplotlib.

J0908-4913, interpulse




Figure 8.4: PSR J0908–4913 interpulse: broad-band polarization properties. As for Fig. 8.3, now showing the interpulse. Figure made with matplotlib.

of known certainties in pulsar observations. This is particularly true when using new broad-band observations that reveal new levels of detail about unexplained phenomena (open question 3). My focus has been on keeping the assumptions underpinning the modelling to be simple and limited in number. This is to enable testing of the extent to which these assumptions are sufficient to describe the data, in the knowledge that the assumptions in question do not encompass the breadth of variety observed in pulse profiles in general. As a result, in my research I have paid almost more attention to the cases where the modelling breaks down, in comparison to the cases where the modelling is sufficient and successful.

Since my goal was to understand the limitations of these modelling applications, I began the modelling with minimal knowledge of the priors applicable to each model in question. For both the RM-DM fitting algorithm (Chapter 6), and for the scattering models (Chapter 7), I therefore began by keeping priors on model parameters uniform within the widest allowable range of the parameter values. Exploration of the results obtained for pulsars with different observable features, combined with tests of the modelling results for simulations, meant that I could then converge upon an understanding of the limitations of the models and the circumstances under which they are applicable. For example, the accuracy of the RM-DM fitting is strongly dependent on the curvature of the PA profile, which itself depends on the pulsar geometry, and the accuracy of scatter modelling depends on the relative sizes of the scattering timescale and the intrinsic width of the modelled pulse profile in question. Such conclusions are almost as important as the tangible results of the modelling process, since they will inform the prior knowledge required for future modelling of these facets of pulsar behaviour.

The application of this modelling philosophy to the scatter modelling has particular relevance in terms of my approach towards understanding the intrinsic pulse profile and its evolution with frequency. My early intentions had been to constrain the frequency evolution of the pulse profile within the modelling process. The logic behind this was that we have prior expectations that pulse profile shapes evolve smoothly with frequency, and that they tend to increase in width as frequency decreases, particularly below 1 GHz. However, the reality is that pulsars show considerable variety in their frequency evolution, and so quantifying this qualitative description is difficult. I therefore decided instead to use a simple, consistent, model for which the intrinsic profile shape could freely converge on the best fit parameters for the pulse profiles at each frequency independently. I then used qualitative prior knowledge to compare the results for the independent profile fits across frequency and hence constrain the cases

for which modelling of both scattering and of the intrinsic pulse profile was successful. Now that my results indicate what to expect from modelling scattered pulsars, based on the sample of single-component pulsars used, this information can be applied to modelling multi-component pulsars, which require more complex descriptions of the intrinsic profile shape.

8.3 Pulsar beam structure

My work on the single pulses of PSR J1136+1551, and on the integrated broad-band profiles of PSRs J1359–6038 and J1056–6258, reveals information about the beam structures of these pulsars, addressing open questions 1a–d. Some measurable outcomes of this research are the estimates of emission heights and height ranges for the observations analysed. As these provide concrete values, they can be compared both to expectations from theory and to the measurements made in other pieces of research. In this section I will explain the origins of the emission heights I calculated for PSR J1136+1551 and how these values depend upon the assumptions made about aspects of the pulsar beam structure. I will compare my results with those measured by others, and discuss how the underlying assumptions made with alternative methods may similarly impact these results. Finally I will discuss how my emission height ranges for PSRs J1359–6038 and J1056–6258 compare with the expectations for variation of such measurements for \dot{E} , and how this might inform investigations going forward.

8.3.1 Emission heights for the hollow cone and fan beam models

Chapter 5 describes my work using single pulses of PSR J1136+1551 to investigate this pulsar’s beam structure. The work concluded that the pulsar beam shape was best described by the fan beam model, with additional evidence of the behaviour of orthogonal polarization modes of emission in the pulsar magnetosphere. This work can be compared with that of Desvignes et al. (2019), who mapped out the sky-projected beam shape for the precessing pulsar PSR J1906+0746 by modelling how its polarization properties evolved and linking this to the beam geometry. Their map shows an emission region that extends radially, which may be taken as evidence for the fan beam model being appropriate to describe the beam structure for that pulsar. Assuming the fan beam model, it is then possible to calculate how the height at which the radio emission is produced varies with frequency. I calculated that to be $r(f) =$




Figure 8.5: Emission heights vs. input footprint parameter ratio for hollow cone and fan beam models.

This figure shows how the emission height range over the observing band 330–470 MHz varies depending on the input value of s_L^{peak} . For a given value of s_L^{peak} the associated range of emission heights is given by the extent of the grey shaded region (hollow cone model) or hatched region (fan beam model). Dashed horizontal lines mark the highest and lowest heights for the observing band for each model, corresponding to 330 MHz and 470 MHz respectively, for $s_L^{peak} = 0.5$ and $s_L^{peak} = 0.8$. These two values of s_L^{peak} are marked with vertical dashed lines to guide the eye. The emission height range for the fan beam model with $s_L^{peak} = 0.5$ is indicated by a black double-headed arrow; this corresponds to the emission heights shown in Fig. 5.11. The emission height calculated by Noutsos et al. (2015) at 400 MHz, using the hollow cone model with $s_L^{peak} = 1$, is marked with a cross. Figure made with matplotlib.

$424 f^{-0.39}$, where r and f are in units of km and GHz respectively (or, when weighting the power law fit by the uncertainties on the emission heights, $r(f) = 447 f^{-0.33}$). However, it is also important to understand how the model assumptions listed in Section 4.1.1 affect the heights calculated, and to what extent information from other sources can constrain our knowledge of pulsar emission heights.

For the single pulse analysis I fixed the footprint parameter ratio for the field line responsible for the average emission peak at the reference frequency, s_L^{peak} , to 0.5 (see Section 5.1). However, qualitatively similar metrics can be obtained using a different value, though doing so would shift the emission heights and height ranges. In other words, the results of the single pulse analysis do not constrain s_L^{peak} , meaning that there is the freedom to change its value given additional information. Fig. 8.5 shows how the heights and height ranges compare for the hollow cone and fan beam models, and how these values change with s_L^{peak} . Previous work modelling the emission heights of the integrated profile of PSR J1136+1551 has used the hollow cone model with the emission fixed to the last open field lines so that $s_L^{peak} = 1$. The emission height that Noutsos et al. (2015) calculated for PSR J1136+1551 is marked on Fig. 8.5. They calculated this for a frequency of 400 MHz, using the component peak separation and their result is in agreement with what my single pulse measurements predict, assuming $s_L^{peak} = 1$.

8.3.2 Emission heights from the delay-radius relation

Since emission heights can be estimated through a few different methods, it is instructive to compare my results with those calculated by others using alternative techniques. As described in Chapter 2, the delay-radius relation is responsible for a phase lag between the intensity and PA profiles at a given frequency, such that $\Delta\phi = 4r/R_{LC}$ (see Blaskiewicz et al. 1991; Johnston & Weisberg 2006; Dyks 2008). Noutsos et al. (2015) measure this phase lag for PSR J1136+1551 at 150 MHz, defining the fiducial point of the intensity profile as the mid-point between the two component peaks and that of the PA as the point where it has the steepest gradient. This gives an emission height of 349_{-150}^{+158} km. In the absence of polarimetric data, and given that I dedispersed and aligned the data such that the means of the subpulse distributions evolved symmetrically with respect to the magnetic axis, I could not directly measure delay-radius relation effects. However, I am able to obtain an equivalent emission height from my modelling work. Doing so requires a change in footprint parameter ratio from $s_L^{peak} = 0.5$ to $s_L^{peak} = 0.8_{-0.1}^{+0.2}$. The height shift that this entails is shown in Fig. 8.5.

Kramer et al. (1997) measured an upper limit on the pulse time-of-arrival difference of $1640 \mu\text{s}$ between 1.41 and 32 GHz, which they translated to a height range of 310 km using $\alpha = 147.0^\circ$ (for $\alpha = 51.3^\circ$ this translates to 276 km). Extrapolation of my heights for the fan beam model to their frequency band, using $r(f) = 424 f^{-0.39}$, results in height ranges below their upper limit for all values of $s_L^{peak} \geq 0.5$.

Hassall et al. (2012) measured no time delay attributable to timing A/R (see definition 2h in Chapter 2) across the lower frequency range of 40–180 MHz, which translates to an upper bound of the height range of 59 km. This is very low compared to my extrapolated height range. In principle, we might expect measurable A/R in their data. The low heights they postulate may be attributed to their technique for determination of the fiducial point and dedispersion, absorbing the effect. My work with simulated profiles in Chapter 6 shows that profile misalignment due to the delay-radius relation is indeed absorbed by the DM and RM, if it is of a comparable magnitude to that predicted by the height-frequency relation $r(f) = 424 f^{-0.39}$.

All these considerations of the effects of the delay-radius relation, and of timing delays due to A/R, so far have assumed a strict one-to-one mapping of radius and frequency. However, it is already clear from my simulations that the single pulse behaviour of PSR J1136+1551 is better explained by there being two emission paths which can be associated with OPMs. A contribution from the frequency-independent X mode would reduce the magnitude of PA–intensity delay measured in comparison to that which would otherwise be expected. Noutsos et al. (2015) refute the possibility that the frequency widening of PSR J1136+1551 may be caused by diverging OPMs, suggesting that the lack of jumps in the PA profile at frequencies below 1400 MHz is possible evidence against this being the widening mechanism. There are, however, emission geometries and X and O mode distributions that could result in the amplitude of one being everywhere smaller than that of the other along the line of sight. This would mean a net widening could be seen without the need for any mode jumps. Furthermore, if emission at a given frequency comes not from a single height but from some range, the frequency dependence of a timing A/R effect would be far less clear than otherwise expected. Polarization data could be used to clarify these issues for this pulsar, and the evidence from the studies of PSRs J1056–6258 and J1359–6038 in Chapter 6 is that it is possible for A/R effects to be present but they may be easily absorbed by the choice of the DM and RM.

8.3.3 Impact of assumptions about emission region shape

My simulation results favour the fan beam model, a description of the emission region that is independently favoured by the work of Wang et al. (2014), in their investigations of precessing pulsars. Unlike the hollow cone model, it does not provide a clear explanation for the symmetry of the two emission regions about the fiducial plane. However, the symmetrical frequency evolution, which implies symmetrical positioning of the emitting regions, is not directly inferred from the data. It is a result of my choice of DM, the reasoning for which has been discussed in Section 4.2. There is therefore some flexibility: the exact positioning and frequency evolution of the sub-pulse distributions do not have to be perfectly symmetrical in the simulation. Indeed, given that the two distributions have different shapes, it might be surprising if they were.

I have addressed above the key implications of this model on the emission heights generated, but only in the case of a one-to-one radius-to-frequency mapping. It seems likely that the region within which a given frequency is produced could extend over at least a small range of heights. The model of Gangadhara (2004), which relates frequency to field line curvature instead of height, presents an alternative avenue of consideration. In such a picture the observable radiation must originate from a range of heights, as the line of sight crosses regions at different perpendicular distances from the magnetic axis. A simulation investigating these considerations would more accurately replicate the fan beam model of Dyks & Rudak (2015), for which each frequency is emitted by an extended region along the beam.

8.3.4 Emission heights in the context of understanding pulsars

My research has shown that emission height measurements are strongly dependent upon a number of underlying assumptions in the process of transforming pulsar measurements to those heights. However, even if the exact heights of emission cannot be unequivocally calculated, comparison of results, both from different methods, and from different pulsars, will still reveal useful information. The emission height ranges I calculated for PSRs J1359–6038 and J1056–6258 were broadly consistent with wider expectations of pulsar emission. A further factor of relevance can be considered is the measurement of the spin-down luminosity \dot{E} for each pulsar (see Section 1.1.2) and how this might be related to emission height ranges.

Karastergiou & Johnston (2007) developed an empirical model of radio pulsar beams to describe the diverse population of observed pulsar profiles. This model proposed different emission height structures for pulsars with different values of \dot{E} . They suggested that pulsars with low \dot{E} emitted over a broad range of emission heights, whereas those with high \dot{E} would have their emission restricted to a much narrower range.

PSR J1359–6038 has an \dot{E} of 1.2×10^{35} ergs s⁻¹, and I estimated its height range to be 120±120 km. PSR J1056–6258 meanwhile has a smaller \dot{E} of 1.9×10^{33} ergs s⁻¹ and larger height range estimate of 430±350 km. My calculations of emission height ranges for these two pulsars are therefore in accordance with this idea, without providing further constraints. Similar calculations for a large survey of pulsars would be necessary to investigate whether the behaviour of the population as a whole reflects this prediction.

8.4 Understanding inhomogeneous regions in the ISM

Understanding pulsars increases the accuracy to which we can use them as tools to probe the ISM (open question 2b). Having developed my knowledge of how broadband pulsar observations behave and vary across frequency, I then used this information to inform my analysis of pulsars scattered by structures in the ISM. In this section I will discuss the results of my analysis within the context of previous work in this area, and how the results of my work could inform the wider practice of using pulsar observations to constrain measurements of scattering.

8.4.1 Scattering parameters in context

Cross-comparing my sample of pulsars with the literature, I found 9 pulsars in our sample that have previous measurements for α , which I have listed in Table 8.1. Some of these have been measured multiple times, giving a total of 13 measurements of α to compare to my results, found in Löhmer et al. (2001), Lewandowski et al. (2013) and Lewandowski et al. (2015b). 6 of the 7 estimates of α given in Lewandowski et al. (2013) are larger than my corresponding estimates, whilst 4 of the 5 estimates given in Löhmer et al. (2001) are smaller than mine. This suggests that differing modelling approaches may tend to give systematic differences in the resultant parameters estimated. Of the 9 pulsars, 5 have at least one measurement consistent with mine to within the uncertainties. The 4 that are entirely inconsistent are all found

Table 8.1: Table of pulsars for which there have been previous measurements of α . It lists the values of α measured in this work (see Table 7.1) and those values found in the literature, with the appropriate citations. The reference codes in the table correspond to the following papers: L2001 is Löhmer et al. (2001); Lew2013 is Lewandowski et al. (2013); Lew2015b is Lewandowski et al. (2015b).

PSRJ	α (this work)	α (literature)	Reference
J1801–2304	4.529±0.003	3.9±0.4	L2001
		4.92±0.11	Lew2013
		3.45±0.09	Lew2015b
J1818–1422	3.787±0.008	3.5+0.5–0.6	L2001
		3.97±0.12	Lew2013
J1820–1346	3.81±0.01	3.3+0.7–0.5	L2001
J1822–1400	3.31±0.08	4.1+0.7–1.2	L2001
		3.96±0.18	Lew2013
J1824–1118	3.77±0.02	3.5+0.7–0.6	L2001
J1825–1446	4.0±0.2	3.77±0.24	Lew2013
J1835–0643	3.77±0.07	4.37±0.3	Lew2013
J1841–0425	3.29±0.07	3.91±0.14	Lew2013
J1857+0143	4.0±0.1	4.69	Lew2013

in Lewandowski et al. (2013) and all have larger values of α quoted there than those presented here.

Fig. 8.6 shows τ calculated at 1 GHz plotted against DM. The values for the final sample are plotted as black circles, and for comparison I have also plotted those values obtained from pulsars subsequently determined to be either not scattered or poorly modelled. The horizontal lines mark the smallest and largest time resolutions for a single bin of a pulse profile. Since all observations have 1024 bins across the profile, these values correspond to $1/1024^{\text{th}}$ of the smallest and largest periods in the sample. I took the equation relating τ and DM fit by Krishnakumar et al. (2015) at 327 MHz, scaled it to 1 GHz using the best fit α of 4.0 and plotted it as a black line. I also plotted the scaled fits for $\pm\sigma$ about the mean, $\alpha = 3.4$ and $\alpha = 4.6$, as dashed lines. The functional form of this equation is

$$\tau_s = 3.6 \times 10^{-6} \text{ DM}^{2.2} (1.0 + 0.00194 \text{ DM}^2) (1000/327)^{-4.0}, \quad (8.1)$$

where τ has units of ms and DM has units $\text{cm}^{-3} \text{ pc}$. It can be seen from Fig. 8.6 that the pulsars in my final sample have high DMs and that they are clustered close to the best fit line of Krishnakumar et al. (2015). There is some spread in the values, meaning that my results are also compatible with previous fits in the literature made by Ramachandran et al. (1997); Löhmer et al. (2004); Bhat et al. (2004) and

Lewandowski et al. (2015a). Since the MeerKAT sample is limited to a narrow DM range, it is not possible to perform an equivalent model fit for the τ -DM relationship using its results alone. It is a point of note that my results are not symmetrically distributed about the Krishnakumar et al. (2015) model. This may be related to a difference in modelling approach: whereas I modelled the intrinsic pulse profile as a Gaussian without constraining its parameters, Krishnakumar et al. (2015) used a high frequency unscattered profile as the template for the intrinsic pulse profile.

The choice of how to model the intrinsic pulse profile necessarily affects the results for the measured scattering timescales and scattering spectral indices. For example, the slower rise time of the thick screen model (Williamson, 1972) may absorb the intrinsic pulse shape into its measurement, or conversely the choice of a Gaussian intrinsic pulse may obscure evidence for a thick screen scattering function. Another example is the comparison of my scattering timescale result for PSR J1316-6232 (see Table 7.1) with the estimate published by Crawford et al. (2001) of $\tau \sim 150$ ms at 1.35 GHz. Scaling my result to this frequency gives a scattering timescale almost twice as large. This difference follows directly from the difference in modelling choices. Whereas Crawford et al. (2001) provided an estimate of the scattering timescale based only on the exponential decay of the profile intensity, my model involves a convolution of that exponential decay with a Gaussian representing the intrinsic profile shape.

This suggests caution should be exercised when directly comparing individual τ values measured using different methods. A systematic shift of τ values caused by the choice of method is likely to have less of an impact on α , provided all the values of τ used to calculate it were estimated using the same method. However, calculations of α made using τ values measured with a variety of methods are likely to be skewed by this. This example highlights the importance of applying the same modelling approach to a large-scale sample of pulsars, such as the TPA survey on MeerKAT, in order to be able to compare the scattering results for different pulsars.

In this context, there is another source of scattering timescales with which it is worth making a comparison. The NE2001 model (Cordes & Lazio, 2002) describes the distribution of free electrons in the Galaxy and can be used to make predictions of expected scattering timescale for a given set of galactic coordinates and either an input DM or distance. Fig. 8.7 is a repeat of Fig. 8.6, now with the NE2001 predictions for scattering timescale overlaid in blue. These scattering timescales were extracted from the model by inputting, for each pulsar, its galactic coordinates and our modelled DM into the python wrapper for the model, PYNE2001 (developed




Figure 8.6: Plot of τ vs. DM.

Plot of scattering timescale τ at 1 GHz vs. DM. Black points: the values for the final filtered sample. Black line: the best fit equation identified by Krishnakumar et al. (2015), scaled to 1 GHz using our best fit α of 4.0. Dashed lines: the same equation, now scaled using $\alpha = 4.0 \pm 0.6$. Pale grey crosses: pulsars rejected from the sample due to either to inaccurate modelling or the pulsars not being scattered. Horizontal dot-dash lines: these mark the minimum and maximum time resolutions of a single bin for the pulsar sample, where all observations had a resolution of 1024 bins across the pulse period. Figure made with matplotlib.




Figure 8.7: Plot of τ vs. DM including results from NE2001. Repeat of Fig. 8.6, now with the corresponding predictions for scattering timescale from the NE2001 model overlaid in blue. The main text describes how these predictions were obtained from the model. Figure made with matplotlib.




Figure 8.8: Histograms of τ residuals for this work and NE2001.

This figure shows the residual in scattering timescale with respect to the scaled τ -DM equation from Krishnakumar et al. (2015), plotted as a histogram, both for this work and for the scattering timescale predictions from the NE2001 model. Grey: scattering timescale residuals for this work. Blue: those for the corresponding values of τ predicted by the NE2001 model for the same galactic coordinates and DMs. A vertical black line marks $\Delta\tau = 0$. Inset: the same histogram, now rebinned and zoomed to the x-range ± 50 ms, which corresponds to the three bins centred on $\Delta\tau = 0$ in the main plot. Figure made with matplotlib.

by Vincent Morello, see <https://github.com/v-morello/pyne2001>). For our successfully modelled pulsars (marked with dots) the NE2001-predicted scattering timescales clearly overlap our measurements, but have a tendency towards smaller values. To inspect this difference more carefully, without the confounding visual difficulty of the log-scale, I plot histograms in Fig. 8.8 of scattering timescale residuals with respect to the scaled functional form from Krishnakumar et al. (2015), as given in equation 8.1. Both histograms peak around $\Delta\tau = 0$, indicating broad consistency with this function. However, my modelled scattering timescales are skewed towards larger scattering timescales (positive $\Delta\tau$), whilst the predictions of NE2001 are skewed towards smaller values (negative $\Delta\tau$). This is clearest in the inset figure, which zooms in on $\Delta\tau$ between ± 50 ms. Investigating the possible origins of this distance, I found that, of the 1143 pulsars whose dispersion measures were input into NE2001, 63 were also in my sample of 84 successfully-modelled pulsars. These DMs were taken from Taylor et al. (1993), the Princeton pulsar catalog and the Parkes Multibeam survey and can be found at the NE2001 source code link (accessible from <https://github.com/v-morello/pyne2001>). I found that they were generally larger than my own measurements, though only on the order of the ΔDM corrections I made through the scatter modelling, as shown in Fig. 7.12. Querying the NE2001 model with a smaller DM at the same galactic coordinates ought in theory to result in a smaller estimated scattering timescale τ . However, it is clear that this DM discrepancy is insufficient to explain the τ differences shown in Fig. 8.8, because if I query the model using my uncorrected input DMs for these pulsars, rather than my corrected values, this results in an increase in τ of order only 1 ms. A possible origin for this discrepancy in scattering timescale is a difference in the method by which (Cordes & Lazio, 2002) estimate scattering timescales as input information for the model, as compared to the method I present in this thesis. As described previously, Crawford et al. (2001) measured a much smaller value of τ for PSR J1316–6232 than that given in this work, due to their use of a different modelling approach. It could be that the discrepancy with NE2001 seen here is a similar effect, scaled to a large number of sources.

Of the values of τ associated with poor modelling shown in Fig. 8.6, there are several at smaller DMs for which τ is significantly larger than would be expected from both the Krishnakumar et al. (2015) model and the prediction from NE2001. This is easily explained based on the results of my investigations into the causes of poor model fits. It is expected that pulsars with smaller DMs will, in general, be less strongly scattered. My work shows that small scattering timescales are likely to be over-estimated, particularly when the timescale size is similar to the width of the

intrinsic pulse profile. These poorly modelled pulsars also have τ estimates close to the temporal resolutions of the pulsars. Indeed, some of these pulsars were concluded not to be scattered at the observing band, which takes the concept of small values of τ being overestimated to the logical extreme. I caution therefore that there may be a wider tendency to overestimate τ at the lower limits of both DM and τ attainable at a given frequency.

This may go some way to explaining Fig. 8.9, which is a plot of α against DM for my values (black circles) and those given in the literature as described in Chapter 7.4 (Cordes et al., 1985; Löhmer et al., 2001, 2004; Lewandowski et al., 2011, 2013, 2015a,b; Krishnakumar et al., 2017, 2019). My best fit of $\alpha = 4.0(\pm 0.6)$, shown as a horizontal black line, is consistent with the results for the literature. On calculating mean α values for 5 bins across the DM range I saw no significant evolution of α with DM, save for a slight increase in α at the highest DM bin, encompassing the three points in Fig. 8.9 at around $1000 \text{ cm}^{-3} \text{ pc}$). More data are required to determine whether this increase is statistically significant. Weighting the averages by the uncertainties in α tends to favour higher values in comparison to the unweighted averages. Inspecting the literature values of α , it can be seen that there is a greater spread in the literature values at lower DMs. As I have described, the over-estimation of smaller values of τ tends to lead to under-estimation of α . The reduced strength of scattering at lower DMs should cause this to have a greater effect at low DMs, which is what can be seen in Fig. 8.9.

8.4.2 Scattering at low and high frequencies

Whereas my modelling explains the distributions of α and τ obtained for pulsar observations at frequencies $\geq 400 \text{ MHz}$, other behaviour is seen in the low frequency studies performed by Kuzmin & Losovsky (2007) and Geyer et al. (2017). Kuzmin & Losovsky (2007) identified a τ -DM relation of $\tau = 60(\text{DM}/100)^{2.2} \text{ ms}$ at 100 MHz , and Geyer et al. (2017), whose results for τ also corresponded well to this equation, measured a distribution of values of α that is systematically shifted to lower values in comparison to mine. An explanation for this may be that low frequency pulsar observations are probing a different scattering environment. One aspect of this is that lower frequency pulsar emission probes a wider region of space than that at higher frequencies since it is scattered more strongly, as shown in fig. 1 of Cordes et al. (2016). Applying the same isotropic model may therefore result in different measurements of α at different frequencies, something not measurable for the frequency range of the data presented here but meriting further exploration with datasets at lower frequencies.




Figure 8.9: Plot of α vs. DM.

Black points: our sample of 84 pulsars. Pale grey crosses: previously published values obtained through time domain scatter modelling. Horizontal line: the mean value of α calculated from a Gaussian fit to the distribution of our results. Figure made with matplotlib.

A further consideration is the effect of distance. Scattering analyses performed at low frequencies focus on nearby pulsars, due to the extreme scatter broadening that occurs in high DM pulsars. It is interesting to consider the effect of anisotropy in this context. In the early stages of the modelling process I compared anisotropic model fits with the isotropic fits that I ultimately presented in Chapter 7 and found them to be generally unsuccessful. I concluded that there was no evidence of anisotropic scattering in the MeerKAT TPA sample of scattered pulsars. For nearby pulsars, for which scattering is predominantly observable and measurable at lower frequencies than the TPA sample, it is likely that the pulsar emission only passes through a single or a very small number of screens. This means that the scattering behaviour observed may bear hallmarks of anisotropy. There is evidence of anisotropic scattering in low frequency surveys, e.g. Geyer et al. (2017), and using an isotropic model to identify the scattering parameters where such anisotropy is evident then results in under-estimation of α , as described by Geyer & Karastergiou (2016). This may be another factor contributing to the smaller values of α at lower DMs seen in Fig. 8.9: pulsars with smaller DMs are likely to be nearby and hence an anisotropic model may be more appropriate. By contrast, the emission from our high DM pulsars is likely to have passed through multiple scattering screens, all of which may have different levels of anisotropy oriented along different axes. The net scattering effect observed will therefore approximate isotropy, as observed in my results.

