

The Variability of Radio Pulsars



Paul Brook
St. Edmund Hall
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2015

For mum.

Acknowledgements

My PhD is all but done, but just before submission,
I'd like to thank the people who deserve some recognition.
To Aris Karastergiou, my savvy supervisor,
Who's only two years older, but at least a decade wiser.
I must thank Aris first of all, to not would be unethical,
For reasons inexhaustible, not simply alphabetical.
His door was always open for me; piles of praise I owe,
So Aris I would like so say a big *ευχαριστώ*.
Cheers too to Simon Johnston, whom I owe a lot of beer,
For super supervision in the Southern Hemisphere.
Something more I'm grateful for and owe Simon a debt:
Allowing me to access an outstanding dataset.
On all my Aussie allies, whopper thanks I must bestow,
This work would not have happened without C.S.I.R.O.
To Duncan and to Maura, I must take this chance to say,
Thanks for warmth and welcoming me to the U.S.A.
And when I'm home in Oxford, thanks for emails to and fro,
Connecting me to Morgantown, without the ten foot snow.
For finding me the money to jet off to places hotter,
I thank the purse-string person, the good-hearted Garret Cotter.
Garret's helped me plenty, always answers when I knock,
Advising me pre-PhD, now into my postdoc.
His sidekick Ashling Morris, has the patience of a saint,
She takes my half-filled travel forms with barely a complaint.
I take my woes to Ashling and she fixes all of them,
From lost receipts to airplane seats, she really is a gem.
A final thanks to Jocelyn, to her I tip my hat,
For helpful chats and guidance but, of course, not only that;
This PhD could never be, if she had never seen,
That piece of scruff, LGM-1, CP nineteen nineteen.

Declaration

I declare that no part of this thesis has been, or is being, submitted for any qualification other than the degree of Doctor of Philosophy at the University of Oxford.

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

Some of this work has been published in the following journal:

Chapter 3: P. R. Brook, A. Karastergiou, S. Buchner, S. J. Roberts, M. J. Keith, S. Johnston, and R. M. Shannon. Evidence of an Asteroid Encountering a Pulsar. *Astrophysical Journal, Letters*, 780:L31, January 2014. doi: 10.1088/2041-8205/780/2/L31.

*Paul Richard Brook
Oxford, March 2015*

“The pursuit of truth and beauty is a sphere of activity in which we are permitted to remain children all our lives.”

Albert Einstein

“Space? Well there’s nowt up there!”

Norma Brook

Abstract

Neutron stars are amongst the most exotic objects known in the universe; more than a solar mass of material is squeezed into an object the size of a city, leading to a density comparable to that of an atomic nucleus. They have a surface magnetic field which is typically around a trillion times stronger than the magnetic field here on Earth, and we have observed them to spin up to around 700 times per second. The existence of neutron stars was first proposed by Baade and Zwicky in 1934 but later graduated from theory to fact in 1967 as the first pulses were detected by Jocelyn Bell-Burnell, a then graduate student at the University of Cambridge. There are now well over 2000 neutron stars whose radio emission beams point at, and have been detected on Earth. We call these objects pulsars. Because of their remarkable properties, pulsars are very useful to physicists, who can employ them as precision timing tools due to the unwavering nature of their emission and of their rotation. Having an array of ultra-accurate clocks scattered throughout our galaxy is very useful for performing astrophysical experiments. In particular, precise pulsar timing measurements and the models that explain them, will permit the direct detection of gravitational radiation; a stochastic background initially, and potentially the individual signals from supermassive black hole binaries.

Our models of pulsar behaviour are so precise that we are now able to notice even slight departures from them; we are starting to see that unmodelled variability in pulsars occurs over a broad range of timescales, both in emission and in rotation. Any unmodelled variability is, of course, detrimental to the pulsar's utility as a precision timing tool, and presents a problem when looking for the faint effects of a passing gravitational wave.

We are hoping that pulsar timing arrays will detect gravitational radiation in the coming decade, but this depends, in part, on our ability to understand and mitigate the effects of the unmodelled intrinsic instabilities that we are observing. One important clue as to the nature of the variability in pulsar emission and rotation, is the emerging relationship between the two; we sometimes observe

correlation on timescales of months and years.

We have been observing pulsars for almost fifty years and our expanding datasets now document decades of pulsar behaviour. This gives us the ability to investigate pulsar variability on a range of timescales and to gain an insight into the physical processes that govern these enigmatic objects.

In this thesis I describe new techniques to detect and analyse the emission and rotational variability of radio pulsars. We have employed these techniques on a 24 year pulsar dataset to unearth a striking new example of a dramatic and simultaneous shift in a pulsar's emission and rotation. We hypothesise that this event was caused by an asteroid interaction, although other explanations are also possible.

Our variability techniques have also been used to analyse data from 168 young, energetic pulsars. In this thesis we present results from the nine most interesting. Of these, we have found some level of correlated variability in seven, one of which displays it very strongly.

We have also assessed the emission stability of the NANOGrav millisecond pulsars and have found differing degrees of variability, due to both instrumental and astrophysical causes.

Finally, we propose a method of probing the relationship between emission and rotation on short-timescales and, using a simulation, we have shown the conditions under which this is possible.

Throughout the work, we address the variability in pulsar emission, rotation and links between the two, with the aim of improving pulsar timing, attaining a consolidated understanding of the diverse variable phenomena observed and elucidating the evolutionary path taken by pulsars.

Contents

1 Pulsars	1
1.1 Neutron star theory	1
1.1.1 Interior and environment	3
1.1.2 Pulsar radio emission	7
1.1.3 Pulsar populations	10
1.2 Pulsar observations	15
1.2.1 Pulsar spectra	15
1.2.2 Pulsar Timing	16
1.2.3 Radio emission	20
1.2.4 Effects of radiation propagation	30
1.3 Long-term monitoring	35
1.3.1 Rotational irregularities	35
1.3.2 Intermittent pulsars	38
1.3.3 State-switching	40
1.3.4 Magnetar variability	43
1.4 Open questions	43
1.4.1 How do pulsar populations evolve with time?	43
1.4.2 Are long and short-term emission variability the same phenomena?	45
1.4.3 What initiates a change in magnetospheric currents?	46
1.4.4 What is the nature of timing noise?	47
1.4.5 How is emission linked to glitching?	47
1.5 This thesis	49
2 Observations and Analysis Techniques	50
2.1 Observations	51
2.1.1 PSR J0738-4042 and Fermi Timing Programme observations	51
2.1.2 NANOGrav Observations	53

2.2	Analysis techniques	55
2.2.1	Pulse profile monitoring technique	55
2.2.2	Spindown monitoring technique	56
2.2.3	Gaussian Process Regression	57
2.2.4	Enhanced pulse profile monitoring technique	57
2.2.5	Enhanced spindown monitoring technique	63
3	The variability of PSR J0738-4042	65
3.1	Emission history	65
3.1.1	Polarisation	68
3.2	The dataset	70
3.3	The 2005 event	71
3.4	Interpretation	73
3.5	Unresolved issues	79
4	Monitoring the variability of young, energetic radio pulsars	82
4.1	Fermi Timing Programme Observations	83
4.2	Modelling the timing residuals	83
4.3	Variability correlation	87
4.4	Notable examples	88
4.4.1	PSR J1830-1059 (B1828-11)	88
4.4.2	PSR J1602-5100 (B1558-50)	89
4.4.3	PSR J0738-4042 (B0736-40)	92
4.4.4	PSR J0742-2822 (B0740-28)	97
4.4.5	PSR J0908-4913 (B0906-49)	99
4.4.6	PSR J0940-5428	102
4.4.7	PSR J1105-6107	102
4.4.8	PSR J1359-6038 (B1356-60)	109
4.4.9	PSR J1600-5044 (B1557-50)	109
4.5	Discussion	114
5	The emission variability of NANOGrav millisecond pulsars	124
5.1	The search for gravitational waves	124
5.1.1	Gravitational waves	124
5.1.2	Pulsar timing arrays	127
5.1.3	NANOGrav	129
5.2	Results from ASP and GASP backends	129

5.2.1	Instrumental issues	129
5.2.2	Astrophysical profile changes	135
5.3	PUPPI and GUPPI results	143
5.3.1	Instrumental issues	143
5.3.2	Astrophysical profile changes	143
5.4	Discussion	146
6	Detecting variable spindown rates in mode-changing and nulling pulsars	155
6.1	Motivation	155
6.2	Simulation and results	157
6.3	Proposal 1: Do state-fractions change over time?	168
6.4	Proposal 2: Continuous monitoring	169
7	Conclusions	171
7.1	New techniques	171
7.2	New findings	173
7.3	Proposed framework for the interpretation of pulsar variability . .	177
	References	184

List of Figures

1.1	The traditional magnetic dipole model of a rotating neutron star and its magnetosphere	2
1.2	Density and mass of a neutron star as a function of radius	4
1.3	Cross-section of a neutron star	5
1.4	Emission beam models	9
1.5	The $P - \dot{P}$ diagram	12
1.6	Pulsar flux density spectra	15
1.7	Timing residuals of PSR B1133+66	18
1.8	Pulse power distribution	20
1.9	Sub-pulse drifting and nulling in pulsars	22
1.10	Mode-changing pulsars	23
1.11	Nulling in pulsars	24
1.12	Individual pulse from PSR B0301+19	25
1.13	Examples of integrated pulse profiles	26
1.14	Pulse profiles at different observing frequencies	28
1.15	The rotating vector model	29
1.16	Pulse dispersion and incoherent dedispersion	32
1.17	Timing noise	36
1.18	Glitches in PSR B1930+22	37
1.19	The rotation and emission behaviour of the intermittent pulsar	39
1.20	Integrated pulse profile changes	41
1.21	The relationship between profile shape and $\dot{\nu}$	42
1.22	The spindown rate and flux-density for XTE J1810-197	44
1.23	the short-lived double-peaked mode of PSR J1119-6127	48
1.24	RRAT-like pulses seen in PSR J1119-6127	48
2.1	Comparison between total power and circular polarisation profile	52
2.2	GP regression models	58

2.3	Pulse profile intensity for 93 observations of PSR J1830-1059 . . .	59
2.4	Illustration of the uses of GP regression	61
3.1	The integrated pulse profile of PSR J0738-4042	66
3.2	71 irregularly spaced observations of PSR J0738-4042	67
3.3	PSR J0738-4042 pulse profiles with and without the transient component	68
3.4	PSR J0738-4042 pulse profiles with polarisation information . . .	69
3.5	Simulations of PSR J0738-4042 pulse profiles	70
3.6	Variations in the profile shape and spindown rate seen in PSR J0738- 4042	72
3.7	The 2005 drifting event seen in PSR J0738-4042	74
4.1	Timing residuals and GP model for PSR J0940-5428 using one covariance kernel	85
4.2	Timing residuals and GP model for PSR J0940-5428 using two covariance kernels	86
4.3	Pulse profile and spindown variability for PSR J1830-1059	90
4.4	Pulse profile variability and correlation with spindown rate for PSR J1830-1059	91
4.5	Pulse profile and spindown variability for PSR J1602-5100	93
4.6	Pulse profile variability and correlation with spindown rate for PSR J1602-5100	94
4.7	Pulse profile and spindown variability for PSR J0738-4042	95
4.8	Pulse profile variability and correlation with spindown rate for PSR J0738-4042	96
4.9	The average pulse profiles of PSR J0742-2822	98
4.10	The pulse profile shape and spindown variability of PSR J0742-2822	99
4.11	Pulse profile and spindown variability for PSR J0742-2822	100
4.12	Pulse profile variability and correlation with spindown rate for PSR J0742-2822	101
4.13	Pulse profile and spindown variability for the main pulse of PSR J0908-4913	103
4.14	Pulse profile and spindown variability for the interpulse of PSR J0908-491	104
4.15	Pulse profile variability and correlation with spindown rate for PSR J0908-4913	105

LIST OF FIGURES

4.16	Pulse profile variability and correlation with spindown rate for PSR J0908-4913	106
4.17	Pulse profile and spindown variability for PSR J0940-5428	107
4.18	Pulse profile variability and correlation with spindown rate for PSR J0940-5428	108
4.19	Pulse profile and spindown variability for PSR J1105-6107	110
4.20	Pulse profile variability and correlation with spindown rate for PSR J1105-6107	111
4.21	Pulse profile and spindown variability for PSR J1359-6038	112
4.22	Pulse profile variability and correlation with spindown rate for PSR J1359-6038	113
4.23	Long-term behaviour of J1600-5044	115
4.24	Pulse profile and spindown variability for PSR J1600-5044	116
4.25	Pulse profile variability and correlation with spindown rate for PSR J1600-5044	117
4.26	Template matching simulation	121
5.1	The orbital decay of binary pulsar system PSR B1913+16	126
5.2	The Hellings and Downs plot	127
5.3	Characteristic strain sensitivity for current and future gravitation wave detector	128
5.4	A deviant pulse profile from PSR J2317+1439	131
5.5	Pulse profile of PSR J1713+0747	132
5.6	Individual frequency channels 1-8 of the PSR J1713+0747 profile .	133
5.7	Individual frequency channels 9-16 of the PSR J1713+0747 profile	134
5.8	Variability maps for the main pulse of PSR B1937+21	136
5.9	Variability maps for the interpulse of PSR B1937+21	137
5.10	Pulse profiles from PSR B1937+21	138
5.11	Variability maps of PSR J1853+1303	139
5.12	Observations of PSR J1853+1303	140
5.13	Variability maps of PSR J1910+1256	141
5.14	Pulse profiles of PSR J1910+1256	142
5.15	Observations of PSR J1713+0747	142
5.16	Deviant observations of three pulsars made on MJD 55305	144
5.17	Pulse profile changes of PSR J1713+0747	145
5.18	Variability maps for the main pulse of B1937+21	147
5.19	Variability maps for the interpulse of B1937+21	148

5.20	Observations of the main pulse and interpulse of PSR B1937+21	149
5.21	Variability maps for PSR J1600-3053 at 1500 MHz	150
5.22	Variability maps for PSR J1600-3053 at 800 MHz	151
5.23	Examples of the pulse profile variability of PSR J1600-3053	152
6.1	Correlation between spindown rate and a state-fraction varying randomly with time	160
6.2	Correlation between spindown rate and a state-fraction varying randomly with time	161
6.3	Correlation between spindown rate and a state-fraction varying linearly with time	162
6.4	Correlation between spindown rate and a state-fraction varying linearly with time	163
6.5	Correlation between spindown rate and a state-fraction varying linearly with time	164
6.6	Correlation between spindown rate and a state-fraction varying linearly with time	165
6.7	Correlation between spindown rate and a state-fraction varying sinusoidally with time	166
7.1	Effects of normalisation on Gaussian functions	176

List of Tables

1.1	Properties of the intermittent pulsars	39
1.2	Properties of six state-switching pulsars with correlated pulse profile and spindown rate variability	41
3.1	40 years of average profiles from PSR J0738–4042	66
5.1	List of NANOGrav millisecond pulsars	130
6.1	M value and correlation coefficient for state-fraction and spindown rate	168

Chapter 1

Pulsars

1.1 Neutron star theory

When a typical star exhausts its sources of internal energy, inward gravitational forces prevail and drive the collapse of the star, bringing its life to an end. The mass of the collapsing star will determine the nature of the corpse that will eventually be left behind. The most massive stars are believed to undergo an unstoppable gravitational collapse resulting in a supernova explosion, and end their lives as a black hole. The least massive become white dwarfs, with more than 95% of stars ending their lives this way. Electron degeneracy pressure within a white dwarf star is able to resist further gravitational collapse.

For stellar remnants with masses in excess of $1.46 M_{\odot}$ the gravitational forces become too great for even electron degeneracy pressure to resist (Chandrasekhar, 1931) and the star will collapse further, to such an extent that the electrons are forced into atomic nuclei, producing neutrons and electron neutrinos,






Figure 1.1: The traditional magnetic dipole model of a rotating neutron star and its magnetosphere. The various components are discussed throughout Section 1.1. Figure from Lorimer & Kramer (2005).

Gravity is unable to overcome the neutron degeneracy pressure against which it now competes; the stellar corpse finds equilibrium as a neutron star, a conclusion to stellar collapse first theorised by Baade & Zwicky (1934).

1.1.1 Interior and environment

Neutron star interiors

Due to the complete neutron degeneracy within a neutron star, the temperature becomes unimportant; the crucial relationship is between pressure and density (the equation of state). Unfortunately, the exotic environment means that the relationship cannot be tested in laboratories. We are instead reliant on theoretical models, constrained by observational parameters such as neutron star masses and radii.

It is possible to obtain the mass of a neutron star by observing it in a binary system; the typical value is $1.4 M_{\odot}$, with theoretical models imposing an upper limit of $3 M_{\odot}$ (Oppenheimer & Volkoff, 1939). The highest mass neutron stars currently known are around $2 M_{\odot}$ (Demorest et al., 2010; Antoniadis et al., 2013). At a mass of $1.4 M_{\odot}$, most equation of state models place the star's radius to be between 10-12 km (see right panel of Figure 1.2). This results in an average mass comparable to that of nuclear matter.

In reality, the density of a neutron star is far from uniform. The solid crystalline crust is thought to consist primarily of iron nuclei and a sea of degenerate electrons at a density of approximately 10^6 g cm^{-3} . The density increases with depth below the surface, until protons and electrons are forced to combine to produce neutrons (Equation 1.1). The number of free neutrons also increases with depth,




Figure 1.2: Left panel: The density distribution as a function of neutron star radius, calculated from a range of equations of state. Right panel: The total neutron star mass as a function of radius, calculated from a range of equations of state. Figures from Wiringa et al. (1988).

as large nuclei become unstable. Above a density of around $2 \times 10^{14} \text{ g cm}^{-3}$, the neutron star is composed of a mixture of around 95% free, superfluid neutrons and 5% superconducting electrons and protons. Near the centre of the star, the density reaches around six times that of nuclear matter, and various theories compete to describe its potentially exotic composition. Interactions between the superfluid core and the crystalline crust are thought to be responsible for rotational irregularities observed in pulsars, which may also be linked to changes in pulsar emission (Section 1.3.1).

Neutron star environments

If the magnetic field in a collapsing progenitor star is conserved, we expect relatively tiny neutron stars to be left with huge surface magnetic flux densities (B). The actual B field is a result of this principle and also of supernova processes which are not well understood. Neutron stars typically have surface magnetic flux densities of the order of 10^{12} G . Analogously, neutron stars also have incredi-




Figure 1.3: Cross-section of a neutron star. Figure from Lyne & Graham-Smith (2012).

ble rates of rotation due to the conservation of angular momentum during their birth; the rotation period of a newly born neutron star is of the order of milliseconds. The pulse period of a pulsar P lengthens as the star loses rotational energy,

$$\dot{E} \equiv -\frac{dE_{rot}}{dt} = -\frac{d(I\Omega^2/2)}{dt} = -I\Omega\dot{\Omega} = 4\pi^2 I \dot{P} P^{-3}, \quad (1.2)$$

where E_{rot} is the rotational kinetic energy, t is time, I is the moment of inertia, Ω is the rotational angular frequency, P is the star's rotation period, and \dot{P} is its first time derivative, or *spindown rate*. In pulsar astronomy, the rotational frequency ν and its first time derivative $\dot{\nu}$ are often used instead of P and \dot{P} .

The forces exerted on particles by the magnetic field of a neutron star (usually modelled as a dipole), totally dwarf those exerted by gravity. The axis of this magnetic field is usually misaligned with the rotational axis of the star (Figure 1.1). The rotating magnetic dipole, therefore, generates an electromagnetic wave which contributes to the energy loss of the system. This is known as *dipole braking*.

Generally, we can model the rotational evolution as a power-law

$$\dot{\Omega} = -\kappa\Omega^n, \quad (1.3)$$

where κ is usually assumed to be a constant, and n is the *braking index*,

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}, \quad (1.4)$$

which is equal to 3 in the case of pure dipole braking.

If we do make the assumptions that braking is *entirely* due to the rotating magnetic dipole and that the spin period at birth is much smaller than the present value, we are able to obtain a value for the pulsar's *characteristic age*,

$$\tau_c \equiv \frac{P}{2\dot{P}} = 15.8 \text{ Myr} \left(\frac{P}{\text{s}}\right) \left(\frac{\dot{P}}{10^{-15}}\right)^{-1} \quad (1.5)$$

Again, assuming that dipole braking is the dominant cause of energy loss, we can calculate the surface magnetic field of the star,

$$B_s \simeq 10^{12} \text{ G} \left(\frac{\dot{P}}{10^{-15}}\right)^{1/2} \left(\frac{P}{\text{s}}\right)^{1/2}, \quad (1.6)$$

The rotating magnetic field also induces an electric field (E), which results in the extraction of charged particles from the interior of the star. The high degree of conductivity both inside and outside the neutron star means that the induced electric field is cancelled out by the relocation of charged particles, leaving a force

free state,

$$\mathbf{E} + \frac{1}{c}(\boldsymbol{\Omega} \times \mathbf{r}) \times \mathbf{B} = 0, \quad (1.7)$$

where Ω is the angular velocity of the magnetic field and c is the speed of light. In the case where the rotation and magnetic axes are aligned, Goldreich & Julian (1969) show that the number density of charged particles once this equilibrium has been reached is

$$n_{GJ} = \frac{\Omega B_s}{2\pi c e}, \quad (1.8)$$

where B_s is the magnetic field strength at the surface and e is the charge of an electron. The result of this equation is a magnetosphere filled with plasma and forced to corotate with the star by the $\mathbf{E} \times \mathbf{B}$ field. The radius of the magnetosphere extends until the corotation reaches the velocity of light. This defines the edge of the *light cylinder* (Figure 1.1).

Although equations 1.5 and 1.6 above are defined with the assumption of pure dipole braking, many examples of pulsars in which $n \neq 3$ (e.g. Lyne et al., 1988, 1996) show that it is not the sole source of rotational energy loss. Other contributions come from a braking torque induced by particles outflowing along the magnetic axes (Section 1.1.2). Any variability in these outflowing currents will have consequences for the emission and rotational properties of a pulsar; this is thought to be the source of the correlated emission and rotation changes seen in a growing number of pulsars (Section 1.3.2 and Section 1.3.3), including those discussed in Chapter 3 and Chapter 4.

1.1.2 Pulsar radio emission

As depicted in Figure 1.1, magnetic field lines can be categorised as either open or closed; the magnetic field lines contained entirely within the light cylinder define

the closed and open field line regions. Charged particles are trapped within closed field lines, but are permitted to flow outwards along open lines. The radio emission mechanism is not well understood, but is thought to be a consequence of the huge electric field which is generated at the magnetic polar cap. Here, particles are accelerated by the electric field and emit curvature radiation in the form of γ -ray photons. These photons subsequently undergo pair-production after interacting with the magnetic field. The resulting electron-positron pairs are then accelerated, produce more photons, and a particle cascade develops. Such cascades produce a plasma above the polar cap from which the radio emission emerges in the form of a beam which is aligned tangentially to the open field lines at the height of emission.

As the emission beam from a spinning pulsar repeatedly sweeps by the Earth, we observe periodic pulses, which accumulate to produce a *pulse profile*. This is an average measure of the flux density received from a pulsar as a function of its period. The stability of a pulse profile means that it can be considered a unique fingerprint for each pulsar at a given observational frequency and is a property which facilitates precise timing of the arriving pulses (see template matching in Section 1.2.2). To what degree and on what timescales pulse profiles are truly stable are open questions; addressing them is a main theme of this thesis.

Beam geometry

The geometry of the pulsar radio beam is generally modelled as a cone, centred on the magnetic axis. Each time we see a pulsar emission cone sweep by, it registers as a pulse. Two competing models of the radio beam geometry of a pulsar are the *nested cone* (Rankin, 1993; Gil et al., 1993) and the *patchy beam* (Lyne & Manchester, 1988) models. The nested cone model is characterised by multiple




Figure 1.4: Two radio emission beam models which complete to explain multi-component pulse profiles. Left panel shows a nested cone structure (Rankin, 1993; Gil et al., 1993), while the right panel shows the patchy beam model developed by Lyne & Manchester (1988). Figure from Lorimer & Kramer (2005).

cones nested around a central core component, centred on the magnetic axis (left panel of Figure 1.4). The line of sight dictates the number of profile components observed. In the patchy beam model, regions of radio emission are distributed unevenly throughout the cone of emission (right panel of Figure 1.4).

More recently, Karastergiou & Johnston (2007) show how a model in which radio emission of different frequencies originates at different heights in the pulsar magnetosphere, is able to explain the diverse phenomena observed in the profiles of radio pulsars.

A number of pulsars show two distinct pulses per rotation, separated by approximately half a rotation period. The most popular interpretation of such an interpulse, is that the rotational and magnetic axes of the neutron star are approximately 90° apart, such that emission from both the north and south magnetic poles can be seen from Earth. This view is supported by Keith et al. (2010); after studying five interpulse pulsars, they show that all have an angle between their magnetic and rotational axes close to 90° .

1.1.3 Pulsar populations

We know of distinct pulsar categories, with quite different physical properties. These properties can provide us with clues regarding their birth and evolution. The pulsar parameters that give an insight into the nature and evolution of pulsar populations are encapsulated by the $P - \dot{P}$ *diagram* (see Figure 1.5).

Ordinary pulsars

PSR B1919+21 was the first pulsar detected. This and the vast majority of those discovered subsequently, typically have a spin period between a few tenths of seconds and a few seconds. This period is lengthening at the rate of between 10^{-17} and 10^{-13} seconds per second as the pulsar loses energy. The youngest pulsars are known to be energetic, having high \dot{E} values and being more likely to be observable in γ -rays (Weltevrede & Johnston, 2008). They are also more likely to glitch and display timing noise (Section 1.3.1). It was shown by Johnston & Weisberg (2006) that the youngest pulsars also have other distinct observational properties. They demonstrate that the pulse profiles of young pulsars are simple, consisting of one or two prominent components. For the two component profiles, the trailing component is almost always brightest and is the only one of the two to show circular polarisation. The linearly polarised fraction is nearly always in excess of 70%. Additionally the position angle swing is generally flat across the leading component and steep across the trailing component (see *Integrated profiles* in Section 1.2.2). By contrast, older pulsars often have more complex pulse profiles with multiple components. Because of their slower rotation rate, they also have light cylinders with larger radii (Section 1.1.1). This means that the magnetic field lines that remain open at the boundary of the light cylinder, will define a

relatively narrow beam. Older pulsars, therefore, have narrow pulse profiles as a fraction of the pulse period (i.e. a small *duty cycle*), with respect to young pulsars with higher rates of rotation. To learn more about the pulse profile stability in young, ordinary pulsars, an analysis of their emission and rotational stability is presented in Chapter 4.

One in ten radio pulsars is in a binary system that has evidently survived the supernova in which the pulsar was created. This situation can give rise to a very different type of pulsar.

Millisecond pulsars

The evolutionary sequence for a millisecond pulsar begins with two main sequence stars in a binary system. The more massive star evolves most rapidly, undergoes a core-collapse supernova and becomes a neutron star. The destructiveness of this event disrupts most binary systems (Radhakrishnan & Shukre, 1985), leaving a normal, isolated pulsar. For a typical pulsar, P is of the order of a second, and \dot{P} is around 10^{-15} .

For those systems which remain, the second star will later reach the end of its life and enter a red giant phase. At a close enough approach, the neutron star can now gravitationally accrete matter from the atmosphere of its red giant companion and becomes a bright source of thermal X-rays. At this stage of evolution, these systems are known as Low Mass X-ray binaries (LMXB) or High Mass X-ray binaries (HMXB) depending on the companion.