8.5 Overall discussion and future work

The common thread running through my work is that conclusions about both pulsar emission and the ISM need to be informed by the relationship between the physical mechanisms taking place and what is observable given the capabilities of the instrument being used. My techniques to investigate this have focused on a simplistic modelling approach, which I have used to test how the observations deviate from theoretical expectations. The broad-band observational capability that new instruments give us reveals both the frequency evolution of pulse shape from both the single pulse and integrated profile perspectives, and how polarization evolves with frequency. How well we understand these factors directly informs how well we can calculate and account for the effects of the ISM, in particular dispersion, Faraday rotation and scattering. Simultaneously, knowledge of how to correct for these effects informs our understanding of the pulsar emission beam structure and hence the processes that bring it into existence.

Future work on understanding the pulsar radio beam should be informed by the results of this work in terms of how best to interpret broad-band features of pulsar emission. Profile evolution with frequency, both in terms of the pulse shape and the polarization features, can be clearly linked to theories about beam structure in the pulsar magnetosphere. The wider the band being addressed, the more important it is to account for the effects of the interstellar medium to enable comparisons to be made across frequency. My work provides a method for doing so, and in the future I intend to apply my techniques to a wide scale survey of pulsars observed with the UWL receiver. The initial results from this are revealing the additional importance of circular polarization. Applying my analysis and considering this additional aspect will reveal information about beam structure and emission heights from a population perspective.

My work on modelling scattered pulsars observed with the MeerKAT telescope showed that including profile evolution improves the information available from scatter modelling, increasing understanding of the ISM. Future work should extend this to include more complex profiles, including their evolution with frequency. The time-variable aspect of scattering is also relevant: a strong understanding of the pulse profile helps converge on the scattering parameters accurately, which would be relevant for measuring any changes of how that profile is scattered over time.

8.6 Summary

My work has demonstrated how data from new broad-band instruments can reveal considerable information about the intrinsic behaviour of pulsar emission. In addition to explaining factors in the emission mechanism, such as the beam structure and emission heights above the pulsar surface, this work has improved understanding of the ISM, as I have shown through my survey of interstellar scattering. This research lays the ground work from which to build, in terms of extending understanding of the broad-band polarimetric properties of pulsars to include circular polarization, and applying the analytical techniques developed here to multiple pulsars, in order to investigate the wide variety of phenomena observed.

Chapter 9

Conclusions

9.1 Intentions and context

Observations of radio pulsars over the last half-century have revealed a wealth of complexity that is not adequately explained by the canonical model. In addition, the presence of the ISM has an impact on every pulsar observation made. The combined efforts of theorists and observers have resulted in good characterization of the features of pulsar radio emission, and a set of theoretical frameworks to explain these characteristics and predict further expectations for pulsar behaviour. New advances in radio instrumentation meant that I could test these theories to a new level of detail, and so address long-standing questions about the broad-band nature of pulsar radio emission. The work presented in this thesis builds upon a range of ideas and techniques developed within this field of research in the following ways.

9.1.1 Context: ideas and techniques

An idea that is central to much of the investigatory work of this thesis is that measurements of the effects of the ISM (dispersion, Faraday rotation and scattering) encode information, not just about the ISM structure, but about the pulsar as well. This idea is evident in the work of Ilie et al. (2019), who identified effects of magnetospheric origin in the phase-resolved RMs that they measured for a large sample of pulsars, which in turn is a style of measurement used first by Noutsos et al. (2009). Karastergiou & Johnston (2006) laid the groundwork for probing the meaning of the DM applied in the context of what this means for both the choice of RM and our understanding of the fiducial point in a pulse profile. Ilie et al. (2019) demonstrated clearly how the measurement of these two values, RM and DM, is covariant, opting,

for each pulsar, to choose the DM which minimized variations in the phase-resolved RM calculated.

When investigating broad-band alignment of pulse profiles I advanced the concept of alignment based on the PA profile as applied by Karastergiou & Johnston (2006), converting it to simultaneous calculation of RM and DM since I had the broad-band capacity necessary to be able to do so. The concept of pulse profile evolution with frequency being a key factor which must be corrected for in pulsar timing is well laid out by Hassall et al. (2012), who showed that the choice of fiducial point in a pulse profile has an impact on the resultant DM and the understanding of pulsar beam structure, particularly with relation to emission heights. Tension between various measurements of emission heights for PSR J1136+1551 in the work of Hassall et al. (2012) and Noutsos et al. (2015) prompted my investigation using observations of that pulsar's single pulses. For this I built upon the geometrical framework laid out by Gangadhara (2004) to build simulations of 3D pulsar beams and the conceptual description by Dyks & Rudak (2015) to understand and test how RFM would manifest observationally from the fan beam model in such a simulation. The results of such investigations demonstrated the importance of incorporating the concept of intrinsic profile shape and evolution in a consistent way into models of scattered pulsars. When performing a survey of how pulsars are scattered by the ISM I was guided by the approach taken by Geyer et al. (2017) to include the effects of profile evolution into forward modelling of pulse profiles. I set up a new framework for applying the model developed by Geyer et al. (2017) and added to our understanding of ISM structure by completing the largest survey of time-domain scattered pulse profiles done using a single instrumental set-up.

9.1.2 Context: comparison with other work

This work has its comparative origins in the broad-band profile width evolution investigated by Thorsett (1991) and the efforts of many since to characterize pulsar profile evolution with frequency in terms of RFM and emission heights, a notable example of which is the work of Mitra & Li (2004). New broad-band observations provide an opportunity to extend the understanding of pulsar polarimetric behaviour, the trends of which have been described in such works as Karastergiou et al. (2005), Johnston & Weisberg (2006), Johnston et al. (2008) and Johnston & Kerr (2018). The assumptions made and parameters constrained in my work modelling pulsar beam structure should be compared with similar work by Hassall et al. (2012) and Noutsos et al. (2015), and with the beam mapping results for a binary pulsar by Desvignes

et al. (2019). Previous efforts to measure the scattering timescale and its frequency evolution for pulse profiles are documented in Chapter 7 (e.g. Krishnakumar et al., 2015; Geyer et al., 2017), along with how differences in modelling approach can result in different conclusions being drawn about the ISM, such as differences in measuring scattering timescale (Crawford et al., 2001) or the effect on α of using isotropic fits to anisotropic scattering (Geyer & Karastergiou, 2016). Further mention should be made of the scattering analysis done through methods other than forward modelling, including deconvolution analysis with the CLEAN algorithm and measurement of scattering parameters via scintillation (e.g. Bhat et al., 2004), and using scintillation analysis to map the shapes of ISM inhomogeneities on the sky (Brisken et al., 2010; Simard et al., 2019; Main et al., 2020). Such techniques provide complementary results that reveal the complexity of structure of the ISM on different distance scales from those probed by my dataset.

9.2 Approach and goals

Upon obtaining broad-band pulsar observations, both integrated profiles and single pulses as described in Chapter 3, the first stage of the analysis process was to identify metrics through which to compare these data to the theoretical frameworks developed by the authors listed above. I tested the metrics and algorithms I developed by creating simulations, in order to understand the limits of what a given modelling process would be able to encompass and explain. The focus of the modelling itself was to understand the extent to which the various pieces of theory could replicate what was observed in the data. This revealed information about the accuracy, uniform applicability and limitations of the theories in question, which could then be interpreted in the context of the open questions being addressed.

In summary, I measured the frequency evolution of single pulse profile shapes, created an algorithm to align PA profiles across frequency by correcting for RM and DM simultaneously, and developed a framework to apply the scattering model created by Geyer et al. (2017) to a large survey of scattered single-component pulse profiles. I investigated the accuracy and validity of my modelling processes by creating simulations of pulsar data and using these as test cases. From the results of these investigations I interpreted broad-band pulsar behaviour in terms of calculating emission heights and height ranges, and characterizing the effects of the inhomogeneous ISM in causing profiles to be scattered. This work has made progress in answering open questions about both pulsar emission and ISM behaviour. It has in turn demonstrated which

aspects of radio pulsar observations are the most important to be studied, in order to continue to understand pulsar radio emission processes.

9.3 Answers and questions

9.3.1 Theories ruled out by my results

The results of this thesis help to constrain understanding of pulsar emission and propagation processes by ruling out certain theories related to these processes. The theory of RFM has previously been tested in observations by fixing emission from integrated profiles to the circular region of the last open field lines. As shown in Fig. 8.5, the choice of both the beam model cross-section and the footprint of the emitting field lines on the pulsar affects the resultant emission heights and height ranges inferred by a significant amount. The behaviour of single pulses of PSR J1136+1551 does not support the circular symmetry commonly assumed, instead favouring a fan beam model and revealing that additional complexity is required. This includes the differential behaviour of orthogonal polarization modes and a full modelling process relating the evolution of the wider emission region to the phase variation along the line-of-sight in a geometrically consistent way. The results of both RFM measurements and the relative misalignments of intensity and PA profiles indicate a relatively narrow range of emission heights over a broad band. This rules out very wide height ranges, or RFM behaviour with a frequency relationship that is strongly distinguishable from the quadratic behaviour of the DM, since it was possible to absorb all PA misalignment into the measurements of DM and RM. Finally the survey of scattering in high DM pulsars did not support the presence of strong anisotropic scattering in the time domain profiles, with the isotropic model fitting much more successfully to the pulsar sample. Since nearby pulsars observed at lower frequencies have exhibited evidence of anisotropic scattering, this implies a variation in scattering environment in the ISM on different distance scales. This is most likely due to there being an increased number of scattering screens over the larger distances, which will result in an apparent average isotropic scattering effect.

9.3.2 How my work has progressed understanding

This research has produced results that answer key questions posed by the results of earlier work. The single pulse work has provided observational evidence in support of the fan beam model, in particular the ways in which RFM is expected to manifest

under that model according to the theoretical description by Dyks & Rudak (2015). The alignment of PA profiles at two discrete frequencies for PSRs J1056–6258 and J1359–6038 by Karastergiou & Johnston (2006) resulted in misalignment of the total intensity, which prompted the conclusion that intrinsic evolution of the PA profile shape might be responsible. Addressing this question with new broad-band data demonstrated that it was possible to obtain alignments for both these pulsars when constrained to using only the DM and RM. Since doing so replicated qualitatively the intensity profile shift with frequency seen in Karastergiou & Johnston (2006), this suggests an origin based instead in the variation of profile emission heights with frequency. Hassall et al. (2012) ruled out RFM as an explanation for the profile evolution of PSR J1136+1551, based on the mismatch between emission height results for the hollow cone model and A/R-based timing delays. This work provides evidence that RFM is applicable to this pulsar provided the effect of the beam cross-section is considered within the modelling set-up. Furthermore, my work emphasizing the impact of DM choice upon understanding of beam emission heights shows that the choice of DM can absorb the misalignment effect of RFM, such that it could easily be underestimated or missed in analysis. My survey of scattered pulsars resulted in a well constrained statistical measurement of α that is consistent with theoretical results, and showed that within my modelling process the effect of inaccurate modelling was a tendency to underestimate α . This suggests that results in the literature with anomalously low values of α , where modelling processes have been similar, may have been affected by the same limitations as those found in this work.

9.3.3 Progress with the open questions

Here I repeat the open questions laid out in section 1.3 and summarize the progress that the work in this thesis has made.

1. How do pulsars produce radio beams?

- (a) **What is the structure of the radio beam?** The fan beam model is favoured for PSR J1136+1551 (see for example Fig. 5.9), and 2D beam cross-section modelling suggests that observations are consistent with emission propagating tangent to magnetic field lines. The importance of orthogonal mode propagation and interaction should be stressed. Similar study of a large number of pulsars would be necessary to understand results on a population level.

- (b) **What dictates pulsar beam width, filling fraction and distribution of active emitting sites?** Unanswered questions remain about the distribution of active pulsar emission sites, but their random distribution and phase connection across frequency are both favoured by the broad-band results for all the pulsars studied so far.
- (c) **How does beam structure vary with frequency?** My modelling process reveals that comparatively simple RFM simulations, with one-to-one frequency-to-height relationships, replicate the data well for PSR J1136+1551. The results of this thesis constrain emission height ranges to be fairly narrow and the heights to be at small fractions of the light cylinder radius (predominantly less than 1% for the frequency ranges studied, with the highest height estimated for PSR J1136+1551, 665km, corresponding to 1.2% of R_{LC}). This is consistent both with others' work and with my estimates for PSRs J1056–6258 and J1359–6038, although these estimates are indicative only, and not strongly constrained. There is scope for more complex height-frequency relationships, for example a given frequency being emitted over a range of heights, and there is clear evidence of the necessity to account for the complexity of orthogonal mode propagation and interaction, in particular for pulsars with high fractional circular polarization.
- (d) **How are single pulses produced and how does this result in the properties of integrated pulse profiles?** Single pulses were found to display similar behaviours to the averaged integrated profile, but with a spread of variability that required a statistical approach. Their behaviour also suggested interaction of orthogonal modes was relevant to the production of the resultant pulses observed. My work used single pulses to address the concept of broad-band emission being traceable to active field lines and found it consistent with observations for PSR J1136+1551.

2. How do observational constraints affect our understanding of the pulsar population?

- (a) **How do the observational properties of pulsars vary across the pulsar population?** Within the context of work at other frequencies, the work on scatter modelling in this thesis reveals the observational effects of scattering to be dependent on distance, where high DM pulsars are consistent with isotropic scattering despite existing evidence that individual

screens in the ISM display strong anisotropy. This can be rationalized by high DM pulsars having their profiles scattered by multiple screens, so that the observed profile appears isotropically scattered. The simple model I used does not converge on accurate results for pulsars with small scattering timescales, nevertheless more strongly scattered pulsars were found to be very well modelled by this theory.

(b) **How does pulsar emission interact with the ISM?** This work demonstrates techniques to understand the effects of the ISM separately from the intrinsic frequency-dependence of pulse profiles. Work identifying RM and DM values consistent with expectations about pulsar emission processes resulted in calculating emission height ranges consistent with theoretical expectations. Applying an understanding of intrinsic profile evolution is important for constraining characterization of scatter modelling in a large survey. The results of this survey revealed the properties of scattered pulsars to be consistent with theory.

(c) **How do the polarization properties of pulsar observations map to the emission and propagation processes undergone?** This work supports theoretical concepts of RFM and emission heights being measurable in polarimetric pulse profiles through the delay-radius relation. However, it puts constraints on emission heights being within narrow ranges such that the delay-radius misalignment can be entirely absorbed by the DM and RM over broad frequency ranges. Calculation of accurate DM and RM values and corresponding understanding of pulsar emission processes therefore requires detailed study of polarization properties of pulsars. Broad-band evolution of pulsar polarization, particularly in relation to circular polarization, points to the interaction of orthogonal modes being a strong contributory factor to the pulse profile behaviour observed. More work needs to be done to understand the role of circular polarization and its relationship to the coherent or disjoint interaction of orthogonal modes.

3. **How can we use new broad-band radio observations to answer these questions?** Broad-band radio pulsar observations reveal the frequency dependence of pulsar emission in unprecedented detail, reducing the degeneracy of models attempting to describe its behaviour. It is important to be aware of the limitations of any given model that cannot describe the data to full precision, for example my neglect of circular polarization in the quantitative analysis of

broad-band polarization in Chapter 6. In this context, an in-depth analysis of both simulations and a small number of pulsars reveals these limitations prior to applying such modelling to large samples. By approaching new broad-band pulsar data from this perspective, it was possible for me to identify and study the individual aspects of radio pulsar observations in ways not previously possible, as described in the questions above. Doing so has contributed to our understanding of pulsar radio emission behaviour as a whole.

9.4 Going forward

The framework for understanding broad-band pulsar polarization and its interaction with the ISM has now been built; what remains is to apply it to wider analysis of many pulsars. There remains complexity to address, in terms of both understanding diverse polarimetric behaviours and how the key observational features intersect in more complex cases, such as scattered pulsars with multiple components. It is at the point at which a model breaks down that its utility for uncovering new behaviour is at its greatest. The approach laid out in this thesis has been propelled by a focus on these limiting scenarios, and how broad-band observations help break the degeneracies that result from the cases where the effects of profile evolution and the ISM intersect. Applying the same approach to address questions about circular polarization, complex pulse profile shapes and time-variable interactions with the ISM will continue to give answers about the behaviour of radio pulsars.

Future work will focus on finding the origins of pulsar polarization and solving pulsar beam structure. For this a detailed understanding is needed of the broad-band polarimetric properties of a large number of pulsars, in addition to understanding the effects of propagation through the ISM on the observations. Using this information will enable me to improve fundamental experiments in physics, through creating broad-band templates to increase the accuracy of pulsar timing. In addition to extending the analysis processes developed in this thesis to multiple pulsars, future work must focus on the role of circular polarization and the broad-band polarization properties of single pulses. The results of this work will also enable me to extend our understanding of the ISM by applying pulsar beam science to models of more complex scattered pulse profiles.

The new broad-band observations made possible with the latest radio telescopes open up a new swathe of possibilities for understanding the radio emission from pulsars. These possibilities will only increase with the arrival of the Square Kilometre

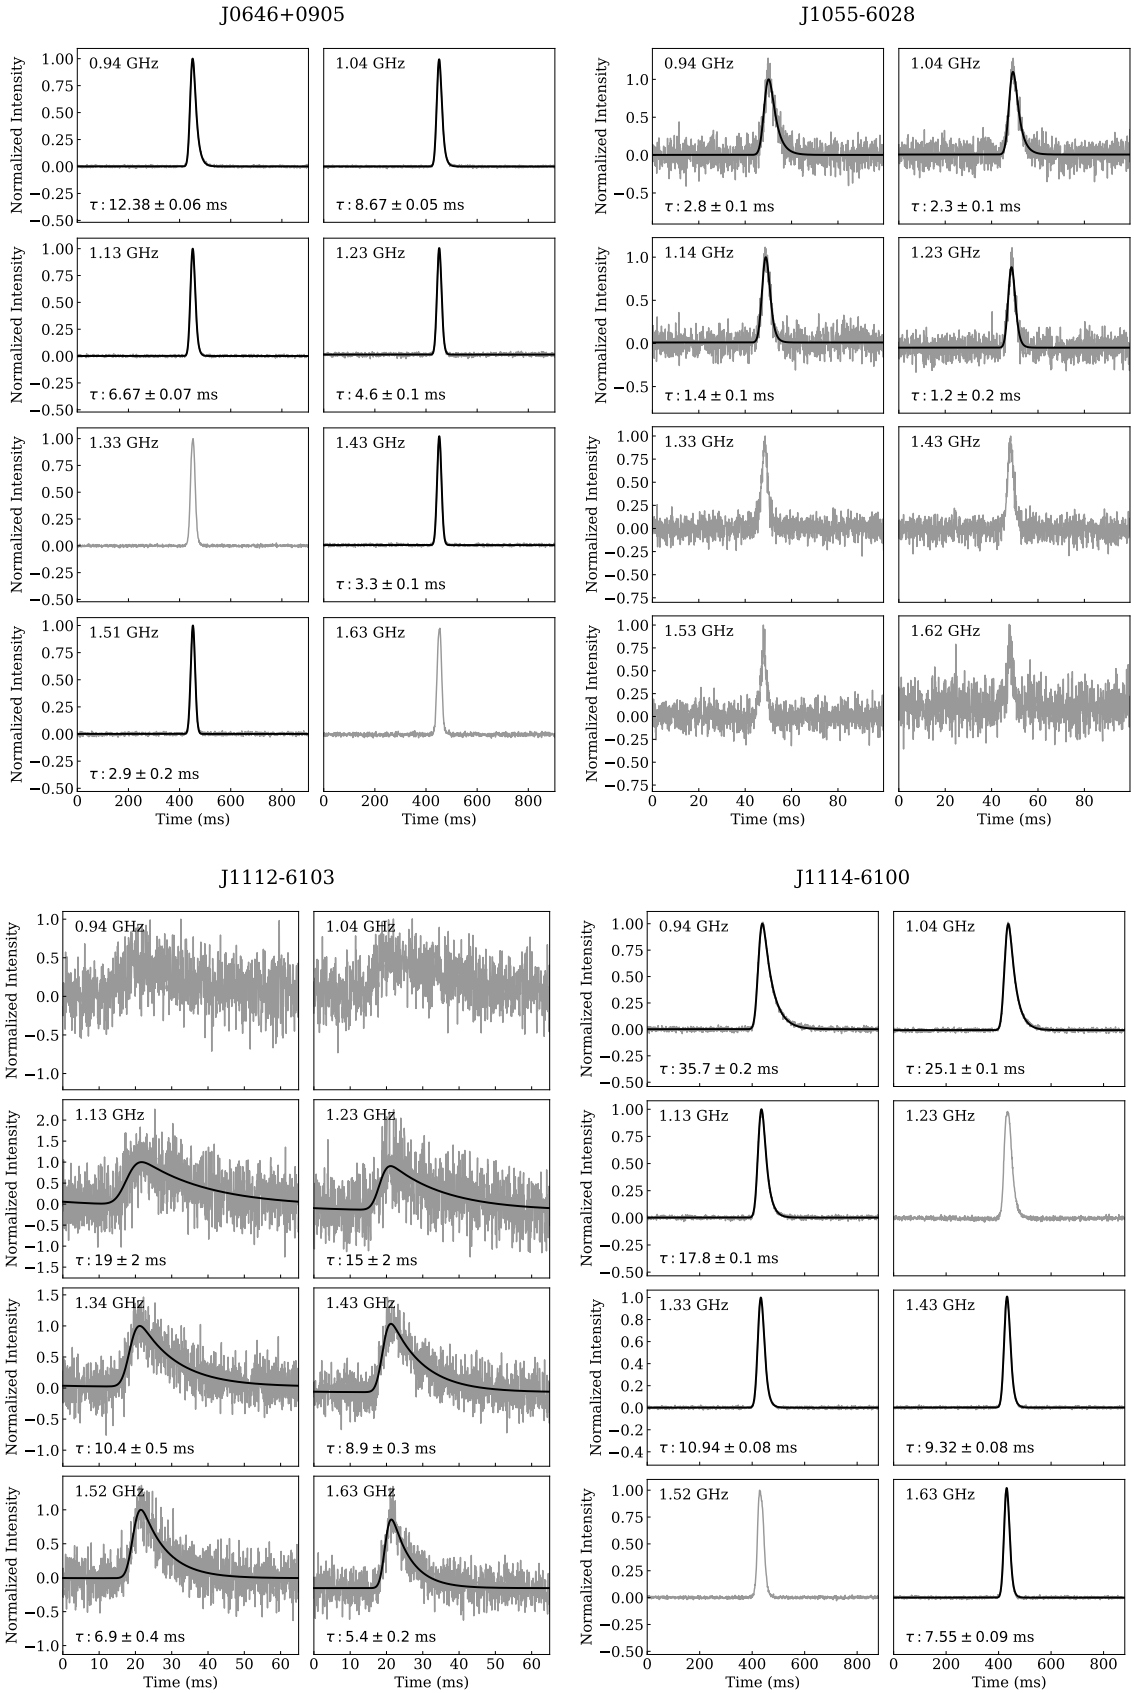
Array telescope, due to begin construction in 2021. This research has the potential to explain not only the extreme physics of neutron stars and the complexities of the ISM, but to answer fundamental questions about the laws of physics and the Universe.