The accretion of material also transfers angular momentum to a pulsar, resulting in an increase of its spin frequency. When the transfer of the red giant atmosphere is complete, the pulsar has a period of a few tens of milliseconds or less. This population of millisecond pulsars (or recycled pulsars) takes up a new position on




Figure 1.5: The $P - \dot{P}$ diagram. Millisecond pulsars are located toward the bottom left. Pulsars known to be part of a binary system are marked by a circle. Young pulsars still associated with supernova remnants are shown as stars. The location of intermittent and state-switching pulsars (Section 1.3.2 and Section 1.3.3) are depicted by red and green diamonds respectively. PSR J0738-4042 (Chapter 3) is marked as a blue triangle and PSR J1602-5100 by a yellow triangle (Section 4.4.2). Examples of millisecond pulsars that show variability in Chapter 5 are marked as peach-coloured circles. Lines of constant magnetic field strength B , characteristic age τ and spin-down luminosity \dot{E} are shown. The grey regions are where pulsar emission mechanisms are predicted to fail. Figure adapted from Lorimer & Kramer (2005).

the $P - \dot{P}$ diagram (see Figure 1.5).

Due to an apparent reduction during the mass transfer process, millisecond pulsars have magnetic fields which are typically three or four orders of magnitude lower than isolated pulsars. This results in a low spindown rate, assuming that the spindown of a pulsar is primarily caused by the generation of rotating dipole radiation (see Section 1.1).

After the mass transfer, the millisecond pulsar is left with either a white dwarf companion, or one sufficiently massive that it may also undergo a supernova collapse. If the system also survives this second supernova, it emerges as a double neutron star binary, containing one millisecond pulsar and one normal, young pulsar. If the system is disrupted, both elements will instead be isolated in space. Although millisecond pulsars share many observational properties with their longer-period counterparts, Kramer et al. (1998) found that they are slightly less luminous and less efficient radio emitters than ordinary pulsars. Kramer et al. also demonstrated that profile changes as a function of observational frequency are slow in millisecond pulsars. This is suggestive of a compact magnetosphere in which different frequencies are emitted over a small range of heights.

In contrast to the slowly rotating ordinary pulsars, millisecond pulsars have narrow light cylinders and, hence, wide beams of emission and a high duty cycle. Although the pulse profile of a millisecond pulsar may be wide as a fraction of a pulse period, the high rate of rotation means that the pulse is still very narrow in comparison to an ordinary pulsar. This characteristic, combined with a small period derivative means that millisecond pulsars are very stable rotators that can be timed to a high level of precision (see Section 1.2.2 and Section 5.1.2).

Although the rotational reliability of millisecond pulsars is demonstrable, an unchanging pulse period is also necessary for optimum pulsar timing (see template

matching in Section 1.2.2), and the degree to which their integrated pulse profiles are stable requires further investigation. The precision timing of millisecond pulsars indicates that their profiles *are* generally robust, but a detailed investigation is presented in Chapter 5.

Magnetars

Magnetars are a subset of pulsars with long rotation periods, high spindown rates and, hence, are inferred to have extremely high magnetic fields (Equation 1.6). They fall into two subcategories: Soft Gamma Ray Repeaters (SGRs) and Anomalous X-ray pulsars (AXPs).

The first SGRs seen were initially categorised as γ -ray bursts. They were later noticed to repeat from the same location in space and subsequently acknowledged as distinct sources. Periodicities between two and eight seconds are seen in SGRs (observed in either X-rays or γ -rays), interpreted as the spin periodicity of their neutron star origin.

AXPs have rotation periods between 6 and 12 seconds, and spindown rates on the order of 10^{-11} , which lead to high magnetic fields (Equation 1.6). The decay of this magnetic field is thought to be the source of energy which heats the star and leads to the soft thermal X-ray emission observed.

AXPs and SGRs share parameter space on the $P - \dot{P}$ diagram (Figure 1.5) and also share some physical characteristics; some SGR are additionally X-ray sources (Kouveliotou et al., 1998), (Kouveliotou et al., 1999) while a number of AXPs have produced SGR-like bursts of X-rays (Gavriil et al., 2002).

Magnetars are known to produce X-rays flares that trigger extremely variable radio emission, as discussed in Section 1.3.4.




Figure 1.6: Flux density spectra for PSR B0329+54 (panel (a)) and PSR B1929+10 (panel (b)). The spectral behaviour of PSR B0329+54 can be fit by a single power law and shows low frequency turnover, while PSR B1929+10 can be fit by a broken power law and also shows a possible high-frequency turn-up. Figure from Lorimer & Kramer (2005).

1.2 Pulsar observations

1.2.1 Pulsar spectra

Although only a small fraction of a pulsar’s energy losses are converted into radio emission, the overwhelming majority of observational data that we have collected from pulsars has arrived as radio waves. The flux density spectra of radio pulsars is generally steep and can be approximated by a power law,

$$S_\nu \propto \nu^\alpha, \quad (1.9)$$

where S_ν is flux density at frequency ν , and α is the spectral index. The average value of α is ~ -1.4 (Bates et al., 2013). Around 5% of pulsars, however, require a two-component power law model to describe their spectra (Maron et al., 2000). As seen in the left panel of Figure 1.6, a low-frequency turnover is also often observed. At low radio frequencies, the propagation effects of scattering and

dispersion dominate observations and so can obscure weak and distant pulsars (see Section 1.2.4). Due to their steep radio spectrum, pulsar observations also become difficult above 3 GHz. Most pulsar searches are, therefore, carried out around 1.4 and 1.7 GHz, at frequencies reserved for HI and OH observations. At 1.4 GHz, measured pulsar flux densities range between $20 \mu\text{Jy}$ and 5Jy , and have typical values of a few mJy.

Despite the obstacles at low frequency, some instruments, such as the Low Frequency Array (LOFAR) are designed to observe at the largely unexplored low-frequency regime (van Haarlem et al., 2013). LOFAR is an interferometric radio telescope array with stations across Europe, concentrated in the Netherlands. It is comprised of low-band antennas that cover the 10-90 MHz range, and high-band antennas that cover 110-240 MHz.

1.2.2 Pulsar Timing

Rotation and timing

The high angular momentum of a rapidly spinning, dense star results in extreme rotational stability and gives pulsars their reputation as cosmic clocks. Our ability to precisely measure the arrival time of periodic pulses enables us to harness the power of pulsar timing.

The time of arrival (TOA) of a pulse is measured by a process of *template matching* (e.g. van Straten, 2006). The stability of pulse profiles at a given frequency permits the cross-correlation of an observed profile with a high signal-to-noise ratio (S/N) template, to provide a TOA of the former. The template is either a model or the average of many previous observations. The cross-correlation may be performed in

either the time or Fourier-transformed frequency domain. In the time-domain, the precision is typically one-tenth of the sampling interval, whereas in the frequency domain, the S/N of the pulse profile is the limiting factor (Taylor, 1992).

Timing models

A *timing model* attempts to comprehensively describe all physical aspects that affect the TOA of a pulse on Earth. This includes fundamental rotational parameters such as P and \dot{P} , estimations of the pulsar's position and proper motion, along with the pulsar's dispersion measure (Section 1.2.4), in order to account for delays incurred by radio pulse propagation. If the pulsar has a binary companion, parameters that describe the system must also be included. Additional parameters can be added to create an increasingly sophisticated and precise timing model, such as \ddot{P} , or details of rotational irregularities which are known to have occurred within the pulsar (Section 1.3.1).

The term *timing residuals* is used to describe the differences between the pulse arrival times expected by a timing model and the arrival times observed. As such, systematic trends in the timing residuals are either the result of an inaccurate or incomplete timing model, timing noise (Section 1.3.1), or another unanticipated process (e.g. Gravitational radiation; Section 5.1). Figure 1.7 demonstrates the consequences of a poor model for the timing residuals. A timing model is usually optimised by using a least-squares-fitting procedure to minimise the residuals. Bright millisecond pulsars with their intrinsically narrow pulses can be timed with the greatest precision. The uncertainty of a pulse time of arrival scales as follows:

$$\sigma_{TOA} \simeq \frac{W}{S/N} \quad (1.10)$$




Figure 1.7: Timing residuals of PSR B1133+66. Panel (a) shows a well-fitting timing model with residuals showing only white noise. Panel (b) shows the results of a timing model which underestimates the spindown rate of the pulsar by 4%. The cyclical nature of the timing residuals in Panel (c) is due to a position offset (here a $1'$ declination error) in the timing model. Panel (d) show the results of omitting the pulsar's proper motion of 380 mas yr^{-1} . Figure from Lorimer & Kramer (2005).

where W is the pulse width and S/N is the signal to noise ratio of the pulse. The most accurately measured millisecond pulsars have stability that is comparable to that of atomic clocks over decades, with TOAs that are measurable to microsecond precision and rotation periods that can be predicted to one part in 10^{15} (McLaughlin, 2013). Consequently, precision pulsar timing has a range of useful applications.

- Extracting pulsar masses from binary systems.
- It allows us to determine the astrometric parameters of pulsars, as it is sensitive to small inaccuracies in, for example, assumed position.
- The distance to nearby pulsars (within around 1 kpc of the Sun) can be determined by measurement of the small cyclic changes in pulse TOAs, induced by an apparent change in pulsar position over the course of the

Earth's orbit.

- We can also test gravitational physics in the strong-field regime; such studies only become possible in the extreme environment of a double neutron star system. The double pulsar system PSR J0737-3039 (Burgay et al., 2003; Lyne et al., 2004) permits tight constraints of its orbital parameters and, therefore, a stringent test of general relativity.
- The first exoplanets were discovered around a neutron star; the level of precision means that the timing data are sensitive to any cycles which would be induced by orbiting bodies.
- We can probe the galaxy and interstellar medium by observation of the various propagation effects experienced by the radio pulses. The speed, direction, polarisation and flux-density of radio waves can all be affected by the environment through which they travel, as discussed in (Section 1.2.4).
- Sudden rotational irregularities in pulsars known as *glitches* (Section 1.3.1) are thought to be due to interactions between the superfluid core and the crystalline crust. Glitch observations via precision pulsar timing have provided a method of probing neutron star interiors, allowing us to create models that describe the composition and behaviour of material inside the star (e.g. Andersson et al., 2012).
- Furthermore, pulsar timing is at the forefront of the race to directly detect gravitational radiation; an application that will be discussed in detail in Chapter 5.

As pulsar timing models become progressively precise, subtle deviations which are due to intrinsic instabilities in the pulsar are unveiled. It is these instabilities




Figure 1.8: Individual pulse power distributions for three pulsars made from 408 MHz observations. Figure from Smith (1973).

which are addressed throughout this thesis, and must be understood in order to advance the field of pulsar timing.

1.2.3 Radio emission

Individual pulses

The individual radio pulses that we receive have a typical width of 1° to 3° of the pulsar's 360° rotation. They vary from pulse-to-pulse in both intensity and location within the window of emission. In some pulsars, the power of an individual pulse has a Gaussian distribution around a mean value. In others, the distribution is asymmetric (Figure 1.8).

Giant pulses

The Crab pulsar is the prototypical emitter of giant pulses; their intensity can

be 1000 times larger than a typical individual pulse. After 25 years of being the sole example, the Crab was joined by other pulsars, including millisecond pulsars, that are also known to produce giant pulses (e.g. Cognard et al., 1996; Romani & Johnston, 2001; Johnston & Romani, 2003). The pulse longitude of the giant pulses seem to be at the same location as the double peaks of high energy emission. Giant pulses have been defined by Johnston & Romani (2004) as having a flux-density distribution that obeys power-law statistics and an association with non-thermal high energy emission. Extremely bright and extremely infrequent pulses demonstrate the existence of mechanism that can produce emission variability on short and long timescales.

Individual pulses have been observed with durations as short as 0.4 nanoseconds and intensities of 2 MJy (Hankins & Eilek, 2007), which implies a source less than a metre in any dimension.

Sub-pulse drifting

The location of individual pulses may favour particular phases within the window of emission, or be ostensibly random. It was noticed by Drake & Craft (1968), that in some cases, the pulses will appear at a phase which steadily drifts across the emission window, as illustrated in Figure 1.9. This phenomenon is known as sub-pulse drifting. The drift can occur in either direction, and is predominantly seen to occur in older pulsars (Weltevrede et al., 2006). One interpretation is that the emission regions within the cone of emission are moving as part of a rotating carousel of sub-beams (Rankin et al., 2005).

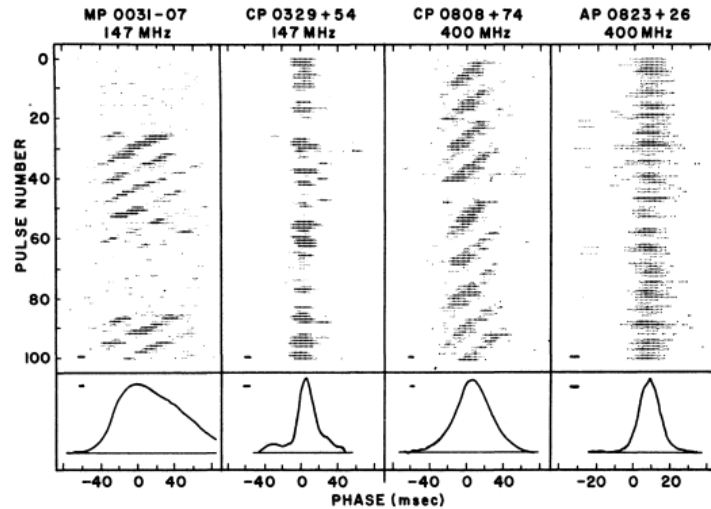


Figure 1.9: Sub-pulse drifting and/or nulling as seen in four pulsars. The positions of each sub-pulse is shown with respect to the integrated pulse profile. Figure from Taylor & Huguenin (1971).

Micropulses

High time resolution observations reveal features *within* individual pulses, known as micropulses, which have durations typically of a few microseconds. Observations taken at various frequencies simultaneously show that these features have a wide bandwidth.

Nulling and mode-changing

In the early 1970s, it was discovered that emission changes occur in pulsars on short timescales, in the forms of *nulling* and *mode-changing* (Backer, 1970a), (Backer, 1970b). Mode-changing is a phenomenon in which pulsars are seen to discretely switch between two or more emission states, each of which produces a different average profile over a sequence of individual pulses (Figure 1.10). When drifting subpulses (Section 1.3.2) are present, the drift rate is observed to change

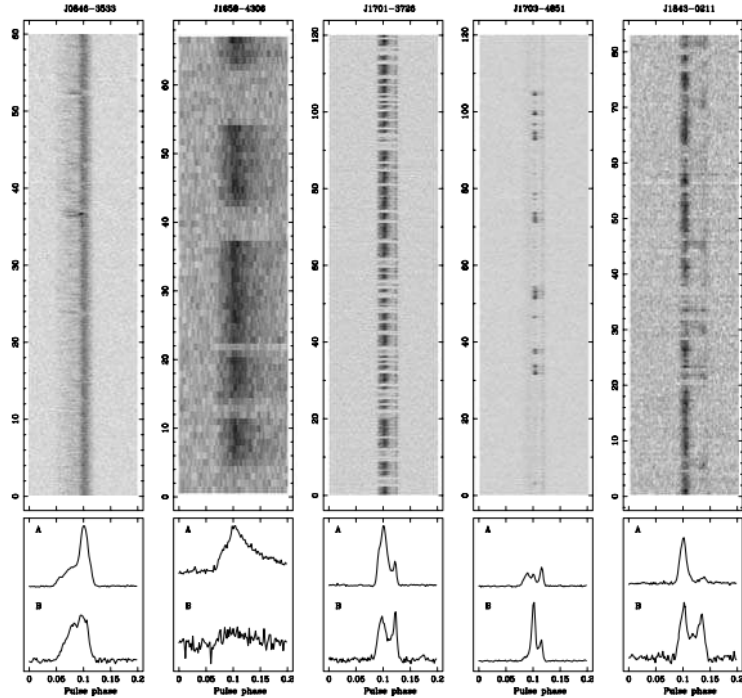


Figure 1.10: Five pulsars which exhibit mode-changing. The grey-scale is linear in intensity from zero to the maximum observed value. The lower panels show mean pulse profiles for the two observed modes. Figure from Wang et al. (2007).

when a mode-change occurs (Redman et al., 2005).

Nulling can be thought of as an extreme form of mode changing, with one state showing no, or low emission (Figure 1.11). The timescale of mode-changing and nulling ranges from a few pulse periods to many hours or even days (Wang et al., 2007). The fraction of time in which the pulsar is in a null state (the *nulling fraction*), also varies from zero to more than 0.5, and has been found to correlate with both age (Ritchings, 1976) and with pulse period (Biggs, 1992). The particle acceleration potential of a pulsar magnetosphere is dependent on P and \dot{P} , and theorised to drop below a critical value once the pulsar is beyond the *deathline* on the $P - \dot{P}$ diagram, a boundary beyond which radio emission is expected to cease. (Figure 1.5). The possibility must be considered, therefore, that nulling




Figure 1.11: Six pulsars that exhibit nulling. The grey-scale is linear in intensity from zero to the maximum observed value. Figure from Wang et al. (2007).

is the manifestation of a failing emission mechanism as the pulsar evolves toward the deathline.

Many pulsars exhibit nulling, mode-changing *and* sub-pulse drifting, suggesting a link between the phenomena. One important issue is whether nulling and mode-changing are related to pulsars that switch between states on much longer timescales (Section 1.3.2 and Section 1.3.3). This is the subject of Chapter 6.

Rotating radio transients

Rotating radio transients (RRATS) are a class of pulsar which are often regarded as extreme examples of nulling pulsars due to the extended periods of time they spend in a quiescent state. RRATs produce detectable emission (bursts typically lasting milliseconds) only sporadically, at irregular and infrequent intervals. More than 70 are now known since their discovery in 2006.

Pulsar searches usually make new discoveries by looking for precisely periodic signals in the data, which rise above the background noise when folded. Because




Figure 1.12: A 22 second series of pulses from PSR B0301+19 as observed by the Arecibo telescope. The lower panels show details of the pulse structure. Figure from Lorimer & Kramer (2005).

of the sparse nature of their pulses, RRATS are discovered via their individual pulses. Analysis of their burst arrival times does reveal underlying regularity of the order of seconds and they have comparable spindown rates to other neutron star classes (McLaughlin et al., 2006). They also show dispersion similar to a normal pulsar.

Integrated pulse profiles

As Figure 1.12 shows, the shape of radio pulses changes substantially from pulse to pulse.

If we take the average of at least a few hundred pulses, however, we start to see a stable shape emerge. As mentioned in Section 1.1.2, a pulse profile is an average measure of the flux density received from a pulsar as a function of its period and considered unique to each pulsar at a given observational frequency. Pulse profiles often are comprised of several (usually less than 5) components, which can usually be modelled as Gaussian or von Mises functions (e.g Kramer et al., 1994; Weltevrede & Johnston, 2008). The width of the pulse profile is measured as a fraction of the pulse period. The profile is usually defined as the window

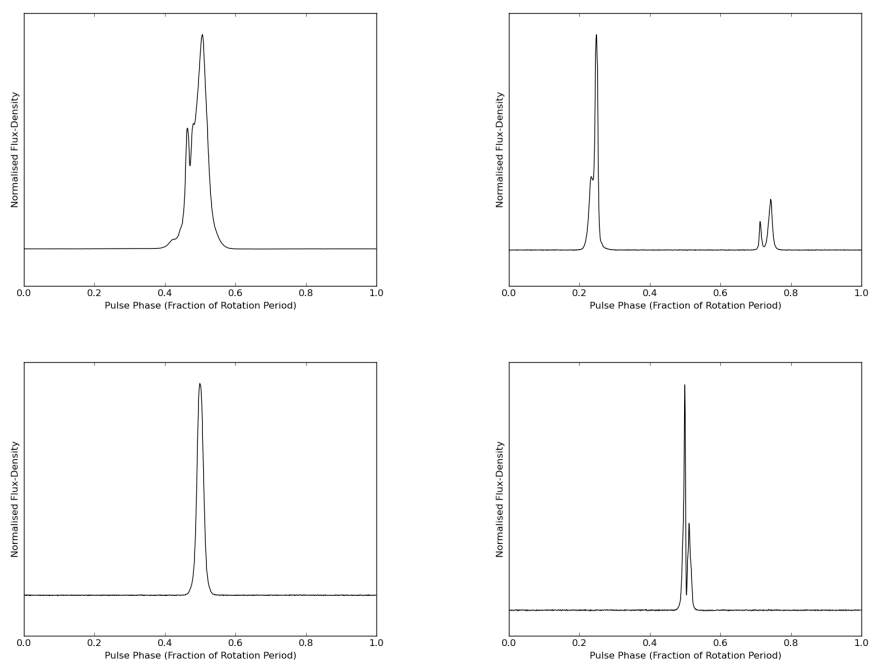


Figure 1.13: Four examples of integrated pulse profiles. Clockwise starting from top-left: PSR J0738-4042, PSR J0908-4913 which has an interpulse, PSR J1602-5100 and PSR J1359-6038.

in which the flux density stays above 10% of the peak (W_{10}). The width of the profile typically spans between 10-20% of the pulse period, but examples can be found in which the profile spans as little as 1%, or almost the whole pulse period. Assuming a circular emission cone, the pulse profile width can be calculated from the geometry of the emission beam,

$$\sin^2\left(\frac{W}{4}\right) = \frac{\sin^2(\rho/2) - \sin^2(\beta/2)}{\sin(\alpha)\sin(\alpha + \beta)} \quad (1.11)$$

(Gil et al., 1984), where α is the angle between the rotational and magnetic axes, β is the angle between the closest approach of the line of sight to the magnetic axis, W is the width of the pulse profile, and ρ is the angular radius of the emission cone.

The shape of the pulse profile is dependent on viewing geometry of the system. As described in Section 1.1, the pulsed radio emission is contained within a conical beam. The pulse profile observed is a result of how our line of sight cuts this emission beam and the number and nature of emission regions it intersects. The components of the pulse profile are each thought to correspond to an emission region in the beam (Rankin, 1993).

The width of a profile is usually seen to increase as the observing frequency drops (e.g. Xilouris et al., 1996). This seems to be due to the expansion of the spacing between individual profile components, and not of the components themselves. This observation has led to the idea that different radio frequencies are emitted at varying heights in the magnetosphere as described by Thorsett (1992). The radius-to-frequency mapping (RFM) model attributes an increasing profile width to the increasing angular width of an emission cone. This angular increase occurs as open magnetic field lines splay as they move away from the neutron star sur-

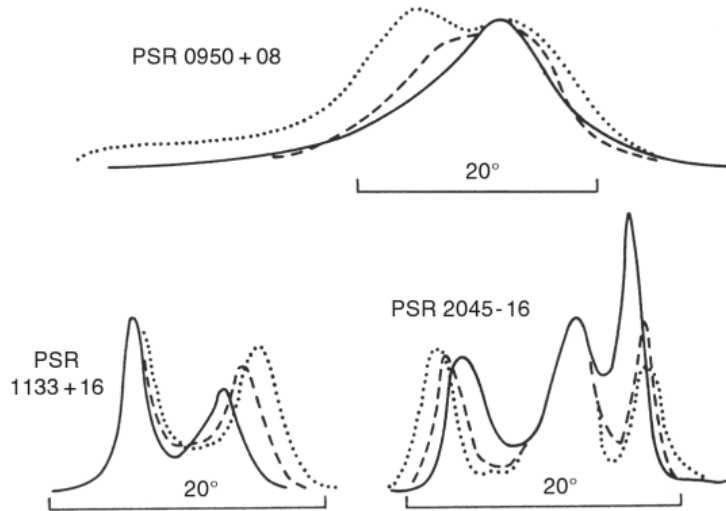


Figure 1.14: Integrated pulse profiles of three pulsars at three frequencies. Dotted line shows the profile at 150 MHz, the dashed line at 240 MHz and the solid line at 610 MHz. Figure from Lyne & Graham-Smith (2012).

face. In this model, lower frequency emission is assumed to occur further from the neutron star surface and to result in wider pulse profiles. Higher frequency radio emission is hypothesised to occur closer to the surface where the magnetic open field lines are more collimated.

The radio emission from many pulsars, especially younger ones, shows a high degree of linear polarisation. The orientation of the magnetic field lines at the emission region dictate the position angle of the linear polarisation. The Rotating Vector Model (RVM; Radhakrishnan & Cooke (1969)) predicts that the orientation of the magnetic field lines and, therefore, the position angle of the linear polarisation will gradually change as the line of sight cuts through the cone of emission. This will be observable as a gradually changing position angle across a pulse profile.

Many pulsars show an S-shaped swing in position angle (Figure 1.15) as predicted by the RVM. Others, however, display considerably more complexity, such as 90°

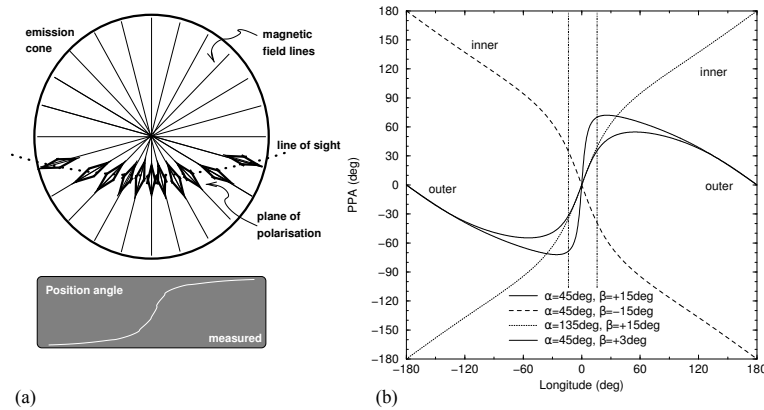


Figure 1.15: The rotating vector model. Panel (a) shows the magnetic axis from directly above. The plane of polarisation of the polar cap emission is dictated by the orientation of magnetic field lines from which it originates. As the line of sight sweeps across the magnetic field lines, the plane of polarisation changes such that the S-shaped curve at the bottom of the panel is produced. The steepest point of the curve is expected to occur when the line of sight passes closest to the axis. Panel (b) shows the shape of the position angle sweep for various impact parameters and inclination angles. The vertical lines denote the typical width of a pulse profile. Figure from Lorimer & Kramer (2005).

steps in position angle and a significant amount of circular polarisation. Such orthogonal jumps in position angle are common in pulsars with medium or low levels of linear polarisation. Karastergiou et al. (2011) demonstrate that this can be explained by the superposition of two entirely polarised orthogonal modes (see also Section 3.1.1).

Individual pulses often show a high degree of polarisation. The amount and orientation of polarisation of successive individual pulses, however, is not necessarily the same. The superposition of individual pulses can, therefore, diminish levels of polarisation, and can explain the lower levels seen in an average pulse profile.

1.2.4 Effects of radiation propagation

Dispersion

The group velocity of an electromagnetic wave depends upon the refractive index μ of the medium through which it is travelling. For the cold ionised plasma that makes up the interstellar medium, μ is frequency dependent. At different observing frequencies, therefore, the velocity of the wave propagation and hence travel time over a particular distance will change. The time delay due to this dispersion

$$t = D \times \frac{DM}{\nu^2}, \quad (1.12)$$

where the dispersion measure DM quantifies the total electron content between pulsar and observer, ν is the radiation frequency, and D is the dispersion constant

$$D \equiv \frac{e^2}{4\pi m_e c} = 4.1488 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}, \quad (1.13)$$

where e is the electron charge, m_e is the electron mass and c is the speed of light. Equation 1.12 shows that the delay of a signal is inversely proportional to the square of the observing frequency. A result of this is that the higher frequencies of a radio pulse will arrive at the observer earlier than the lower frequencies; a broadband pulse of intrinsically short duration will be stretched out. It is possible to correct for this using a technique called dedispersion (see below). A consequence of having a value for the DM of a pulsar along with an independent measure of its distance (e.g. via parallax), is a value for mean electron density in that direction of the sky.