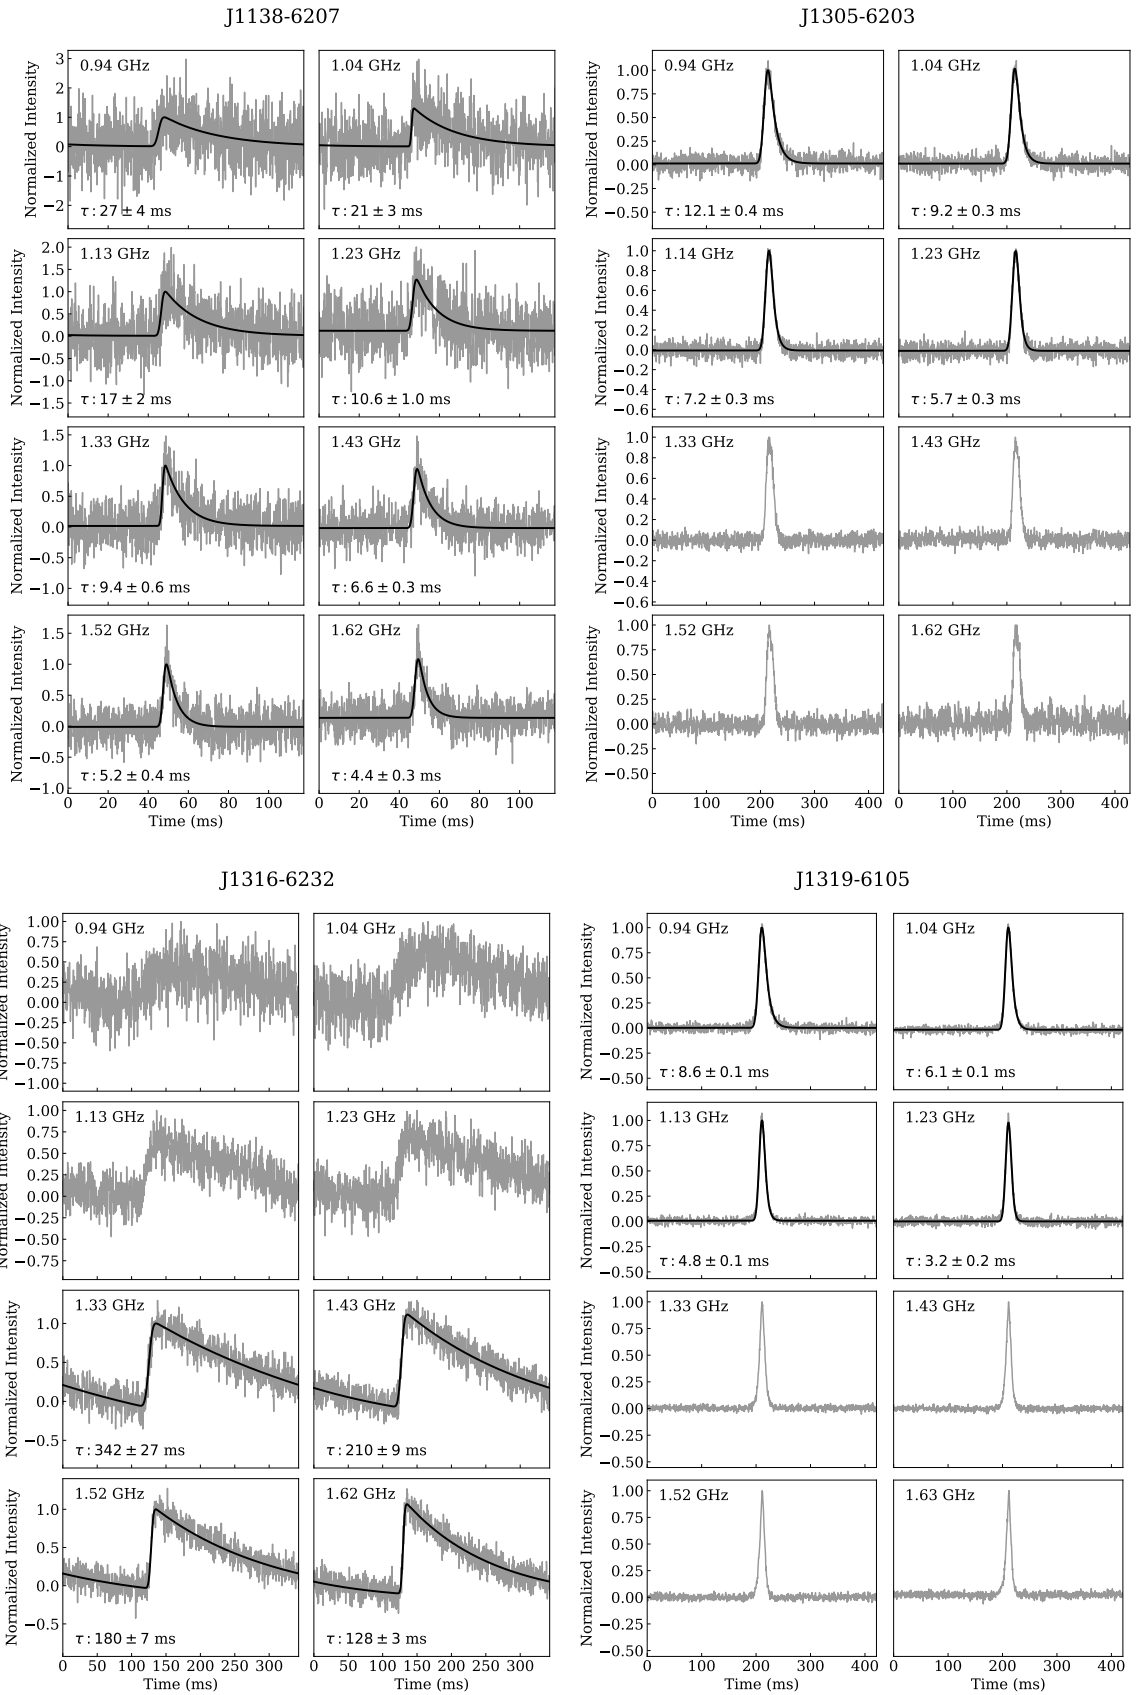
Appendix A

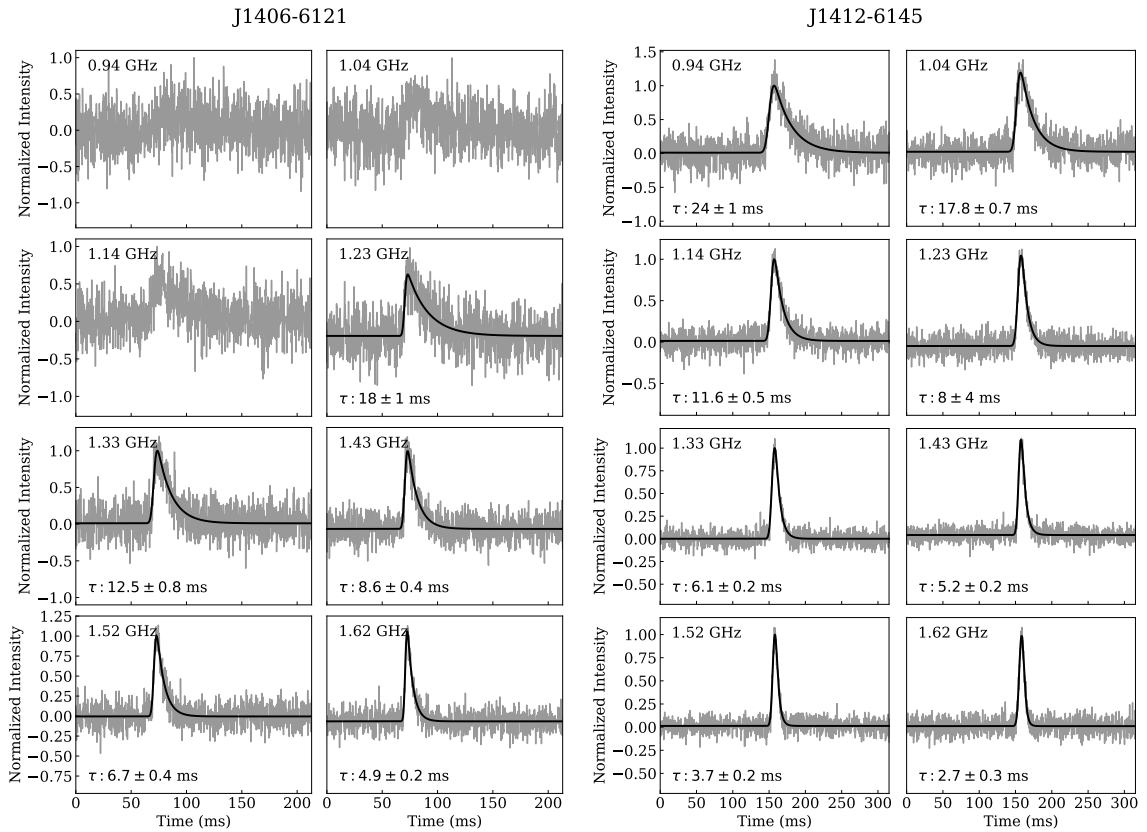
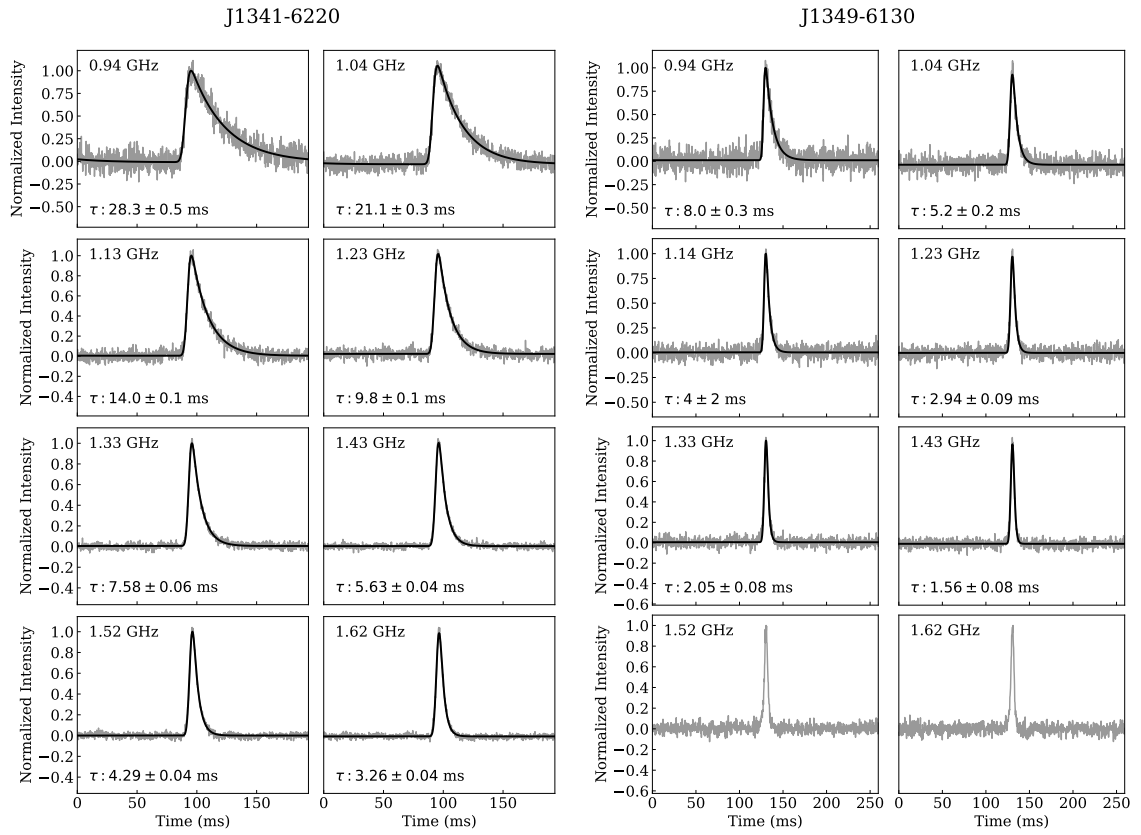
Scatter modelling results: pulse profiles with best fit models

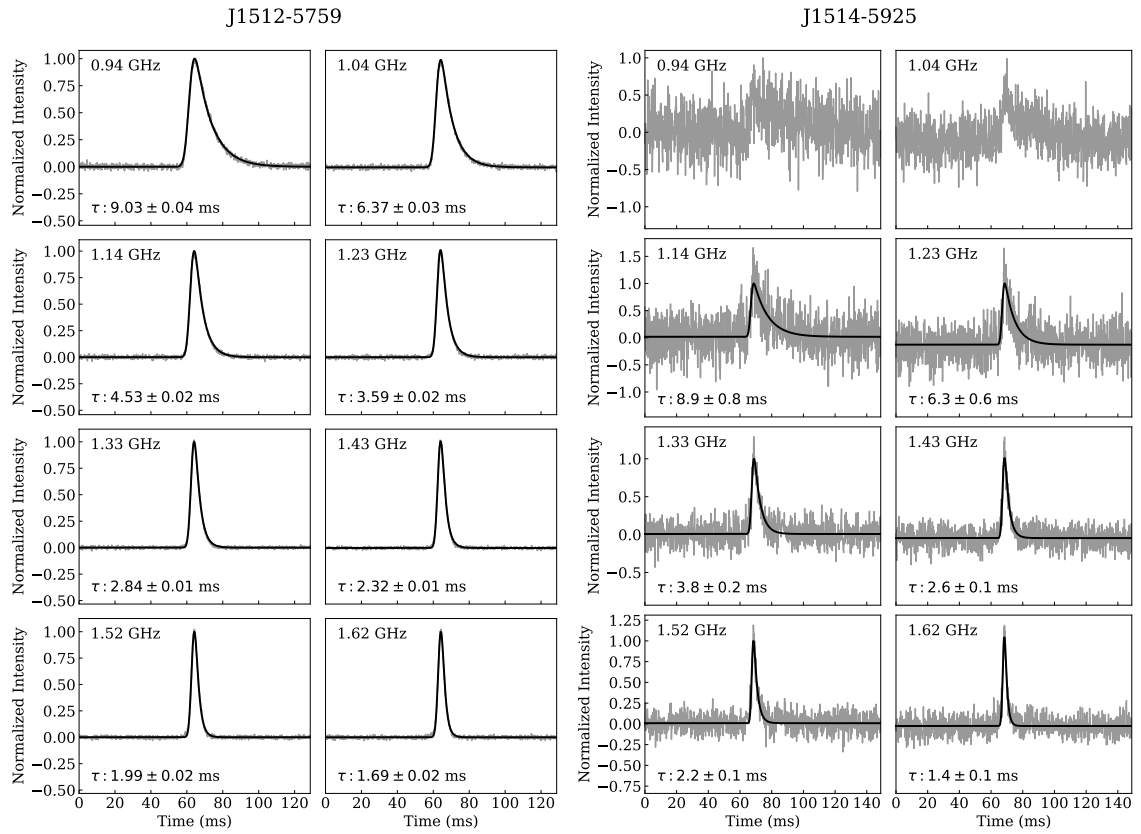
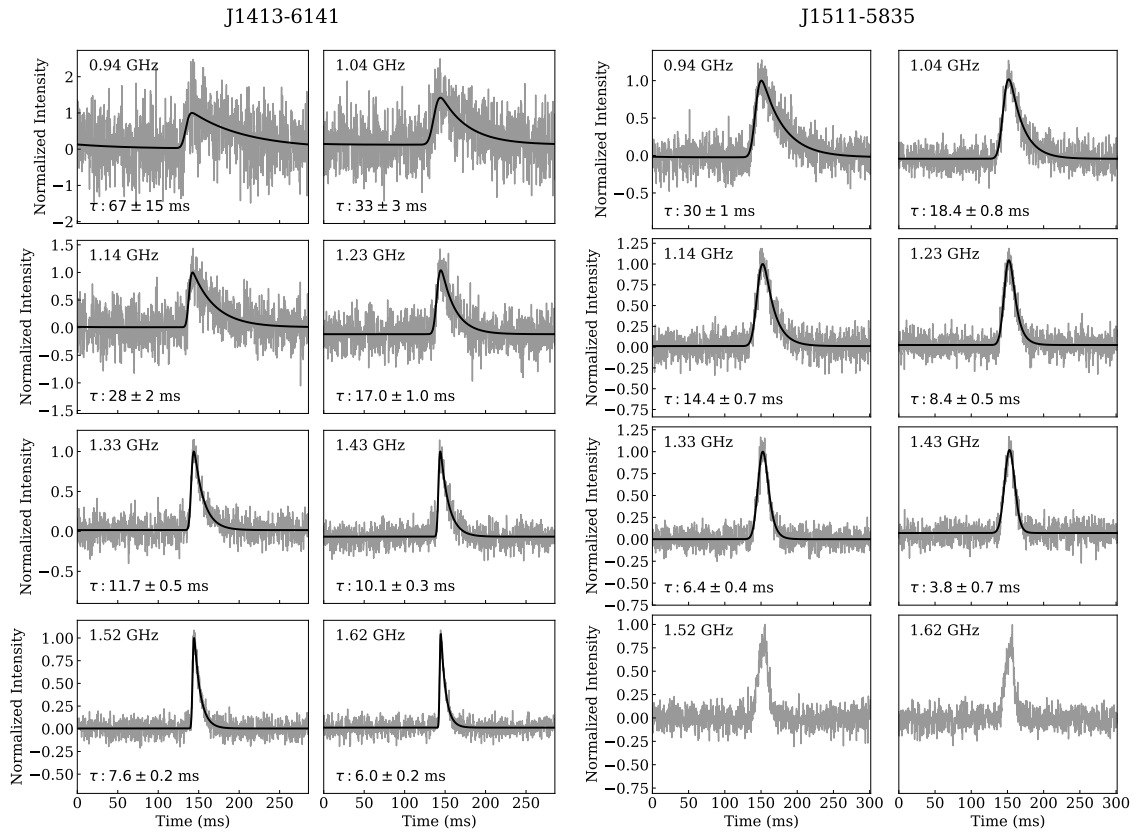
The figures in this appendix show plots of pulse profiles for all 84 successfully modelled scattered pulsars detailed in Chapter 7 and Table 7.1. The profiles are shown with successful model fits overlaid, as in Fig. 7.2. For profiles at discrete frequencies where model fitting was unsuccessful, the profile alone is shown. Each pulsar has at least four successful model fits out of the eight pulse profiles that span the observing band.

Figure A.1:











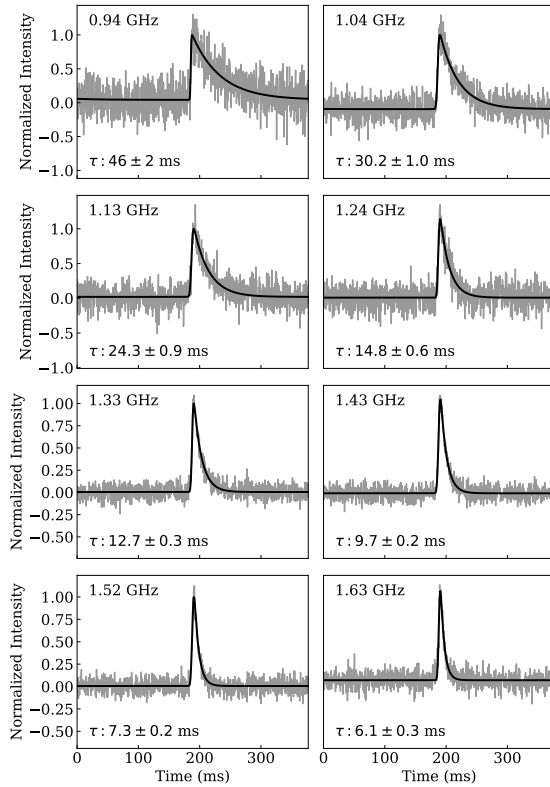
J1519-5734




J1538-5551

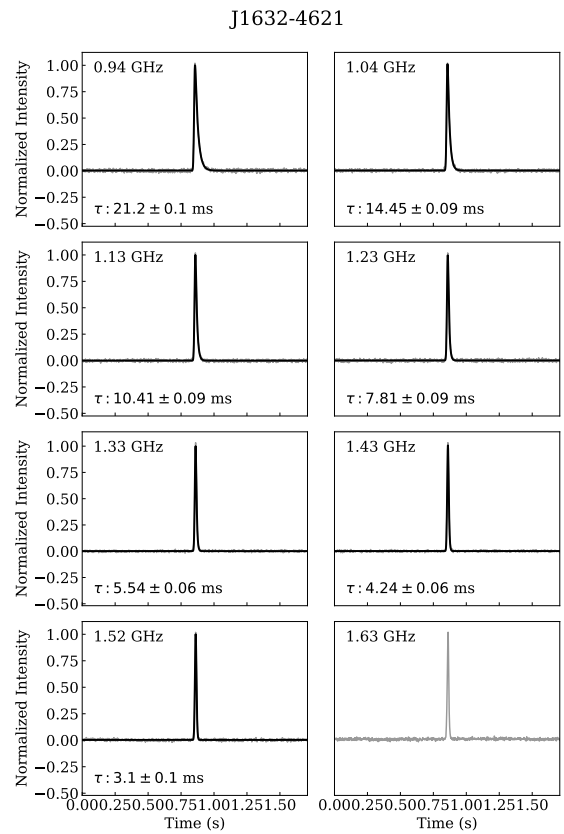
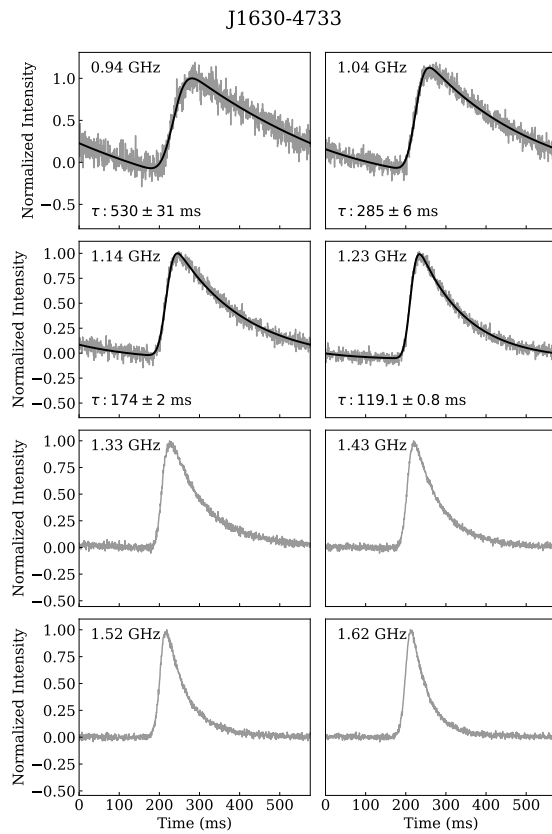

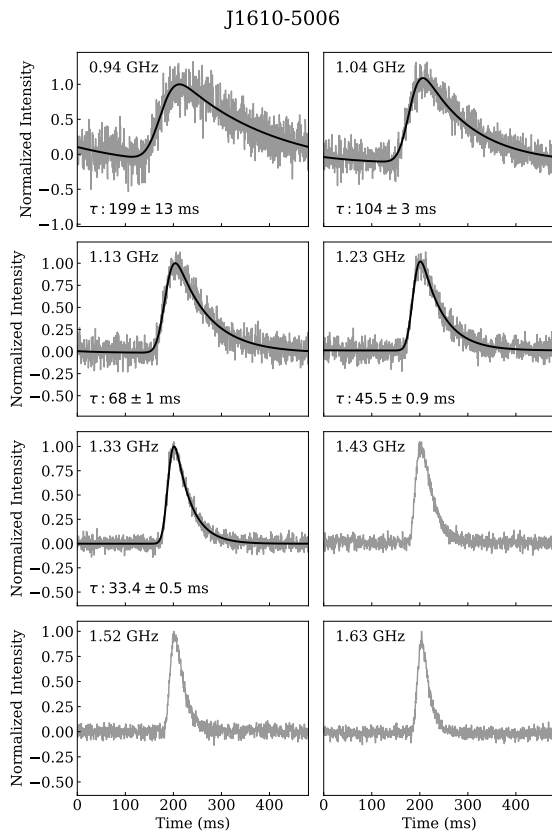


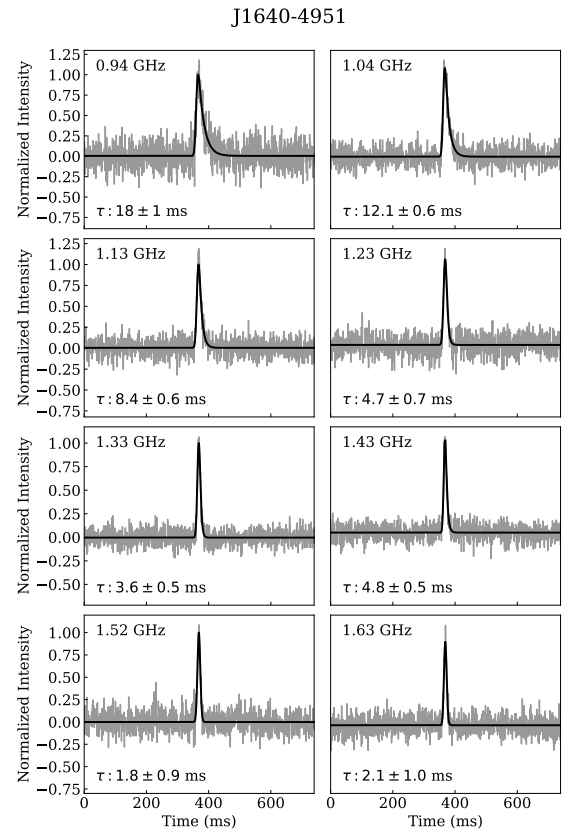
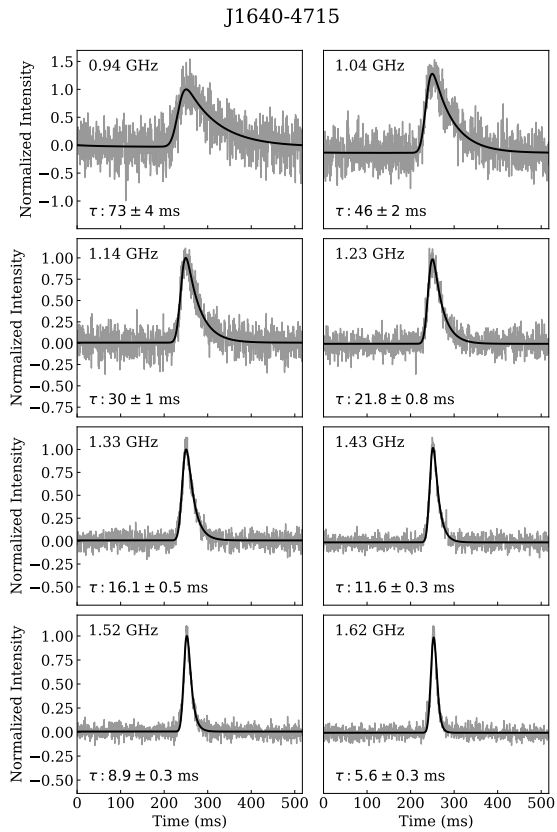
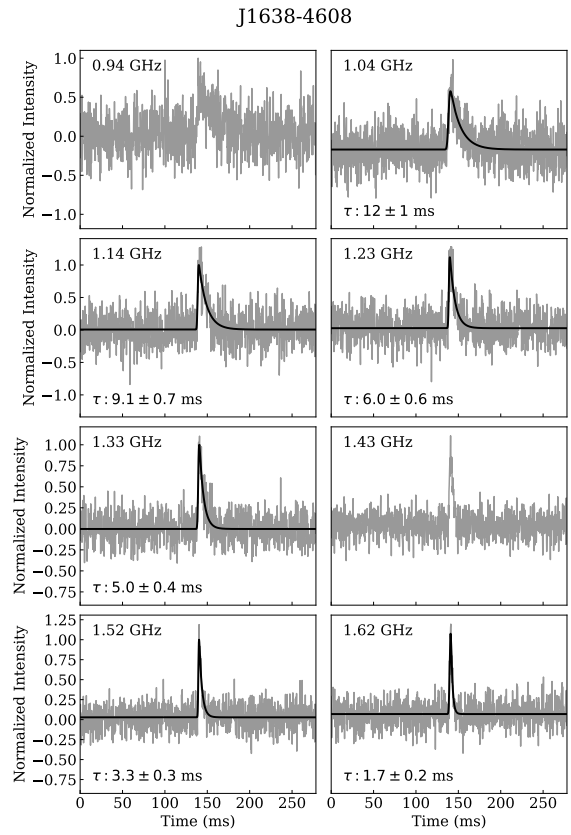
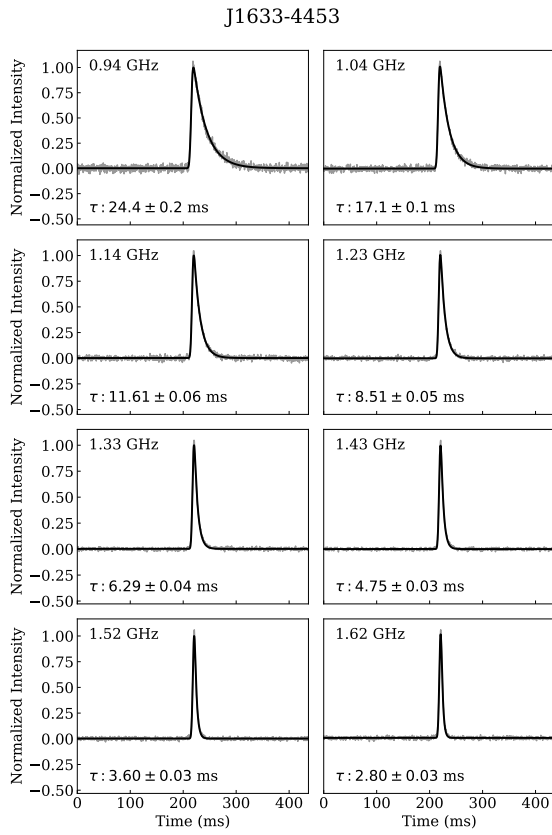
J1543-5459

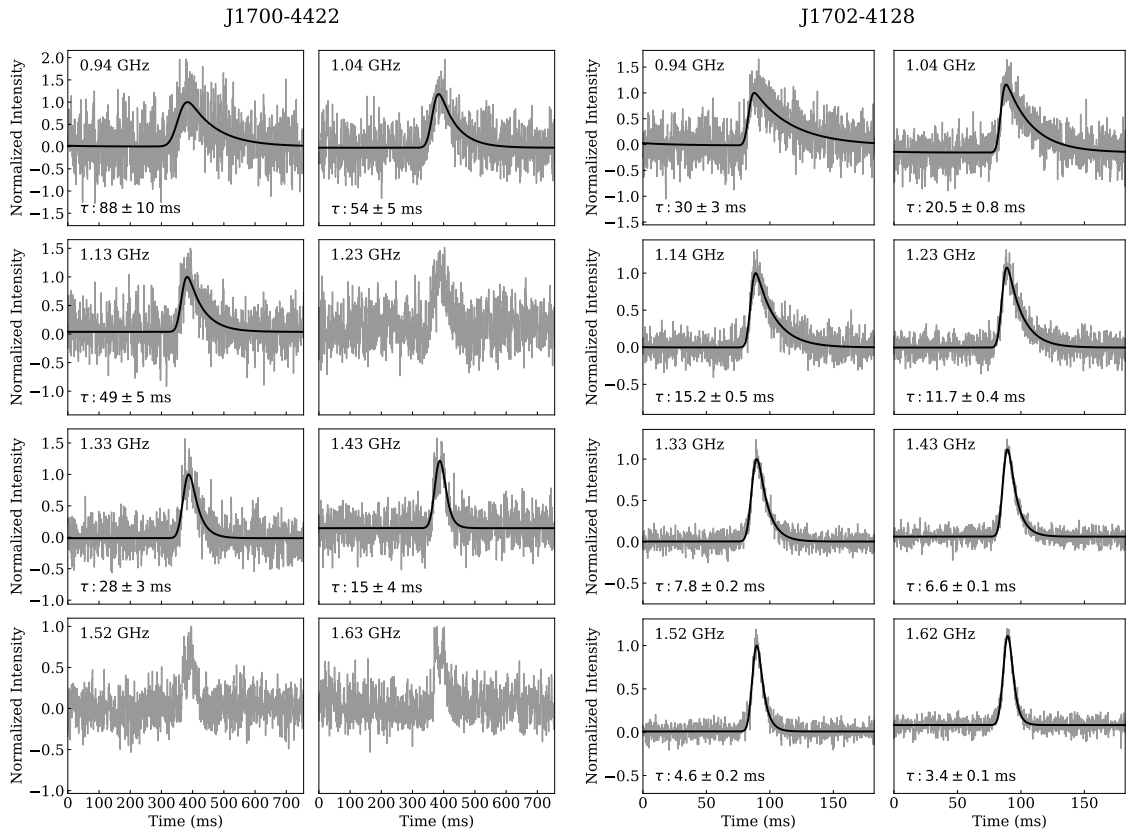
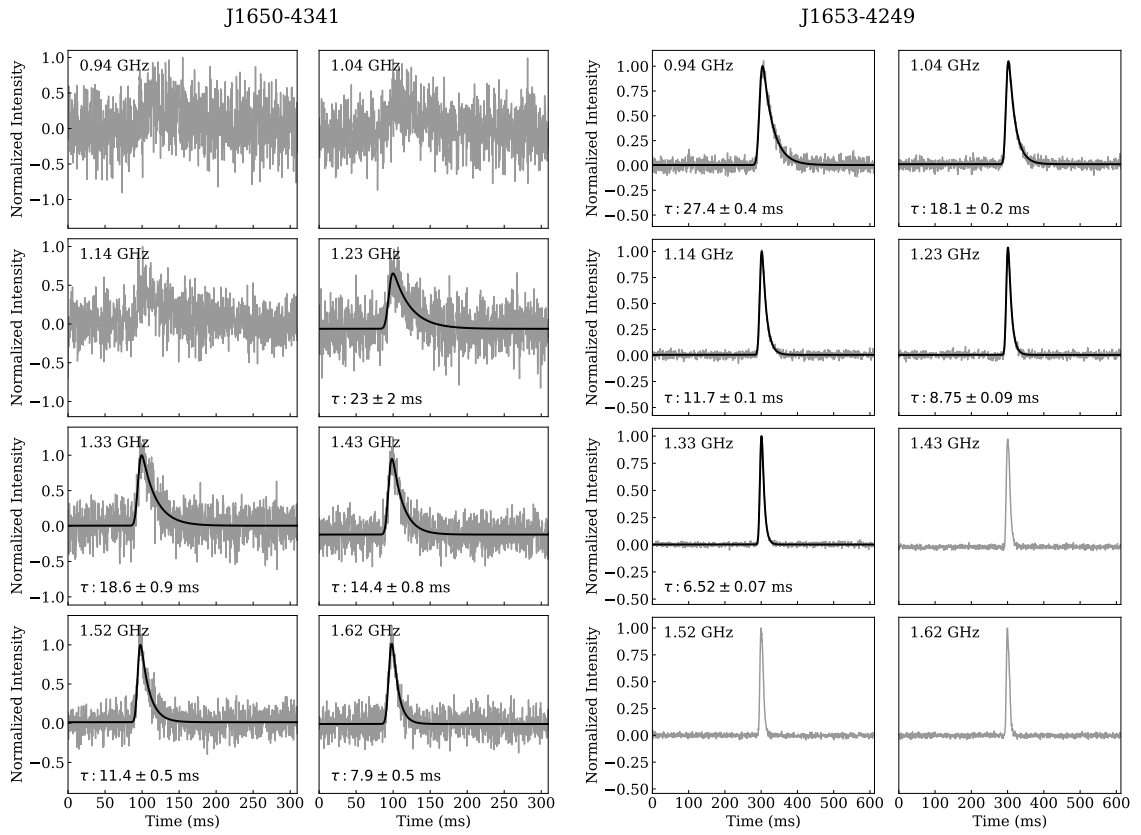


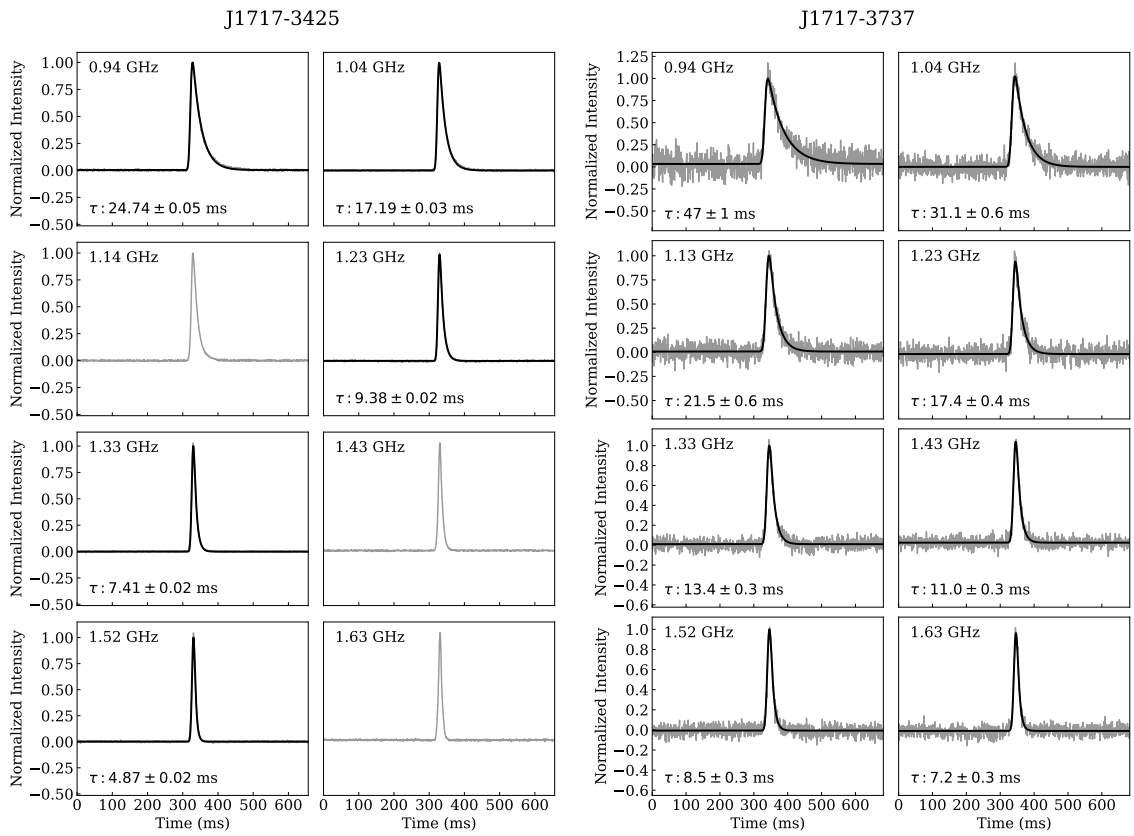
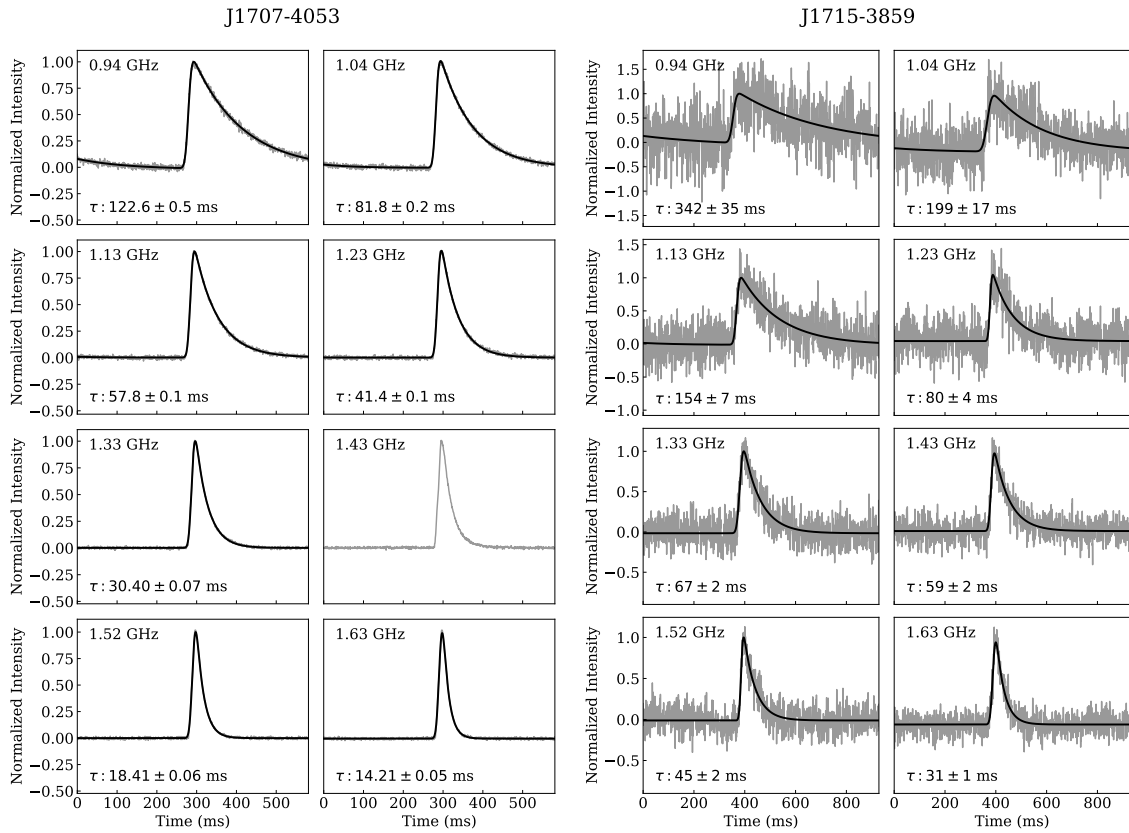
J1551-5310

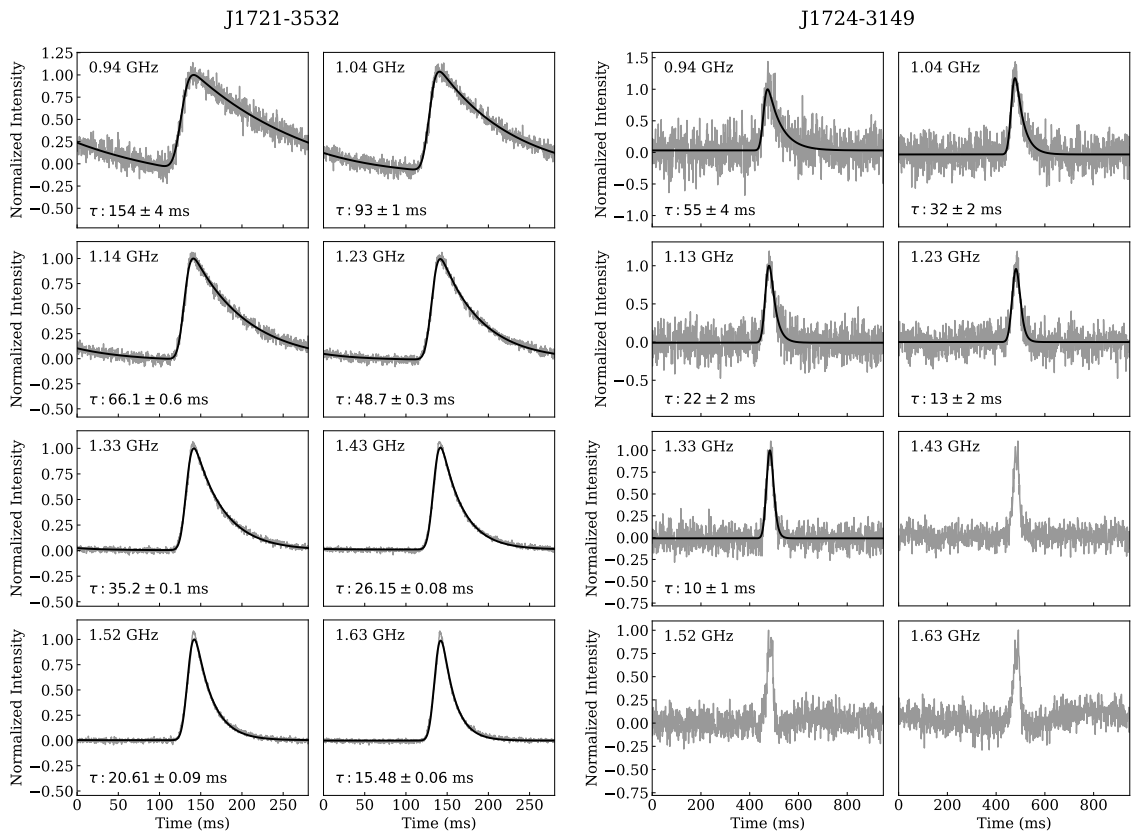
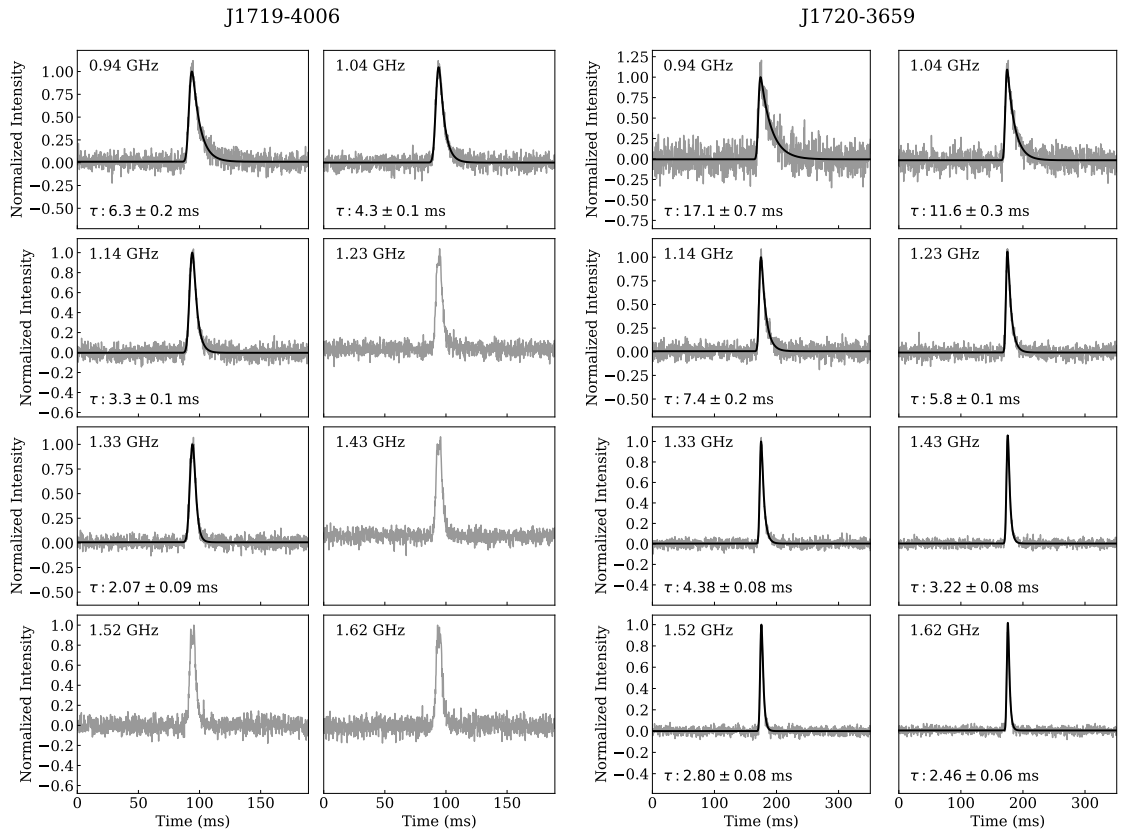


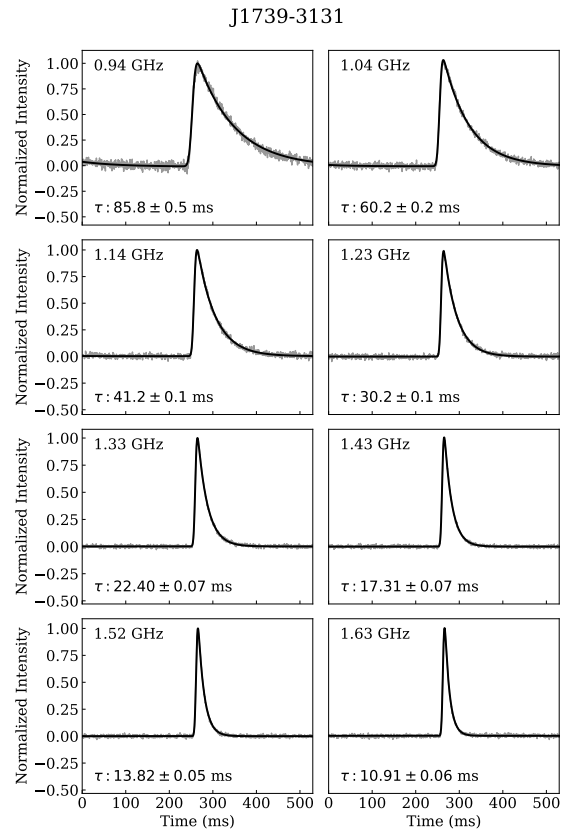

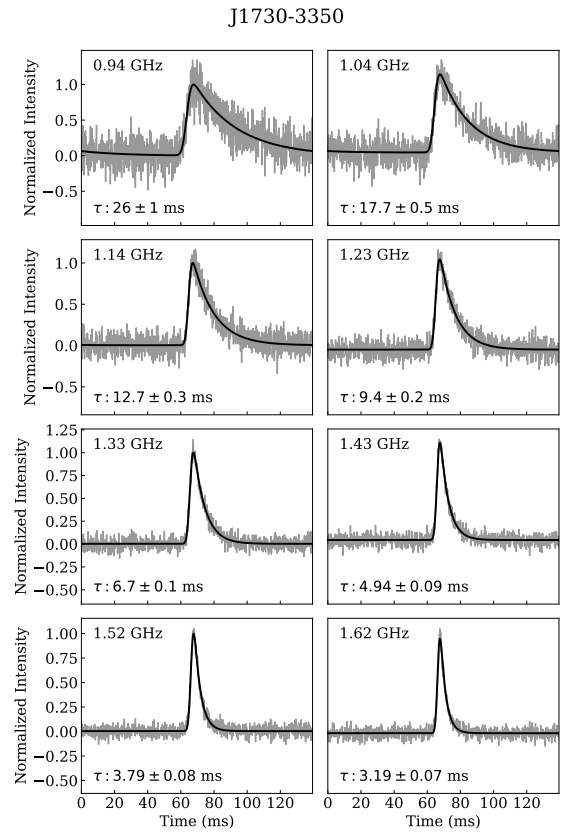
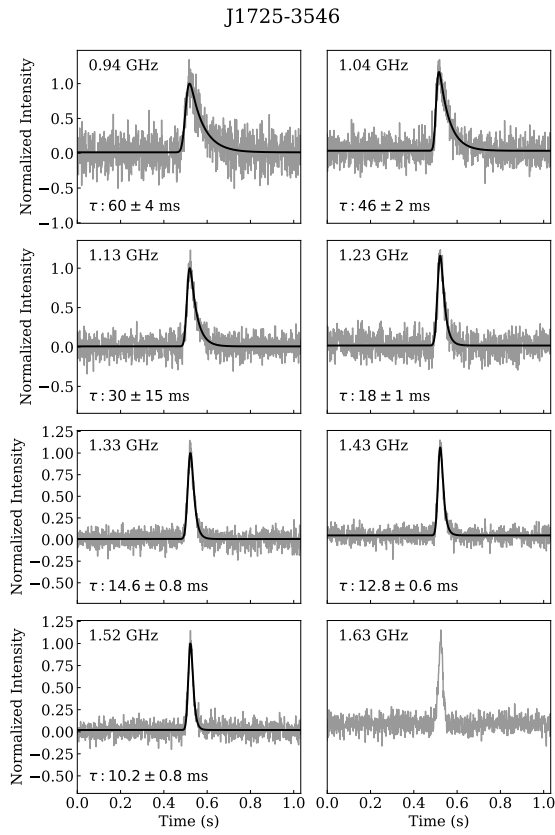