In reality, the electron density in the ISM is inhomogeneous on various length

scales, and the DM for a pulsar can be subject to fluctuation. These fluctuations, if not accounted for in the timing model, can be a source of timing noise (Section 1.3.1). Such timing deviations can be misinterpreted as rotational changes intrinsic to the pulsar. Significant changes in electron density along the line of sight to a pulsar must, therefore, be considered when contemplating the sources of pulsar variability.

Dedispersion

It is possible to compensate for interstellar dispersion by splitting the observing frequency band into discrete channels, and applying a time delay to each. The magnitude of the delay in each channel is dependent on the DM. This is known as *incoherent dedispersion* (Figure 1.16). Of course, dispersion will occur within channels as well as across them, and so the channels must be chosen to be narrow enough to negate intra-channel effects.

A more sophisticated technique is known as *coherent dedispersion*, which results in better timing and pulse profile precision. In this process, the effects of dispersion are corrected for while the signal is still a raw voltage induced in the telescope by the electromagnetic signal. A frequency dependent phase delay is applied to the Fourier transform of this voltage. The delay is calculated to negate the effect of phase rotations induced by the ISM on the pulsar signal. A transformation back into the time domain is the final step in obtaining the required dedispersed data.

Scattering and scintillation

Another consequence of inhomogeneities in the ISM, is that electromagnetic waves travelling through an irregular medium are scattered and follow different paths to




Figure 1.16: Dispersion of pulses is shown in the top two panels. The bottom two panels shows how dividing the frequency band and applying the appropriate delay to each of the smaller channels can reduce pulse broadening and increase the S/N. This is incoherent dedispersion. Figure from Lorimer & Kramer (2005).

the observer. This necessarily gives rise to variations in wave amplitude. Relative motion between observer, pulsar and the intervening matter produce intensity fluctuations on various timescales. This is known as interstellar scintillation

$$I(t) \propto e^{-\Delta t/\tau_s}, \quad (1.14)$$

where Δt is the geometric time delay, and the scattering timescale $\tau_s \propto \nu^{-4}d^2$, where ν is the radiation frequency and d is the distance from the pulsar to the Earth.

The intensity fluctuations are produced by interference between waves which have travelled different paths, but this interference can only occur if the phases of the waves do not differ by more than approximately one radian such that

$$2\pi\Delta\nu\tau_s \sim 1, \quad (1.15)$$

where $\Delta\nu$ is the scintillation bandwidth. Waves within this frequency band can contribute to the interference. We see, therefore, that $\Delta\nu \propto 1/\tau_s \propto \nu^4$ results from a simple model which assumes a uniform scale for ISM inhomogeneities.

Scintillation can be either classed as *weak* or *strong*. Weak scintillation is expected only at high radio frequencies, observing objects at small distances, and so most pulsar observations are made in the strong scintillation regime. Strong interstellar scintillation is further divided into diffractive and refractive interference subcategories. Diffractive scintillation results from the interference, as described above. The intensity variations due to refractive scintillation occur on much longer timescales, typically weeks and months (e.g. Stinebring et al., 2000), and are a result of large scale focusing and defocussing of the radiation.

When observing long-term variations in flux density, it is not easy to disentangle the changes due to intrinsic emission instability in the pulsar, and those induced by refractive scintillation. The state-switching pulsars (Section 1.3.3), as well as examples in Chapter 4 seem to indicate that intrinsic changes have a tendency to alter the shape of a pulse profile, and not just its brightness.

Faraday Rotation

An additional dispersive phenomenon experienced by electromagnetic waves is *Faraday rotation*, in which the radiation undergoes frequency dependent changes in its plane of polarisation. The rotation in radians is given by

$$\phi = R\lambda^2, \tag{1.16}$$

where R is the *rotation measure* and λ is the observation wavelength. The rotation measure is dependent on electron density n_e and the magnetic field component along the line of sight $B \cos \theta$ such that

$$R = 0.81 \int n_e B \cos \theta dl, \tag{1.17}$$

where dl is an incremental distance to the source. A measurement of Faraday rotation along with an independent value of DM gives a direct measurement of the intervening magnetic field.

Variability in Faraday rotation could result in subtle pulse profile variability in poorly calibrated polarisation data.

1.3 Long-term monitoring

As we approach half a century of pulsar observations, we are learning more about how they behave over long timescales. There are many pulsars for which observations now span many years, and some with datasets spanning decades. These are helping us to realise that pulsars do not always maintain the stability with which they are synonymous; they show many departures from their dependability. The primary focus of this thesis is to learn more about these departures in order to elucidate their nature and their causes. This is done by discovering and scrutinising new cases of pulsar instability, along with the existing examples discussed in this section.

1.3.1 Rotational irregularities

Timing noise

Timing noise is the term given to the unexplained, systematic deviation from the modelled rotational behaviour of a pulsar. Although modest on short timescales, its amplitude can be large over months and years. It is younger pulsars that often show the most timing noise, and correlation is seen between timing noise activity and spindown rate (Hobbs et al., 2010). Timing noise is thought to often result from a pulsar switching between two distinct spindown rates (Lyne et al., 2010). In Chapter 4 of this thesis, we broaden this idea and explore the possibility that all timing noise is primarily due to changes in the spindown rate of a pulsar.

Glitches

Whereas timing noise has a fairly continuous and long-term nature, a more discrete




Figure 1.17: Examples of timing noise. The timing residuals shown are relative to simple timing models of ν and $\dot{\nu}$. Figure from Lyne et al. (2010).




Figure 1.18: Glitches in PSR B1930+22. From top to bottom: frequency, frequency residuals relative to a simple spindown model, and $\dot{\nu}$, each plotted as functions of time. Figure from Espinoza et al. (2011).

timing irregularity occurs when the rotation rate of a pulsar suddenly increases. This is known as a glitch (Figure 1.18).

Observed by the early arrival rate of pulses, a glitch is often followed by an exponential relaxation toward the pre-glitch rotation rate. The cause of glitches is thought to be a result of the transfer of angular momentum between the neutron star crust and the inner superfluid (Section 1.1). The crust is continuously being slowed by the electromagnetic torques provided by the magnetosphere. Due to low friction, however, the superfluid is not affected by this torque, and the disparity in rotation speed between crust and superfluid grows. A sudden increase in the coupling between these two components leads to an abrupt transfer of momentum and a discrete increase is observed in the pulse frequency.

As with timing noise, glitches are also predominantly observed in younger pulsars (Espinoza et al., 2011). They have also occasionally been observed to apparently trigger emission changes in pulsars, as discussed in Section 1.4.5.

1.3.2 Intermittent pulsars

A group known as *intermittent pulsars* go through a quasi-periodic cycle between phases in which radio emission is and is not detected (Kramer et al., 2006; Camilo et al., 2012; Lorimer et al., 2012). The timescale of this behaviour in our three current examples, ranges from months to years. In these objects, each of their two states is associated with a different spindown rate (c.f. the state-switching pulsars, Section 1.3.3). Because the energy associated with radio emission is only a fraction of a pulsar’s spindown energy, it is hypothesised that a global change in magnetospheric particle currents is taking place (Kramer et al., 2006). The changes in spindown are attributed to varying torque induced on the neutron star by changing currents. The radio emission would be simultaneously effected by the presence or absence of charged particles at the polar cap. State-switches in pulsars (Section 1.3.3) are similarly thought to be explained by changing magnetospheric currents.

The intermittent pulsar discovered by Kramer et al. (2006) was the first example of a pulsar showing emission changes that were strongly linked to rotational behaviour. Since this discovery, the long-term relationship between emission and rotation has been observed in other intermittent and ordinary pulsars.

The phenomenon displayed by intermittent pulsars is often considered an example of long timescale nulling. There is, however, nothing remarkable about the P and \dot{P} values of the intermittent pulsars discovered, and they are not in the region of the $P - \dot{P}$ diagram in which pulsar emission is expected to fail (Figure 1.5). The cause for the sudden change in magnetospheric current, therefore, remains unknown.

There are currently three published examples of intermittent pulsars; their prop-

Pulsar Name	J2000 Name	ν (Hz)	$\dot{\nu}_{off}$ (Hz s ⁻¹⁵)	$\dot{\nu}_{on}$ (Hz s ⁻¹⁵)	$\Delta\dot{\nu}/\dot{\nu}$ (%)	Typical <i>Off</i> Length (Days)	Typical <i>On</i> Length (Days)
B1931+24	J1933+2421	1.229	-10.8	-16.3	51	5-10	25-35
J1841-0500	J1841-0500	1.095	-16.7	-41.7	150	>350	580
J1832+0029	J1832+0029	1.873	-3.1	-5.4	77	1500	650-850

Table 1.1: Properties of the intermittent pulsars. $\Delta\dot{\nu}/\dot{\nu}$ is the fractional change in spindown rate (Kramer et al., 2006; Camilo et al., 2012; Lorimer et al., 2012).

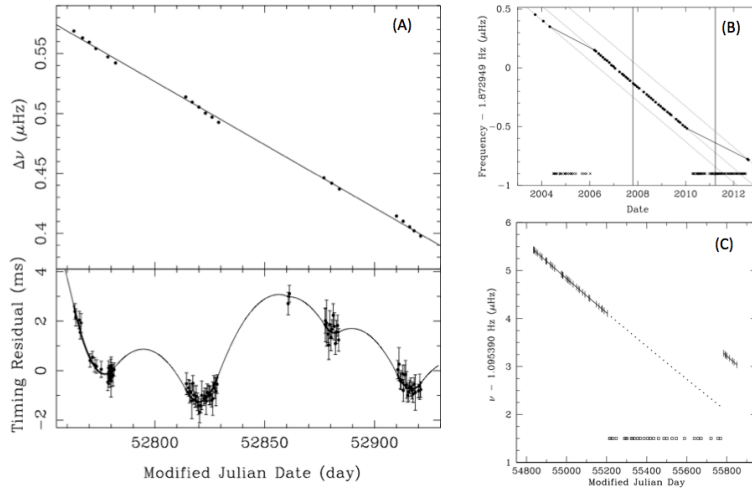


Figure 1.19: The rotation and emission behaviour of the intermittent pulsars. Each panel shows the decreasing rotational frequency of the pulsar over time. The large gaps between groups of data are the inactive emission states. The lower section of Panel (A) shows the pulsar’s timing residuals with respect to a simple timing model. A model with a distinct rate of spindown for each emission state is the solid line fit. Panel (A), (B) and (C) show data from PSR B1931+24 (Kramer et al., 2006), PSR J1841-0500 (Lorimer et al., 2012), and PSR J1832+0029 (Camilo et al., 2012) respectively.

erties are summarised in Table 1.1. The rotational history for each is plotted in Figure 1.19. It is worth noting that on MJD 55176 during the on phase, PSR J1841-0500 was detected as usual in a 300 second observation. It was then undetected in two subsequent observations, before being observed again later in the day; the pulsar appears to have been off for between 10 minutes and 2.7 hours. As well as its intermittency, this pulsar is unusual in a number of other ways. Firstly, it is located within only 4’ of another neutron star, the very young magnetar 1E 1841-045, which is within supernova remnant Kes 73. PSR J1841-0500 was actually discovered during a search for radio emission from the magnetar.

Secondly, the pulse profile experiences a huge amount of scatter. Furthermore, the magnitude of the rotation measure is the largest known at -3000 rad m^{-2} . Camilo et al. (2012) propose discrete intervening sources (e.g. HII region or supernova remnant) to explain the high scatter and RM observed.

For both PSR J1841-0500 and PSR J1832+0029, X-ray observations were made with the Chandra X-ray Observatory, but no emission was detected.

1.3.3 State-switching

In 2010, Lyne et al., showed six pulsars for which timing noise is correlated with changes in pulse shape over long timescales. The pulsars switch between two emission states, each with a distinct pulse profile (Figure 1.20). Lyne et al. characterise the profiles by a shape parameter, which can be seen in Figure 1.21 to either correlate or anti-correlate with the spindown rate of the pulsar. PSR J0742-2822 shows the most rapid changes, switching on a timescale of around 100 days, while PSR B2035+36 showed only 1 switch in 19 years of observation.

The pulsars featured in Lyne et al. were the first examples to show long-term correlated changes in pulse profile and rotational behaviour. Pulsars such as these may give valuable insight when trying to understand the root causes of pulsar variability, and motivate much of this work. New examples are discovered and discussed in Chapter 3 and Chapter 4 of this thesis. In the latter, we follow Lyne et al. in attributing timing noise to a changing spindown rate. We have further analysed the relationship between pulse profile and spindown variability in 168 young pulsars, nine of which are presented in Chapter 4.

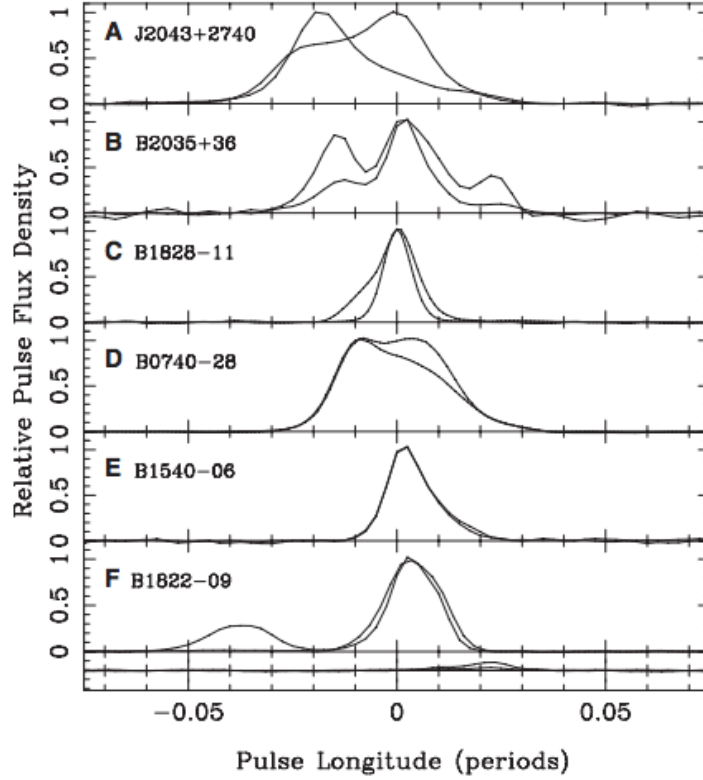


Figure 1.20: Integrated pulse profile changes of six pulsars observed at 1400 MHz. For each pulsar, the two traces show the most extreme profiles observed. Figure from Lyne et al. (2010).

Pulsar Name	J2000 Name	ν (Hz)	$\dot{\nu}$ (Hz s ⁻¹⁵)	$\Delta\dot{\nu}/\dot{\nu}$
B0740-28	J0742-2822	5.996	-604.36	0.66
B1540-06	J1543-0620	1.410	-1.75	1.71
B1822-09	J1825-0935	1.300	-88.31	3.28
B1828-11	J1830-1059	2.469	-365.68	0.71
B2035+36	J2037+3621	1.616	-12.05	13.28
J2043+2740	J2043+2740	10.40	-135.36	5.91

Table 1.2: Properties of six state-switching pulsars with correlated pulse profile and spindown rate variability. $\Delta\dot{\nu}/\dot{\nu}$ is the peak-to-peak fractional change in the spindown rate seen in Figure 1.21. (Lyne et al., 2010).




Figure 1.21: The relationship between profile shape and $\dot{\nu}$. For each of the six pulsars, the upper trace tracks a pulse shape parameter while the lower trace is $\dot{\nu}$. W_{10} , W_{50} and W_{75} are the full widths of the profile at 10%, 50% and 75% of peak amplitude respectively. W_{eq} is the pulse equivalent width (the ratio of the area under the pulse to the peak amplitude). A_{pc}/A_{mp} is the ratio of the amplitudes of the precursor and main pulse. Figure from Lyne et al. (2010).

1.3.4 Magnetar variability

Six magnetars for which the timing has been monitored over at least a few months all seem to show substantial timing instability (Kaspi et al., 2001; Gotthelf et al., 2002; Gavriil & Kaspi, 2002, 2004; Woods et al., 2002; Kaspi & Gavriil, 2003). Two of the six were SGRs, and show the most timing noise, with $\dot{\nu}$ varying by up to a factor of four. One magnetar, XTE J1810-197, was discovered in 2003 when its X-ray flux surged to a value around 100 times larger than had been seen over the previous quarter of a century (Camilo et al., 2007). Previously undetected positionally coincident radio emission was also observed after the X-ray discovery. The magnetar shows correlation between radio flux density and spindown rate. During a nine month observation period, the spindown rate decreases by $\sim 60\%$, while large fluctuations in flux density and flux shape are seen, driving a downward trend in flux density (Figure 1.22). It is noteworthy that a glitch is believed to have occurred at the time of the X-ray outburst in 2003.

1.4 Open questions

1.4.1 How do pulsar populations evolve with time?

The rotation period of a pulsar and its time derivative can be measured with astonishing accuracy. This gives us some insight into the rotational evolution of the neutron star. We expect that pulsars are born onto the top left of the $P - \dot{P}$ diagram (Figure 1.5) (Faucher-Giguère & Kaspi, 2006) with short periods and high rates of spindown. From here, energy loss guides their evolution through the central body of pulsars, before reaching the death-line. At this point, the electric field generated by the spinning magnetic dipole is not strong enough to




Figure 1.22: The spindown rate and flux-density at 1.4 GHz for XTE J1810-197 over a nine month period. The vertical dashed line denotes July 2006, at which time $\dot{\nu}$ increased at a large rate, the flux densities reached a local minimum, and the pulse profiles changed in character. Figure from Camilo et al. (2007)

accelerate particles as required for radio emission. As described in Section 1.1.3, the evolutionary path from the graveyard to the millisecond pulsar population is thought to be understood, but the track taken by magnetars and RRATs is less apparent.

Magnetars are thought to be young, as we find them close to the galactic plane. They are also expected to be short-lived due to their high spindown rates. It is not clear whether magnetars can evolve from ordinary pulsars or if they have a different origin from which they gain their high magnetic field. Some radio pulsars have physical properties which overlap those of AXPs, lending credence to the former scenario.

Similarly, it is not clear whether RRATs form a distinct population or if they can be evolutionarily linked to other classes of neutron star (Karastergiou et al., 2009). If the former is true, then we must ask why pulsars with similar spin properties are not also observed as RRATs.

We have seen that intermittent pulsars can be inactive for hundreds of days at a time, and could, therefore, belong to a population that remains largely undetected. Could it be possible that all pulsars undergo periods of intermittent activity or other emission variability if observed for long enough, and how would this affect their evolution?

1.4.2 Are long and short-term emission variability the same phenomena?

A natural question to ask when addressing the emerging examples of long-term emission variability, is how they relate to the short-term phenomena of nulling and mode-changing. Pulsar intermittency is now observed on a very wide range

of timescales. Are RRATs just pulsars which are exhibiting an extreme form of nulling, with nulling fractions reaching around 99%? Do the intermittent pulsars belong to this family or is intermittency produced by a completely distinct phenomena as suggested by Burke-Spolaor et al. (2011). Similarly, are state-switching and mode-changing the same phenomena on different timescales?

We could take a step towards unifying the long and short timescale phenomena by demonstrating that nulling and mode-changing behaviour is correlated with rotational changes, as long-term emission variability has been shown to be. Because they occur on timescales much shorter than the duration over which \dot{P} can be measured, however, we currently remain agnostic regarding rotational behaviour during the emission variability. A resolution to this problem is suggested in Chapter 6.

1.4.3 What initiates a change in magnetospheric currents?

Kramer et al. (2006), were the first to offer an explanation for the relationship between radio emission and rotation spindown rate seen in intermittent pulsar PSR B1931+24 (Section 1.3.2), citing variable charged particle currents in the pulsar magnetosphere. Changing plasma levels are expected to modify both the material outflow along open field lines at the polar cap, and the subsequent emission produced. Plasma variations would also alter the braking torque on the pulsar, and we would expect to see a change in spindown accompanying any significant change in emission. The cause of any shift in magnetospheric currents, however, remains unknown. Similarly, if changing magnetospheric currents also causes the quasi-periodic state switching, what causes the switch and dictates the timescales involved? We provide one hypothesis in Chapter 3.

1.4.4 What is the nature of timing noise?

If varying magnetospheric currents are indeed responsible for changes in emission and braking torque, as hypothesised, then correlation between spindown rate and emission variability follows. In some pulsars at least, switching between two rates of spindown is known to be the source of timing noise. Could it be that the ubiquitous phenomenon of timing noise is always the result of a mismodelled spindown rate? Is it possible, therefore, to account for timing deviations via the analysis of a changing pulse profile? If the answer is yes, continuous observation of a pulsar's emission profile would allow mitigation of any timing noise and vastly improve the pulsar's utility as a precision timing tool. We will investigate these questions in Chapter 4.

1.4.5 How is emission linked to glitching?

How does glitching fit into the paradigm of pulsar variability? Why do we occasionally observe emission changes after a glitch has occurred? As with timing noise, glitch activity has been shown to decrease with a pulsar's age (Espinoza et al., 2011), and both timing noise and glitches have been linked to pulse profile variability. In 2011, Weltevrede et al. presented various types of variable behaviour in PSR J1119-6127. The usually single-peaked profile showed a clear but short-lived double peak during 2007 (Figure 1.23). Additionally, a handful of strong RRAT-like pulses were observed from PSR J1119-6127 around the same time (Figure 1.24). Both the transient component and the RRAT-like pulses seen in PSR J1119-6127 were preceded by a large amplitude glitch in the spindown parameters.

As well as the above example, we have discussed glitches that are linked with




Figure 1.23: An observation showing the short-lived double-peaked mode of PSR J1119-6127. The left panel shows that dispersion produces the same effect in both components. The right panel shows the three minute period over which the double-peaked mode was observed. The intensity of both panels is shown in grey-scale. Figure from Weltevrede et al. (2011).




Figure 1.24: RRAT-like pulses seen in PSR J1119-6127. The top panel shows the average pulse profile from a 22 minute observation. The bottom panel shows strong individual pulses that are offset with respect to the average profile. Figure from Weltevrede et al. (2011).

X-ray emission in magnetars (Section 1.3.4) and radio emission in PSR B0740-28 (Section 1.3.3).

A relationship between glitching and emission variations has also been noted in RRATs; it has been suggested that in RRAT J1819-1458, significant increases in the rate of pulse detection and in the radio pulse energy, as well as a long-term decrease in spindown rate follow glitching events (Lyne et al., 2009).

1.5 This thesis

In the pursuit of answers to these questions, we have looked at long data sets of millisecond pulsars and ordinary pulsars. We have developed new techniques to discover long-term emission and rotational variability, and to probe the relationship between the two. We have also proposed how a similar relationship in short-term variability could potentially be discovered. Through work of this nature, we ultimately hope to fully understand the nature of pulsar evolution, to develop a more coherent picture of the apparently disparate pulsar classes, and to improve precision pulsar timing, aiding the detection of gravitational radiation.

Chapter 2

Observations and Analysis

Techniques

This chapter describes pulsar observations made to collect the data used in this thesis. It also details the subsequent analysis techniques used to investigate pulsar variability. The following is most directly relevant to work discussed in Chapters 3, 4 and 5. Chapter 3 deals with the variability of PSR J0738-4042, Chapter 4 with the variability pulsars from the Fermi Timing Programme and Chapter 5 discusses the variability of millisecond pulsars used by the NANOGrav pulsar timing array.

2.1 Observations

2.1.1 PSR J0738-4042 and Fermi Timing Programme observations

In 2014, astronomers from the University of Oxford, the Commonwealth Scientific and Industrial Research Organisation (CSIRO) in Australia and the Hartebeesthoek Radio Astronomy Observatory (HartRAO) in South Africa published a paper regarding the variability of PSR J0738-4042 (Brook et al., 2014). This work is the subject of Chapter 3. In order to monitor exactly how the profile and rotation of the pulsar have changed, we analysed 24 years of data, collected from two southern hemisphere radio observatories.

Hartebeesthoek Radio Astronomy Observatory

The data from PSR J0738-4042 used in Brook et al. (2014) were collected from September 1988 to September 2012 using the 26 m antenna at HartRAO. Observations were made at intervals from 1 to 14 days using receivers at 1600 MHz (1664 and 1668 MHz) or 2300MHz (2270 and 2273 MHz). We employed the 1600 MHz data for the profile stability analysis, whereas both datasets were used to obtain timing results. The observations were made of a single polarisation: left-hand circular. However, the degree of circular polarisation in this pulsar is low (Karastergiou et al., 2011), resulting in a negligible difference between true total power profiles and the HartRAO data (Figure 2.1).

Over the dataset, the backend provided a single frequency channel of 10 MHz (until 2003 April) and then of 8 MHz at 1600 MHz and of 16 MHz at 2300 MHz. Dispersion due to the interstellar medium is limited to ~ 3 ms across a 10 MHz




Figure 2.1: Comparison between a total power profile received using a full polarisation system and a single circular polarisation profile. As the degree of circular polarisation is low, the two profiles are comparable.

band centered at 1600 MHz. Observations usually consisted of three consecutive 15 minute (2400 pulsar period) integrations, each resulting in a single integrated profile. There is a gap in the regular coverage from April 1999 to August 2000 (MJD 51295-51775) due to an equipment upgrade. In order to determine the pulse time of arrival at each epoch, an analytic template consisting of three Gaussian components was fitted to the integrated profiles. The arrival times were then processed using the timing software *TEMPO2* (Hobbs et al., 2006).

Parkes Radio Telescope

The Parkes radio telescope is a 64 metre diameter parabolic dish, located in New South Wales, Australia. The Parkes telescope has been operational since 1961, and first observed a pulsar in 1968. Since then, it has gone on to discover more pulsars than any other telescope. These discoveries include PSR J2144-3933, the most slowly spinning radio pulsar known, with a period of 8 seconds; three double neutron star binaries, including the first double pulsar system; the first intermittent pulsar; dozens of millisecond pulsars and the first RRAT.

Since 2007, PSR J0738-4042 has been observed on a roughly monthly basis at

1369 MHz with the Parkes Radio Telescope in Australia, as part of the Fermi timing program (Weltevrede et al., 2010). The data were recorded with one of the Parkes Digital Filterbank systems (PDFB1/2/3/4) with a total bandwidth of 256 MHz in 1024 frequency channels. RFI was removed using median-filtering in the frequency domain and then manually excising bad sub-integrations. Flux densities have been calibrated by comparison to the continuum radio source 3C 218. The data were polarisation-calibrated for both differential gain and phase, and for cross coupling of the receiver. After this calibration, profiles were formed of total intensity (Stokes parameter I), averaged over time and frequency. The profiles were cross-correlated with high S/N templates to obtain times of arrival, using standard techniques for pulsar timing (Hobbs et al., 2006).

2.1.2 NANOGrav Observations

NANOGrav is a collaboration of researchers using the Green Bank Telescope and the Arecibo Observatory to detect and study gravitational radiation via pulsar timing (Chapter 5).

Green Bank Telescope

The Robert C. Byrd Green Bank Telescope is located at the National Radio Astronomy Observatory in Green Bank, West Virginia, and is one of two used by the NANOGrav collaboration as part of a pulsar timing array (Chapter 5). The dish is 100 metres in diameter, is fully steerable and can see all sky above its elevation limit of 5° .

The Green Bank Astronomical Signal Processor (GASP) was the previous pulsar backend system (Demorest et al., 2013) at Green Bank, which performed real-time

coherent dedispersion, full-Stokes detection, and pulse period folding. The maximum bandwidth used was 64 MHz, and two widely separated radio frequencies are used to allow for the correction of propagation effects (Section 1.2.4). The data were dispersed within 4 MHz sub-bands. The timing analysis was performed by cross-correlating observation profiles with noise-free templates to measure the pulse TOAs. The current backend system is the *Green Bank Ultimate Pulsar Processing Instrument* (GUPPI), an FPGA-based spectrometer capable of processing up to 800 MHz of bandwidth (DuPlain et al., 2008). A bandwidth of 200 MHz is used for observations at a centre frequency of around 800 MHz and a bandwidth of 800 MHz is used for observations at a centre frequency around 1500 MHz. The data are coherently dedispersed in 1.5625 MHz sub-bands.

Arecibo Observatory

The William E. Gordon telescope is a spherical reflector, located at the Arecibo Observatory in Puerto Rico and is also used by the NANOGrav collaboration as part of a pulsar timing array. At 305 metres in diameter, it is the largest radio telescope in the world, and can detect extremely weak radio signals, such as those from millisecond pulsars. The large dish, however, cannot be moved around and so the parts of the sky that Arecibo can access are limited; it is capable of seeing the sky directly overhead and 20° to each side.

The *Astronomical Signal Processor* (ASP) was the previous pulsar backend system (Demorest et al., 2013) at Arecibo, performed real-time coherent dedispersion (Section 1.2.4), full-Stokes detection, and pulse period folding in the same manner as GASP. The current backend system is the *Puerto Rico Ultimate Pulsar Processing Instrument* (PUPPI), operating in the same way as the equivalent GUPPI backend at Green Bank. At Arecibo, however, the bandwidths are 40 MHz, 700

MHz and 600 MHz at centre frequencies around 400 MHz, 1400 MHz and 2300 MHz respectively.

2.2 Analysis techniques

In the following section, we describe how we have used the TOA data and the average pulse profile of pulsar observations to study pulsar variability. Two versions of this analysis have been employed. The first was in our study of PSR J0738-4042 data up to 2012 (Brook et al., 2014). In subsequent work, the analysis techniques were improved as described below.

2.2.1 Pulse profile monitoring technique

This sub-section details how the average pulse profile of pulsar observations at the Parkes Telescope and HartRAO were used to study the long-term emission variability of PSR J0738-4042, as described in Chapter 3.

In general, the intervals between pulsar observations are not strictly regular and we, therefore, developed a technique that provides a useful visualisation of the variability in sparsely sampled data.

Firstly, profiles with an especially low S/N or a very malformed pulse profile, suggestive of instrumental error, are removed. Low S/N observations are removed for a number of reasons. Firstly, in noisy observations it is difficult to observe anything other than striking pulse profile deviations. Secondly, a low S/N relative to other observations of the same pulsar with the same instrument, is primarily the result of a shorter observation time; the profile may not contain enough integrated pulses to stabilise sufficiently. Thirdly, the average of all accepted profiles is later used as a template with which to compare individual observations; noisy profiles

could distort this template and reduce its S/N. We mitigate this possibility by making the template the median of all observations, and not the mean. Finally, noisy observations tend to saturate the colour-scale of the variability maps, which are discussed shortly.

The observed pulse profiles that remain, are normalised to their peak to facilitate comparison. In further iterations of the pulse profile monitoring technique, described in Section 2.2.4, we have considered both flux calibrated profiles and those normalised to the peak, allowing us to trace both shape and flux density changes. If the observations are deemed to have a S/N that is too low and will mask profile changes, then the average of multiple consecutive profiles can be taken as an observation. To emphasise their variability, we subtract a high S/N template from each observed profile to leave only the *emission residuals*. Plotting these resulting profile residuals produces a *variability map*, which highlights the regions in pulse phase and in time, where observations deviate most from the template. An example of a variability map can be seen in Panel F of Figure 3.6.

2.2.2 Spindown monitoring technique

For the timing analysis of PSR J0738-4042 performed in Brook et al. (2014), the pulsar TOAs were computed by a standard technique of cross-correlating the observed profile with a template (Section 1.2.2). Values of $\dot{\nu}$ were then determined from the TOAs at various observing frequencies: 1600 MHz and 2300 MHz at HartRAO and 1369 MHz at Parkes. The values were calculated using the *glitch* plugin to *TEMPO2*. Overlapping regions of 150 days were selected at 25 day intervals and a timing model of ν and $\dot{\nu}$ was fitted within each region.

2.2.3 Gaussian Process Regression

Gaussian process (GP) regression is a method which allows us to recover an underlying process from noisy observations. Modelling a process allows us to subsequently sample it at any interval, making it ideal when working with sparse and irregularly sampled pulsar datasets. Figure 2.2 shows how data constrain the possible models. As evident in the figure, the uncertainty is a measure of the range of models that could describe the data.

A GP generates data located throughout some domain such that any finite subset of that range follows a multivariate Gaussian distribution (Rasmussen & Williams, 2005; Roberts et al., 2012). It is a non-parametric technique, i.e. no functional form is imposed, giving more influence to the data themselves. Some prior knowledge is expressed, however, via the covariance function, which encodes our assumption about the function we wish to generate by describing how observations relate to one another.

There are many standard covariance functions to choose from, depending on the nature of the physical processes that one is trying to model. For additional flexibility, it is also possible to combine and modify existing covariance functions (Section 4.2).

GP regression forms the basis of our next generation monitoring techniques and is used to model variability in both the pulse profile shape and in the rotational behaviour of the pulsar, described in the following sections.

2.2.4 Enhanced pulse profile monitoring technique

We improved the profile modification technique used in Section 2.2.1 to produce improved variability maps and describe the process below.

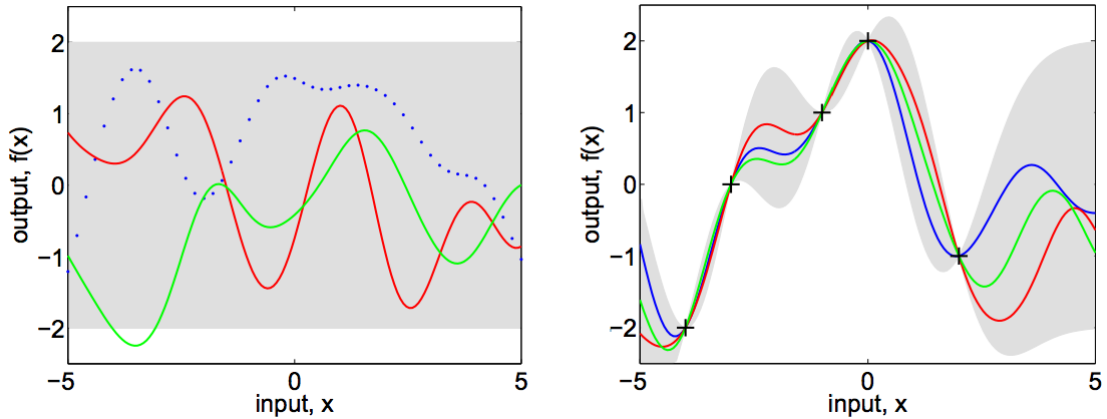


Figure 2.2: The left panel shows three functions that represent samples drawn from the Gaussian process prior distribution. The blue dotted line indicates y values that were generated for the given x values. The solid lines join a large number of evaluated points. In the right panel, five noise free observations have been added, which constrain the models. The solid lines shows three random samples drawn from the posterior distribution. In both plots, the shaded region shows twice the standard deviation at each x value. Figure from Rasmussen & Williams (2005).

Treatment of observed pulse profiles

As before, some observations are considered unreliable and excluded from further analysis. A pulse profile is excluded if the standard deviation of the off-pulse region is more than a factor of two larger than the median value taken from the off-pulse regions across all epochs. The reasons for removing these low S/N observations is described in Section 2.2.1. Observations are also removed manually if they show isolated extreme profile deviations which can likely be attributed to instrumental issues.

Pulse profiles originally consist of n phase bins across the pulse period. If S/N is low for any pulsar, then pulse profile variations can become difficult to detect above the noise. In cases where the observation with the highest S/N in a pulsar data set, has a peak value less than 20 times the standard deviation of the off-pulse noise, the number of phase bins is reduced to $n/8$, leading to poorer




Figure 2.3: Pulse profile intensity for 93 observations of PSR J1830-1059. The dark region for each observation is the peak of the pulse, which wanders in phase due to timing noise. The observations are unevenly spaced in time.

temporal resolution but higher S/N. The individual profiles are aligned by cross-correlation with the average over all epochs. Using the timing information to align the profiles is not possible, given the amount of timing noise in the data. Figure 2.3 demonstrates the phase shift of the unaligned pulse profile as a result of timing noise.

When analysing flux calibrated data, we are interested in using variability maps to monitor two types of profile variability: changes in the flux density across the whole pulse profile, i.e. brightness variations, and changes in the relative flux density of profile components, i.e. shape variations. To observe the latter, we

normalise all observations to the peak of their pulse profile. To identify this peak in the presence of noise, we take the median of 11 pulse phase bins; 5 either side of the maximum. Because of this, the profile peak still shows some jitter despite it being the feature to which the observations are normalised.

Two median profiles were calculated for each pulsar dataset; one for the normalised data described above and one for the non-normalised data with absolute flux calibration. The relevant profile was then subtracted from each observation, leaving a profile residual.

Gaussian Process Regression for pulse profile monitoring

To infer the behaviour of the emission residuals at any point in time, and not just on the days of observation, we must produce a continuous model, guided by the data. For each of the pulse phase bins, a GP regression was performed in order to produce a non-parametric function which best describes the emission residuals in that phase bin (Figure 2.4).

The covariance function chosen for this analysis employs a kernel from the Matern class,

$$k(x_i, x_j) = \sigma_f^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}d}{\lambda} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}d}{\lambda} \right), \quad (2.1)$$

where K_ν is a modified Bessel function, d is the distance $|x_i - x_j|$ between any two data points (training points), σ_f^2 is the maximum allowable covariance, and λ is the characteristic length scale, i.e. a parameter which reflects how significantly the distance between x_i and x_j affects $k(x_i, x_j)$. The positive covariance parameter ν was chosen to be $3/2$ to provide a level of smoothness and flexibility to the

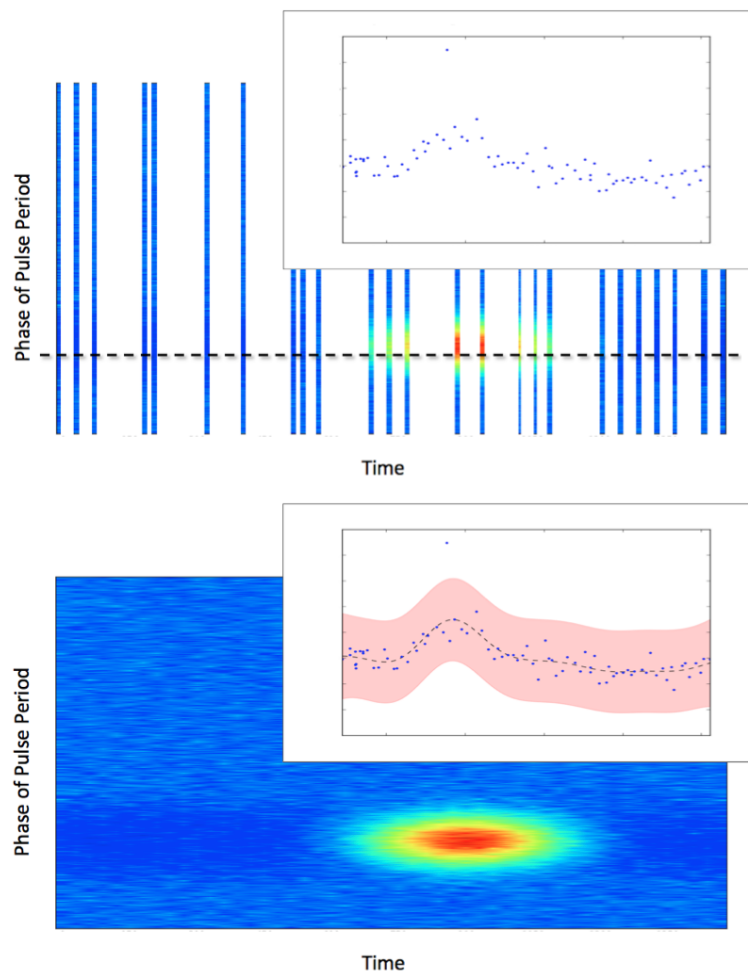


Figure 2.4: Illustration of the uses of GP regression. The top plot depicts unevenly space observations, that show a variability feature around the phase bin marked with a dashed line. The variability in the bin is represented in the inset. In the inset of the bottom plot, a continuous function has been fitted to the data via GP regression. Repeating this for all phase bins produces the solid colour variability map.

covariance function that is suitable for the kind of trends we are trying to model.

$${}_{3/2}k(x_i, x_j) = \sigma_f^2 \left(1 + \frac{\sqrt{3}d}{\lambda} \right) \exp \left(-\frac{\sqrt{3}d}{\lambda} \right) \quad (2.2)$$

This Matern covariance kernel was combined with a Gaussian noise kernel $\sigma_n^2 \delta_{x_i x_j}$ to model the uncertainty in the training points, where $\delta_{x_i x_j}$ is the Kronecker delta function and σ_n is the standard deviation of the noise in the training points. The choice of covariance parameters θ (σ_f , λ , σ_n) employed by the GP is made by maximising $\log p(\mathbf{y}|\mathbf{x}, \theta)$. If desired, we can impose boundaries on the covariance parameters to reflect the nature of the process we are trying to model.

The GP takes the training points (x_i and corresponding y_i) and calculates test points, i.e. the most likely value y_* for any value x_* , and its variance.

$$y_* = K_* K_{ij}^{-1} \mathbf{y} \quad (2.3)$$

$$\text{var}(y_*) = K_{**} - K_* K_{ij}^{-1} K_*^T \quad (2.4)$$

where K_{ij} is a covariance matrix with components $k(x_i, x_j)$ over all training points, K_* is a matrix which reflects covariance between a test point and the training points and has components $k(x_*, x_i)$. The covariance of a test point $K_{**} = k(x_*, x_*)$.

For each pulse phase bin, the GP regression produces a function which models the emission residuals. All phase bins can then be combined to produce a continuous variability map that highlights deviations across the pulse profile and across the data span (e.g. top and middle panels of Figure 4.3).

The strength of the GP regression technique used, is in modelling trends that are coherent over at least a few observations. This necessarily means that isolated

data points that buck a trend, may not be well modelled by GP regression. If, however, an outlying data point is far enough removed (e.g. a badly calibrated pulse profile), then in trying to accommodate it, the GP regression will flag up the location of the rogue data point on the variability map. In this way, the variability maps produced using GP regression techniques are a useful way to spot potentially problematic pulse profiles.

2.2.5 Enhanced spindown monitoring technique

Irregular sampling again raises difficulties in the calculation of $\dot{\nu}$, which we also address using GP non-parametric modelling. In doing so, we have developed an entirely new technique for measuring the spindown rate of a pulsar: employing timing residuals to model the second derivative of the timing residuals using GP regression. This analytical method is in contrast to the numerical methods often employed to determine variable spindown rates (e.g. Keith et al., 2013).

The timing model

When $\dot{\nu}$ deviates from its value in the pulsar timing model, non-Gaussian, systematic noise is introduced into the timing residuals. Under the assumption that this noise is entirely due to variations in $\dot{\nu}$, it is possible to estimate the $\dot{\nu}$ changes by first of all taking the second derivative of the timing residuals. In previous work (e.g. Keith et al. (2013)) the second numerical derivative of a smooth timing residual curve has been taken after interpolating between observations. We use a new technique which employs GP regression to analytically model the second derivative of the timing residuals, allowing us to produce a continuous function representing $\dot{\nu}$ for each pulsar.

Gaussian Process Regression for spindown monitoring

The covariance kernel used for this is the squared exponential kernel, as is it infinitely differentiable.

$$k(x_i, x_j) = \sigma_f^2 \exp\left(\frac{-d^2}{2\lambda^2}\right), \quad (2.5)$$

along with the Gaussian noise kernel. We begin by optimising the parameters θ , as described in Section 2.2.4. Following Holsclaw et al. (2013), the second derivative process of the GP can be estimated as

$$\frac{dy_*}{dx} = K_*'' K_{ij}^{-1} \mathbf{y} \quad (2.6)$$

where the second derivative of K_*

$$K_*'' = \frac{\sigma_f^2}{\lambda^2} \exp\left(\frac{-(x_i - x_*)^2}{2\lambda^2}\right) \left(1 - \frac{(x_i - x_*)^2}{\lambda^2}\right). \quad (2.7)$$

The variance of this method is given by

$$\text{var}\left(\frac{dy_*}{dx}\right) = K_*'' K_{ij}^{-1} K_*'' \quad (2.8)$$

Which gives:

$$\frac{3\sigma_f^2}{\lambda^4} \quad (2.9)$$

We have used the early versions of the techniques featured in this chapter when analysing data from PSR J0738-4042 in Chapter 3. The advanced versions are used for subsequent analysis in Chapter 4 and Chapter 5.

Chapter 3

The variability of PSR J0738-4042

PSR J0738-4042 is a bright, radio-emitting neutron star with rotational properties similar to the main population of middle-aged, isolated, radio pulsars. It has ν and $\dot{\nu}$ values of approximately 2.667 s^{-1} and $-1.15 \times 10^{-14} \text{ s}^{-2}$, respectively (Wang et al., 2001). In 2005 PSR J0738-4042 became a new and explicit example of correlated emission and rotation variability in pulsars.

3.1 Emission history

As a bright pulsar and one of the earliest to be discovered, PSR J0738-4042 has been observed frequently over recent decades. It was noticed by Karastergiou et al. (2011) that the average pulse profile in 2006 was significantly different to the profile observed two years earlier (Figure 3.1). Specifically, the leading edge of the pulse profile showed a prominent new feature that leads the profile peak by $\sim 15^\circ$. The feature was present over a broad frequency range in 2006, but was




Figure 3.1: The integrated pulse profile of PSR J0738-4042 as observed by the Parkes radio telescope at 50 cm, 20 cm and 10 cm wavelengths. The thin line traces the profile in the first half of 2004 and the thick line traces the profile in the latter half of 2006, in which the transient feature can be seen to appear at all wavelengths. Figure from Karastergiou et al. (2011).

Date	Frequency	Component at -15	Reference
<1970	1720 MHz	Strong and discrete	Komesaroff et al. (1970)
<1975	1400 MHz	Strong and discrete	Backer (1976)
<1977	631 MHz	Shoulder to main pulse	McCulloch et al. (1978)
<1977	1612 MHz	Strong and discrete	Manchester et al. (1980)
1979	950 MHz	Strong and discrete	van Ommen et al. (1997)
1990	950 MHz	Weak shoulder to main pulse	van Ommen et al. (1997)
1991	800 MHz	Absent	van Ommen et al. (1997)
1996	1375 MHz	Absent	unpublished
1997	1375 MHz	Absent	unpublished
2004	1375 MHz	Absent	Karastergiou & Johnston (2006)
2004	3100 MHz	Absent	Karastergiou & Johnston (2006)
2005	8400 MHz	Absent	Johnston et al. (2006)
2005	3100 MHz	Absent	Johnston et al. (2007)
2006	1369 MHz	Strong and discrete	Noutsos et al. (2009)

Table 3.1: 40 years of average profiles from PSR J0738-4042 (Karastergiou et al., 2011).

absent in 2004.

This finding precipitated a literature search in order to track the pulse profile over the history of observations. The available data are summarised in Table 3.1; the transient feature has been observed ever since the pulsar’s discovery, until the span between 1991 and 2005 when it was absent.

In 2005, the profile appears again in Parkes observations (Figure 3.2). A clear comparison of average profiles with and without the transient feature is shown in Figure 3.3. The transient feature persists until the current day, although the intensity of the component has not been steady. The biggest change since its latest




Figure 3.2: 71 irregularly spaced observations of PSR J0738-4042 between 2003 and 2011 with the Parkes radio telescope. The transient feature can be seen to appear at profile 23. Figure from Karastergiou et al. (2011).




Figure 3.3: The contrast between pulse profiles with and without the transient feature. The solid red line shows how the pulse profile of PSR J0738-4042 looked around 25 years ago. The dashed blue line shows the profile post-2005.

appearance came in 2010 when there was a discrete increase in the component’s flux density (Section 4.4.3).

3.1.1 Polarisation

Figure 3.4 shows that the appearance of the transient component in the profile of J0738-4042 is accompanied by a $\sim 90^\circ$ jump in the angle of polarisation, along with a drop in levels of linear polarisation. As discussed in Section 1.2.2, these changes can be shown to be the consequence of two competing orthogonal modes, each with entirely polarised radio emission (Karastergiou et al., 2011).

Karastergiou et al. were also able to show that the 2006 and 2010 pulse profiles (containing the transient feature) can be accurately modelled by combining the 2004 profile (without the feature) with a 100% polarised Gaussian component. The polarisation of this simulated component is orthogonal to that of the

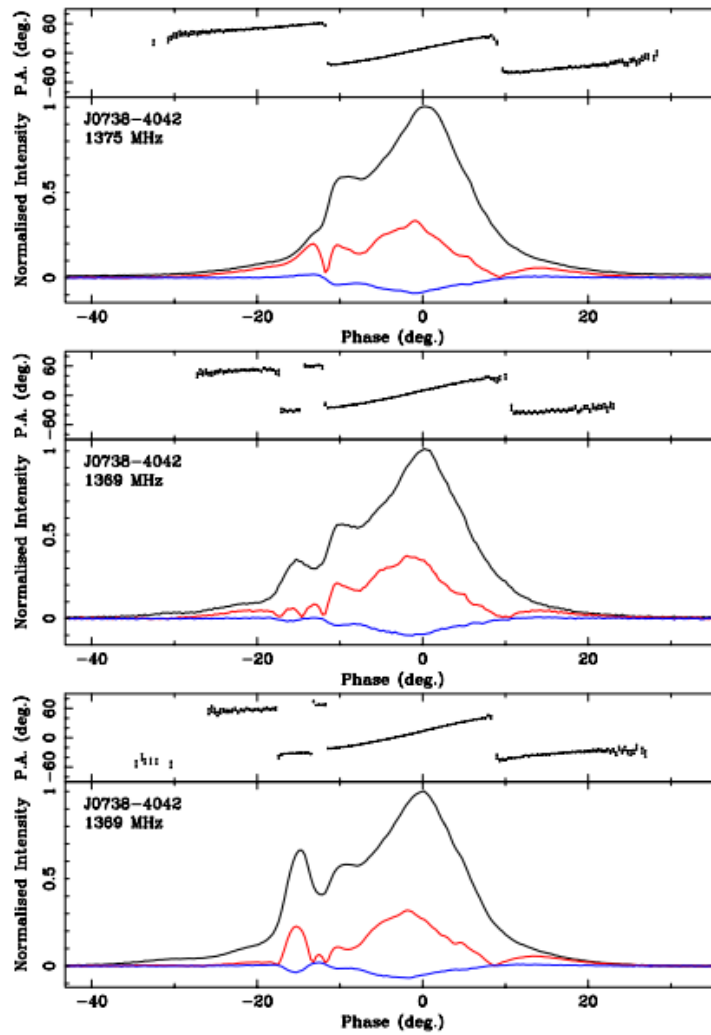


Figure 3.4: Integrated pulse profiles with polarisation information in 2004, 2006 and 2010 (from top to bottom). The black trace is the total intensity, the red and blue traces are the linear and circular polarisation respectively. The panels above each profile shows the position angle (PA) of the polarisation. Figure from Karastergiou et al. (2011).

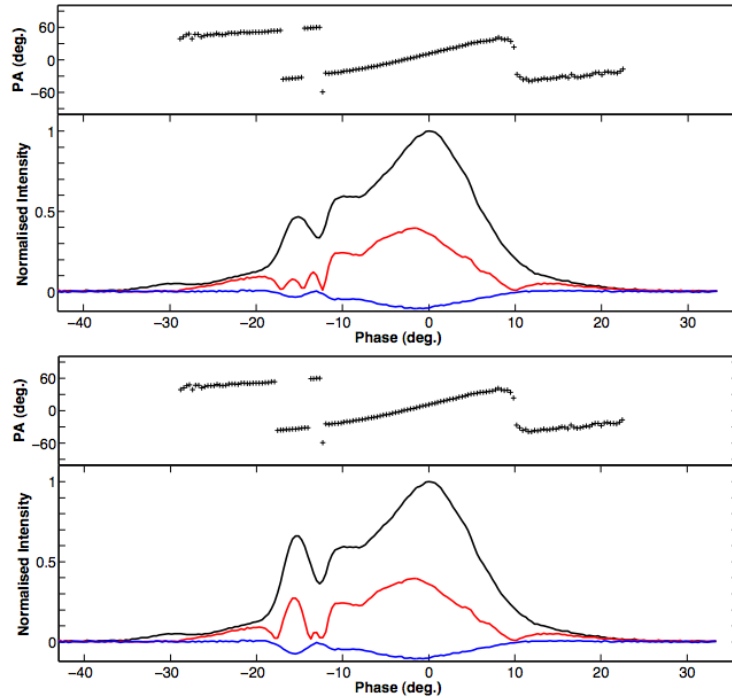


Figure 3.5: Simulations of the 2006 and 2010 profiles from Figure 3.4, created by adding a fixed-width, variable amplitude Gaussian component to the 2004 profile in which the transient feature is not present. Figure from Karastergiou et al. (2011).

corresponding 2004 pulse profile phase.

3.2 The dataset

As detailed in Section 2.1.1 and in Brook et al. (2014), data from HartROA and the Parkes Telescope was combined to produce a unique 24 year dataset with which to analyse the long-term variability of PSR J0738-4042. The radio emission and timing history of the pulsar, between 1988 and 2012 are presented in Figure 3.6. In order to identify changes in the pulse profile, the relatively low S/N of the HartROA observations was improved by combining observations (Section 2.2.1). To obtain profiles evenly spaced in time, we averaged eight profiles per 10 day interval. If a given 10 day interval contained fewer than eight observations, it was

extended by 10 days until eight profiles could be averaged. This average profile was then used for both 10 day intervals from which the data were taken. Because of the higher sensitivity and improved S/N, the averaging process was not necessary for the Parkes data. However, the data were still divided into 10 day windows, to maintain the same sampling as for HartRAO. If multiple observations existed in a single window, the profile with the highest S/N was used. If no observations existed within a given 10 day window, the profile from the previous window was carried over.

We have divided the data into five intervals with distinct profile shapes, shown in panels (A)(E). Panel (F) shows the residual between the data and the model as a function of epoch and rotational pulse phase, centered around the pulse peak and with a temporal resolution of ~ 1.46 ms. After 2008, the observations have a much higher S/N, reflecting the higher sensitivity of the Parkes Telescope.

The history of $\dot{\nu}$, computed at 25 day intervals (Section 2.2.2), is plotted in panel (G) of Figure 3.6. The HartRAO template used to determine the TOAs did not include the transient component, while it did feature on the Parkes template (see template matching in Section 1.2.2). Values of $\dot{\nu}$ with uncertainties in excess of 10^{-16} s $^{-2}$, were considered unreliable and not included in Figure 3.6.