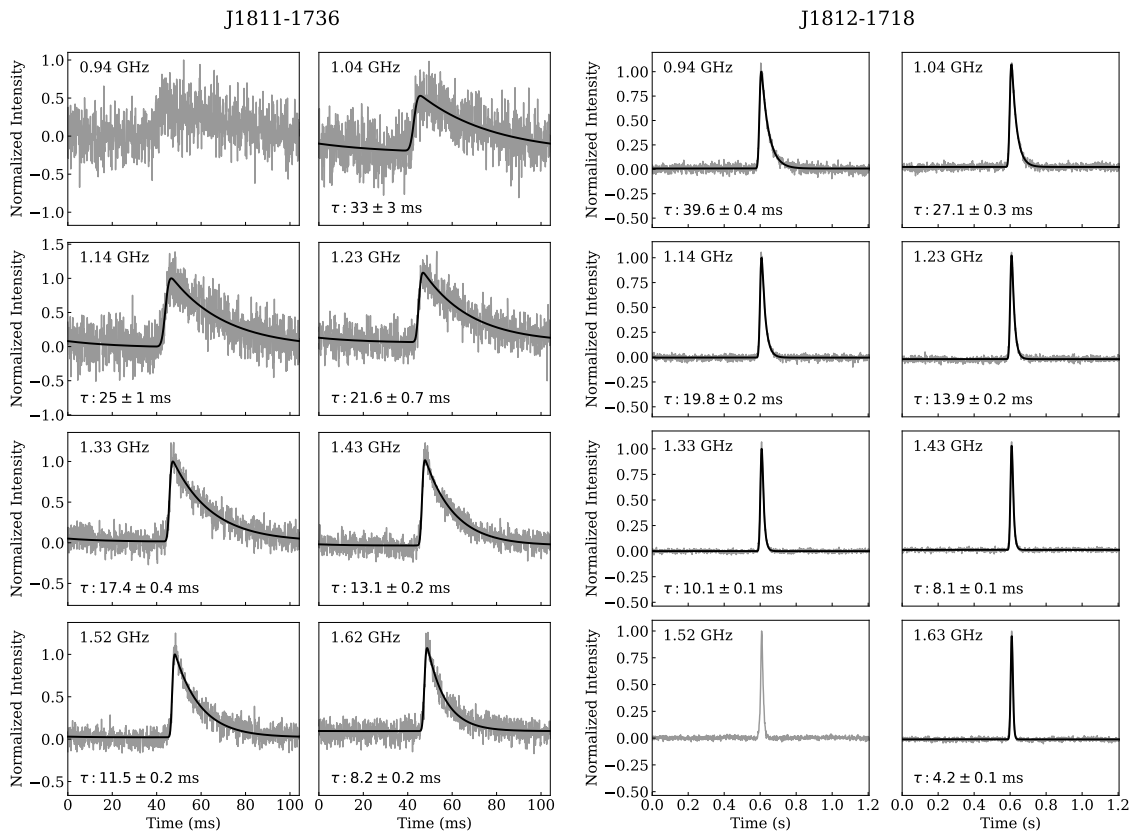
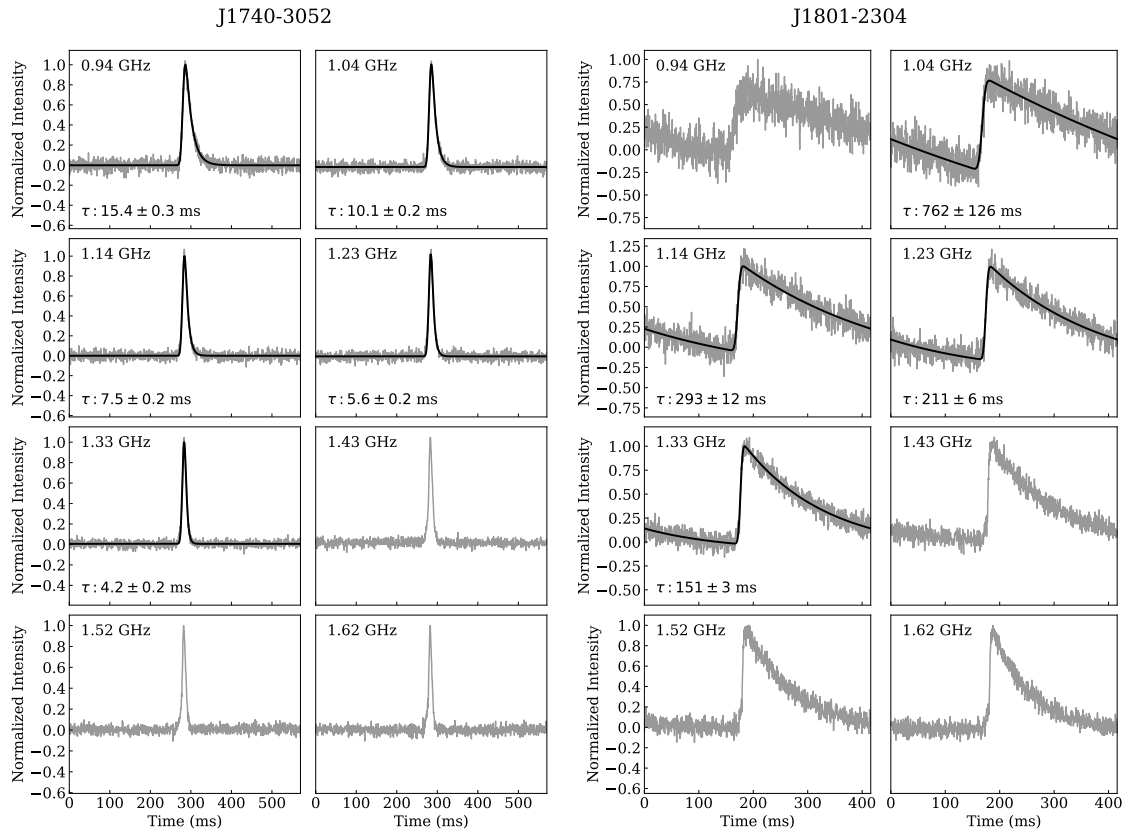


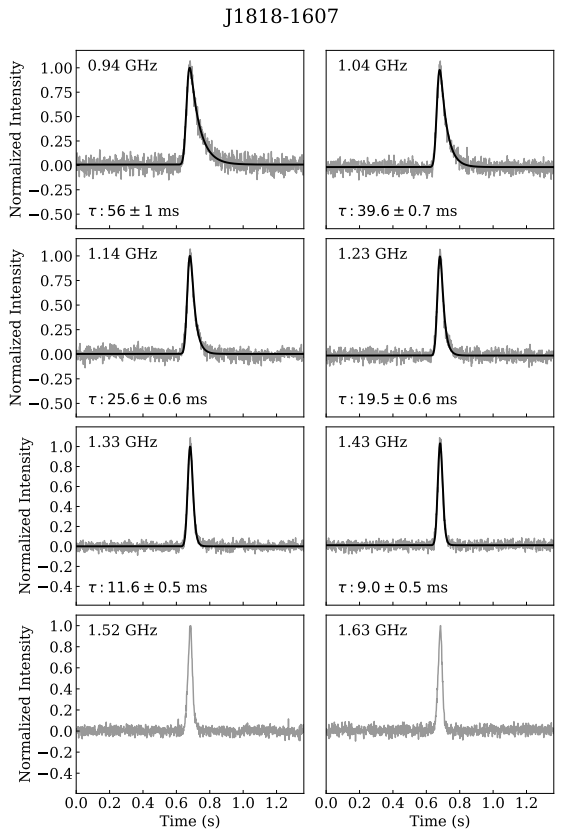
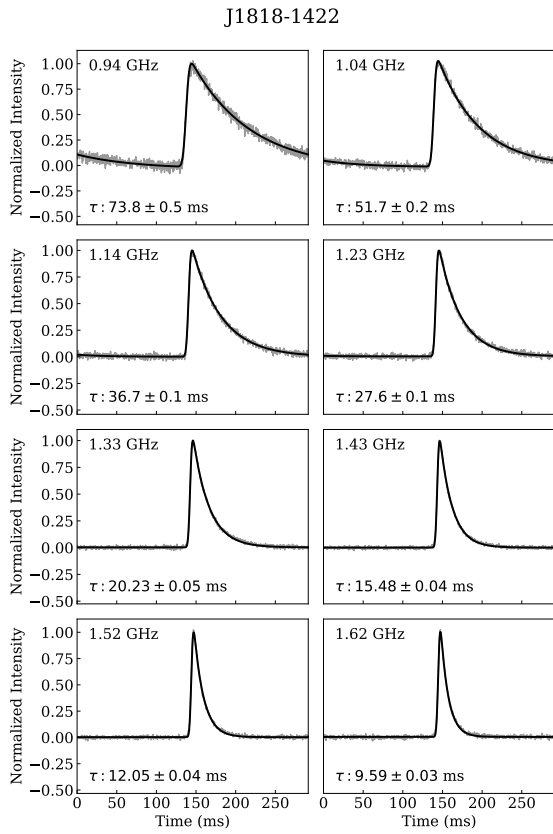
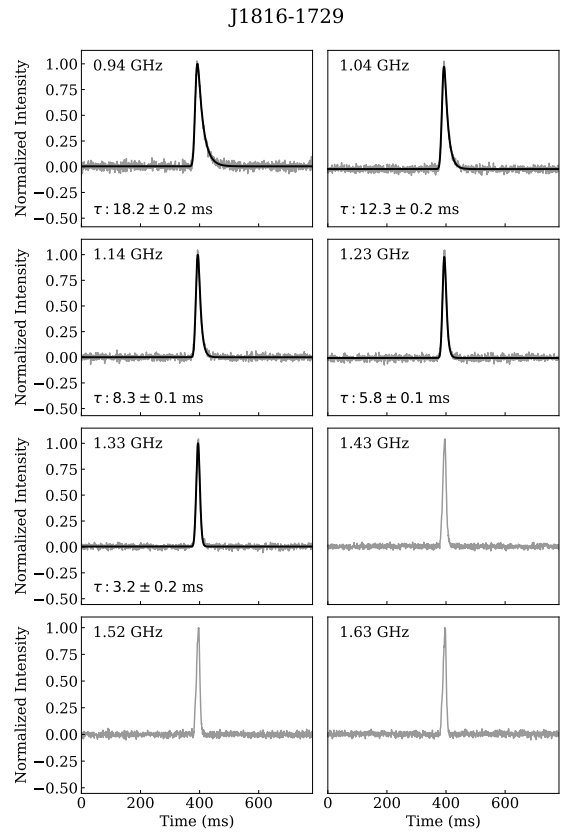
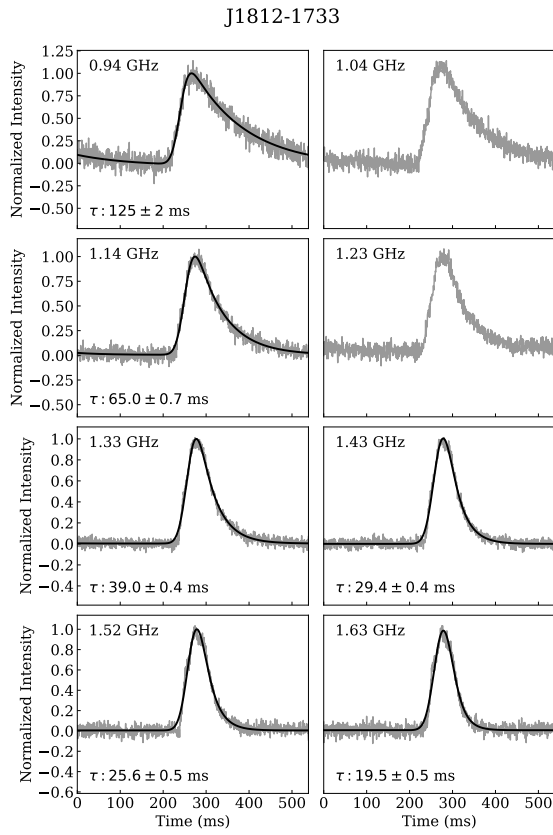


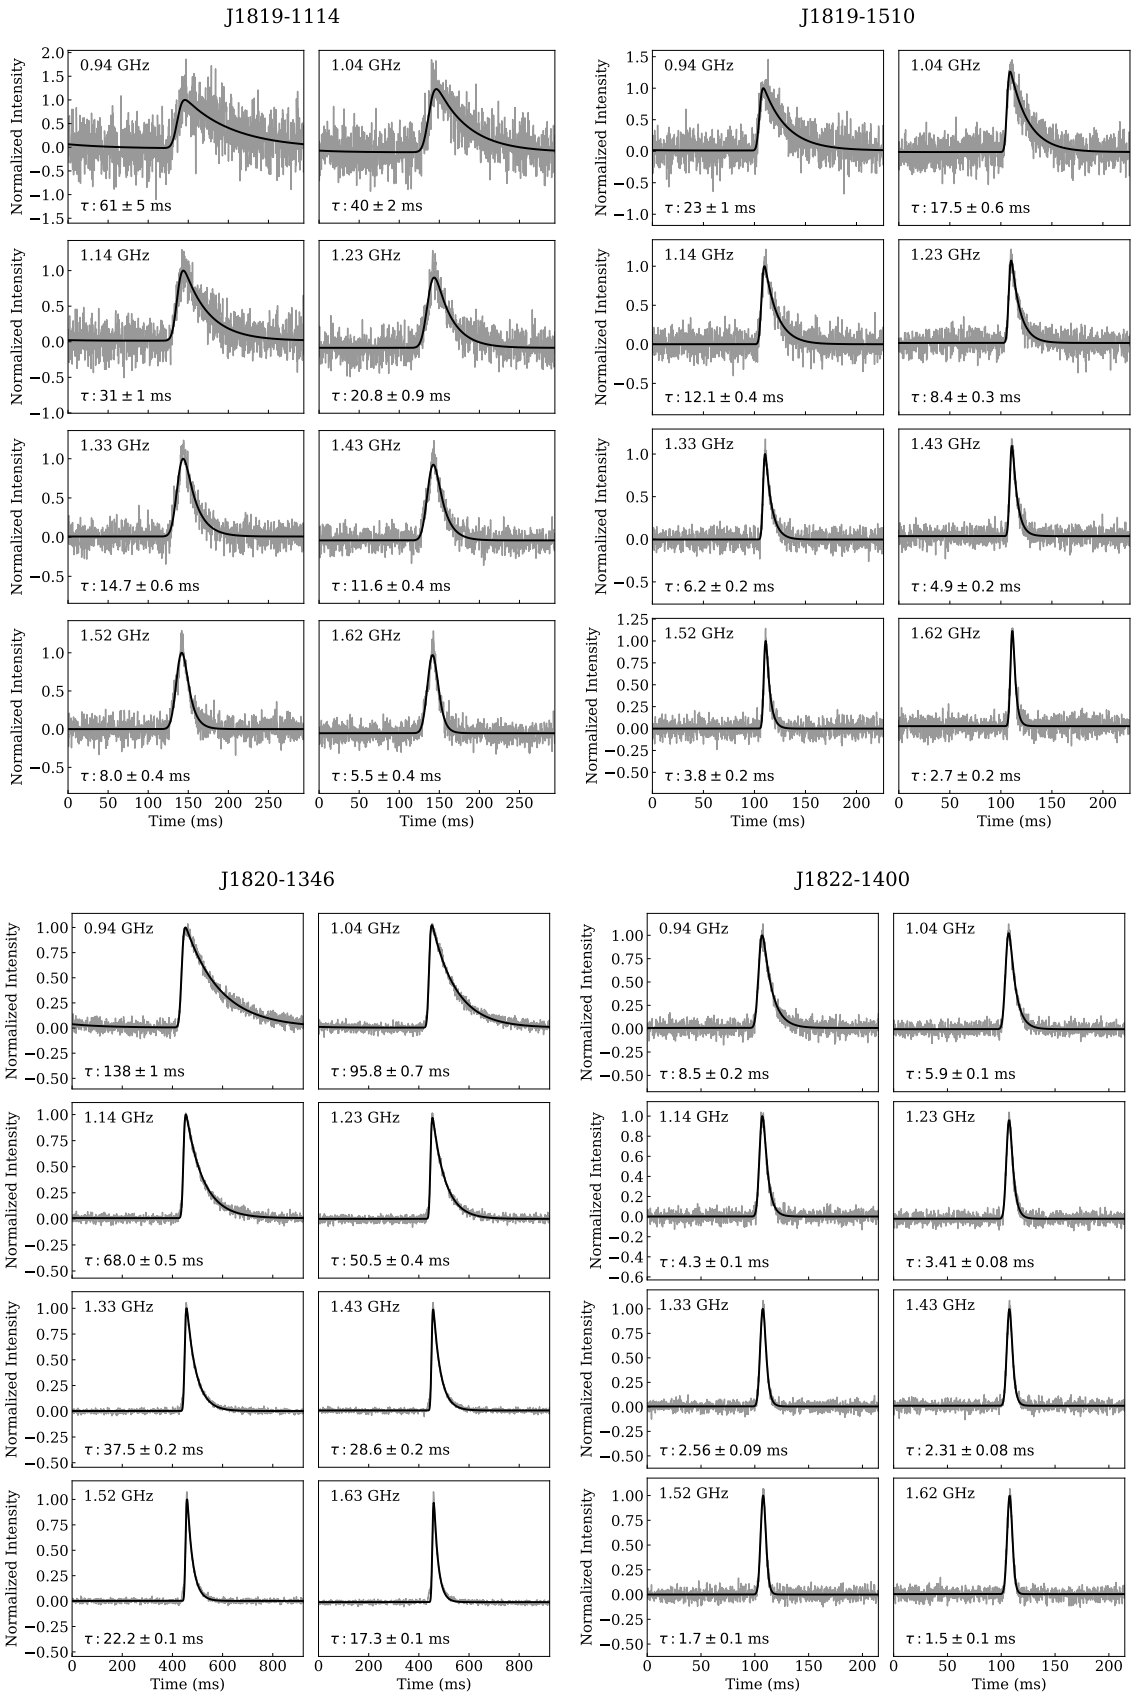



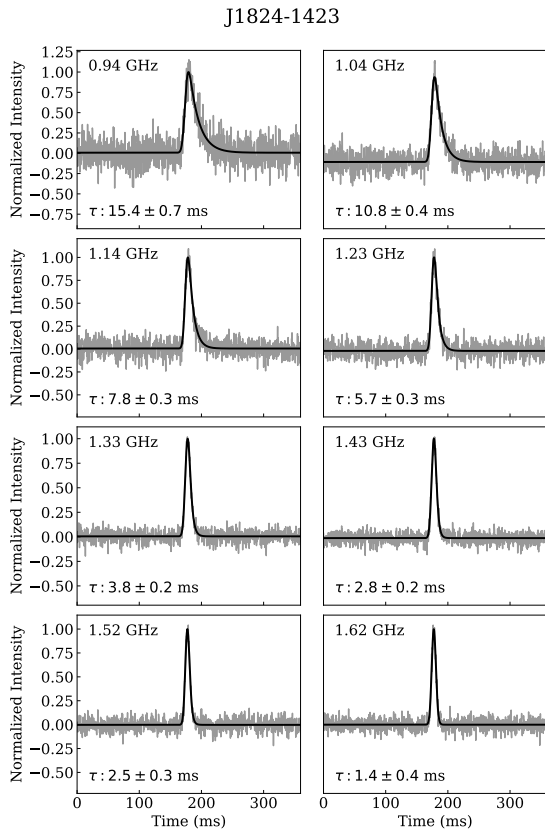
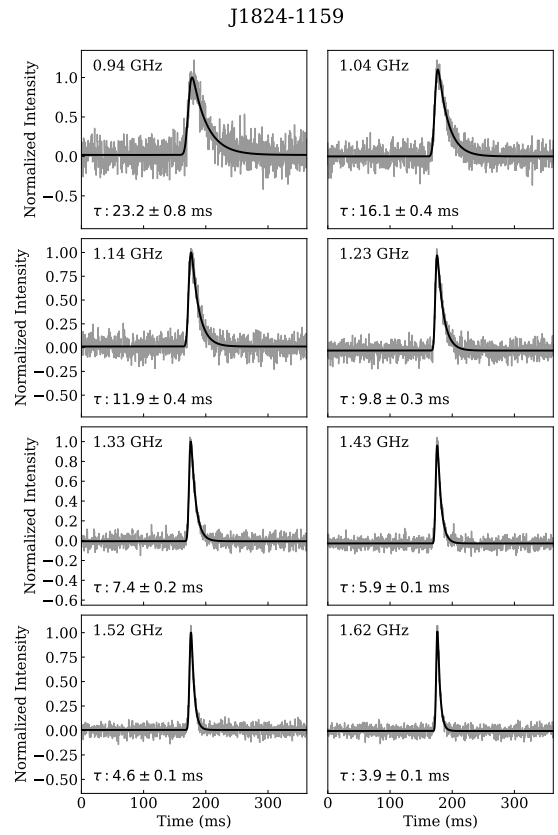
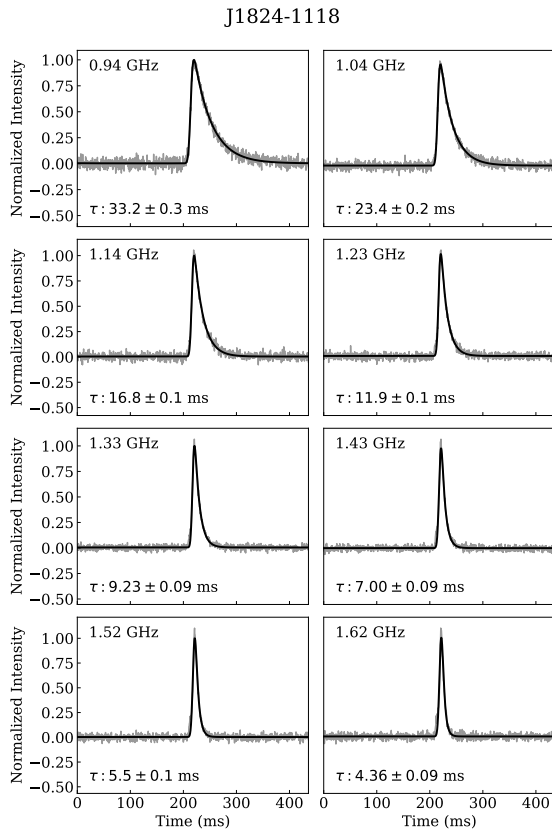


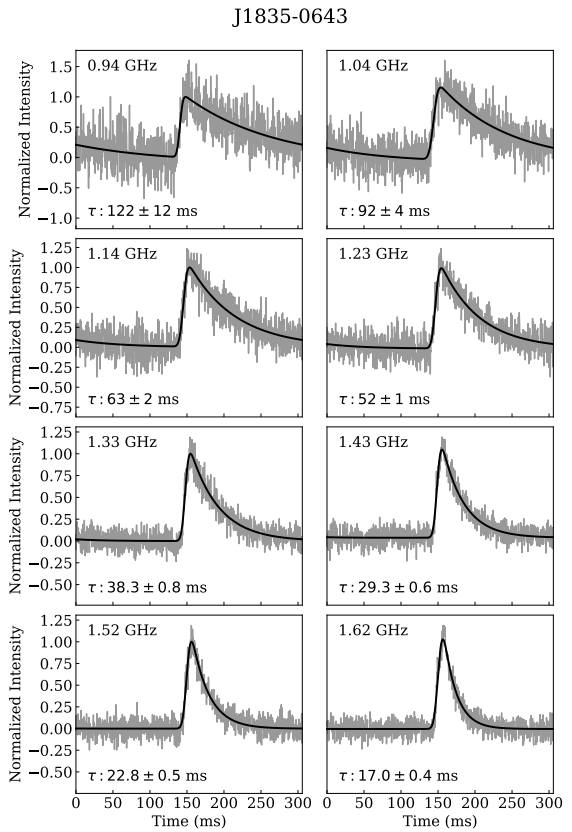
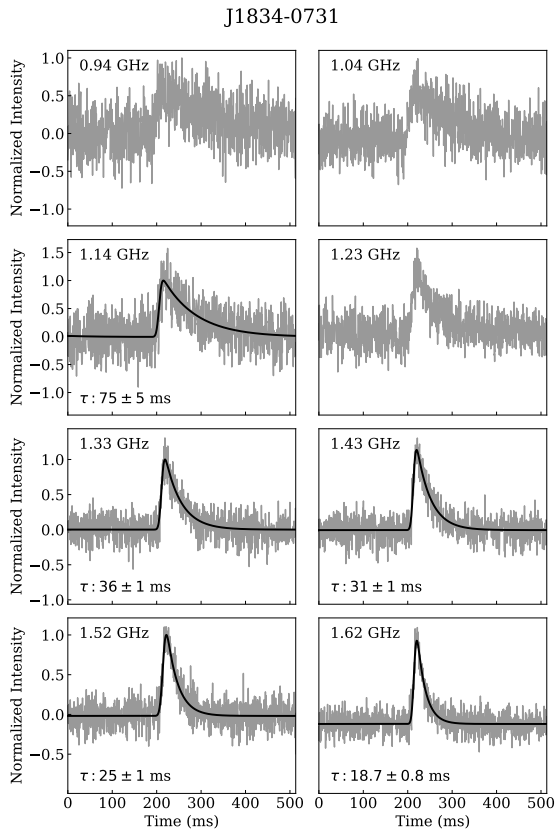

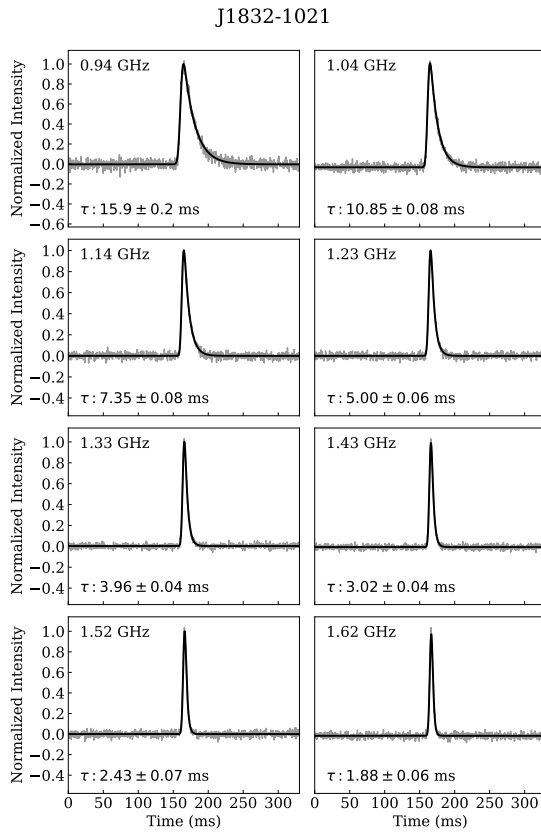


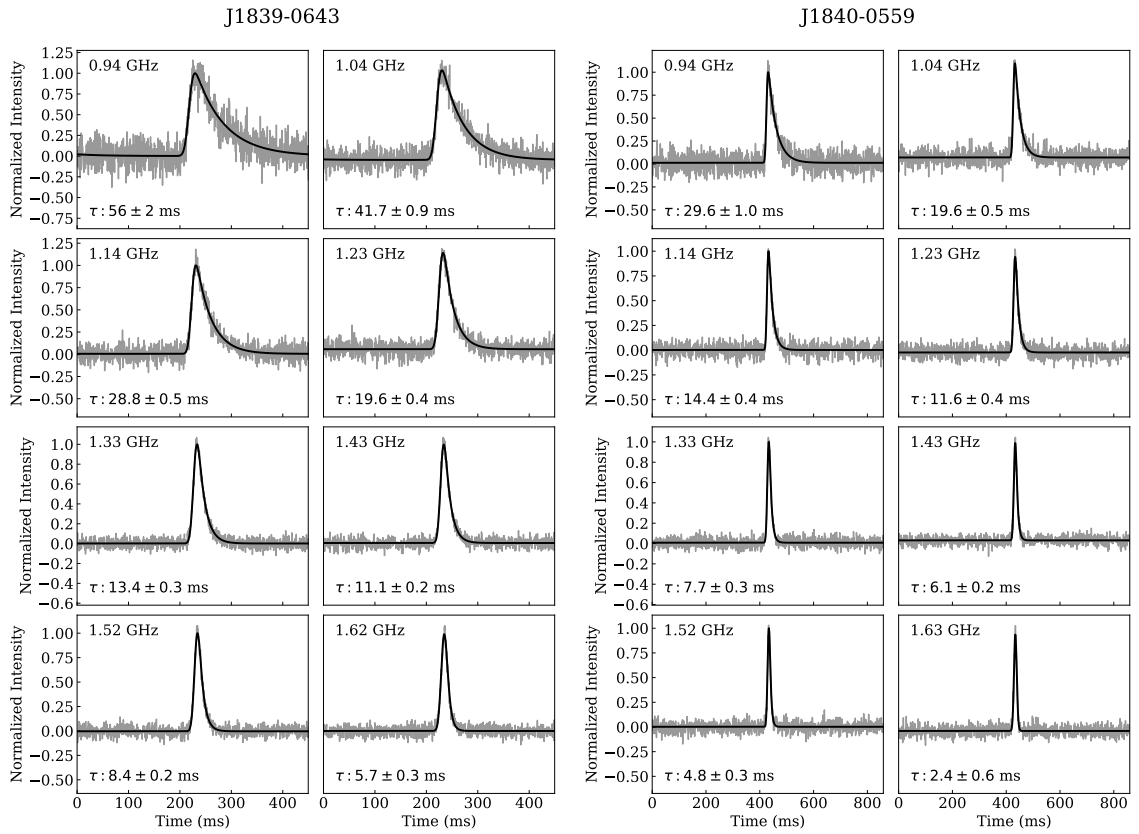
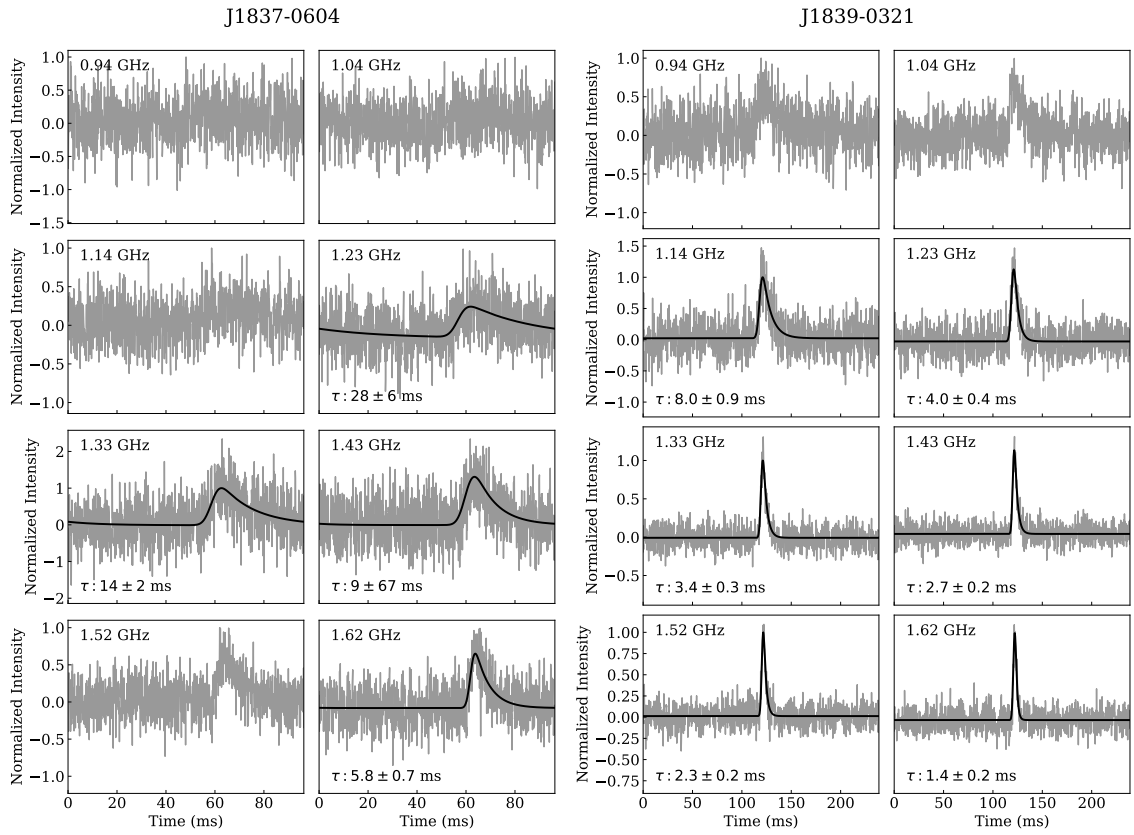