3.3 The 2005 event

In September 2005, a dramatic change in pulse profile is seen to occur simultaneously with an abrupt change in rotational torque. The profile changes begin as a feature which drifts with respect to the rest of the pulse profile, before becoming a new and enduring profile component. The drifting feature and its relationship with $\dot{\nu}$ are shown in high contrast in Figure 3.7. Here, $\dot{\nu}$ is modelled using GP

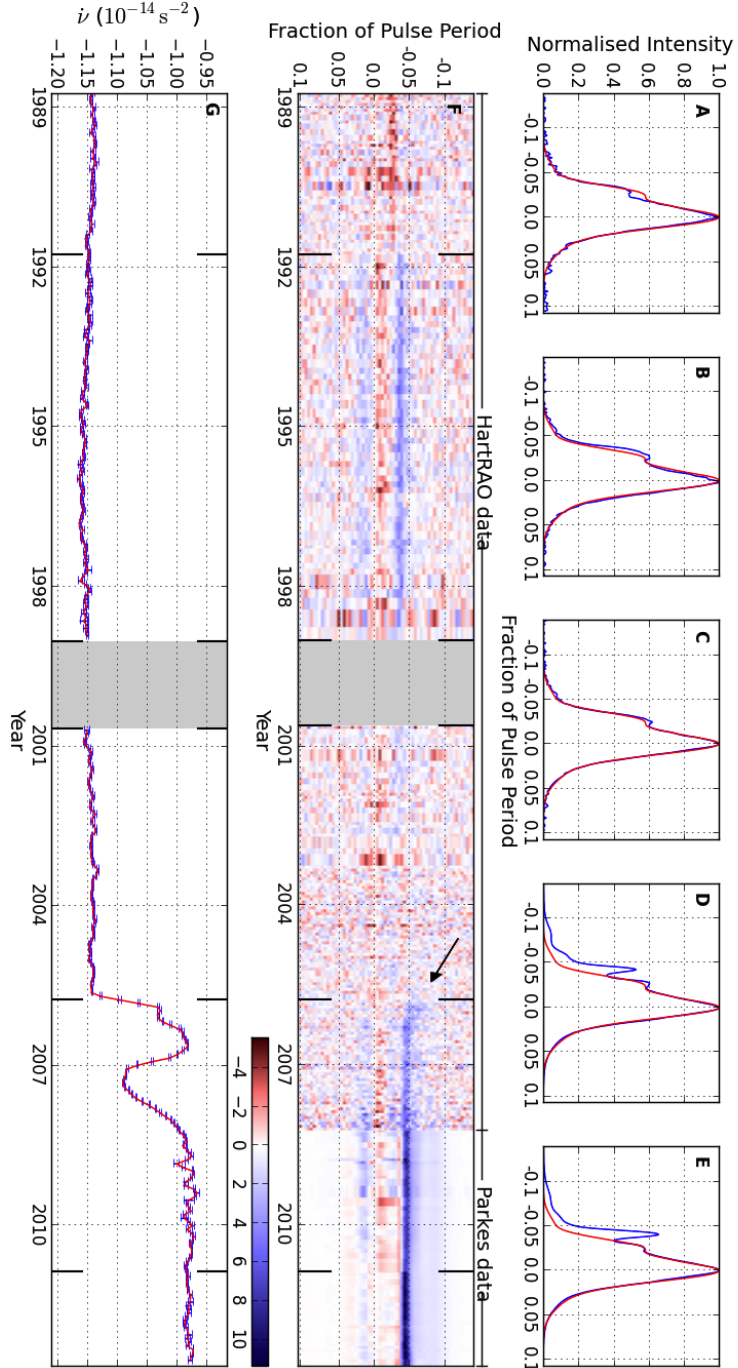


Figure 3.6: Variations in the profile shape and spindown rate seen in PSR J0738-4042. Profiles are observed at 1600 MHz with HartRAO and at 1369 MHz with the Parkes Telescope. Panels (A) to (E): the blue trace denotes the median pulse profile for each of five intervals over the 24 yr dataset, which are demarcated in panels (F) and (G). The red trace in each plot is a constant model profile which represents the median profile for all of the HartRAO data. Panel (F): map showing the difference between data and the constant model, in units of the HartRAO off-pulse standard deviation. The epochs at which data were collected from both telescopes were used to normalise the Parkes data to the HartRAO scale. The arrow points to drifting emission changes which precede the emergence of a new persistent profile component. Panel (G): the pulsar spindown rate as a function of time.

regression, as described throughout Section 2.2. The technique is employed much more extensively in the pulsar data analysed in chapters 4 and 5.

The drift occurs over ~ 0.02 of the pulse period and has a duration of ~ 100 days. As it begins, a pronounced change in $\dot{\nu}$ is simultaneously seen in the interpolated curve. The value of $\dot{\nu}$ is relatively stable both before the 2005 event, at $\sim -1.14 \times 10^{-14} \text{ s}^{-2}$ and ~ 1000 days later at $\sim -0.98 \times 10^{-14} \text{ s}^{-2}$.

Almost a decade after its sudden jump, the spindown rate of PSR J0738-4042 is yet to return to its pre-2005 rate. What the spindown rate does next, could provide important clues about the evolutionary path of this, and other pulsars. If the $\sim 15\%$ drop in $\dot{\nu}$ repeats without switching back to the higher value, then a continuous drop down the $P - \dot{P}$ diagram is inevitable. This behaviour would also set PSR J0738-4042 apart from the pulsars in which the spindown rate is seen to switch in both directions.

3.4 Interpretation

As discussed in Section 1.3.2, simultaneous changes in spindown rate and emission have been attributed to changing currents in the pulsar magnetosphere. The cause for the sudden change, however, is not clear. As no intrinsic pulsar process is known that can cause such sudden and dramatic alterations, the possibility that PSR J0738-4042 encountered an external body must be explored.

Material around pulsars

Evidence for planetary and disk systems around neutron stars is described in Shannon et al. (2013); they show that the timing of PSR B1937+21 is consistent with the presence of an asteroid belt. They also refer to other examples




Figure 3.7: Panel (a): $\dot{\nu}$ as a function of time, as computed over a 1700 day period referenced to 29th June, 2005. The curve is interpolated from data points on which it is overlaid. Panel (b): the profile residuals as they appear over the same time period. High contrast is used in order to emphasise the drifting feature.

of planetary and disk systems around neutron stars, such as the planets around PSR B1257+12 (Wolszczan & Frail, 1992) and PSR B1620-26 (Thorsett et al., 1999); the dust disk around magnetar 4U 0142+61 (Wang et al., 2006) and the unusual γ -ray burst GRB 101225A, thought to be due to a minor body falling onto a neutron star (Campana et al., 2011). Also of interest is their reference to circumstantial evidence for asteroid belts around white dwarfs, suggesting that rocky bodies may exist around post main sequence stars (Debes & Sigurdsson, 2002; Koester & Wilken, 2006; Farihi et al., 2011).

Mechanics of an asteroid encounter

In their 2008 paper, Cordes & Shannon describe the process by which asteroids, formed from supernova fallback material, may enter the magnetosphere of a pulsar and affect both the pulse profile and rate of spindown. The interaction may be initiated by an asteroid's migration due to collisions, orbital perturbations and the Yarkovsky effect, or by direct injection from eccentric orbits. A small body, falling toward a pulsar, is evaporated and ionised via pulsar radiation. The remaining charges, which are electrically captured by gap regions, can perturb particle acceleration in various ways. When accelerated to relativistic energies, charges can produce γ -rays, which give rise to a pair-production cascade. A quiescent region of the magnetosphere can be activated in this way, leading to new observable emission components. The injected charged particles may also diminish the electric field of a gap region, consequently attenuating an existing pair-production cascade. Additionally, the rearrangement of current within the magnetosphere would affect the braking torque and, therefore, the rate of spindown of the pulsar.

Drift height

As discussed in Section 1.2.2, radio emission generated at higher altitudes on a set of dipolar magnetic field lines will be beamed at larger angles with respect to the magnetic axis than emission at lower altitudes. Two interpretations of the drift shown in Figure 3.7 are either azimuthal movement of an emitting region at a given height, or an emission region moving in height along particular magnetic field lines. In the context of the former, the low rate of phase drift does not correspond to any process known, or seen previously in other pulsars (Weltevrede et al., 2006). Additionally, Cordes & Shannon note that a change in pair-production can result in a change in emission altitude for a given frequency, due to the plasma frequency dependence on height

$$\omega_p \propto n_{\pm}^{1/2}, \quad (3.1)$$

where n_{\pm} is the e^-p^+ pair density, and

$$\frac{\delta r}{r} = -\frac{1}{3} \frac{\delta n_{\pm}}{\delta n_{\pm}} \quad (3.2)$$

relates a change in altitude δr to a change in pair density. An increase in e^-p^+ pair density requires an increase in emission altitude to produce radiation at the same frequency.

In the case of a dipolar magnetic field emitting over multiple heights, there is a relationship between the angular radius of the field lines ρ at a given height and the observed pulse phase of the emission ϕ ; the lower the altitude, the closer the emission component will be to the center of the profile. This relationship can be derived using

$$\sin^2\left(\frac{W}{4}\right) = \frac{\sin^2(\rho/2) - \sin^2(\beta/2)}{\sin \alpha \sin(\alpha + \beta)}, \quad (3.3)$$

where α is the angle of the magnetic axis with respect to the rotation axis, β is the closest approach of the line of sight to the magnetic axis and W is the total width of the pulse profile (Gil et al., 1984). The observed pulse phase ϕ , measured from the peak of the profile, can be substituted directly for $W/2$. The values of ϕ at which the drift begins and ends are measured in degrees to be $\sim 23.4^\circ$ and $\sim 16.2^\circ$ respectively (Figure 3.7). In order to obtain a value for ρ , at the beginning and end of the drift (ρ_{begin} and ρ_{end}), α and β for the pulsar are needed. A value for β is obtained via the following equation:

$$\beta = \sin^{-1} \left(\frac{\sin \alpha}{|d\chi/d\phi|_{max}} \right), \quad (3.4)$$

where $|d\chi/d\phi|_{max}$ is the maximum rate of change of the polarisation position angle occurring around the center of the pulse profile (Rankin, 1993). For PSR J0738-4042, $|d\chi/d\phi|_{max}$ is ~ 3 (Karastergiou et al., 2011). The value for α is not easily constrained. Through calculations for various possible values of α , we find that ρ_{begin}/ρ_{end} is largely independent of α ; as α varies from 5° to 90° , equations 3.3 and 3.4 restrict ρ_{begin}/ρ_{end} between 1.23 and 1.20 for the observed drift positions. For emission originating close to the magnetic axis, ρ is related to the height H in the following way (Karastergiou & Johnston, 2007):

$$\rho \sim \sqrt{\frac{9\pi H}{2cP}} \quad (3.5)$$

and therefore,

$$\frac{H_{begin}}{H_{end}} = \left(\frac{\rho_{begin}}{\rho_{end}} \right)^2. \quad (3.6)$$

As ρ_{begin}/ρ_{end} is insensitive to α , so too is H_{begin}/H_{end} which has a value of ~ 1.5 . Therefore, if we do interpret the drift as an emission region moving in height

along particular field lines, we conclude that emission region decreases in height over the drift, and the change is around half of its final height of emission.

Magnitude of the magnetospheric current change

A reconfiguration in current within a pulsar magnetosphere would simultaneously affect the braking torque and, therefore, the spindown rate of the pulsar. The 2005 change in $\dot{\nu}$ can be interpreted as a reduction in the total outflowing plasma above the polar caps (Section 1.1.2). The magnitude of the current change, can be inferred from the change in $\dot{\nu}$ following Kramer et al. (2006). The difference between the two extreme $\dot{\nu}$ values corresponds to a reduction in the charge density ϵ of $\sim 7 \times 10^{-9} \text{ C cm}^{-3}$, where

$$\epsilon = \frac{3I\Delta\dot{\nu}}{R_{pc}^4 B_0}, \quad (3.7)$$

the moment of inertia I is taken to be 10^{45} g cm^2 , the magnetic field $B_0 = 3.2 \times 10^{19} \sqrt{-\dot{\nu}/\nu^3}$ gauss, the polar cap radius

$$R_{pc} = \sqrt{2\pi R^3 \nu / c} \quad (3.8)$$

and where the neutron star radius R is taken to be 10^6 cm . We can relate the difference in charge density associated with the two spindown states to mass supplied to the pulsar, by multiplying it by the speed of light, the polar cap area and the duration of the new spindown state. Between 2005 and 2014, this amounts to $\sim 10^{15} \text{ g}$, which lies within the range of known solar system asteroid masses, and is consistent with the mass range of asteroids around neutron stars proposed by Cordes & Shannon (2008). It should be noted however, that additional outflowing

material at the polar cap would be thought to increase a pulsar's rate of spindown and not decrease it as we see in PSR J0738-4042.

3.5 Unresolved issues

The multiple profile changes of PSR J0738-4042

If we have indeed witnessed an encounter between the pulsar and an asteroid, the question arises as to whether and why an event would be unique. Although the largest emission and spindown changes in PSR J0738-4042 occur in 2005, Panel (F) of Figure 3.6 shows a similar but less pronounced emission increase in 1992 along the leading edge of the pulsar. The figure also shows that in 2010, the new component experiences significant and sudden growth, seen clearly in panels (D) and (E). Neither the 1992 or the 2010 emission changes are accompanied by significant $\dot{\nu}$ changes. We also note the multiple profile modes of PSR J0738-4042, as opposed to the seemingly bimodal nature of the state-switching pulsars in Lyne et al. (2010) (discussed in Section 1.3.3). The 1992 and 2010 emission changes could also be caused by material entering the magnetosphere, but the smaller effect on pulse profile and the apparent stability of the spindown rate during the emission changes may suggest smaller amounts of infalling matter.

If the profile changes in PSR J0738-4042, other than the 2005 event, *are* also due to infalling material, then we must ask why the process repeats quasi-periodically and what dictates the decadal timescales involved. One possibility, is that profile changes could be due to a large orbiting body, such as a planet, which periodically disrupts debris in the fallback disk and precipitates inward migration. As an initial test of this hypothesis, we have simulated the perturbation to timing

measurements of this pulsar by an orbiting planet. We performed a periodicity analysis of the residuals in 1000 day segments, using the Cholesky pre-whitening method of Coles et al. (2011) and fitting for sinusoids with frequencies linearly spaced from 0.001 per day to 10 per day. We then computed a weighted mean of the resulting periodograms. Using a 5σ threshold we found no significant signals at long periods, and only two significant periodic signals, at exactly one per day and two per day, the origin of which is unclear. By injecting a simulated planetary signature with a circular orbit of radius 10^{11} cm (typical gravitational tidal radius; Cordes & Shannon (2008)) we rule out any such planet with mass greater than 6×10^{28} g. A planet of smaller mass may, therefore, exist without imprinting its signature on the timing data of this pulsar.

The relationship with other forms of pulsar variability

The behaviour of PSR J0738-4042 can be considered similar to some of the state switching pulsars discussed in Section 1.3.3. In particular, PSR B2035+36 has a comparable timescale of variability, shown in Lyne et al. (2010) to only switch state once in more than 8 years. Figure 1.21 shows that the spindown signature of PSR B2035+36 is also very similar to that of PSR J0738-4042, but switching from a lower to a higher rate of spindown in this case. The profile change in PSR B2035+36, seen in Figure 1.21, also bears some resemblance to the change in PSR J0738-4042. In both cases, the growth of a leading edge component is the primary change. The PSR J0738-4042 data show that, the rate of spindown is less when the the leading edge feature is present. This is also true of PSR B2035+36. Could the changes seen in PSR B2035+36 also be due to the introduction of material in the magnetosphere? Could external material explain the behaviour of the other state-switching pulsars too? Could it even explain those with quasi-periodic

variability occurring on shorter timescales, e.g. PSR B0740-28 with its switches on 100 day timescales? Could the regular input of material change the pulse profiles and spindown rates of these pulsars in such a similar way each cycle? We can also ask the same question of PSR J0738-4042. We see the transient feature appearing at various points throughout the pulsar's history. If we attribute its presence to a magnetospheric interaction with external material, could each encounter, potentially separated by decades, produce such similar profile changes each time?

The links between the various classes of variability have been discussed in Chapter 1. Is it possible that other forms of intrinsic variability are also triggered by interaction with external material? In some ways, the transient feature of PSR J0738-4042 can be compared to the appearance and disappearance of emission seen in the intermittent pulsars; both are accompanied by changes in spindown rate. All three current examples of intermittent pulsars, however, show an increase in spindown rate when the pulsar is active. This is in contrast to the drop in spindown rate which accompanied the transient feature of PSR J0738-4042 in 2005.

Chapter 4

Monitoring the variability of young, energetic radio pulsars

The striking example of long-term correlated emission and rotational variability observed in PSR J0738-4042 inspired further investigation of other pulsar datasets and the development of the variability analysis techniques created to study it. This chapter features a variability analysis of pulsars that have been observed since 2007 on a monthly basis by the Parkes radio telescope (Section 2.1.1). We have applied the technique to 168 pulsars and present results from nine sources with obvious variability. In light of the burgeoning relationship between emission and rotation variability, the primary goal was to investigate the following hypothesis: all timing noise (Section 1.3.1) can be attributed to unmodelled variations in spindown rate, caused by magnetospheric processes that also change the average emission profile. This idea is tested by computing the variability in $\dot{\nu}$ and in pulse profile, and identifying any correlation between the two.

4.1 Fermi Timing Programme Observations

On June 11th 2008, the Fermi Gamma-Ray Space Telescope was launched, carrying the Large Area Telescope (FermiLAT). The study of pulsars at γ -ray wavelengths is one of its key science projects; those potentially detectable by FermiLAT are young, high \dot{E} objects. Since its launch, FermiLAT has detected and studied the properties of almost 170 pulsars (You et al., 2007).

Unlike radio pulsars, γ -ray pulsars cannot be detected in real-time. Because of their dearth, γ -ray photons must instead be collected over weeks, months or even years in order to build a pulse profile. For all but a handful of the brightest sources, these profiles are produced by folding the photons using an accurate radio timing ephemeris. In order to facilitate such γ -ray observations, the Parkes radio telescope began timing 168 pulsars in 2007 as part of the *Fermi Timing Programme*. The majority of the pulsars have an \dot{E} in excess of 10^{34} erg s $^{-1}$ and are, therefore, prime candidates to be detected by FermiLAT. These young, energetic pulsars often show a high degree of timing noise, and must be observed monthly in order to reach the timing precision required.

4.2 Modelling the timing residuals

As mentioned at the beginning of the chapter, we are testing the hypothesis that all systematic noise in pulsar timing is the product of unmodelled changes in $\dot{\nu}$. Under this assumption we can infer the behaviour of $\dot{\nu}$ from the nature of the timing residuals, as described in Section 2.2.5. Before observing the timing residuals, we must first ensure that our pulsar timing model is optimised. As discussed in Section 1.2.2, this is done by using a least-squares-fitting procedure to minimise

the timing residuals. An assumption of this technique is that the residuals contain only uncorrelated (white) noise. In reality, correlated non-Gaussian noise is often also present (Section 1.3.1), leading to inaccurate parameter estimates in the timing model. Usually, this problem is best addressed using the Cholesky method (Coles et al., 2011), which allows us to find a linear transformation that whitens and normalises the residuals. The transformation can be found by estimating the covariance function of the residuals, and is applied to both the residuals and the timing model prior to performing the least-squares-fitting procedure. However, as we are interested in examining whether the systematic, non-Gaussian is the result of a changing $\dot{\nu}$ value, we do not wish to remove it. In this analysis, therefore, the timing models and residuals are calculated without the Cholesky method and the resulting removal of non-Gaussian noise.

When this process was implemented, however, we often saw a yearly sinusoidal structure in the timing residuals, similar to panels (c) and (d) of Figure 1.7. This indicated that our position and/or proper motion parameters for the pulsar were not optimised. To rectify this, we employed the Cholesky method to more accurately optimise the right ascension, declination and proper motion parameters only. Doing this allowed us to eliminate any yearly cycles, without removing any red noise in the timing residuals.

Using multiple kernels

When using GP regression to model the timing residuals (Section 2.2.4), we noticed that choosing a single squared exponential term for the covariance kernel often resulted in a poor fit, leaving an unmodelled short-timescale feature. Adding a second squared exponential kernel to the covariance function, with a short length-scale, modelled the timing residuals more accurately in the majority of cases.




Figure 4.1: The timing residuals and GP model for PSR J0940-5428 using one kernel in the covariance function. Top panel: The red points are the timing residuals. The black trace shows the GP model, which has a covariance function that employs one kernel with a lengthscale of 276 days. The shaded 2σ uncertainty region indicates the range of GP models that can describe the data. Bottom panel: The timing residuals minus the GP model at the epochs of the observations. The structure seen here implies an ill-fitting model that does not account for the short-term periodic behaviour seen in the data. The uncertainty in bottom panel is that of the timing residuals.

Figures 4.1 and 4.2 show the timing residuals of PSR J0940-5428 fitted with a GP using one and two kernels respectively. The discrepancies between the data and the model are also shown and justify the use of two kernels in this case.

With the timing residuals accurately modelled by the GP regression, the spindown rate can be modelled as described in Section 2.2.5.




Figure 4.2: The timing residuals and GP model for PSR J0940-5428. As Figure 4.1, but employing two kernels in the covariance function; the lengthscales are 60 days and 471 days. The lack of structure in the bottom panel suggests a well-fitting model and justifies the number of kernels and optimised parameters used in the covariance function.

4.3 Variability correlation

In order to formally test the relationship between a pulsar’s spindown rate and the change in pulse profile, we look at the correlation between $\dot{\nu}$ and the profile residuals in each normalised phase bin. Specifically, we find the Spearman’s rank correlation coefficient (r) in each case, which is a measure of statistical dependence between two sets of data. The two sets are assigned rank values p_i and q_i , and then $d_i = p_i - q_i$. The correlation coefficient is given by

$$r_{p,q} = 1 - \frac{6 \sum_i d_i^2}{n^3 - n}, \quad (4.1)$$

where n is the sample size.

The correlation between every phase bin and the spindown rate was calculated this way. Furthermore, this was done with lags of up to ± 500 days between the two timeseries. The resulting *correlation maps* can be seen in the bottom panels of figures 4.4, 4.6, 4.8, 4.12, 4.15, 4.16, 4.18, 4.20, 4.22 and 4.25. In these maps, we can identify regions of the pulse profile that have an interesting relationship with $\dot{\nu}$, and gain information about the temporal relationship between the changes in $\dot{\nu}$ and in pulse profile. The Spearman’s rank correlation coefficient is only calculated for the pulse phase bins that lie within the on-pulse region. We have defined this to be a region where the median pulse profile of a dataset is greater than 1/30 of its peak. In each correlation map, a negative lag means that the spindown timeseries has been shifted forward to lead the phase bin timeseries. A correlation at -100 days lag, for example, would mean that features in the phase bins occurred 100 days before the matching features in $\dot{\nu}$.

4.4 Notable examples

All 168 pulsars were initially analysed, with the aim of detecting variability in their pulse profiles. The rotational variability analysis was more challenging; a first look revealed lots of artificial yearly cycles, indicating that the positional parameters of the pulsar were not fit accurately enough. In order to remove the cycle and uncover the intrinsic nature of the timing residuals, a Cholesky analysis would have to be individually performed on each pulsar (Section 4.2). Instead, pulse profile variability maps were created for all the pulsars, and those with interesting pulse profile changes were earmarked to have their rotational behaviour scrutinised. The majority of the observations in the Fermi Timing Programme dataset were too noisy to clearly see profile variability below a certain level. This leaves the possibility that some variable pulsars in the dataset will remain hidden until higher S/N observations can be obtained. Below are examples of the pulsars which showed the most interesting behaviour.

4.4.1 PSR J1830-1059 (B1828-11)

Long-term variability in both emission and spindown rate are well established in PSR J1830-1059, along with correlation between the two (Lyne et al., 2010). The variability has been attributed to free precession (Stairs et al., 2000) or the effects of an orbiting quark planet (Liu et al., 2007). Quasi-periodic profile shape changes can be seen clearly in the middle panel of Figure 4.3. The pulse profile switches back and forth between two distinct emission states; the contrast in the pulse profile shape of the two states can be seen in Figure 4.4. The value of $\dot{\nu}$ is closely correlated with the emission state of the pulsar; $\dot{\nu}$ is at its greatest when the pulsar is in the brightest mode (Figure 4.3). In the pre-normalised,

flux calibrated observations, one state has a peak brightness that is a factor of ~ 2.5 greater than that of the other, with the brightest peak being typically ~ 250 mJy. Figure 4.4 shows the correlation between $\dot{\nu}$ and pulse profile variability at each pulse phase, as a function of the lag between the two timeseries. As the flux density variability is known to be highly correlated with $\dot{\nu}$ in this pulsar, we see the strongest correlation occurring around zero lag. We also see that the timeseries are correlated when $\dot{\nu}$ has a ± 500 day lag. This is due to the fact that the spindown rate appears to have a clear cycle of around 500 days, as seen in Figure 4.3.

4.4.2 PSR J1602-5100 (B1558-50)

PSR J1602-5100 has the lowest spin frequency of the nine pulsars featured in this chapter, and consequently, only PSR J0738-4042 has a lower spindown luminosity \dot{E} . When brightness variations have been removed by normalisation of the observed profile peaks, the pulsar shows a dramatic profile shape change which occurs over ~ 600 days, beginning at \sim MJD 54700 (Middle panel of Figure 4.5). The flux calibrated observations show that the appearance of the new peak at the trailing edge of the smaller profile component, coincides with a drop in flux density at the primary pulse component. The extent of the profile shape change can be seen in Figure 4.6. A drop in $\dot{\nu}$ of $\sim 5\%$ can also be seen, which is correlated with the shape change; both sets of variations seem to begin, peak and end at approximately the same time. The correlation between $\dot{\nu}$ and flux density variability is strongest at zero lag, around the pulse profile phases at which the transient component appears, and at the leading edge of the main pulse (Figure 4.6). The fact that the features on this correlation map extend further along




Figure 4.3: Pulse profile and spindown variability for PSR J1830-1059. Top panel: Map showing the difference between the flux calibrated observations and the average profile across the dataset. The units are the average level of the off-peak standard deviation of observations. The solid horizontal line highlights the pulse peak and the vertical dashed lines show the dates of the included observations. Middle panel: As top panel, but the observations are first normalised to the pulse peak. Bottom panel: Value of $\dot{\nu}$ on observation dates.

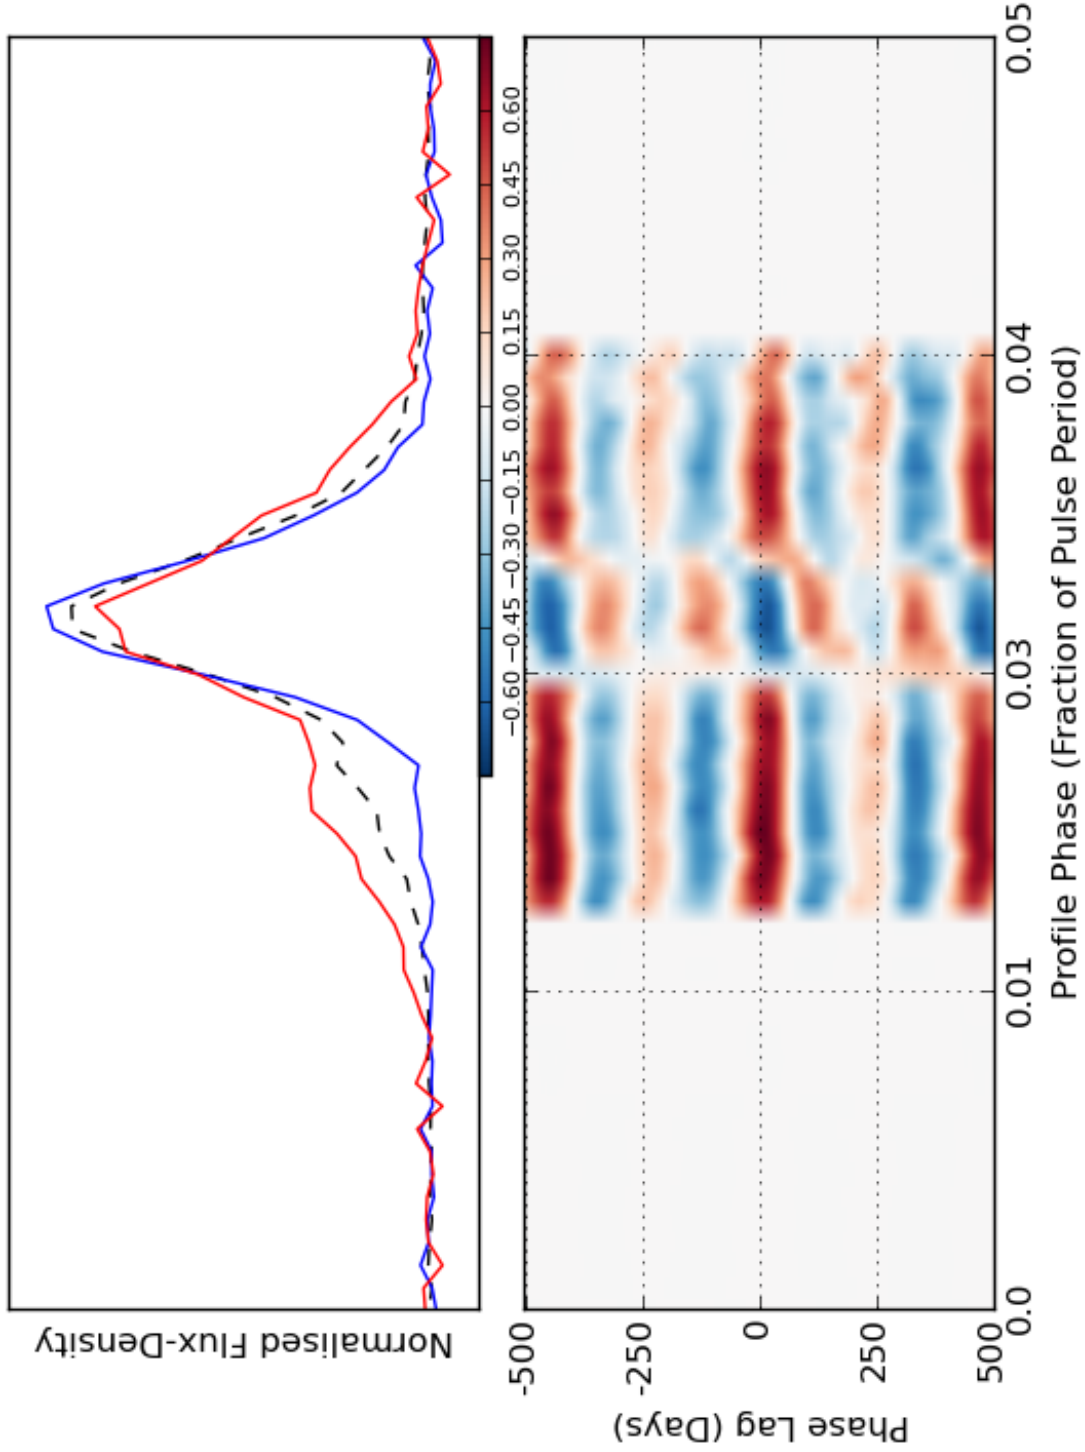


Figure 4.4: Top: Pulse profile variability for PSR J1830-1059. Black dashed line traces the average of normalised pulse profiles across all observations. The blue and red lines are examples of profiles which show the extent of shape changes. The blue profile was observed on MJD 54353, red on MJD 55699. Bottom: Spearman's rank correlation coefficient for ν and flux density variability as a function of the pulse phase and lag between the two timeseries. A negative lag means that flux-density variability is lagging ν

the y-axis than those in the correlation map for PSR J1830-1059, is a reflection of the broad nature of the spindown rate feature as seen in Figure 4.5.

4.4.3 PSR J0738-4042 (B0736-40)

As the subject of Chapter 3, we know PSR J0738-4042 to show a dramatic change in both pulse profile and $\dot{\nu}$, beginning in 2005. This despite the fact that it is otherwise unremarkable, with rotational properties typical of the main population of middle-aged, isolated, radio pulsars. As such, PSR J0738-4042 was not originally a candidate for the Fermi Timing Programme, but is now included in the observations due to its peculiar nature. Regular Parkes observations of the pulsar began in March 2008. Since then, the most prominent change in profile occurred in November 2010 (\sim MJD 55525), when the relative size of the new component increased significantly, and has shown a trend of gradual recession ever since. Preceding this component growth, both chronologically (by \sim 250 days) and in phase (by \sim 0.015 of a pulse period), a drop in flux density of a small profile component is seen. These two changes are illustrated in Figure 4.7 and Figure 4.8. The flux calibrated observations show typical brightness fluctuations of \sim 5% around the median. The value of $\dot{\nu}$ doesn't display any unusual behaviour around MJD 55525 when the primary change in pulse profile shape occurs. Figure 4.8 shows that features in the correlation map for PSR J0738-4042 seem to drift in profile phase. This is reminiscent of the behaviour seen in Figure 3.7. The correlation map also shows that as we move in phase across the leading edge of the pulsar, the correlation alternates between positive and negative values. This seems to show that adjacent phases of the pulse profile are somewhat anti-correlated with each other.




Figure 4.5: Pulse profile and spindown variability for PSR J1602-5100. As Figure 4.3.

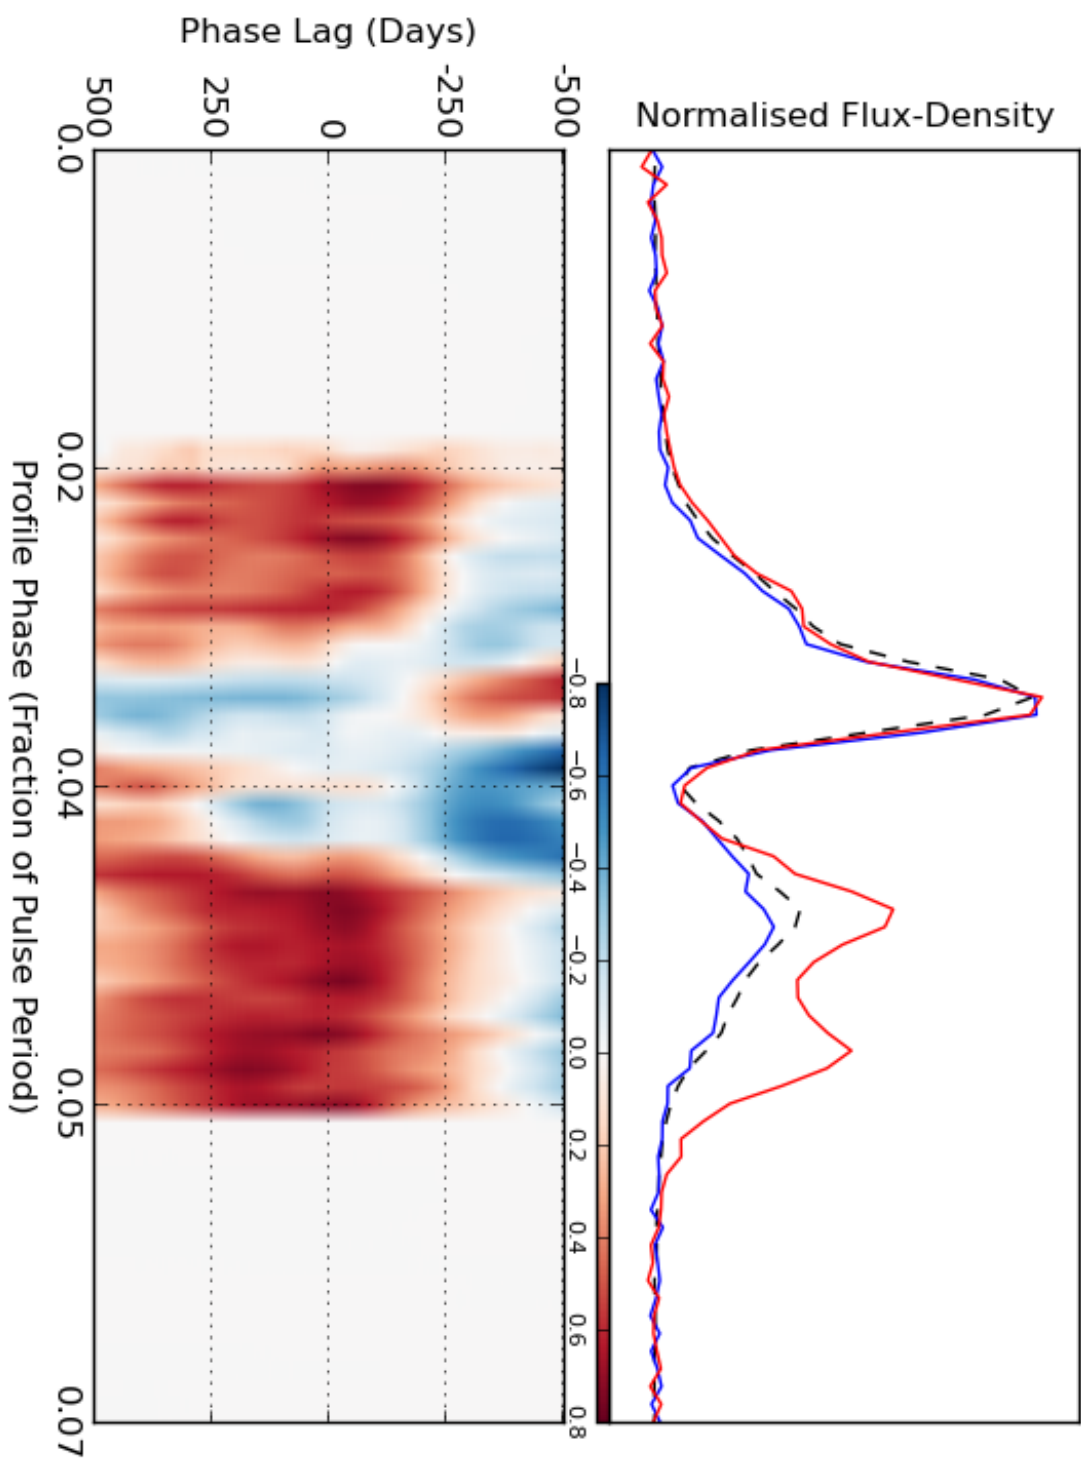


Figure 4.6: Top: Pulse profile variability for PSR J1602-5100. The blue profile was observed on MJD 54306, red on MJD 55072. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.




Figure 4.7: Pulse profile and spindown variability for PSR J0738-4042. As Figure 4.3.




Figure 4.8: Top: Pulse profile variability for PSR J0738-4042. The blue profile was observed on MJD 54548, red on MJD 56002. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

4.4.4 PSR J0742-2822 (B0740-28)

PSR J0742-2822 is known to display profile changes that correlate with $\dot{\nu}$ variations (Keith et al., 2013), as seen in Figure 4.9 and Figure 4.10. It also exhibits the most rapid changes of the six state-switching pulsars analysed in Lyne et al. (2010). The data analysed here show that the flux density across the whole pulse profile varies by $\sim 50\%$ from the median. These variations are shown in the top panel of Figure 4.11. The changes in profile shape are somewhat washed out by this brightness variability, but can be seen more clearly in the normalised observations (middle panel of Figure 4.11); the relative size of two pulse profile components changes quasi-periodically (Figure 4.12). This is most pronounced between MJD 55000 and MJD 55500, and it is in this period when the shape changes can be seen to most closely correlate with $\dot{\nu}$. Keith et al., summarise the profile changes by introducing a shape parameter. They show that the shape parameter correlates particularly well with $\dot{\nu}$ after a glitch (included in the timing model) occurs at MJD 55022.

Two correlation maps are produced for PSR J0742-2822 (Figure 4.12): one using pre-glitch $\dot{\nu}$ data and one using post-glitch $\dot{\nu}$ data. When the pre-glitch data is analysed, $\dot{\nu}$ appears to be anti-correlated with the two profile peaks, and correlated with the central trough. In contrast, when the post-glitch data is analysed, $\dot{\nu}$ appears to be correlated with the peaks, and anti-correlated with the trough. It follows, and can be seen, that certain phases of the pulse profile are correlated and anti-correlation with others.




Figure 4.9: The average pulse profiles of PSR J0742-2822 at 1369 MHz and 3100 MHz. In panels (b) and (d), the black and grey lines denote the two different profile states. The thick lines show total intensity, the thin lines show linear polarisation and the dotted lines show circular polarisation. Panels (a) and (c) show the polarisation position angle as a function of pulse phase. Figure from Keith et al. (2013).




Figure 4.10: The pulse profile shape and spindown variability of PSR J0742-2822. The $\dot{\nu}$ values are plotted in excess of the best-fitting timing model. The circles show a shape parameter for each observation; the parameter is the ratio of the peak height for the two profile components. The dotted line shows the averaged value of the shape parameters in a 60 day sliding window. The vertical dotted line shows where the glitch occurred, after which the correlation between the two timeseries is stronger. Figure from Keith et al. (2013).

4.4.5 PSR J0908-4913 (B0906-49)

PSR J0908-4913 is associated with a pulsar wind nebula (Gaensler et al., 1998). The nebula is considered unusual because it has a low luminosity, steep spectrum and the pulsar is older than any other known to power a pulsar wind nebula. The top panel of Figure 4.13 and Figure 4.14 shows that the emission received from PSR J0908-4913 varies quasi-periodically across the dataset. The pulsar shows interpulse emission and the brightness variations can be seen in both the main pulse and the interpulse. The flux calibrated profiles vary by up to $\sim 50\%$ from the median of around 1200 mJy. The changes in the shape of the pulse profile, however, are slight and gradual across the dataset; a precursor to the main pulse and both interpulse components steadily grow with respect to the main pulse. This trend is drowned out in the flux calibrated observations, but is clearly seen in the middle panel of Figures 4.13 and Figure 4.14. The value of $\dot{\nu}$ shows quasi-periodicity; we see systematic variations which seem to repeat on a timescale of ~ 1000 days, but there is no obvious correlation with the shape changes. The cor-




Figure 4.11: Pulse profile and spindown variability for PSR J0742-2822. As Figure 4.3.

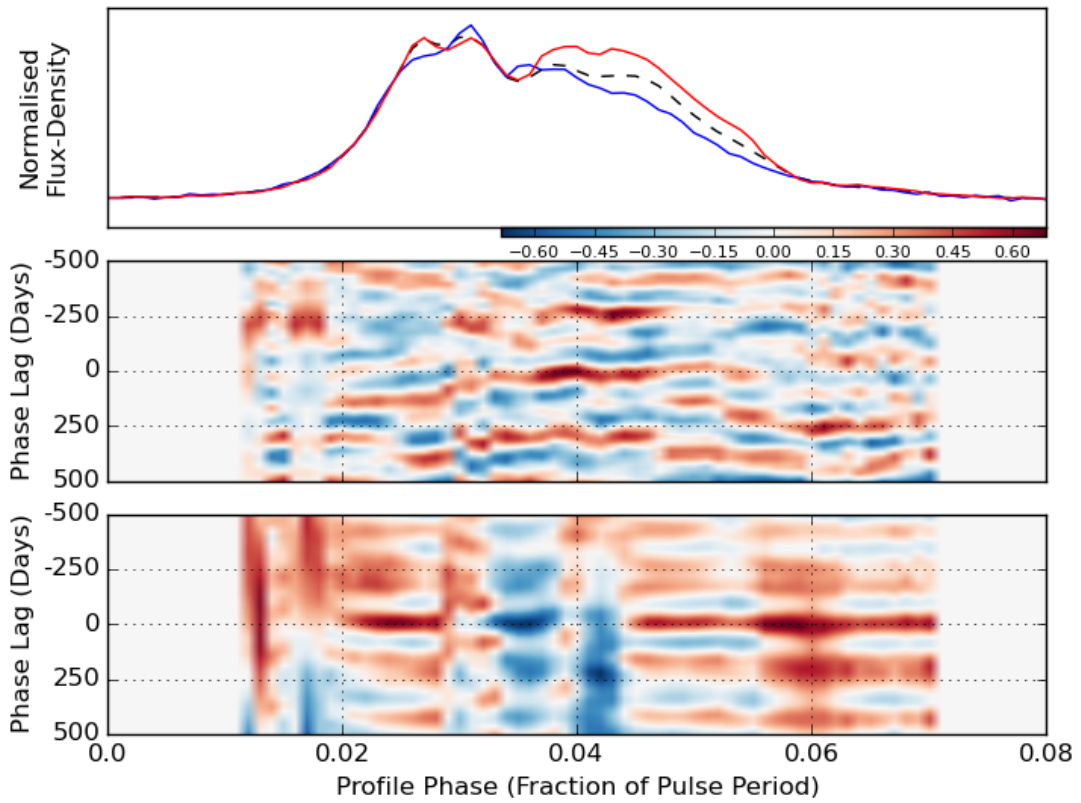


Figure 4.12: Top: Pulse profile variability for PSR J0742-2822. The blue profile was observed on MJD 54881, red on MJD 55358. Middle: Correlation map of $\dot{\nu}$ and flux density, composed of data preceding a glitch on MJD 55022. Bottom: Correlation map composed only of data after the glitch. Panels otherwise as Figure 4.4.

relation maps (Figure 4.15 and Figure 4.16) show no strong relationship between profile and spindown variability in either the main pulse or the interpulse. Only chance correlation in isolated phase bins is seen, rather than coherent correlation features over multiple phase bins.

4.4.6 PSR J0940-5428

PSR J0940-5428 is a faint pulsar that shows only low significance variations in brightness and profile shape (top and middle panels of Figure 4.17 respectively). We feature this pulsar, however, as an example of a pulsar showing $\dot{\nu}$ changes in the absence of significant pulse profile variability; the bottom panel of Figure 4.17, shows systematic changes of $\dot{\nu}$ on a timescale of ~ 200 days. As anticipated, Figure 4.17 shows no correlation features between the pulse profile shape and spindown rate.

4.4.7 PSR J1105-6107

PSR J1105-6107 has a spin period of 63 ms and characteristic age of only 63 kyr. The pulsar has a possible association with a nearby supernova remnant (Kaspi et al., 1997). Throughout our observations, the brightness of this pulsar shows some systematic variation and there are also some changes in profile shape, specifically the relative size of the two profile components (Figure 4.20). One significant shape change that occurs over several observations, beginning \sim MJD 56500 coincides with an increase in $\dot{\nu}$ (Figure 4.19). As seen in Figure 4.20, the strongest correlation between $\dot{\nu}$ and profile shape variability occurs around zero lag in the




Figure 4.13: Pulse profile and spindown variability for the main pulse of PSR J0908-4913. As Figure 4.3. The interpulse variability is shown in Figure 4.14

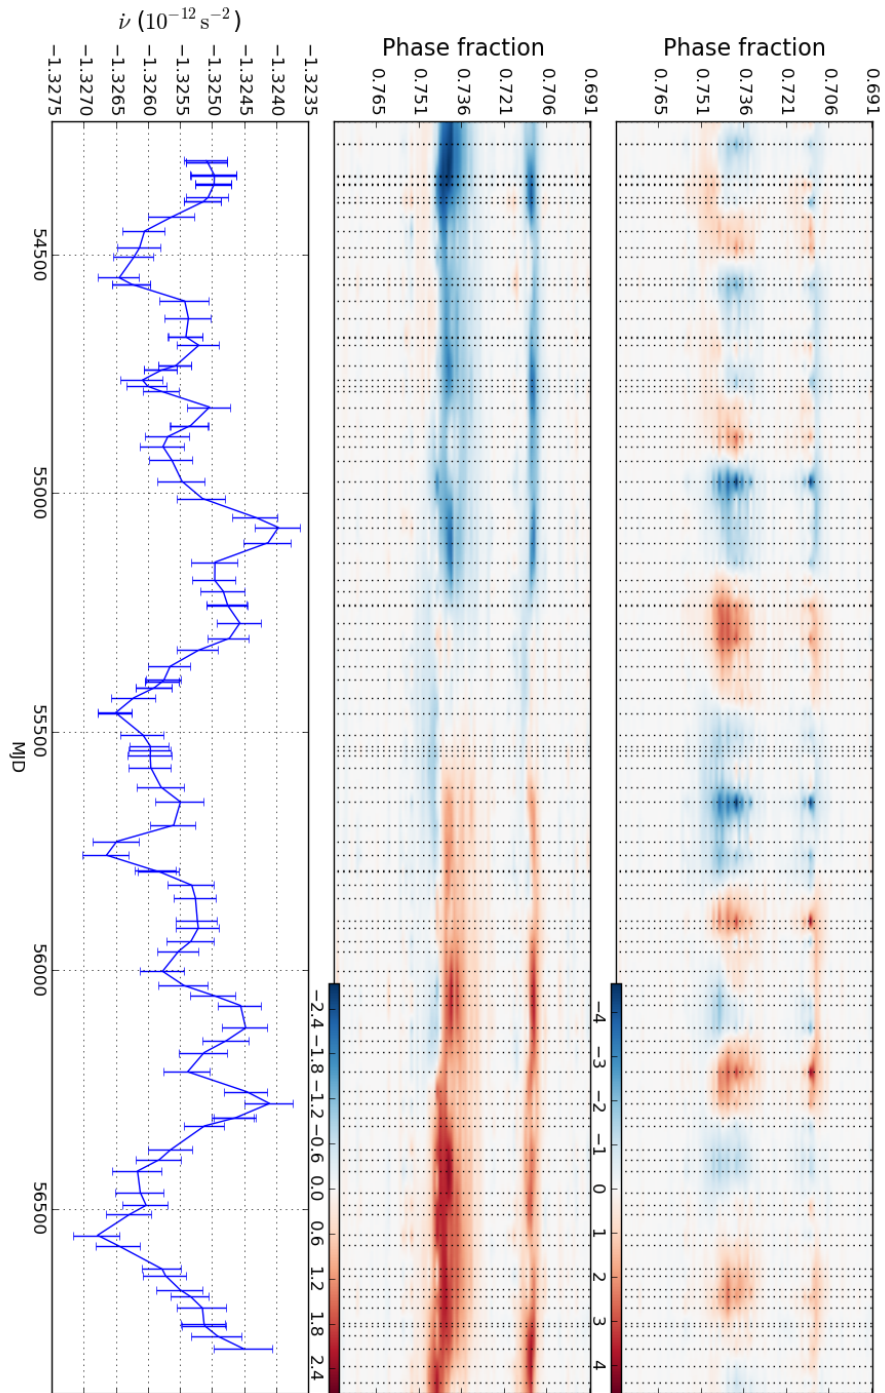


Figure 4.14: Pulse profile and spindown variability for the interpulse of PSR J0908-4913. As Figure 4.3. The main pulse variability is shown in Figure 4.13

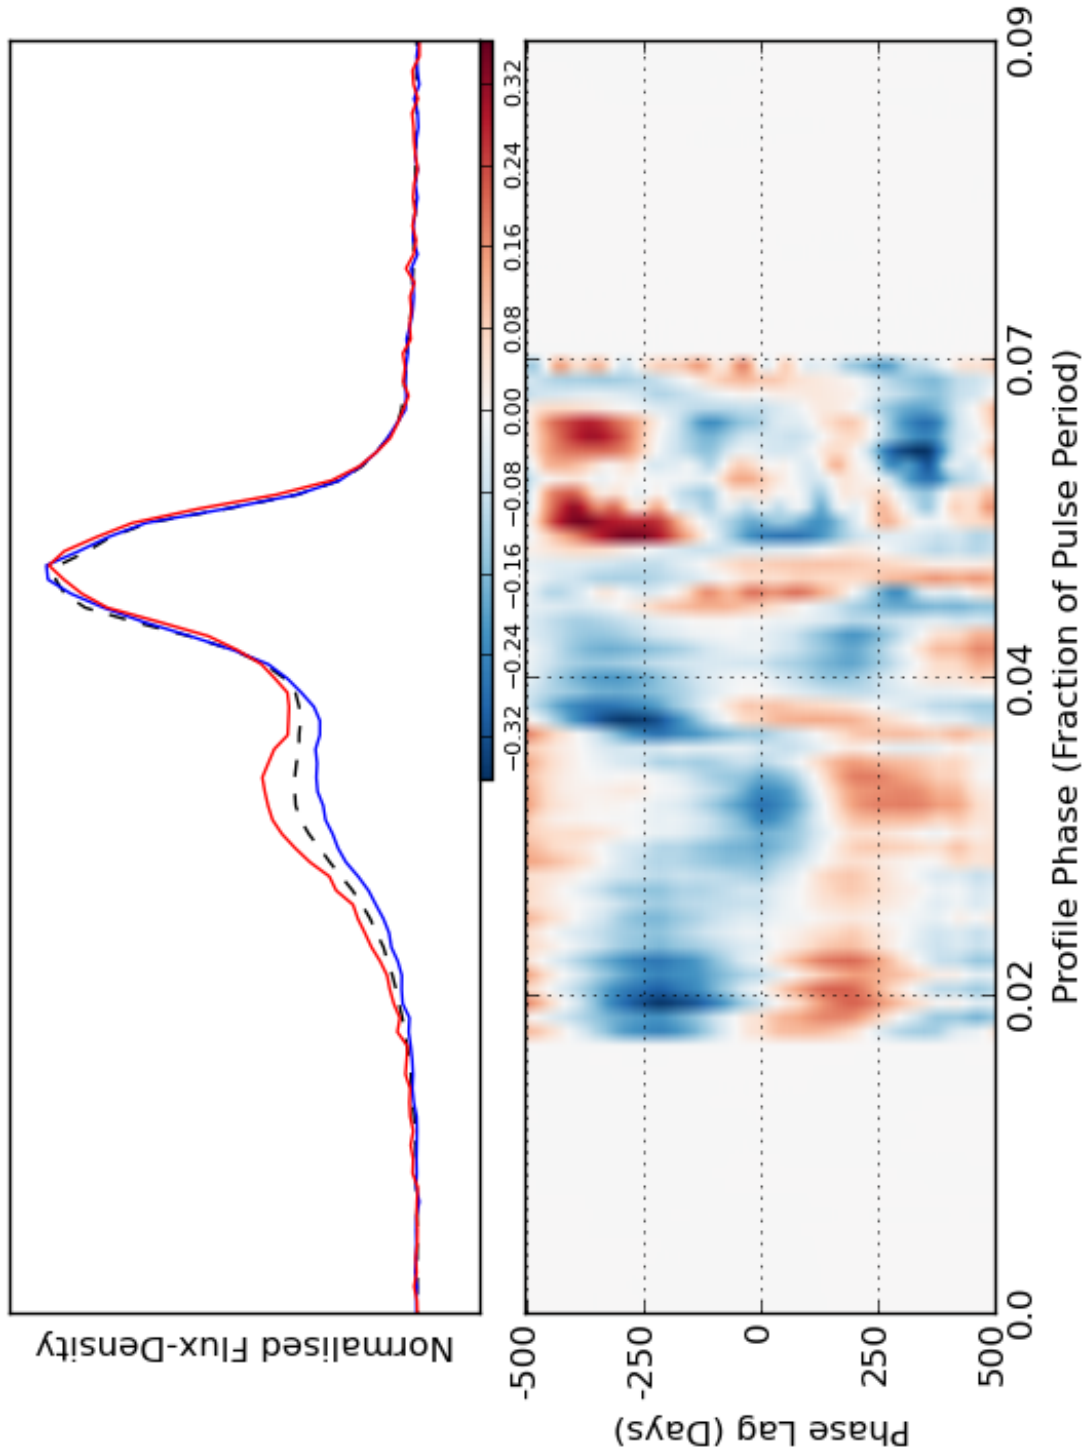


Figure 4.15: Top: Pulse profile variability for the main pulse of PSR J0908-4913. The blue profile was observed on MJD 54220, red on MJD 56746. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