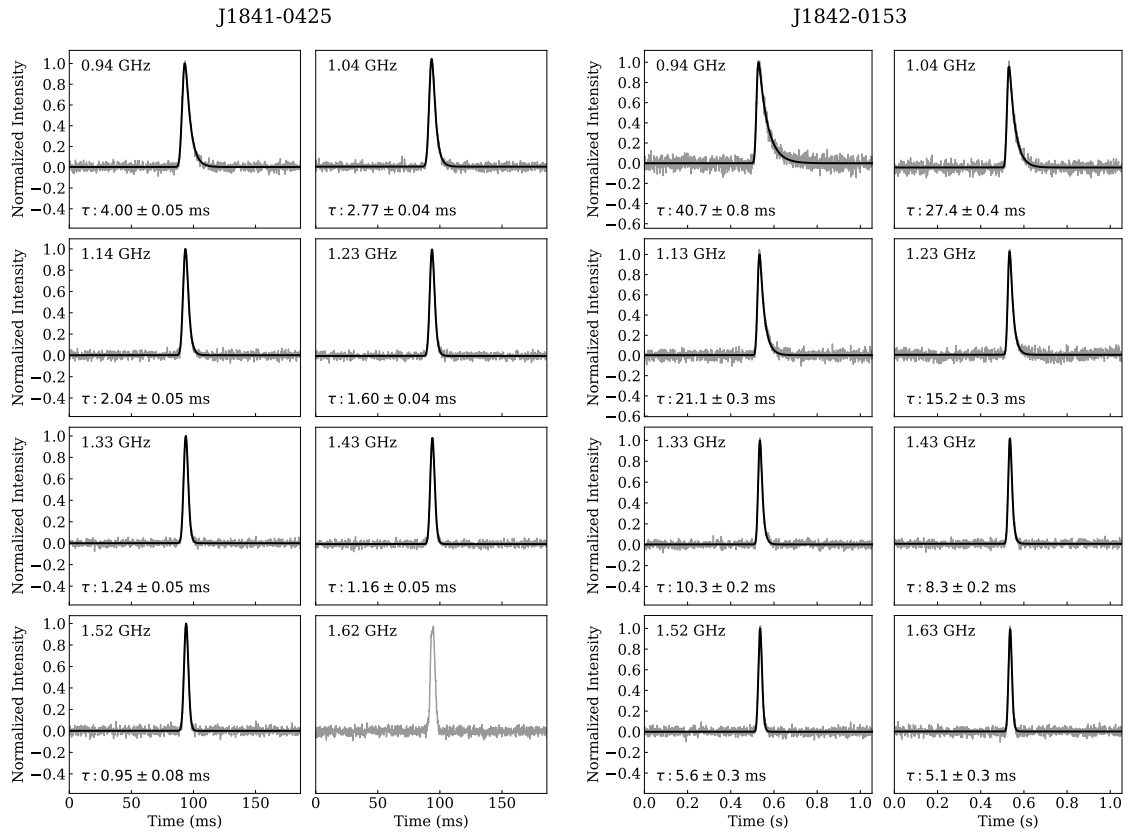


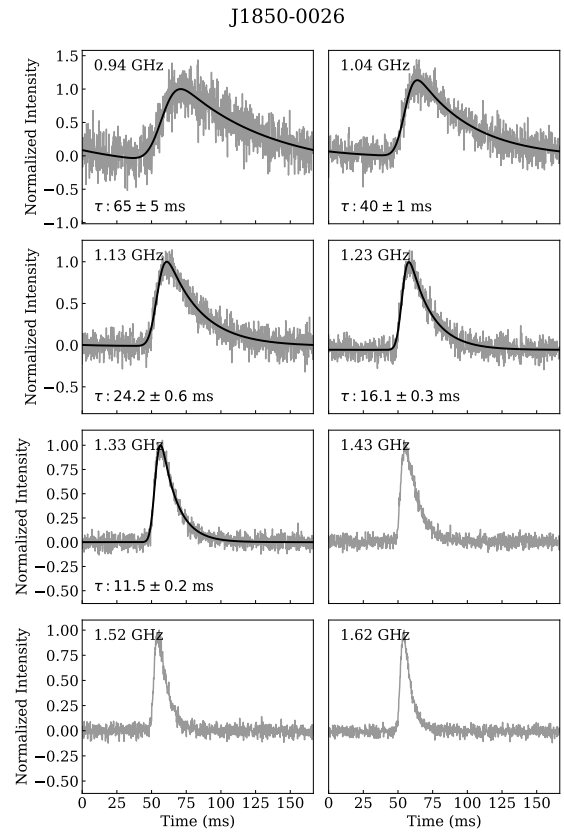
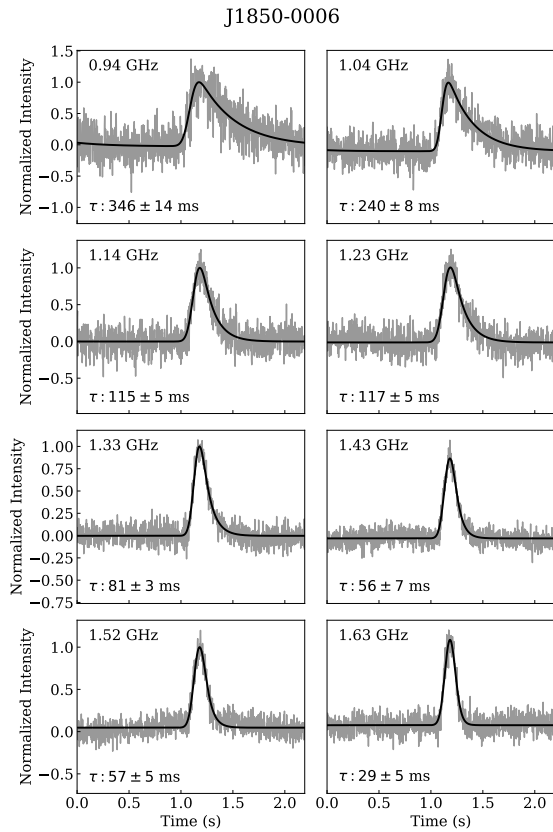
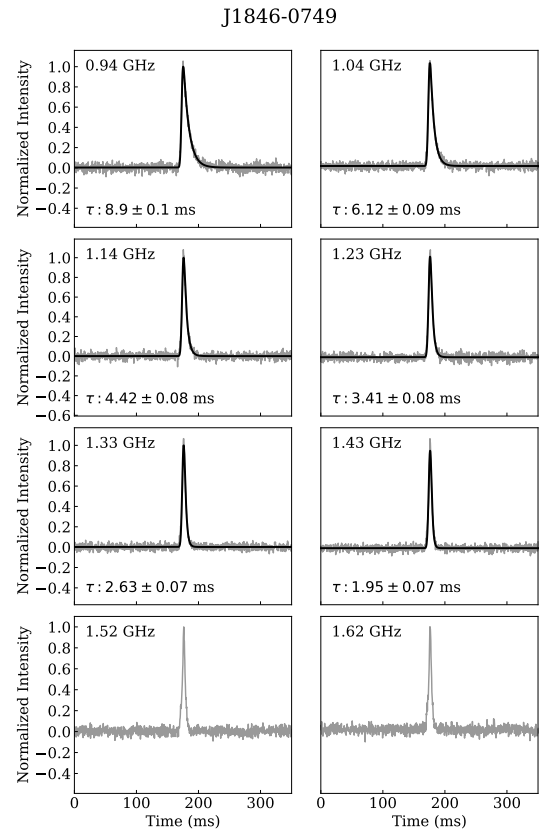
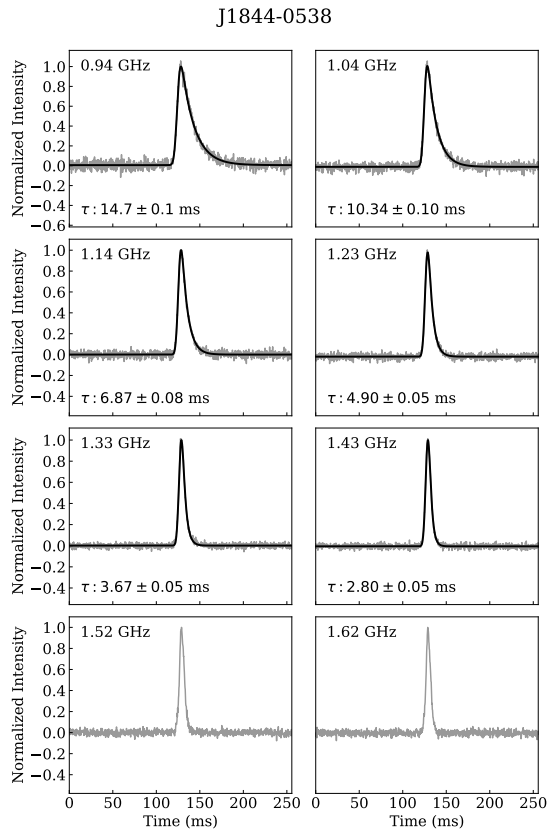


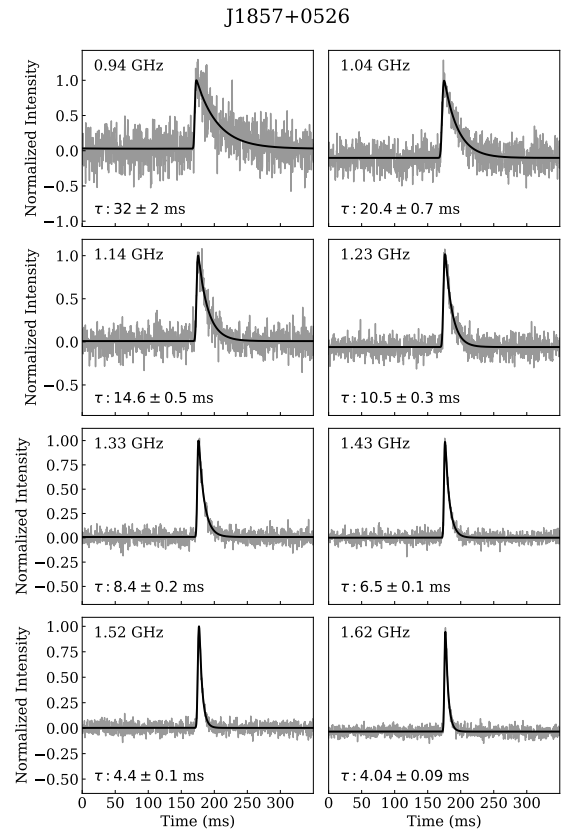

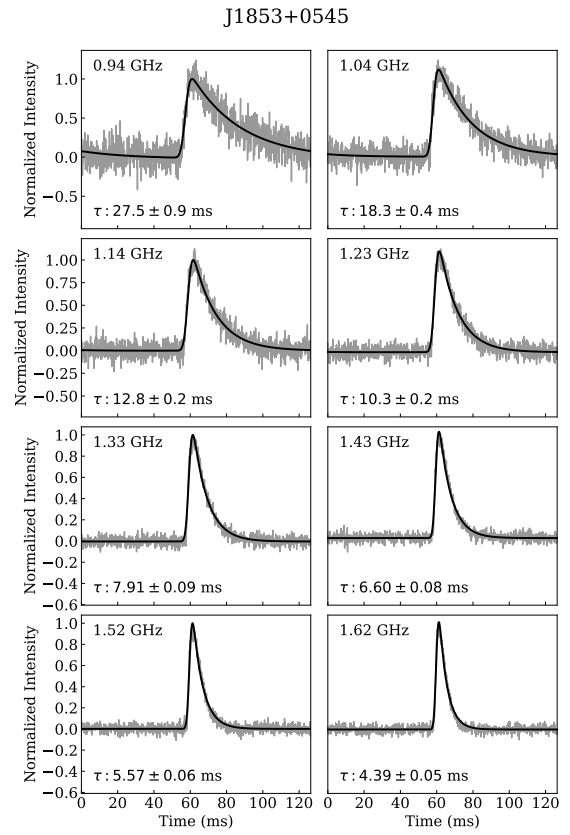
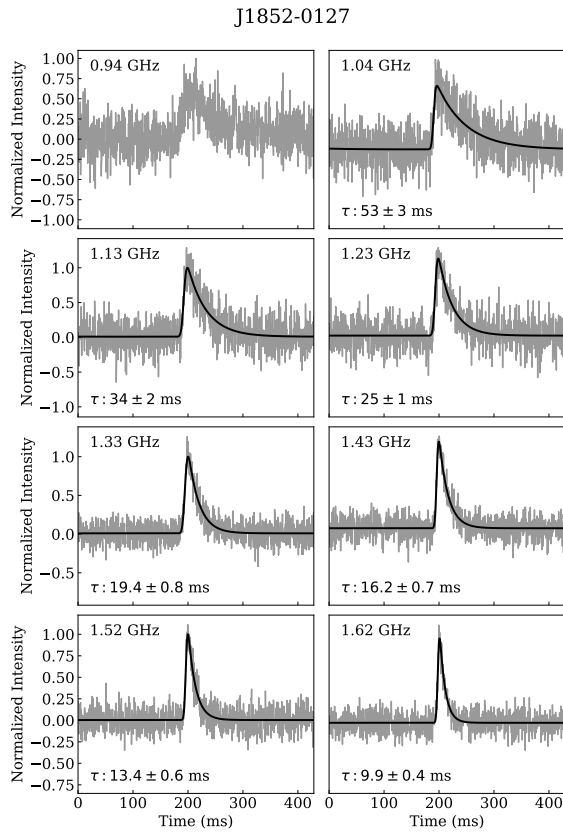