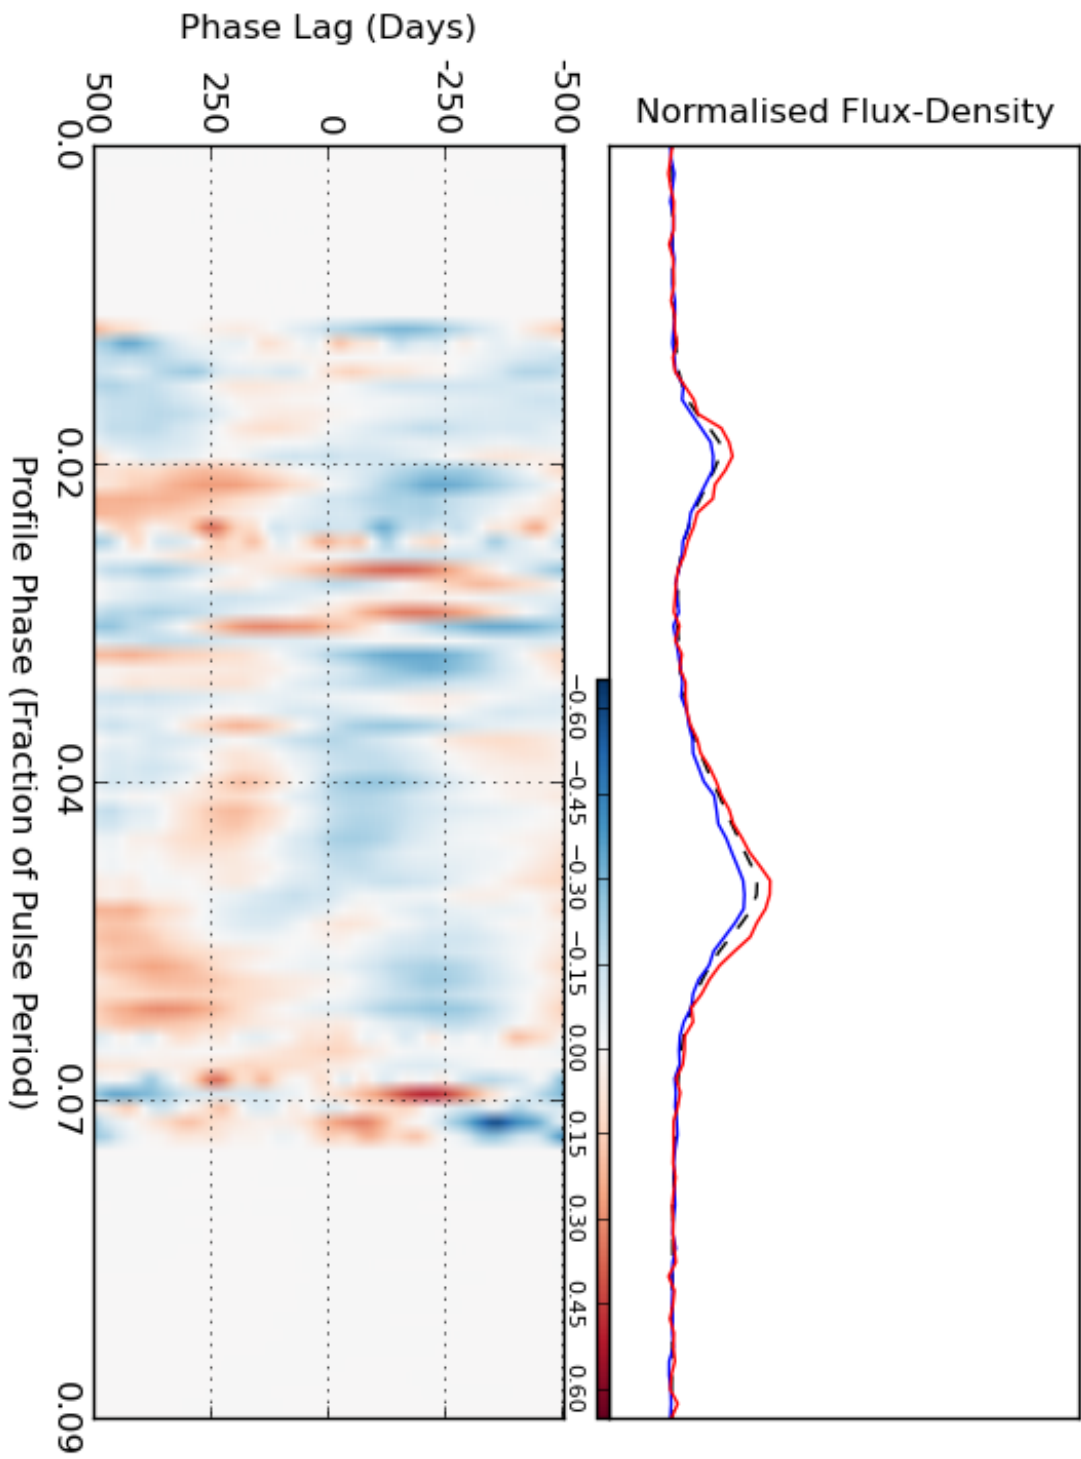


Figure 4.16: Top: Pulse profile variability for the main pulse of PSR J0908-4913. The blue profile was observed on MJD 54220, red on MJD 56746. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

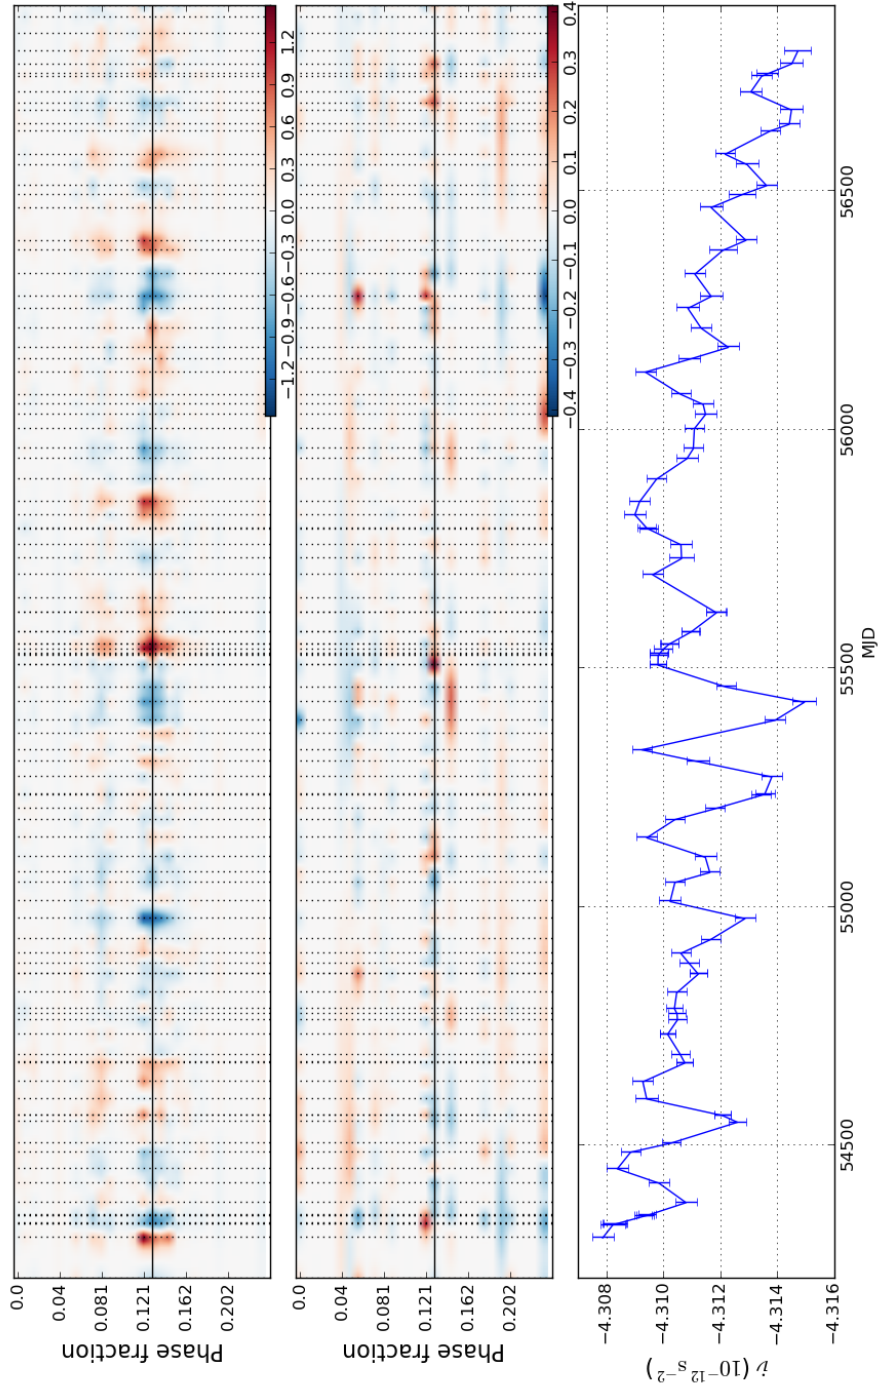


Figure 4.17: Pulse profile and spindown variability for PSR J0940-5428. As Figure 4.3.




Figure 4.18: Top: Pulse profile variability for PSR J0940-5428. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

leftmost of the two components. The peak of the right-hand component is the value to which all observations are normalised, and so only small changes in flux density can occur. A correlation feature also seems to appear towards the trough between the two profile peaks. The correlation occurs when the spindown rate leads the phase bin flux-density. In other words, the flux density changes that take place in the profile trough, follow the behaviour of the rate of spindown (and the peak to its left) around 500 days before.

4.4.8 PSR J1359-6038 (B1356-60)

The brightness of this pulsar systematically varies by $\sim 10\%$ around the median of ~ 550 mJy. The only significant changes in pulse shape occur on three observation days: MJD 56512, 56513 and 56531 (middle panel of Figure 4.21); an example of such a profile is seen in red in Figure 4.22. The corresponding flux calibrated profiles show that the peak flux drops on these observation days. The profile changes coincide with a drop in spindown rate (lower panel of Figure 4.21). A slight correlation feature can be seen at the edges of Figure 4.22. This is predominantly due to the change which occurs in both pulse profile shape and in spindown rate around MJD 56500 as discussed above.

4.4.9 PSR J1600-5044 (B1557-50)

In 2002, a 13-year data set from Hartebeesthoek Radio Astronomy Observatory showed evidence of cyclic variations in both the DM and $\dot{\nu}$ of PSR J1600-5044 (Chukwude, 2002), and the two appear anti-correlated (Figure 4.23). Free pre-




Figure 4.19: Pulse profile and spindown variability for PSR J1105-6107. As Figure 4.3.




Figure 4.20: Top: Pulse profile variability for PSR J1105-6107. The blue profile was observed on MJD 56670, red on MJD 55304. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.




Figure 4.21: Pulse profile and spindown variability for PSR J1359-6038. As Figure 4.3.

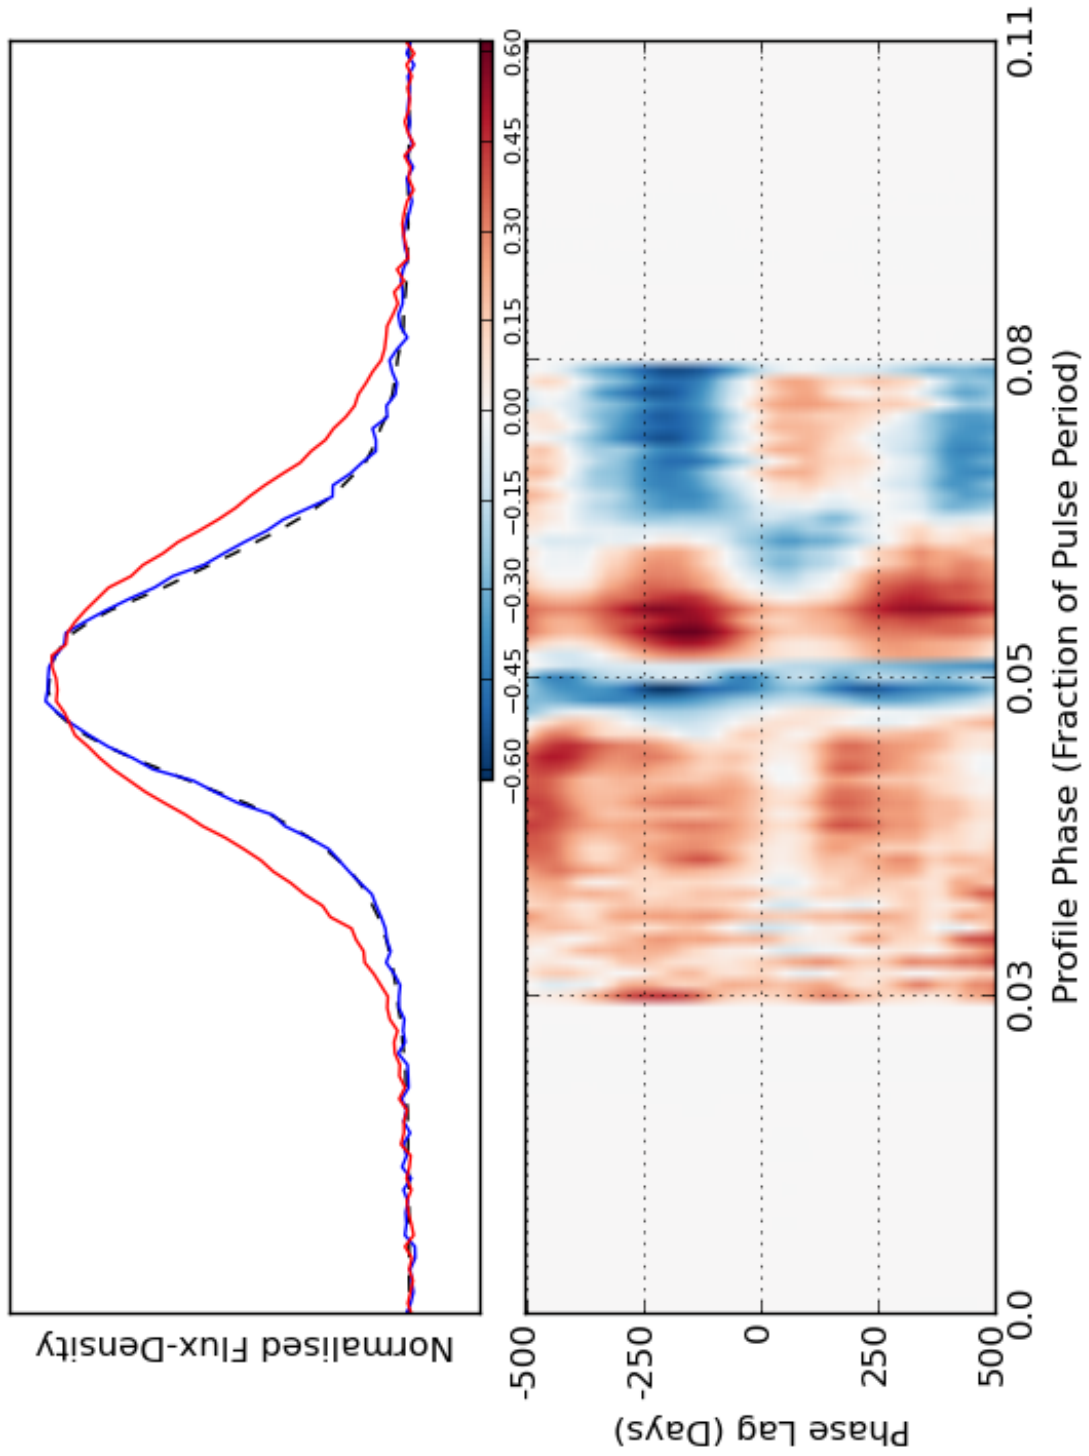


Figure 4.22: Top: Pulse profile variability for PSR J1359-6038. The blue profile was observed on MJD 54786, red on MJD 56531. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

cession of the pulsar is put forward as the most plausible explanation for the signature of the timing residuals. It does not, however, easily explain the DM variations seen. In the Parkes dataset, flux calibrated observations of PSR J1600-5044 show significant variability; the peak of the median profile is ~ 800 mJy, which increases systematically to ~ 1500 mJy (MJD 54920) before dropping back to the level of to the median value. The process occurs over ~ 120 days (top panel of Figure 4.24). When these brightness variations due to refractive scintillation are eliminated by normalising the observations, no significant variability is seen in pulse shape or in $\dot{\nu}$ (middle and bottom panel of Figure 4.24 respectively). Low significance correlation is indicated in the trailing edge of the pulse profile (Figure 4.25).

4.5 Discussion

We have analysed 168 pulsar datasets, spanning up to eight years in length and presented details results from nine. Below are the main conclusions:

Variability in the dataset

When analysing the profile variability of 168 pulsars observed for the Fermi Timing Programme, we see only 11 pulsars showing significant profile shape changes. PSR B1822-09 and PSR B1259-63 were also seen to have emission variability, but were not discussed, as they are well observed examples of variable pulsars (Lyne et al., 2010; Johnston et al., 1994). Due to the level of variability in individual phase bins, the accuracy with which we can measure changes in pulse profile is much lower than the accuracy of our $\dot{\nu}$ measurements. Subtle profile changes, therefore, are easily hidden in pulsars with low S/N; we expect more profile vari-




Figure 4.23: Long-term behaviour of J1600-5044. Panels (a) to (d) shows variations in the timing residuals, dispersion measure, spin frequency and spindown rate respectively. Figure from Chukwude (2002)




Figure 4.24: Pulse profile and spindown variability for PSR J1600-5044. As Figure 4.3.




Figure 4.25: Top: Pulse profile variability for PSR J1600-5044. Bottom: Correlation map of ν and flux density. Both panels otherwise as Figure 4.4.

ability to be seen with longer integration times and more sensitive instruments. As discussed in Section 4.2, the spindown rates of the pulsars often showed cyclic structures which are indicative of imprecise positional parameters in the pulsar timing model and, therefore, it is difficult to say how much spindown variability there is in the Fermi Timing Programme pulsars without further analysis.

New technique for measuring spindown rate

We have developed a new technique for measuring $\dot{\nu}$ by modelling the second derivative of the timing residuals using GP regression. This method is analytical, in contrast to the numerical methods often employed to determine $\dot{\nu}$.

Known variability reproduced

Using the above technique to measure $\dot{\nu}$, along with our method to map pulse profile shape changes we have been able to reproduce the variability already observed in PSR J0742-2422 (Lyne et al., 2010; Keith et al., 2013), PSR J1830-1059 (Keith et al., 2013), PSR B1822-09 (Lyne et al., 2010) and PSR B1259-63 (Johnston et al., 1994).

Discovery of strongly correlated variability in PSR J1602-5100

We have uncovered a new example of correlated $\dot{\nu}$ and profile shape variability in PSR J1602-5100. This pulsar exhibits a dramatic change in profile shape over ~ 600 days, with a simultaneous reduction in $\dot{\nu}$. Such sudden and dramatic changes have previously been attributed to exterior material entering the pulsar magnetosphere (Cordes & Shannon, 2008; Shannon et al., 2013; Brook et al., 2014)), as discussed in Section 3.4. The reconfiguration in current within the magnetosphere,

induced by the introduction of external material would simultaneously affect the braking torque and hence $\dot{\nu}$. The value of $\dot{\nu}$ of PSR J1602-5100 is observed to reduce with the appearance of a new profile component. This is comparable to the 2005 event seen in PSR J0738-4042, hypothesised to be caused an asteroid encounter (Brook et al., 2014). The changes observed in PSR J0738-4042, have persisted ever since, whereas PSR J1602-5100 returned to its previous state after ~ 600 days. The change in $\dot{\nu}$ over this period can be approximated as a step-function jumping between two rates, and interpreted as a reduction in the total outflowing plasma above the polar caps. The magnitude of the current change, can be inferred from the change in $\dot{\nu}$ following Kramer et al. (2006). Following the method detailed in Section 3.4, the difference between the pre- and post-step $\dot{\nu}$ values corresponds to a reduction in the charge density ρ of $\sim 9 \times 10^{-9} \text{ C cm}^{-3}$. We can relate the difference in charge density associated with the two $\dot{\nu}$ states to mass supplied to the pulsar, by multiplying it by the speed of light, the polar cap area and the duration of the new spindown state. Over 600 days, this amounts to $\sim 10^{14}$ g, which lies within the range of known solar system asteroid masses, and is consistent with the mass range of asteroids around neutron stars proposed by Cordes & Shannon (2008).

As TOA measurements are based on template matching (Section 1.2.2), we must ensure that the changing pulse profiles themselves are not responsible for inducing the perceived $\dot{\nu}$ variations that accompany them. Figure 4.26 shows a simulation of pulse profile variations, and how they affect the spindown rate of a pulsar. The simulated pulsar profile evolves in a similar manner to that of PSR J1602-5100; it initially consists of one Gaussian component and later develops a transient Gaussian feature on its trailing edge which lasts for around 600 days. Noise is added to the simulated profiles so that their S/N mimics that of the PSR J1602-5100

observations. To test how the changing observation shape affects the perceived observed spindown rate of the pulsar, we simulate a template matching process. Each observation is correlated with the template. The template is the average of all pulse profile in the dataset. As the shape of the profile changes, so too does their alignment with respect to the template. A shift in alignment can be translated into a fraction of a pulse period (equal to the period of PSR J1602-5100 in this case) and into a timing residual for each observation. The spindown rate is then numerically calculated from the timing residuals.

Figure 4.26 shows that the change in spindown rate due to the template matching process is only around 0.1%, whereas the change seen in PSR J1602-5100 is more than 6%. Additionally, a trailing edge component induces an increase in the perceived spindown rate of the pulsar, in contrast to the decrease seen in PSR J1602-5100. The changing profile shape of PSR J1602-5100 is not responsible for the large, simultaneous change in spindown rate that it experiences.

Two timescales needed to model most timing noise

When fitting a non-parametric function to the timing residuals, all but one of the nine pulsars featured in this paper were best fit by a covariance function which contained two squared exponential kernels plus a noise model. One of the kernels had a lengthscale bounded between 30 and 100 days, and one between 100 and 1000 days, resulting in a function that varies on a short and on a long timescale. As this provides the best fit to the data, this suggests that at least two processes, operating on different timescales, are responsible for the timing noise that we observe.




Figure 4.26: The results of a template matching simulation. Top panel: The simulated pulse profile as a function of time. Middle panel: The profile residuals after the subtraction of the average pulse profile. Bottom panel: How the spindown rate is affected solely due to a mismatch between the average profile shape, and the profile shape of individual observations.

Correlation between pulse profile and $\dot{\nu}$ variations

Of the seven notable examples of pulse profile variability featured in this chapter, five also show some degree of $\dot{\nu}$ variation which could be considered correlated. The profile changes seen, occur to various degrees and on various timescales. It is particularly interesting that no change in $\dot{\nu}$ is seen to accompany the dramatic profile shape change seen in PSR J0738-4042. The present interpretation of the correlated long-term variations observed in emission and rotation, involve changing charged particle currents in the pulsar magnetosphere (Kramer et al., 2006). Distinct pulsar states can be explained by differing levels of plasma present in the magnetosphere. Changing plasma levels are expected to modify both the material outflow along open field lines at the polar cap, and the subsequent emission produced. Plasma variations would also vary the braking torque on the pulsar, and we would expect to see a change in $\dot{\nu}$ accompanying any significant change in emission. Correlated pulse flux density and rotation variability, therefore, hints at a process intrinsic to the pulsar. As well as changes in the shape of the pulse profile, the flux calibrated observations show that large variations in pulse profile flux density, due to refractive scintillation, are ubiquitous. The pulse profile varies as a whole, and its shape is maintained. For some pulsars, these brightness variations are coherent over multiple observations, i.e., on a timescale of hundreds of days. As these brightness variations are primarily due to effects of propagation, we expect and find that pulse profile flux density and $\dot{\nu}$ are not well correlated.

Caveats

When analysing the rotational variations, we have made the hypothesis that timing noise is due to changes in the braking torque on the pulsar, observable as

changes in $\dot{\nu}$. Other possible sources of timing noise are inadequate calibration of the raw observations (e.g. van Straten, 2006), and failure to correct for variations in the interstellar dispersion (e.g. You et al., 2007)).

When testing the correlation between the variability of flux density in individual (or small groups of) pulse phase bins and the spindown rate, the latter is measured much more accurately than the former. This is due to the levels of S/N in our observations. We expect this to be reason why so few of the pulsars in this chapter display profile variability, and potentially why we don't see stronger correlation with $\dot{\nu}$, in those that do.

It is worth noting that the GP regression technique is predominantly sensitive to systematic trends. Therefore, when it is employed to model the flux density variability seen in each pulse phase bin, any profile features that occur in single observations only, will have little effect on the model and, therefore, on the emission variability map overall. This is desirable if the single observation has produced a spurious pulse profile due to instrumental issues, but conversely, any genuine profile deviations that occur in single observations may not feature in the final emission variability map.

Finally, as may be the case for PSR J0742-2822 (Section 4.4.4), if correlation between profile and spindown changes occurs only for fraction of the time spanned by the dataset, the correlation maps may not clearly show any correlation.

Chapter 5

The emission variability of NANOGrav millisecond pulsars

The stability of millisecond pulsars is not well studied, although we expect them to be more stable than ordinary pulsars (Section 1.1.3). This is because millisecond pulsars are spinning faster, and losing energy at a lower rate due to their weaker magnetic fields. In this chapter, we analyse the pulse profile stability of pulsars that are used as part of a pulsar timing array by the NANOGrav collaboration. The results are divided between the old and new generation of backend instruments used by NANOGrav.

5.1 The search for gravitational waves

5.1.1 Gravitational waves

Gravitational waves (or gravitational radiation) are ripples in spacetime, predicted by general relativity and created by accelerating masses in certain circumstances. Currently, our information about the distant universe has been delivered via elec-

tromagnetic radiation and sub-atomic particles. A direct detection of gravitational waves will open an entirely new window on our Universe and will enable us to learn about objects that do not emit electromagnetic radiation, or those which are hard to detect in the electromagnetic regime. Additionally, gravitation waves can carry information to us from before the epoch of reionisation.

It is anticipated that gravitational waves will allow us to learn about how the most massive black holes in the Universe form, how galaxies merge and grow throughout cosmic history, and how gravity behaves at the limit of our understanding. It is also hoped that this new era of astronomy will present new discoveries not yet envisaged.

Gravitational radiation is expected to be very weak because large masses or high acceleration are required to produce significant gravitational signals. Because of this, no direct detection of gravitational radiation signal has been made to date, despite the efforts of teams using ground-based and space-based detectors. Indirect detection, however, was obtained by measuring the orbital period decay in the PSR B1913+16 binary pulsar system (Taylor & Weisberg, 1982; Hulse & Taylor, 1975). Despite the technical difficulties associated with uncovering gravitational wave signals, there are two fields of astrophysics that hope to make a direct detection by the end of the decade. One is the development of second generation km-scale laser interferometers, such as Advanced LIGO (Harry & LIGO Scientific Collaboration, 2010), Advanced Virgo (Acernese et al., 2015) and KAGRA Somiya (2012). These ground-based detectors work at hundreds of hertz and will detect gravitation radiation from inspiralling binary star systems, especially those containing neutron stars or stellar mass black holes. The second field hoping to make an imminent discovery is that of pulsar timing arrays.