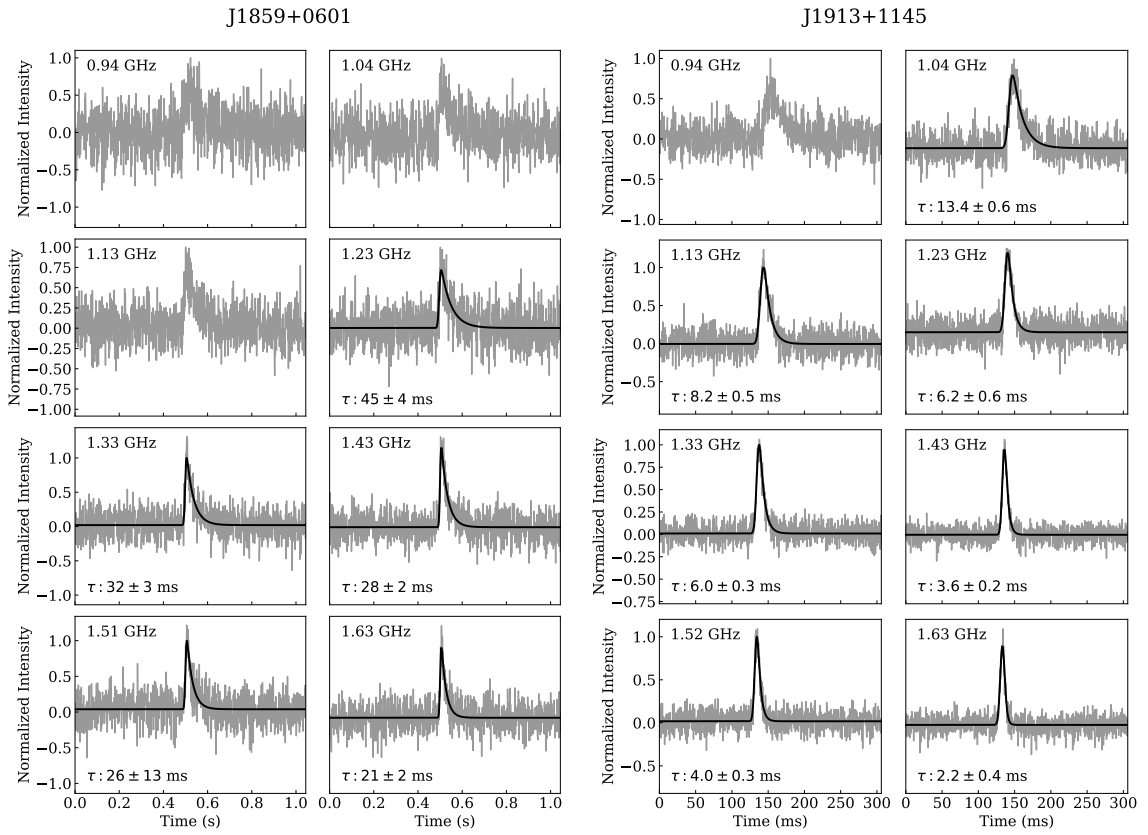










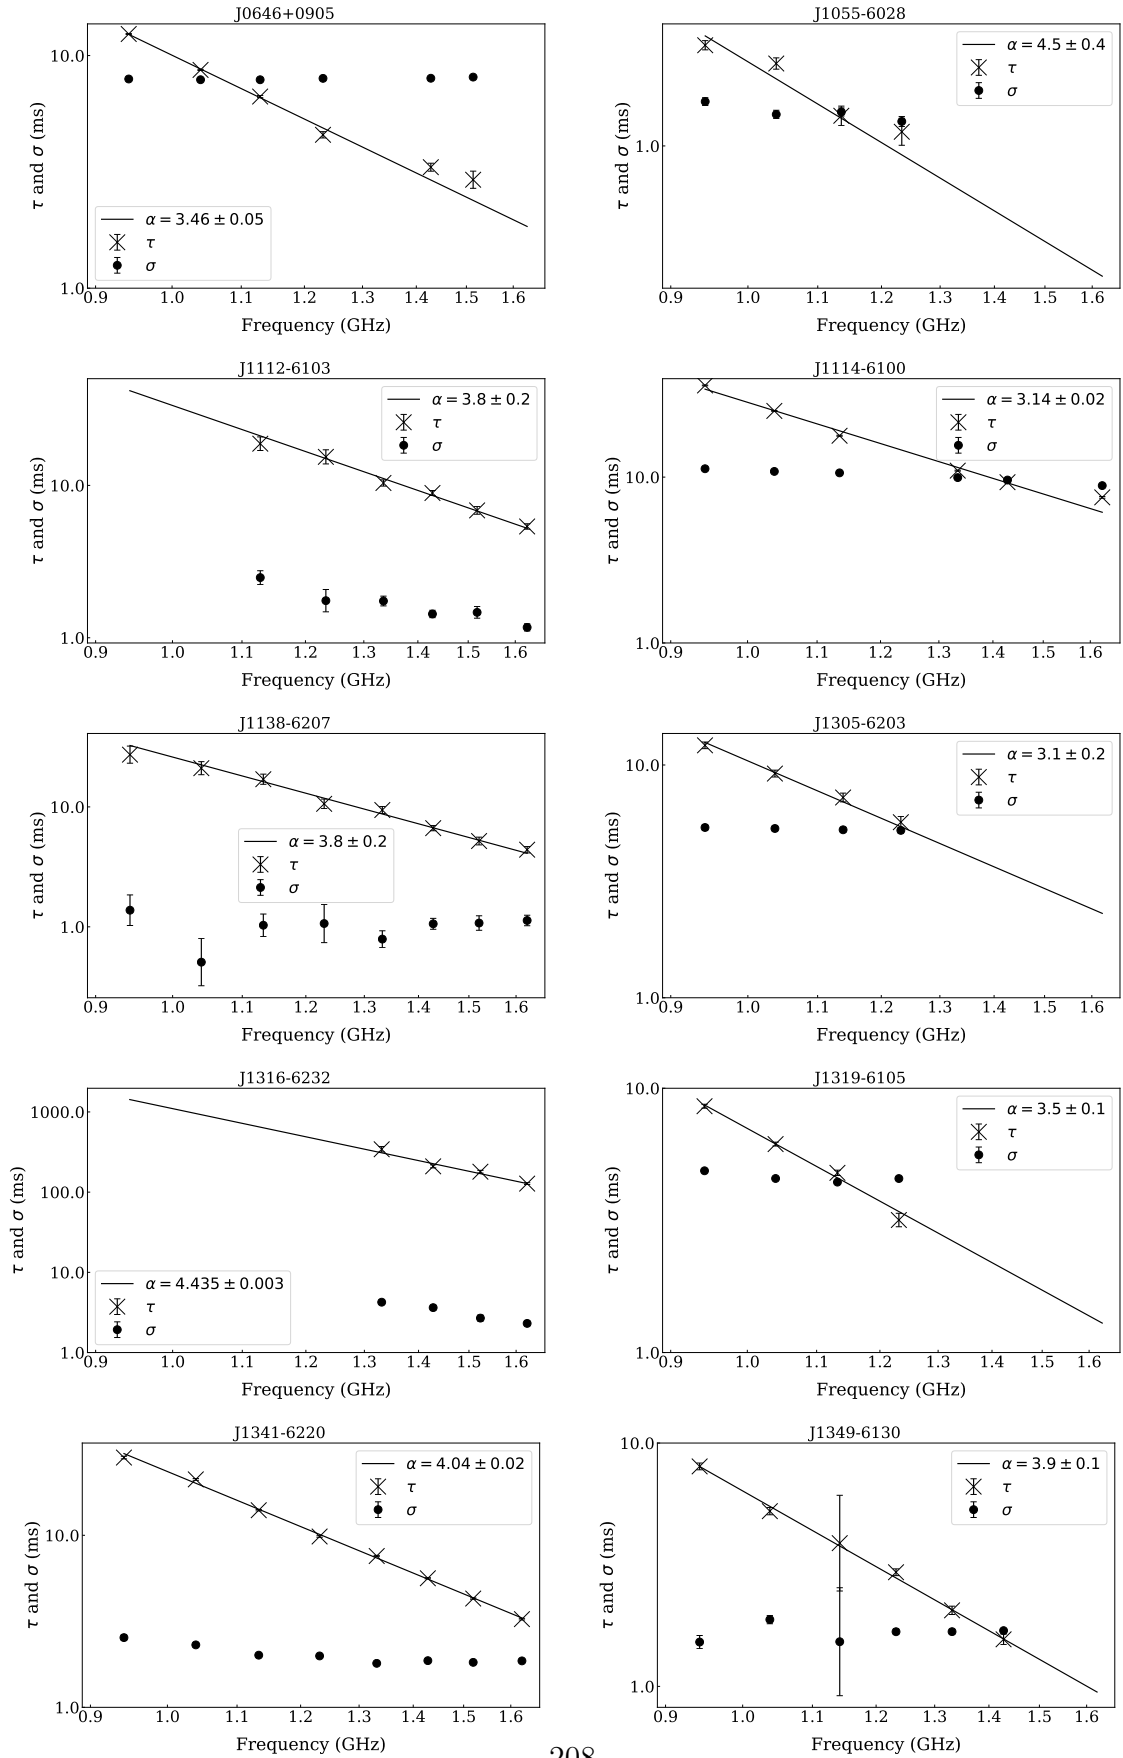


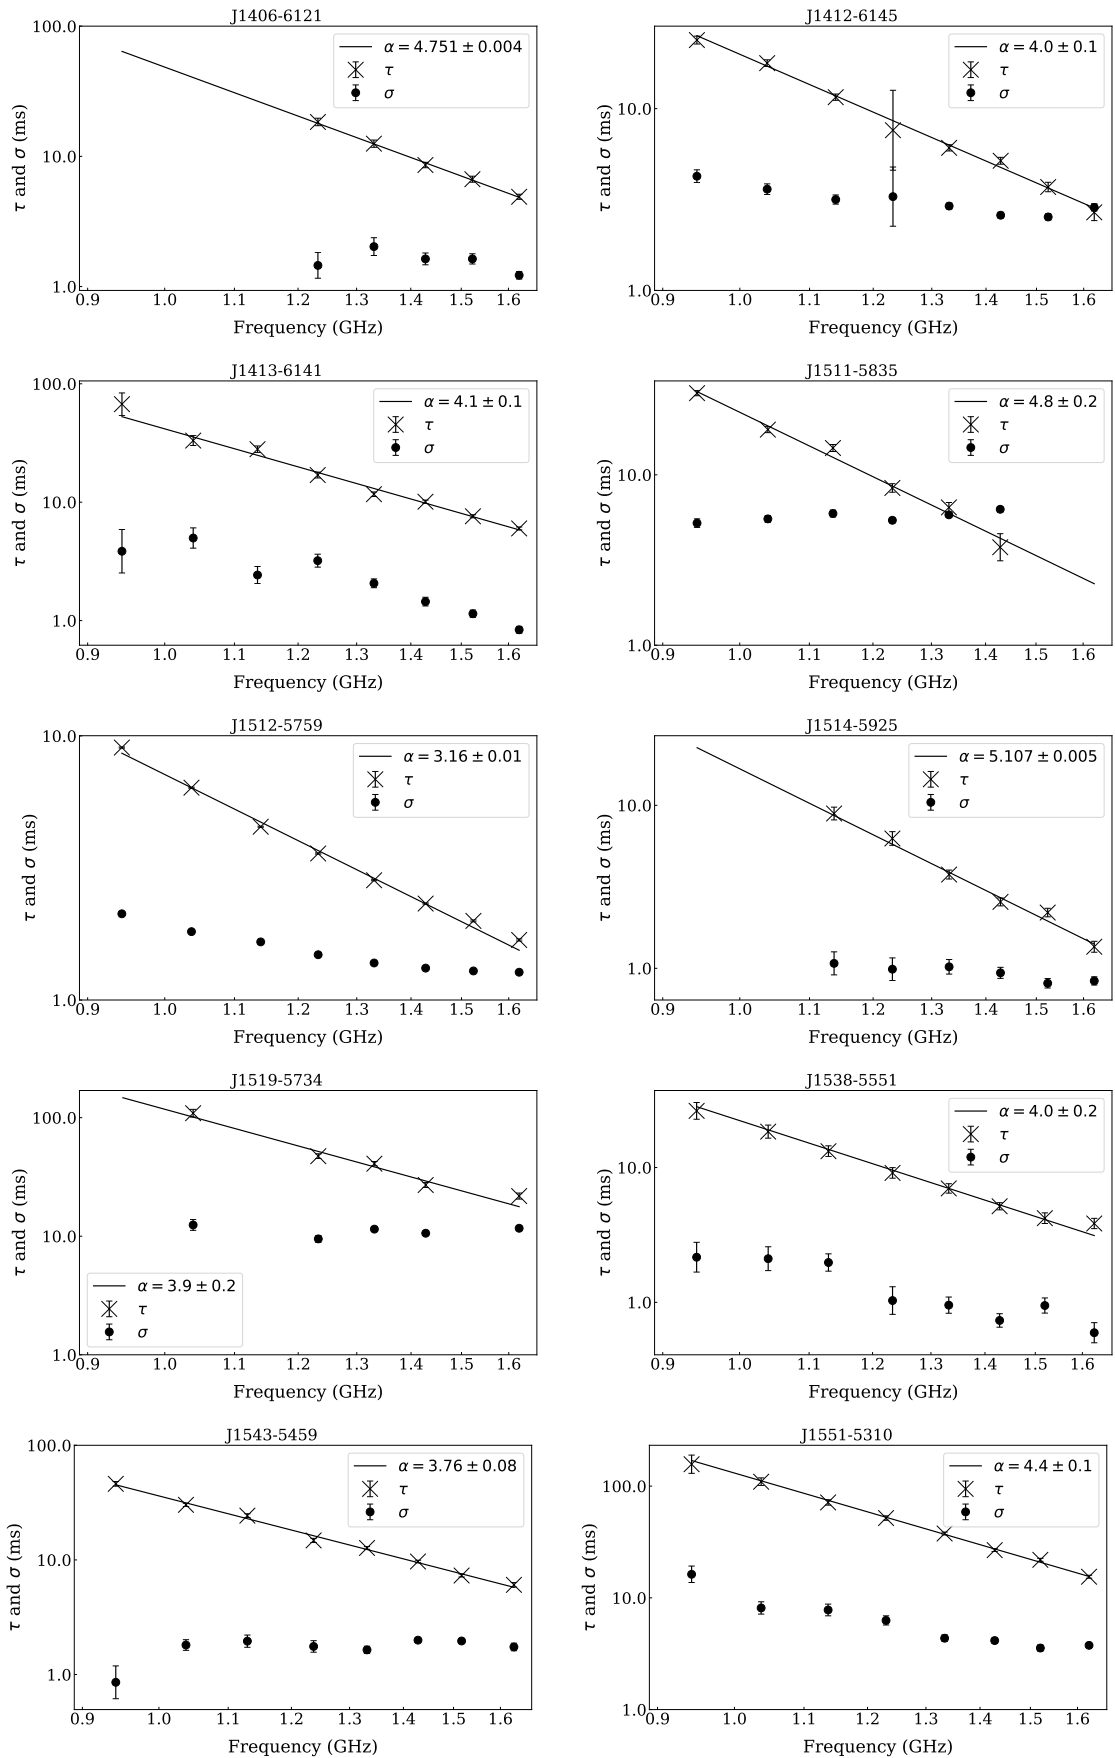
Appendix B

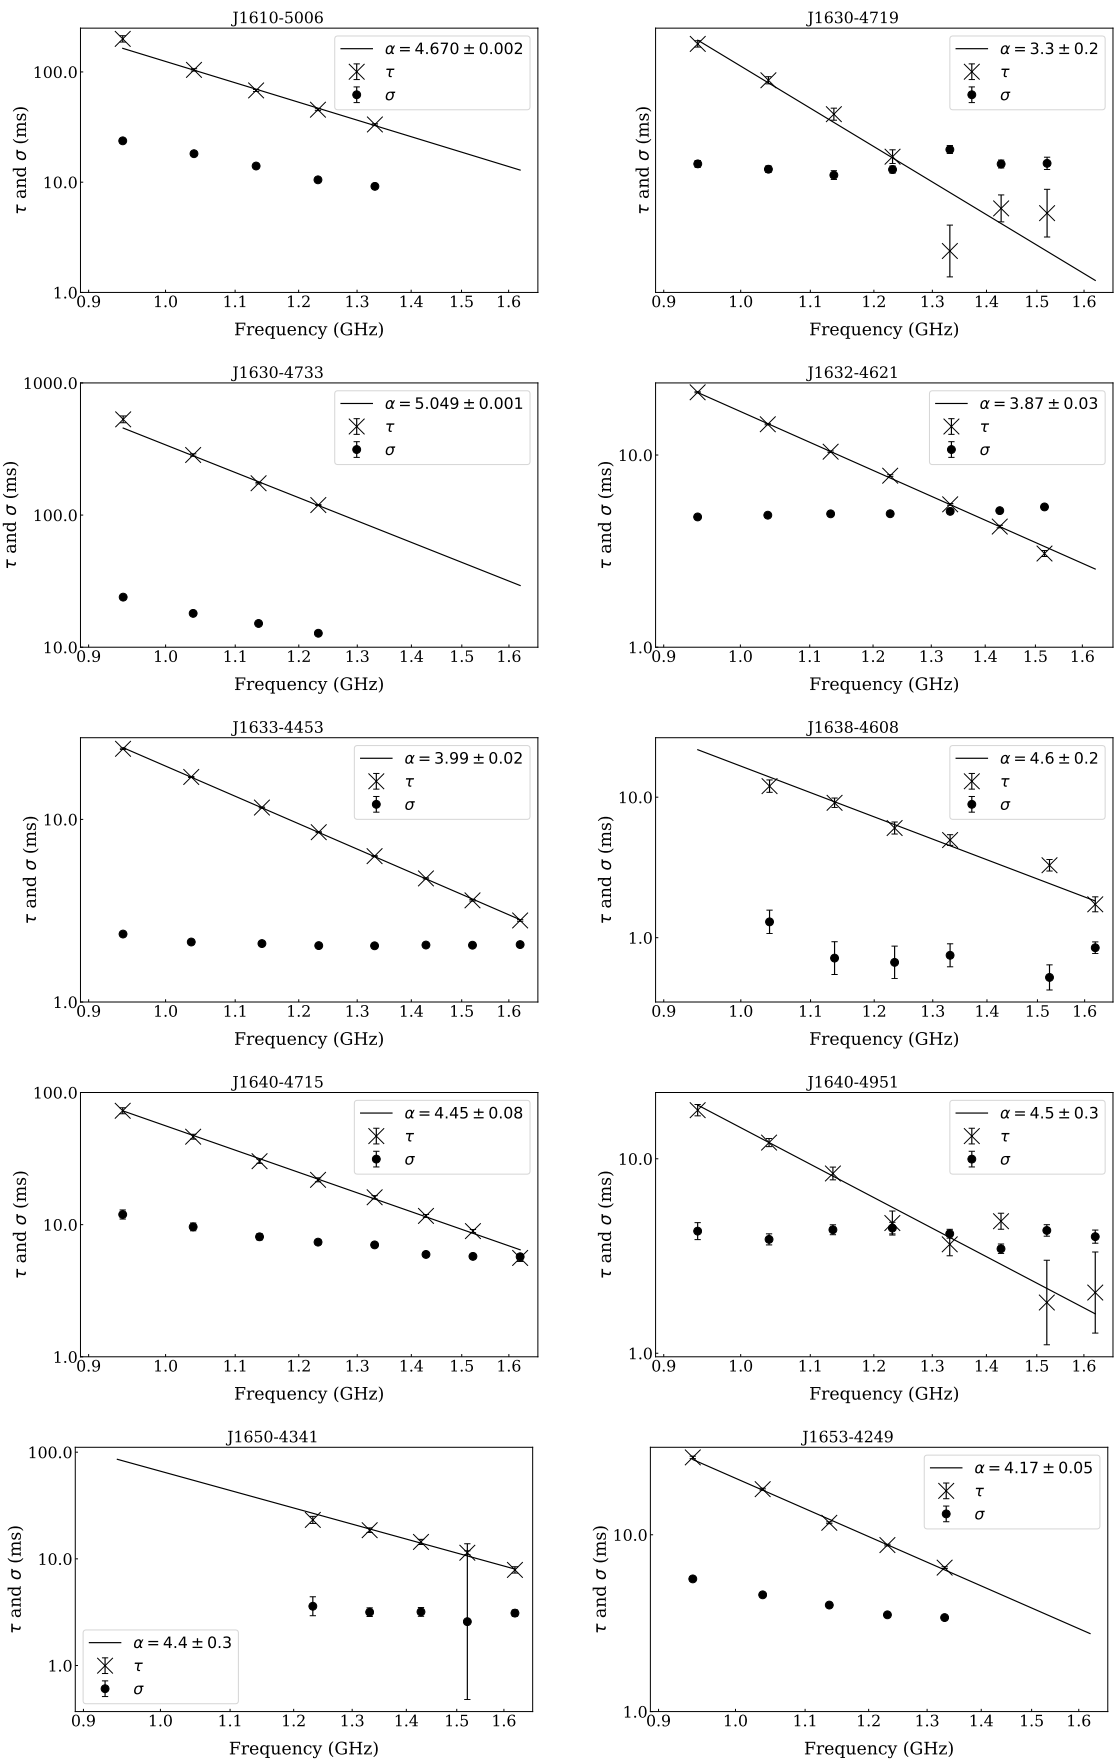
Scatter modelling results: frequency dependence of τ and σ

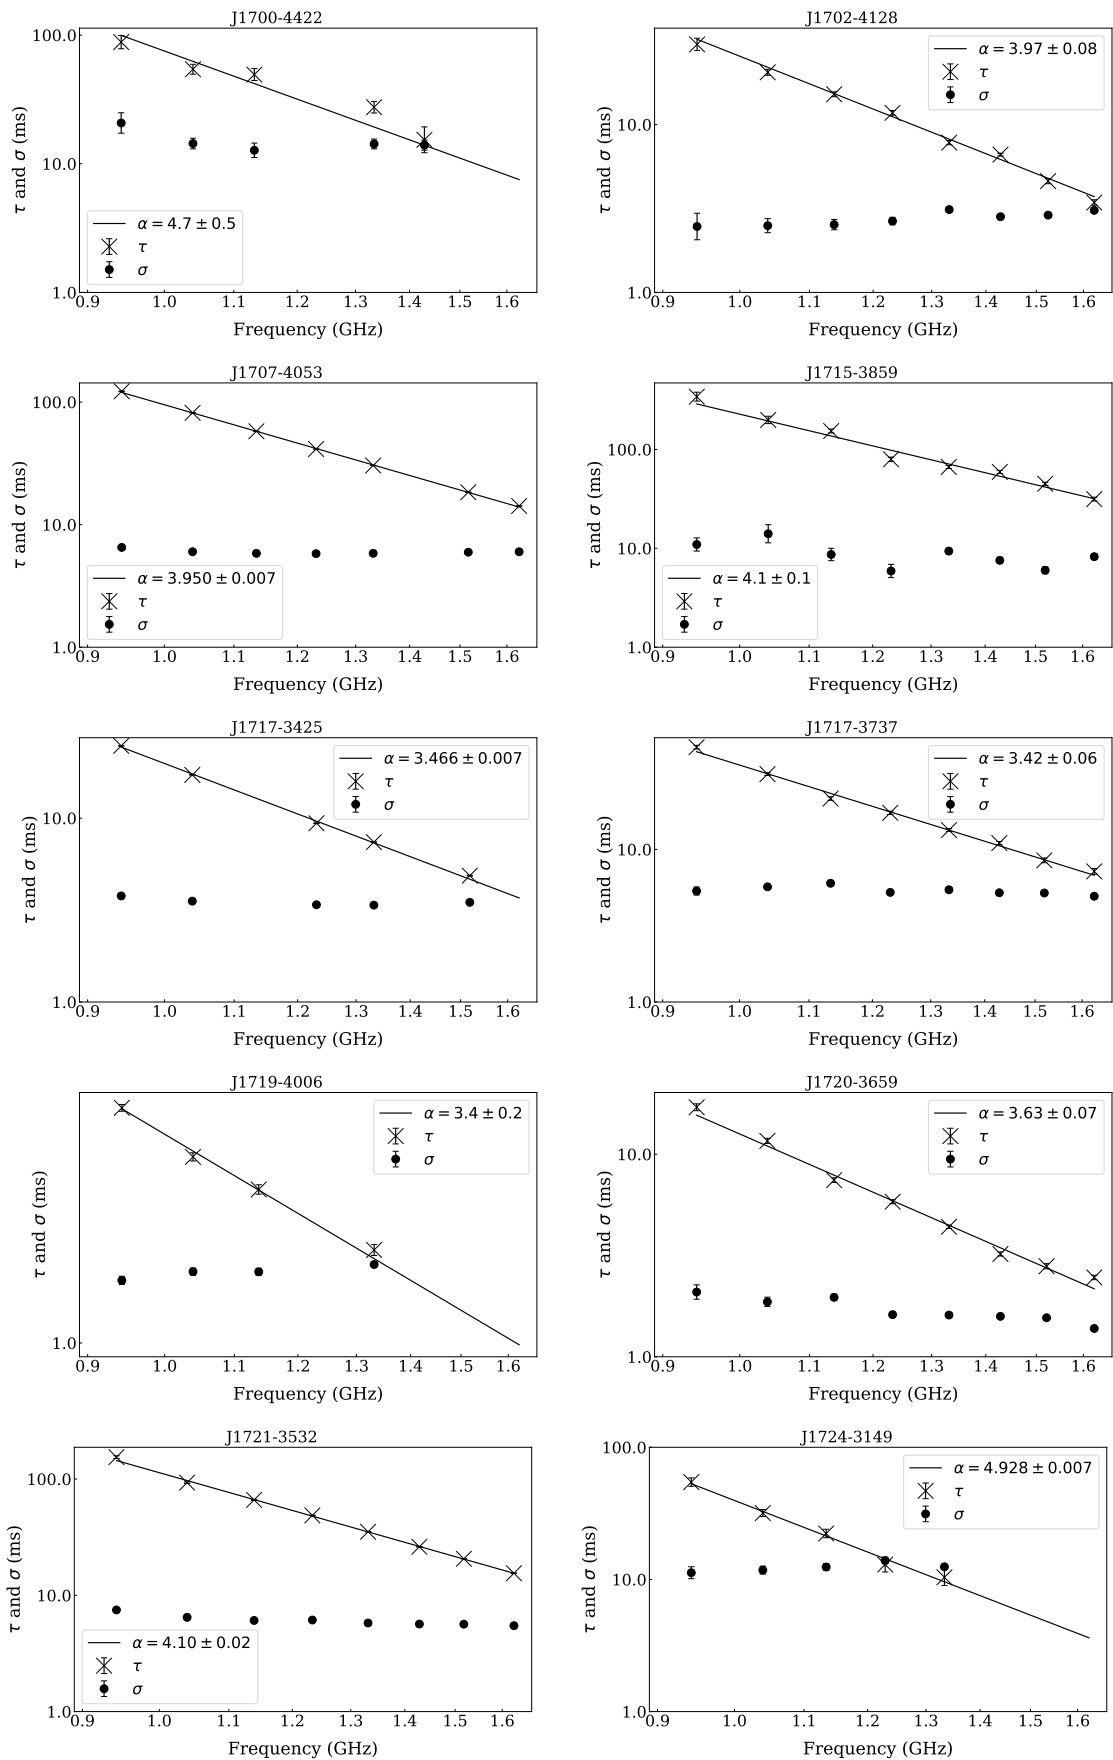
The figures in the following pages show log-log plots of how the modelled scattering timescale τ and the Gaussian standard deviation of the intrinsic pulse σ both vary with frequency, formatted in the same way as Fig. 7.5. This is shown for the 84 pulsars for which I obtained successful scatter modelling fits as described in Chapter 7. Each profile has at least four successfully obtained values of τ and σ , with a maximum of eight across the observing band.

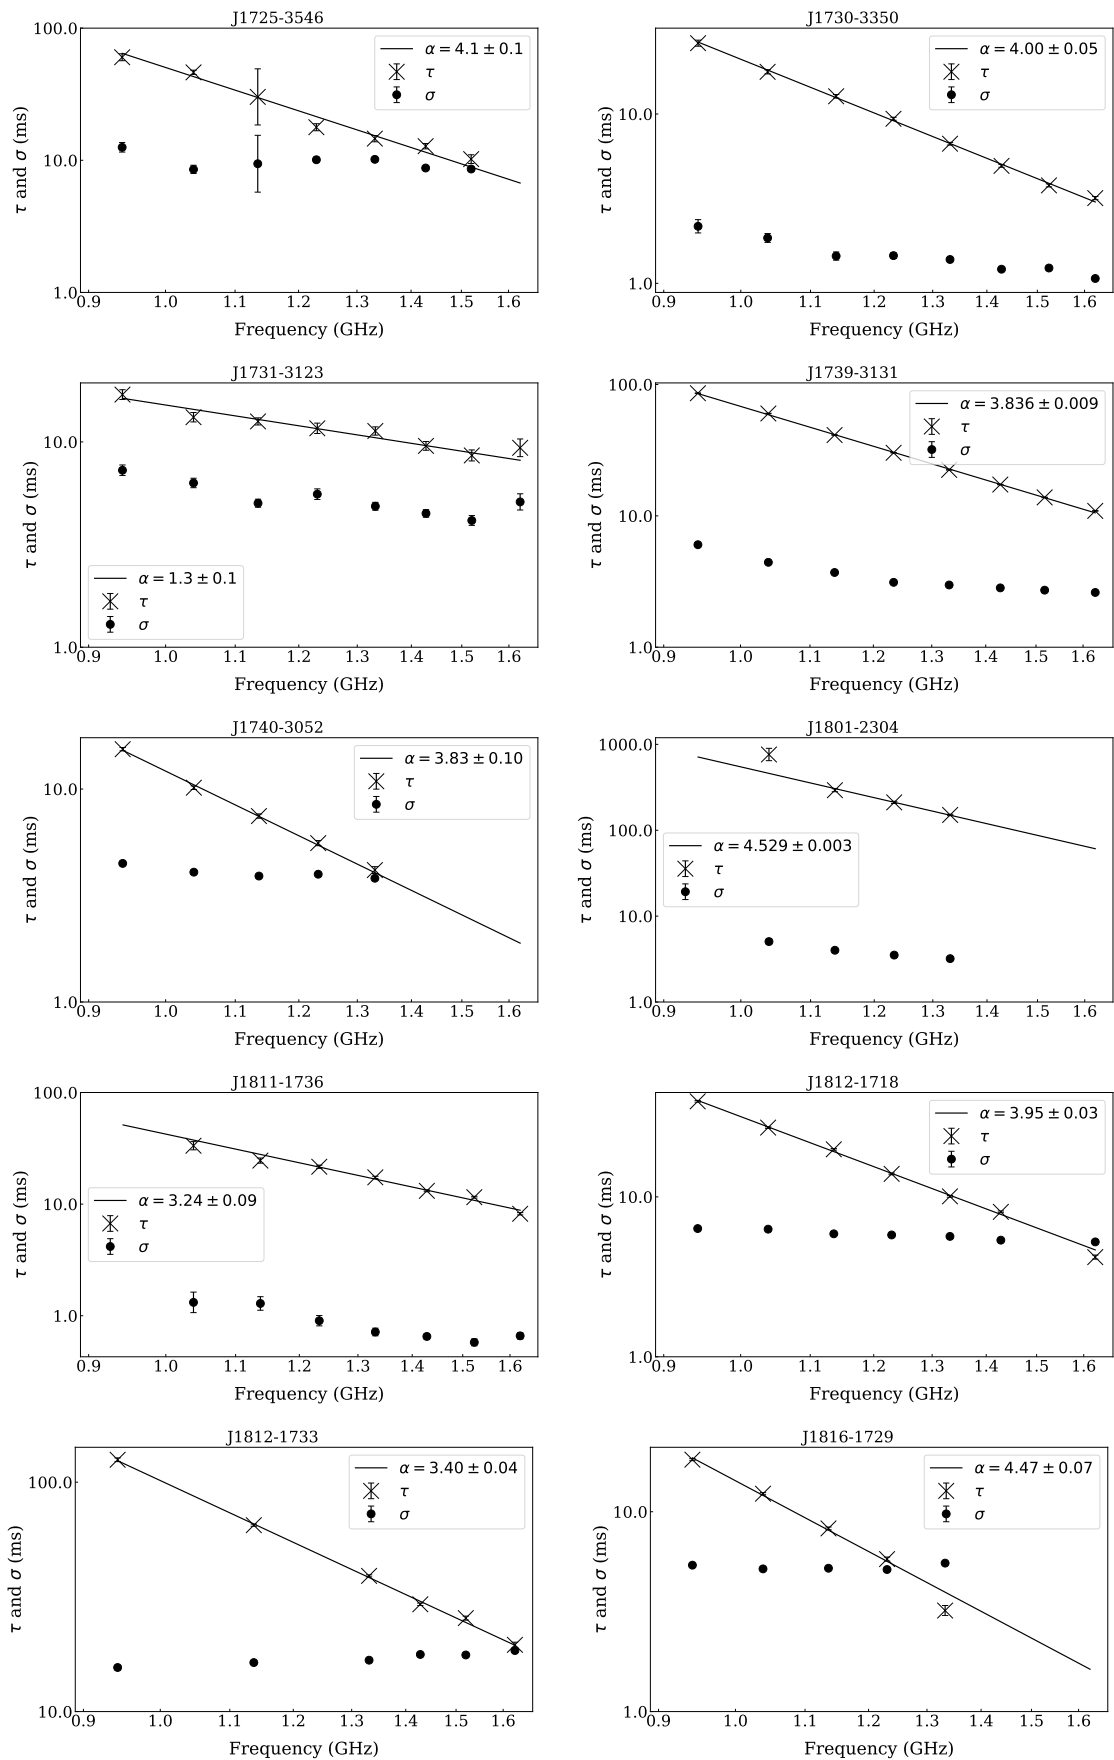
Figure B.1:

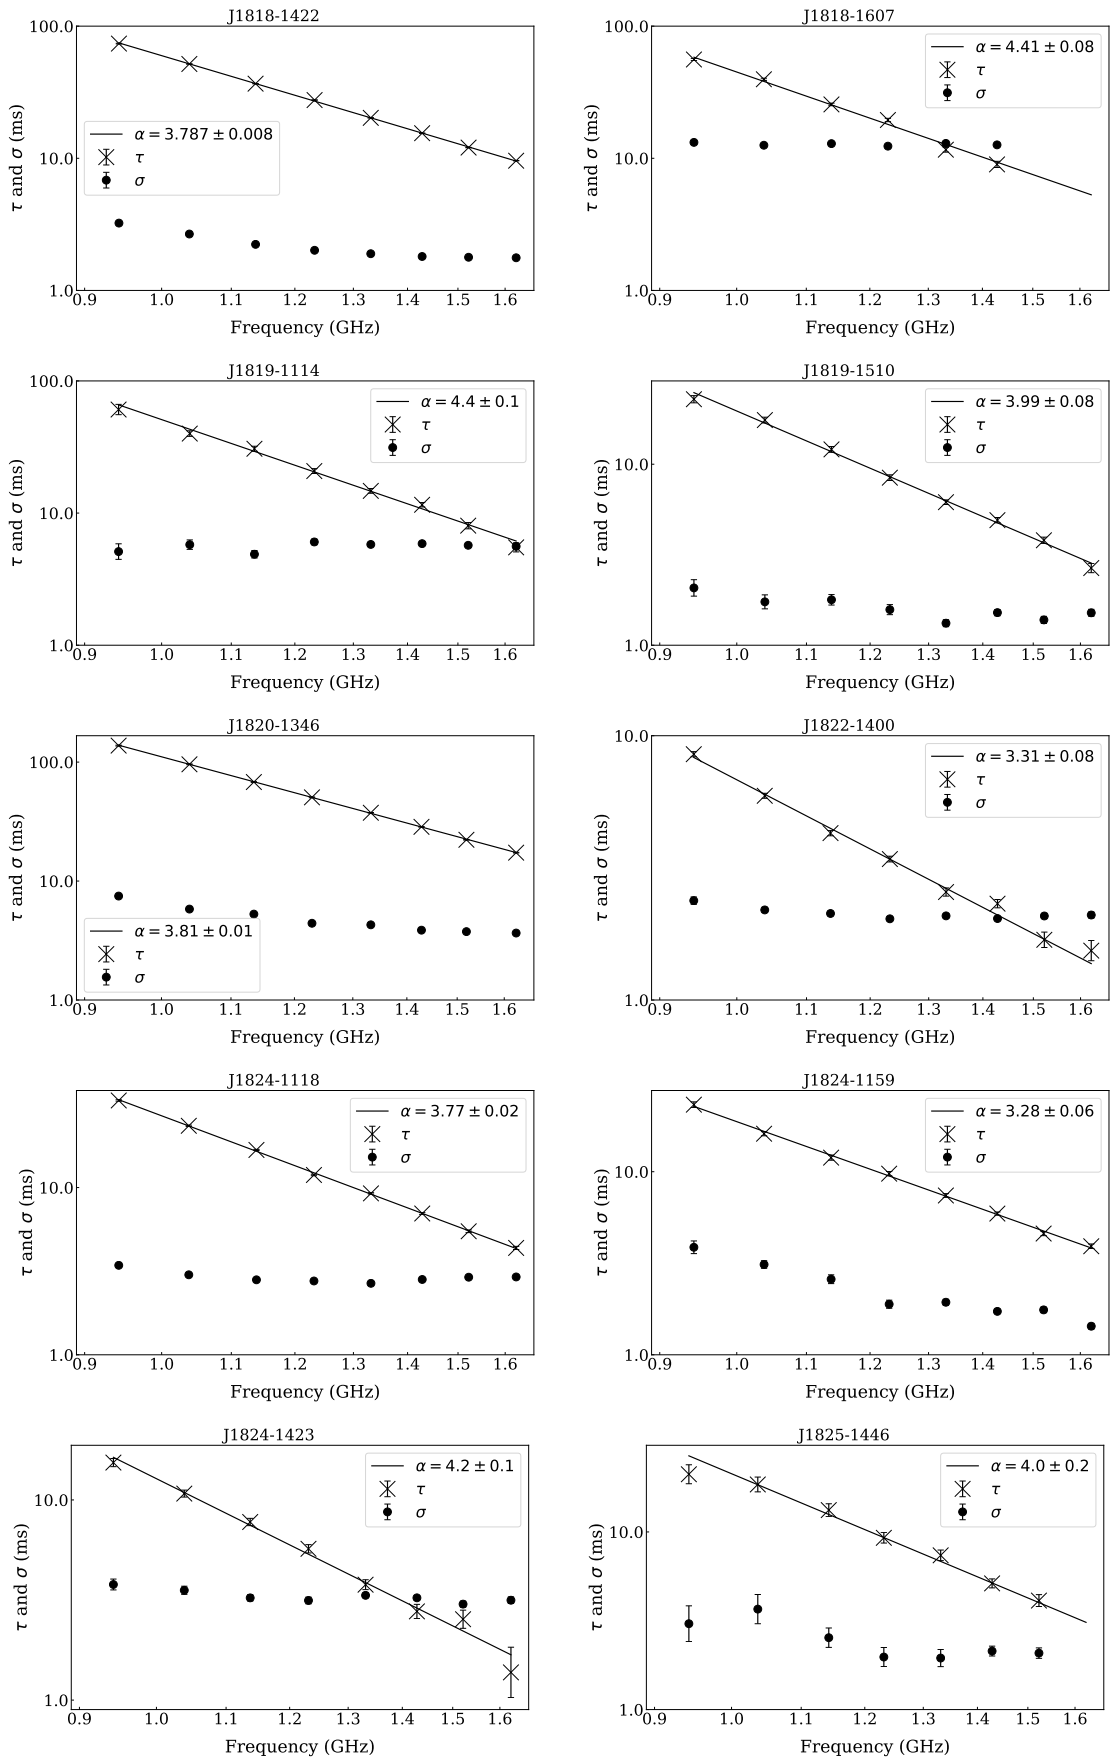














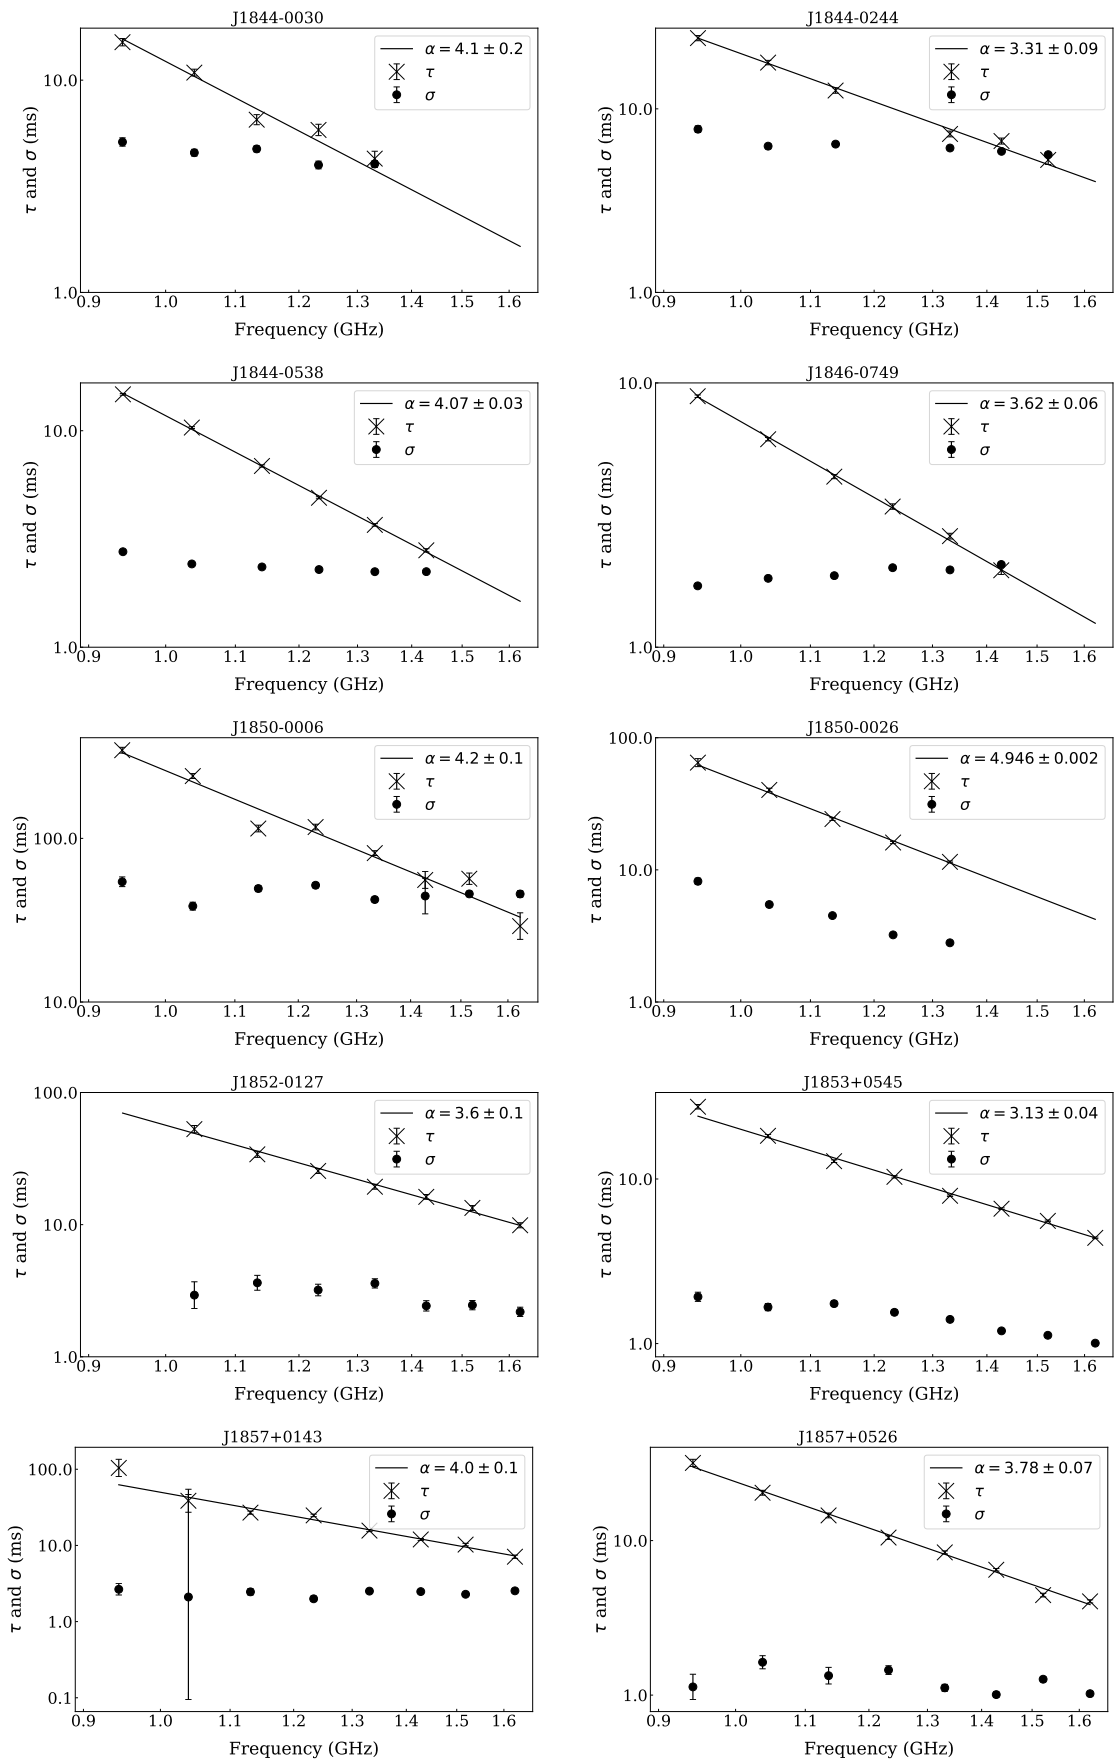


















Bibliography

- Abbate F., et al., 2020, MNRAS, 498, 875
- Anderson P. W., Itoh N., 1975, Nature, 256, 25
- Bailes M., et al., 2016, in MeerKAT Science: On the Pathway to the SKA. p. 11
- Bailes M., et al., 2020, PASA, 37
- Barnard J. J., Arons J., 1986, ApJ, 302, 138
- Bhat N. D. R., Cordes J. M., Chatterjee S., 2003, ApJ, 584, 782
- Bhat N. D. R., Cordes J. M., Camilo F., Nice D. J., Lorimer D. R., 2004, ApJ, 605, 759
- Bhat N. D. R., Gupta Y., Kramer M., Karastergiou A., Lyne A. G., Johnston S., 2007, A&A, 462, 257
- Blaskiewicz M., Cordes J. M., Wasserman I., 1991, ApJ, 370, 643
- Bondonneau L., et al., 2020, arXiv e-prints
- Brentjens M. A., De Bruyn A. G., 2005, A&A, 441, 1217
- Briskin W. F., Macquart J. P., Gao J. J., Rickett B. J., Coles W. A., Deller A. T., Tingay S. J., West C. J., 2010, ApJ, 708, 232
- Cordes J. M., 1978, ApJ, 222, 1006
- Cordes J. M., Lazio T. J. W., 2001, ApJ, 549, 997
- Cordes J. M., Lazio T. J. W., 2002, arXiv e-prints
- Cordes J. M., Rickett B. J., 1998, ApJ, 507, 846
- Cordes J. M., Weisberg J. M., Boriakoff V., 1985, ApJ, 288, 221

Cordes J. M., Shannon R. M., Stinebring D. R., 2016, ApJ, 817, 16

Costa M. E., McCulloch P. M., Hamilton P. A., 1991, MNRAS, 252, 13

Craig H. A., Romani R. W., 2012, ApJ, 755, 137

Crawford F., Manchester R. N., Kaspi V. M., 2001, AJ, 122, 2001

Cronyn W. M., 1970, Science, 168, 1453

D'Alessandro F., McCulloch P. M., King E. A., Hamilton P. A., McConnell D., 1993, MNRAS, 261, 883

Dai S., et al., 2019, ApJ Lett., 874, L14

Davies J. G., Horton P. W., Lyne A. G., Rickett B. J., Smith F. G., 1968, Nature, 217, 910

Desvignes G., et al., 2019, Science, 365, 1013

Dyks J., 2008, MNRAS, 391, 859

Dyks J., 2019, MNRAS, 488, 2018

Dyks J., 2020, MNRAS: Letters, 495, L118

Dyks J., Rudak B., 2015, MNRAS, 446, 2505

Dyks J., Rudak B., Demorest P., 2010, MNRAS, 401, 1781

Edwards R. T., Stappers B. W., 2004, A&A, 421, 681

Edwards R. T., Hobbs G. B., Manchester R. N., 2006, MNRAS, 372, 1549

Everett J. E., Weisberg J. M., 2002, ApJ, 553, 341

Foot C. J., 2005, Atomic physics. Oxford master series in atomic, optical, and laser physics Vol. 135, Oxford University Press

Foreman-Mackey D., 2016, JOSS, 1, 24

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Gangadhara R. T., 2004, ApJ, 609, 335

Gangadhara R. T., Gupta Y., 2001, ApJ, 555, 31

Geyer M., Karastergiou A., 2016, MNRAS, 462, 2587

Geyer M., et al., 2017, MNRAS, 470, 2659

Gold T., 1968, Nature, 218, 731

Goldreich P., Julian W. H., 1969, ApJ, 157

Gupta Y., Gangadhara R. T., 2003, ApJ, 584, 418

Gupta Y., Gothoskar P., Ramesh Bhat N., 2002, IAU Symp., 199, 369

Gupta Y., et al., 2017, Current Science, 113, 707

Hamaker J. P., Bregman J. D., Sault R. J., 1996, A&AS, 117, 137

Han J. L., Manchester R. N., Qiao G. J., 1999, MNRAS, 306, 371

Hankins T. H., Eilek J. A., Jones G., 2016, ApJ, 833, 47

Hassall T. E., et al., 2012, A&A, 543, A66

Hessels J. W., Ransom S. M., Stairs I. H., Freire P. C., Kaspi V. M., Camilo F., 2006, Science, 311, 1901

Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, Nature, 217, 709

Hibschman J. A., Arons J., 2001, ApJ, 546, 382

Hobbs G., Edwards R., Manchester R., 2006, ChJAA, 6, 189

Hobbs G., et al., 2020a, PASA, 37, 16

Hobbs G., et al., 2020b, MNRAS, 491, 5951

Högbom J., 1974, Astron. Astrophys. Suppl, 15, 417

Hotan A. W., Van Straten W., Manchester R. N., 2004, PASA, 21, 302

Hunter J. D., 2007, Computing in Science and Engineering, 9, 90

Ilie C. D., Johnston S., Weltevrede P., 2019, MNRAS, 483, 2778

Johnston S., Kerr M., 2018, MNRAS, 474, 4629

Johnston S., Weisberg J. M., 2006, MNRAS, 368, 1856

Johnston S., Karastergiou A., Mitra D., Gupta Y., 2008, MNRAS, 388, 261

Johnston S., et al., 2020, MNRAS, 493, 3608

Johnston S., et al., 2021, MNRAS, 502, 1253

Jonas J. L., 2009, Proc. IEEE, 97, 1522

Jones P. B., 2013, MNRAS: Letters, 435, 11

Kalapothisarakos C., Wadiasingh Z., Harding A. K., Kazanas D., 2020, ApJ, 907, 63

Karastergiou A., 2009, MNRAS: Letters, 392, 60

Karastergiou A., Johnston S., 2006, MNRAS, 365, 353

Karastergiou A., Johnston S., 2007, MNRAS, 380, 1678

Karastergiou A., et al., 2001, A&A, 379, 270

Karastergiou A., Kramer M., Johnston S., Lyne A. G., Bhat N. D. R., Gupta Y., 2002, A&A, 391, 274

Karastergiou A., Johnston S., Kramer M., 2003, A&A, 404, 325

Karastergiou A., Johnston S., Manchester R. N., 2005, MNRAS, 359, 481

Karastergiou A., Roberts S. J., Johnston S., Lee H., Weltevrede P., Kramer M., 2011, MNRAS, 415, 251

Karuppusamy R., Stappers B. W., Serylak M., 2011, A&A, 525

Kazantsev N. A., Potapov V. A., 2015, ATsir, 1628, 1

Kerr M., et al., 2020, PASA, 37, e020

Kirsten F., Bhat N. D., Meyers B. W., Macquart J. P., Tremblay S. E., Ord S. M., 2019, ApJ, 874, 179

Komesaroff M. M., 1970, Nature, 225, 612

Kramer M., Xilouris K. M., Jessner A., Lorimer D. R., Wielebinski R., Lyne A. G., 1997, A&A, 322, 846

- Kramer M., Karastergiou A., Gupta Y., Johnston S., Bhat N., Lyne A., 2003, *A&A*, 668, 15
- Kramer M., et al., 2006, *Science*, 314, 97
- Kramer M., et al., 2021, *MNRAS*, 000, 1
- Krishnakumar M. A., Mitra D., Naidu A., Joshi B. C., Manoharan P. K., 2015, *ApJ*, 804
- Krishnakumar M. A., Joshi B. C., Manoharan P. K., 2017, *ApJ*, 846, 104
- Krishnakumar M. A., Maan Y., Joshi B. C., Manoharan P. K., 2019, *ApJ*, 878, 130
- Krishnakumar M. A., et al., 2021, *A&A*, 651, A5
- Krzyszowski K., Maron O., Słowikowska A., Dyks J., Jessner A., 2014, *MNRAS*, 440, 457
- Kuzmin A. D., Losovsky B. Y., 2007, *AApTr*, 26, 597
- Lang K. R., 1971, *ApJ*, 164, 249
- Lazarus P., Karuppusamy R., Graikou E., Caballero R. N., Champion D. J., Lee K. J., Verbiest J. P., Kramer M., 2016, *MNRAS*, 458, 868
- Lee L. C., Jokiipii J. R., 1976, *ApJ*, 206, 735
- Lewandowski W., Kijak J., Gupta Y., Krzyszowski K., 2011, *A&A*, 534
- Lewandowski W., Dembska M., Kijak J., Kowalinska M., 2013, *MNRAS*, 434, 69
- Lewandowski W., Kowalińska M., Kijak J., 2015a, *MNRAS*, 449, 1570
- Lewandowski W., Rozko K., Kijak J., Bhattacharyya B., Roy J., 2015b, *MNRAS*, 454, 2517
- Löhmer O., Kramer M., Mitra D., Lorimer D. R., Lyne A. G., 2001, *ApJ*, 562, L157
- Löhmer O., Mitra D., Gupta Y., Kramer M., Ahuja A., 2004, *A&A*, 425, 569
- Lorimer D. R., Kramer M., 2005, *Handbook of Pulsar Astronomy*. Cambridge Observing Handbooks for Research Astronomers, Cambridge University Press, <https://books.google.co.uk/books?id=0Z8tdN6qJcsC>

Lu J. G., et al., 2016, ApJ, 816, 76

Lyne A. G., Manchester R. N., 1988, MNRAS, 234, 477

Lyubarskii Y. E., Petrova S. A., 1998, Astrophys. Space Sci., 262, 379

MacKay D. J. C., 2003, Information Theory, Inference and Learning Algorithms.
<http://www.inference.org.uk/itprnn/book.pdf>

Main R. A., et al., 2020, MNRAS, 499, 1468

Manchester R. N., 2000, PASA, 18, 1

Manchester R. N., Taylor J. H., Huguenin G. R., 1975, ApJ, 196, 83

Manchester R. N., Lyne A. G., Taylor J. H., Durdin J. M., Large M. I., Little A. G.,
1978, MNRAS, 185, 409

Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, ApJ, 129

McKinney W., 2010, in Proceedings of the 9th Python in Science Conference. pp
56–61, doi:10.25080/majora-92bf1922-00a

McKinnon M. M., Stinebring D. R., 2000, ApJ, 529, 435

Melrose D. B., 2017, Rev. Mod. Plasma Phys., 1, 5

Melrose D. B., Rafat M. Z., Mastrano A., 2021, MNRAS, 500, 4549

Michel F. C., 1987, ApJ, 322, 822

Mitra D., Li X. H., 2004, A&A, 421, 215

Mitra D., Rankin J. M., 2002, ApJ, 577, 322

Noutsos A., Johnston S., Kramer M., Karastergiou A., 2008, MNRAS, 386, 1881

Noutsos A., Karastergiou A., Kramer M., Johnston S., Stappers B. W., 2009, MN-
RAS, 396, 1559

Noutsos A., et al., 2015, A&A, 576, A62

Oswald L., Karastergiou A., Johnston S., 2019, MNRAS, 489, 310

Oswald L., Karastergiou A., Johnston S., 2020, MNRAS, 496, 1418

Oswald L. S., et al., 2021, MNRAS, 504, 1115

Parthasarathy A., et al., 2020, MNRAS, 494, 2012

Pen U. L., Macquart J. P., Deller A. T., Brisken W., 2014, MNRAS: Letters, 440, 36

Pennucci T. T., 2019, ApJ, 871, 34

Pétri J., Mitra D., 2020, MNRAS, 491, 80

Petroff E., Keith M. J., Johnston S., van Straten W., Shannon R. M., 2013, MNRAS, 435, 1610

Petrova S. A., 2001, A&A, 378, 883

Price D. C., 2016, arXiv e-prints, arXiv:1607.03579

Radhakrishnan V., Cooke D. J., 1969, ApJ Lett., 3, 225

Ramachandran R., Mitra D., Deshpande A. A., McConnell D. M., Ables J. G., 1997, MNRAS, 290, 260

Rankin J. M., 1983, ApJ, 274, 359

Ransom S. M., et al., 2009, AAS, 214, 605.08

Ransom S. M., et al., 2014, Nature, 505, 520

Rickett B. J., 1977, ARA&A, 15, 479

Rickett B., Johnston S., Tomlinson T., Reynolds J., 2009, MNRAS, 395, 1391

Ridolfi A., et al., 2021, arXiv e-prints

Rohlfs K., Wilson T. L., 1996, Tools of radio astronomy. Springer, doi:10.5860/choice.35-2078, <https://ui.adsabs.harvard.edu/abs/1996tra.book.....R/abstract>

Romani R. W., Narayan R., Blandford R., 1986, MNRAS, 220, 19

Rookyard S. C., Weltevrede P., Johnston S., 2015, MNRAS, 446, 3367

Ruderman M. A., Sutherland P., 1975, ApJ, 196, 51

Sagan C., Sagan L. S., Drake F., 1972, Science, 175, 881

- Serylak M., et al., 2020, arXiv e-prints
- Sidoli L., Israel G. L., Esposito P., Rodríguez Castillo G. A., Postnov K., 2017, MNRAS, 469, 3056
- Simard D., Pen U. L., Marthi V. R., Brisken W., 2019, MNRAS, 488, 4963
- Smith D. A., et al., 2008, A&A, 492, 923
- Smits J. M., Stappers B. W., Edwards R. T., Kuijpers J., Ramachandran R., 2006, A&A, 448, 1139
- Sobey C., et al., 2019, MNRAS, 484, 3646
- Song X., et al., 2020, arXiv e-prints
- Stinebring D. R., McLaughlin M. A., Cordes J. M., Becker K. M., Goodman J. E. E., Kramer M. A., Sheckard J. L., Smith C. T., 2001, ApJ, 549, L97
- Swarup G., 1991, in IAU Colloq. 131: Radio Interferometry. Theory, Techniques, and Applications. pp 376–380
- Szary A., Gil J., Zhang B., Haberl F., Melikidze G. I., Geppert U., Mitra D., Xu R. X., 2017, ApJ, 835, 178
- Tan C. M., et al., 2018, ApJ, 866, 54
- Tange O., 2011, The USENIX Magazine, pp 42–47
- Taylor J. H., Manchester R. N., Lyne A. G., 1993, ApJS, 88, 529
- Thorsett S. E., 1991, ApJ, 377, 263
- Van Der Walt S., Colbert S. C., Varoquaux G., 2011, Computing in Science and Engineering, 13, 22
- Van Rooy P., Whitlow D., Seymour A., 2017, AAS, 229, 137.05
- Verbiest J. P. W., et al., 2016, MNRAS, 458, 17
- Virtanen P., et al., 2020, Nature Methods, 17, 261
- Walker M. A., Melrose D. B., Stinebring D. R., Zhang C. M., 2004, MNRAS, 354, 43

- Wang N., Manchester R. N., Johnston S., Rickett B., Zhang J., Yusup A., Chen M., 2005, MNRAS, 358, 270
- Wang H. G., et al., 2014, ApJ, 789, 73
- Wang P. F., Wang C., Han J. L., 2015, MNRAS, 448, 771
- Weltevrede P., et al., 2010, PASA, 27, 64
- Williamson I. P., 1972, MNRAS, 157, 55
- Xilouris K. M., Kramer M., Jessner A., Wielebinski R., Timofeev M., 1996, A&A, 309, 481
- van Straten W., Bailes M., 2011, PASA, 28, 1
- van Straten W., Manchester R. N., Johnston S., Reynolds J. E., van Straten W., Manchester R. N., Johnston S., Reynolds J. E., 2010, PASA, 27, 104
- van Straten W., Demorest P., Osłowski S., 2012, Astronomical Research and Technology, 9, 237
- van Straten W., van Straten W. 2013, ApJS, 204, 13