Figure 5.1: The orbital decay of binary pulsar system PSR B1913+16. The data points track the change in time of periastron, while the solid lines shows general relativity's predict for a system emitting gravitation radiation. Figure from Taylor & Weisberg (1982).




Figure 5.2: The Hellings and Downs plot. The solid line shows the expected correlation in the timing residuals of pairs of pulsars as a function of their angular separation. The points are simulated pulsar datasets. Figure from Hobbs et al. (2009).

5.1.2 Pulsar timing arrays

The passage of gravitational waves between us and the pulsars we observe, will leave an imprint on the timing residuals. In 1983, Hellings and Downs showed that a signal of gravitational radiation will produce correlated timing fluctuations between pairs of pulsars, where the degree of correlation is a known function of the angular separation (Figure 5.2). As discussed in Section 1.1.3, millisecond pulsars are incredibly stable rotators, and consequently we can predict their TOAs with microsecond precision. They are, therefore, prime candidates for precision timing experiments.

It has been shown that approximately 20 pulsars with timing residuals of around 100 nanoseconds or better, over five years of weekly observations, are needed to make a significant detection of a background signal of gravitational radiation (Sesana et al., 2008)

Pulsar timing arrays are sensitive to gravitational radiation around the nanohertz




Figure 5.3: Characteristic strain sensitivity for current and future gravitation wave detectors. Also shown are predicted backgrounds from various astrophysical phenomena. The coloured dashed lines are expected backgrounds from different cosmological models. Black dotted lines are different levels of gravitational wave energy density content. Figure from Janssen et al. (2015).

frequency regime (Figure 5.3). The strongest expected source in this band is the background signal from supermassive black hole binaries, which result from galaxy mergers (Sathyaprakash & Schutz, 2009). There are currently three pulsar timing arrays: European Pulsar Timing Array (EPTA), which combines data from the Effelsberg, Jodrell Bank, Nancay and Westerbork telescopes; the Parkes Pulsar Timing Array (PPTA) which uses data from the Parkes Radio Telescope (Section 2.1.1); and the North American Nanohertz Observatory for Gravitational Waves (NANOGrav), which combines data from the Arecibo Observatory and the Green Bank Telescope. Since 2008, the three pulsar timing arrays have been collaborating as members of the International Pulsar Timing Array (Manchester

& IPTA, 2013).

5.1.3 NANOGrav

NANOGrav was founded in October 2007 and has since grown to over 60 members at over a dozen institutions. The group currently conducts high-precision timing observations of 36 millisecond pulsars every two weeks with the William E. Gordon Telescope at the Arecibo Observatory and the Robert C. Byrd Green Bank Telescope (Section 2.1.2).

5.2 Results from ASP and GASP backends

All plots in the remainder of this chapter are either variability maps or individual pulse profile plots. The former are as described in Section 2.2.4. Regarding the latter, any red profiles show the average of the whole pulsar dataset, while blue profiles show individual observations.

5.2.1 Instrumental issues

Some observations result in pulse profiles that show substantial deviation from their expected shape (e.g. Figure 5.4). Very dramatic changes that occur across the whole profile and only appear in isolated observations only, are immediately suspected to be the result of instrumental issues rather than a change intrinsic to the pulsar. This suspicion can be vindicated if the MJD on which deviant profiles were observed is noted on multiple occasions for different pulsars. In the NANOGrav data, problematic observation days seem to be preferentially close to the beginning of a dataset, indicating that the instrumental set up was not functioning optimally at the early stages.

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRAV
MILLISECOND PULSARS

Pulsar Name	Period (ms)	DM (pc cm ⁻³)	Observatory	Frequencies (MHz)
J0023+0923	3.05	14.3	AO	430/1410
J0030+0451	4.87	4.3	AO	430/1410
J0340+4130	3.30	49.6	GBT	820/1500
J0613-0200	3.06	38.8	GBT	820/1500
J0645+5158	8.85	18.2	GBT	820/1500
J0931-1902	4.64	41.5	GBT	820/1500
J1012+5307	5.26	9.0	GBT	820/1500
J1024-0719	5.16	6.5	GBT	820/1500
J1455-3330	7.99	13.6	GBT	820/1500
J1600-3053	3.60	52.3	GBT	820/1500
J1614-2230	3.15	34.5	GBT	820/1500
J1640+2224	4.62	62.4	AO	430/1410
J1643-1224	4.62	62.4	GBT	820/1500
J1713+0747	4.57	16.0	AO	1410/2030
J1738+0333	5.85	33.8	AO	1410/2030
J1741+1351	3.75	24.0	AO	430/1410
J1744-1134	4.07	3.1	GBT	820/1500
J1747-4036	1.65	152.9	GBT	820/1500
J1853+1303	4.09	30.6	AO	430/1410
B1855+09	5.36	13.3	AO	430/1410
J1903+0327	2.15	297.5	AO	1410/2030
J1909-3744	2.95	10.4	GBT	820/1500
J1910+1256	4.98	38.1	AO	1410/2030
J1918-0642	7.65	26.6	GBT	820/1500
J1923+2515	3.78	18.9	AO	430/1410
B1937+21	1.56	71.0	AO	1410/2030
J1944+0907	5.19	24.3	AO	430/1410
J1949+3106	13.14	164.1	AO	1410/2030
B1953+29	6.13	104.5	AO	430/1410
J2010-1323	5.22	22.2	GBT	820/1500
J2017+0603	2.90	23.9	AO	1410/2030
J2043+1711	2.38	20.7	AO	430/1410
J2145-0750	16.05	9.0	GBT	820/1500
J2214+3000	3.12	22.6	AO	1410/2030
J2302+4442	5.19	13.8	GBT	820/1500
J2317+1439	3.45	21.9	AO	430/1410

Table 5.1: List of NANOGrav millisecond pulsars, their rotational period, dispersion measure, telescope at which the pulsar is timed and the centre observational frequency. AO is the Arecibo Observatory and GBT is the Green Bank Telescope.




Figure 5.4: An observation of PSR J2317+1439 made at Arecibo on MJD 54629 at 400 MHz. The pulse profile deviates markedly from its expected shape, almost certainly due to instrumental issues. Red trace shows the median profile of the dataset and blue trace shows the individual observation.

Finding certain days on which deviant profiles of various pulsars are observed is one way to identify the cause as instrumental. Another way is to look at the pulse profiles produced by each individual frequency channel of the observation band. If a deviant profile feature does not appear in all channels, then we can assume that the instrumental problem exist within specific frequency channels only. Many examples of this were found in deviant profiles that were observed at Arecibo and processed by ASP. An example, showing an instrumental deviation only in some frequency channels for PSR J1713+0747 is given in figures 5.5, 5.6 and 5.7.

In this and other cases, the frequency channels are often divided; for a profile containing a deviant feature, half of the contributing channels show it and the other half do not (e.g. figures 5.6 and 5.7). The reason for the split is that in the ASP backend, groups of up to 8 channels were distributed to the cluster for processing via different hardware interfaces through different data server computers. The




Figure 5.5: Pulse profile of PSR J1713+0747 observed at Arecibo at an observation frequency of 1400 MHz. The red profile is the average of all observations in the dataset. The blue profile is the observation on MJD 55108. This profile shows an amplified component on the leading edge of the pulse.

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRV MILLISECOND PULSARS

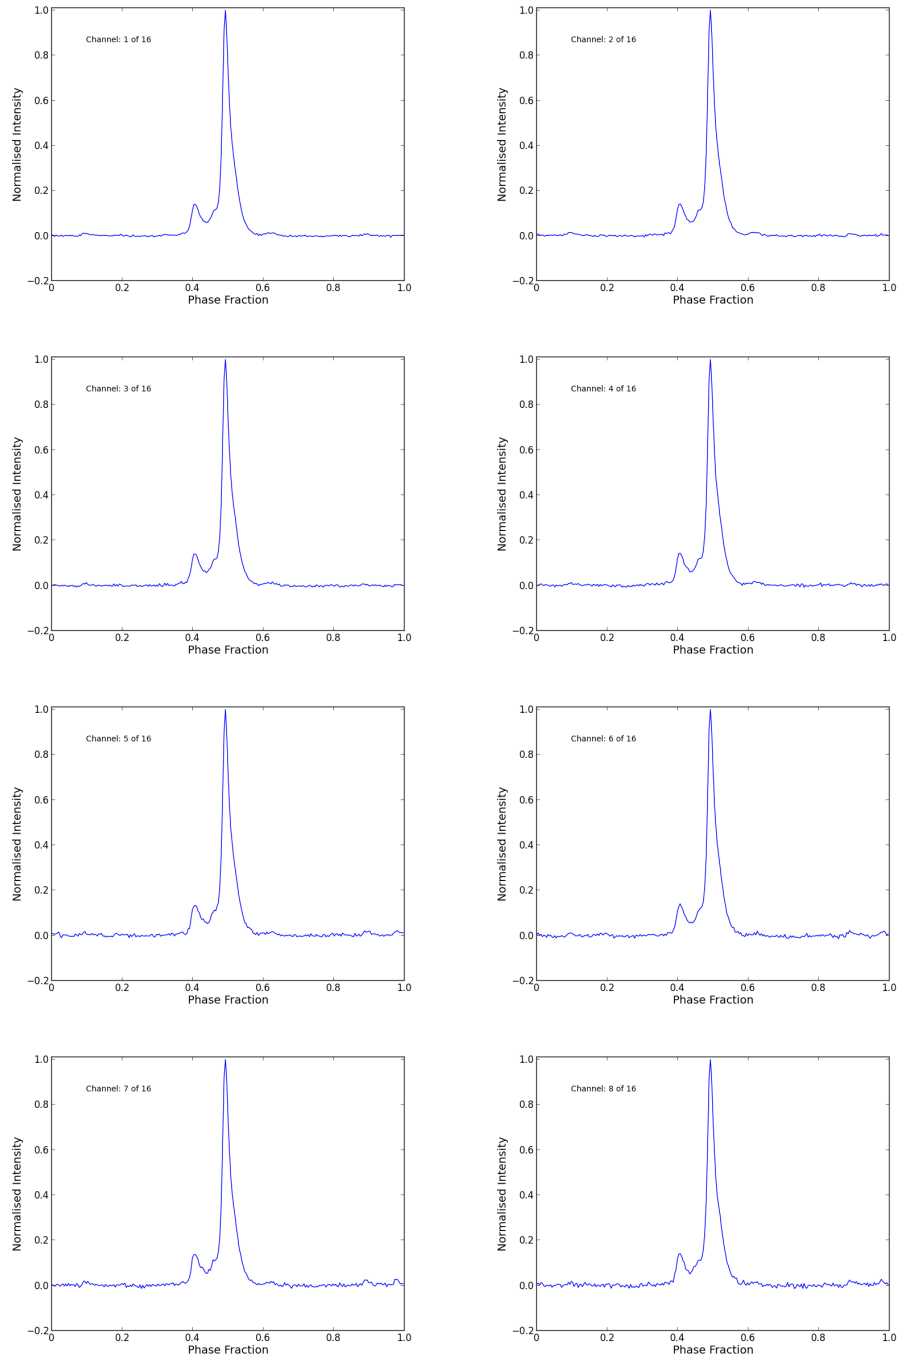


Figure 5.6: Individual frequency channels 1-8 of the PSR J1713+0747 profile seen in Figure 5.5. The deviation can be seen at the leading edge of the profile (at around phase fraction 0.4). It is present in the first 8 frequency channels but not the channels 9-16 (see Figure 5.7).

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRV MILLISECOND PULSARS

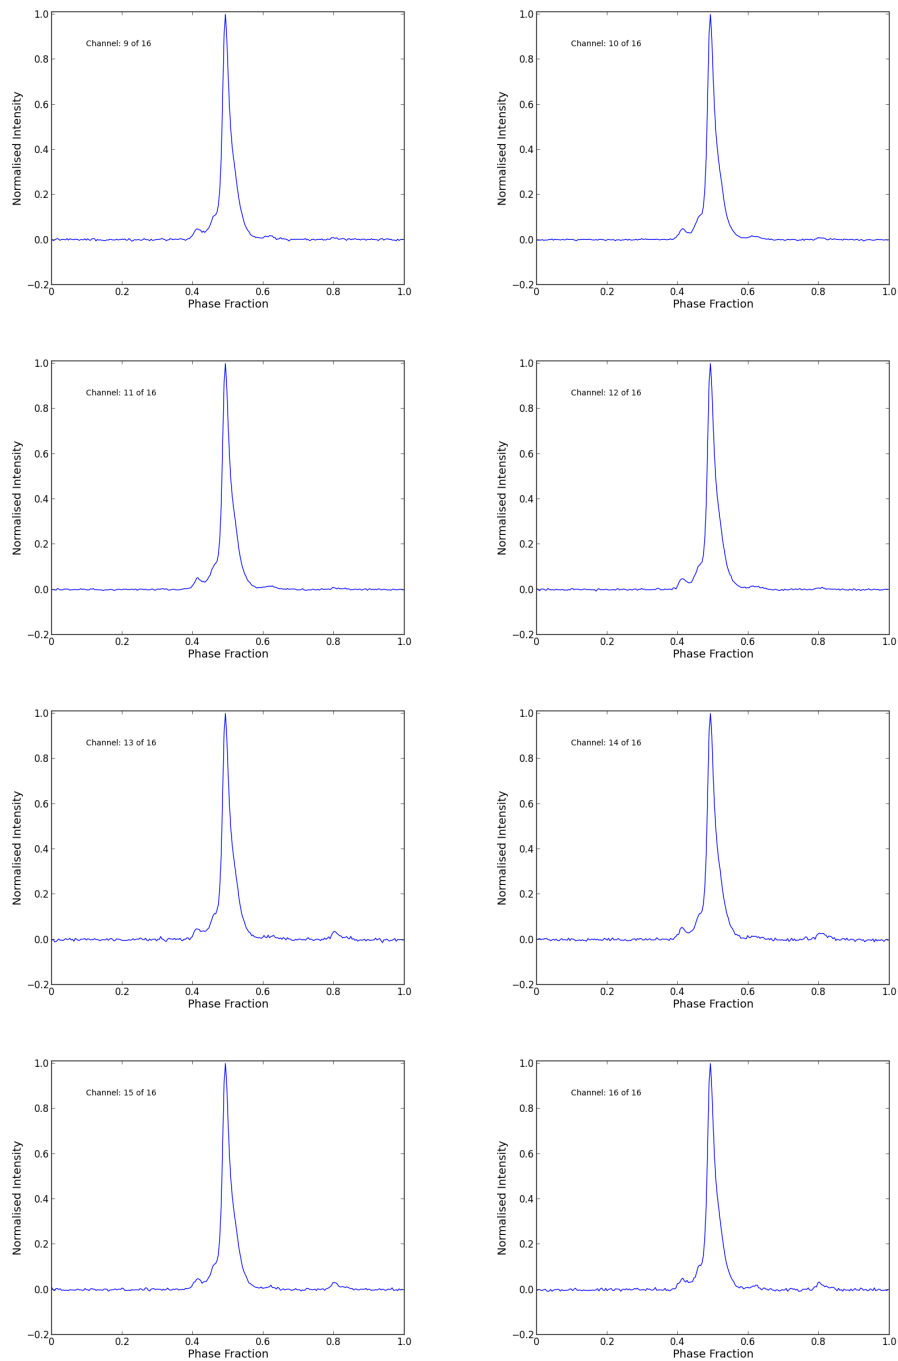


Figure 5.7: Individual frequency channels 9-16 of the PSR J1713+0747 profile seen in Figure 5.5. The deviation at the leading edge of the profile is absent. A smaller deviation can also be seen in the trailing edge (around phase 0.8) of the last four channels.

most common setup was 16 channels with each group of 8 coming from a different data server. If one of the two data servers had a problem during the observation this could lead to the situation where 8 of the channels would be bad. Although they are not the result of astrophysical effects, finding these deviant profiles is important, as their inclusion in a dataset is detrimental to the precision timing of the pulsar. As discussed in Section 1.2.2, a pulsar's TOAs are determined via template matching, and therefore a misshapen observation will provide imprecise pulse TOAs that will contribute to a reduction in the sensitivity of the pulsar timing array.

5.2.2 Astrophysical profile changes

PSR B1937+21

PSR B1937+21 was the first millisecond pulsar to be discovered, has a rotational frequency of 642 Hz and has an interpulse (see *beam geometry* in Section 1.1.2). The bright profile of PSR B1937+21 is seen to deviate at multiple observation frequencies and at both telescopes. The brightness of this pulsar can be seen to vary in the flux calibrated variability maps, which is an expected result of refractive scintillation (Section 1.2.4). When the pulse profiles are normalised to the peak, profile shape variations are additionally present. In particular, the flux density in the interpulse relative to the main peak is seen to fluctuate. Long-term trends can be seen in the Figure 5.9.

PSR J1853+1303

PSR J1853+1303 is a millisecond pulsar in a wide-orbit binary system (Faulkner

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRAV MILLISECOND PULSARS

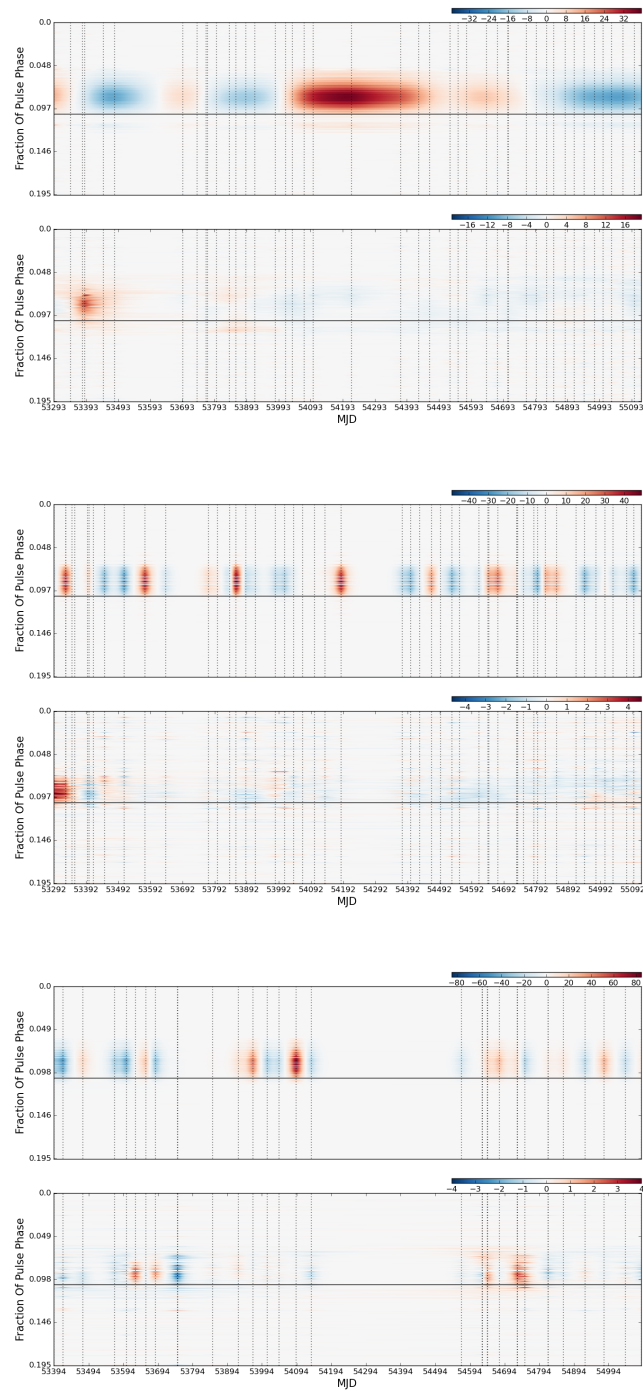


Figure 5.8: Variability maps for the main pulse of PSR B1937+21. Top pair: 800 MHz observations made with the Green Bank Telescope. Middle pair: 1400 MHz observations made with the Green Bank Telescope. Bottom pair: 1400 MHz observations made at Arecibo. Note that the starting date for this plot is slightly different to the others. Long-term changes in brightness can be seen in the flux calibrated Green Bank Telescope observations at 800 MHz only. Other variability occurs on short timescales.

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRV MILLISECOND PULSARS

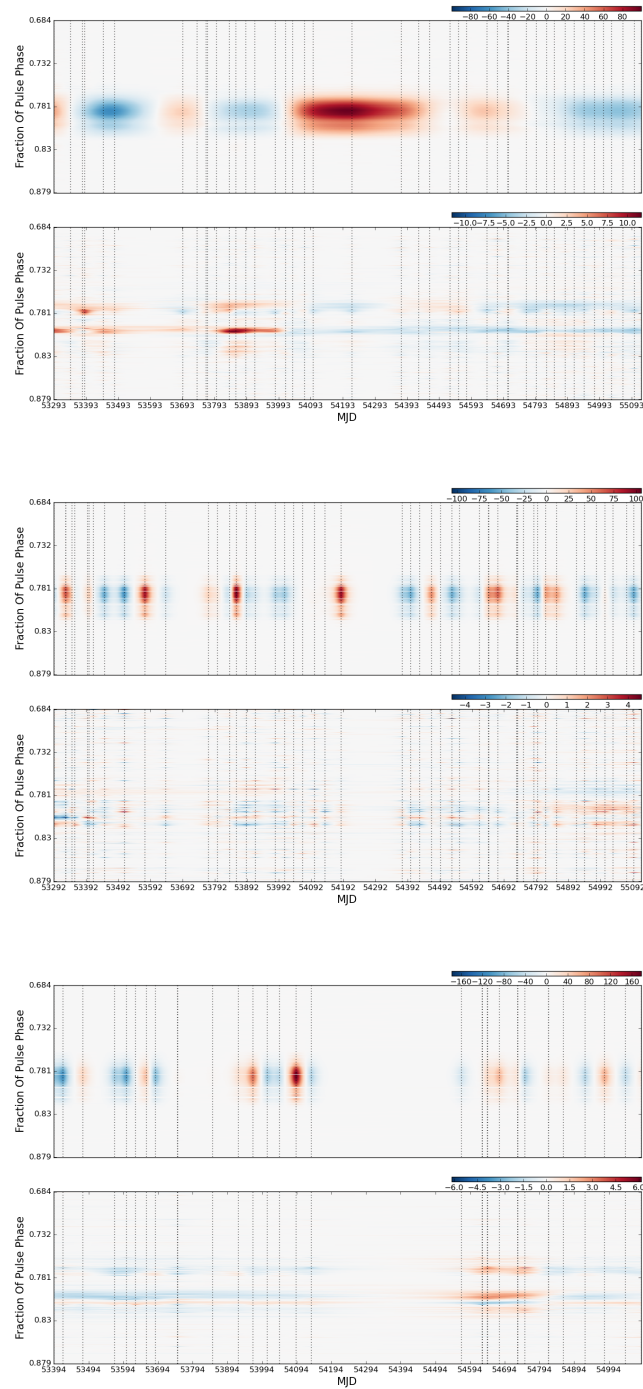


Figure 5.9: Variability maps for the interpulse pulse of PSR B1937+21, as Figure 5.8. Long-term brightness variability is again seen in the flux calibrated Green Bank Telescope observations at 800 MHz. Strong long-term profile shape variability are also clearly seen in the top and bottom panels, but appear anti-correlated.




Figure 5.10: Pulse profiles from PSR B1937+21 from MJD 53629. Top panels: Main pulse and interpulse as observed at 1400 MHz at Arecibo. Bottom panels: Observations at Arecibo at 2300 MHz. Both observations show an increase in interpulse flux density with respect to the main peak.




Figure 5.11: Variability maps of PSR J1853+1303, made with the Green Bank Telescope at 1400 MHz. The individual profile shape deviations are seen in the three observations shown in Figure 5.12

et al., 2004). As can be seen from the lower panel variability map of Figure 5.11, there are a number of profiles around MJD 53600 in which the flux density drops relative to the profile peak. The upper panel shows that the absolute flux of the pulsar around this time is seen to increase. The three profiles showing the change are shown in Figure 5.12.

PSR J1910+1256

Like PSR J1853+1303, PSR J1910+1256 is a wide-orbit binary millisecond pulsar, and was also discovered at the same time (Faulkner et al., 2004). More similarities can be seen in their observations. In the upper panel of Figure 5.13 a region of increased brightness can be seen around MJD 53650, while the lower panel shows a shape change in the corresponding profiles; the pulse narrows slightly as can be




Figure 5.12: Observations of PSR J1853+1303, made with the Green Bank Telescope at 1400 MHz. From left to right, the observation days are MJD 53601, MJD 53661 and MJD 53687.




Figure 5.13: Variability maps of PSR J1910+1256, made with the Green Bank Telescope at 1400 MHz. The upper panel highlights a series of observations in which the pulsar is much brighter than usual. The shape of the profile in these bright observations also changes (Figure 5.14.

seen in Figure 5.14. Reasons for the similarities are discussed in Section 5.4.

PSR J1713+0747

PSR J1713+0747 is a binary millisecond pulsar, discovered in 1993 (Foster et al., 1993). With a timing stability of around 100 nanoseconds on a five year timescale (Verbiest et al., 2009), it is considered one of the highest precision objects in the sky. PSR J1713+0747 shows occasional small profile deviations that are not immediately obvious as instrumental issues (e.g. any deviations are evident in all frequency channels), but appear in isolated observations and so do not appear as a coherent long-term trend on a variability map (Section 4.5).

Figure 5.15 shows two similar profile deviations seen on two separate observations

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRAV MILLISECOND PULSARS




Figure 5.14: Pulse profiles of PSR J1910+1256, made with the Green Bank Telescope at 1400 MHz on MJD 53601. As the brightness of the flux calibrated observations increases, the width of the normalised profiles seems to narrow.




Figure 5.15: Observations of PSR J1713+0747 made at 1400 MHz at Arecibo. The left profile was observed on MJD 53798 and shows a dip in relative flux density either side of the pulse peak. The right panel shows a similar effect on MJD 54519.

at 1400 MHz at Arecibo. When the observation was made on MJD 53798, observations at 2300 MHz (with low S/N) and at 800 MHz with the Green Bank Telescope were also made, but no profile deviation is evident.

5.3 PUPPI and GUPPI results

5.3.1 Instrumental issues

The data processed by backends PUPPI and GUPPI, show far fewer profile deviations that are the result of instrumental issues, yet they are still apparently present. Figure 5.16 shows three deviant profiles from MJD 55305 for three different NANOGrav pulsars. A handful of MJDs are seen to recur when cataloging likely instrumental issues.

5.3.2 Astrophysical profile changes

PSR J1713+0747

As described in Section 5.2.2, the changes in pulse profile seen in this pulsar occur on timescales around, or less than, the span between observations. As discussed in Section 4.5, such variability is not well modelled by our GP, but is demonstrated in Figure 5.17.

Figure 5.17 highlights the more extreme profile deviations. Observations on MJD 56598 are made at three frequencies by two different telescopes; all three observations show abnormal deviation from the average pulse profile. The observations made at Arecibo at 2000 MHz show a drop in flux-density at each side of the pulse peak, whereas the 1400 MHz observations at Arecibo and the 1500 MHz observations made at Green Bank show a rise.

PSR B1937+21

The long-term variability of PSR B1937+21 seen in the ASP and GASP data,




Figure 5.16: Observations made with the Green Bank Telescope at 1500 MHz on MJD 55305. From left to right, the pulsars are PSR J1012+5307, PSR J1455-3330, PSR J1600-3053.

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRV MILLISECOND PULSARS

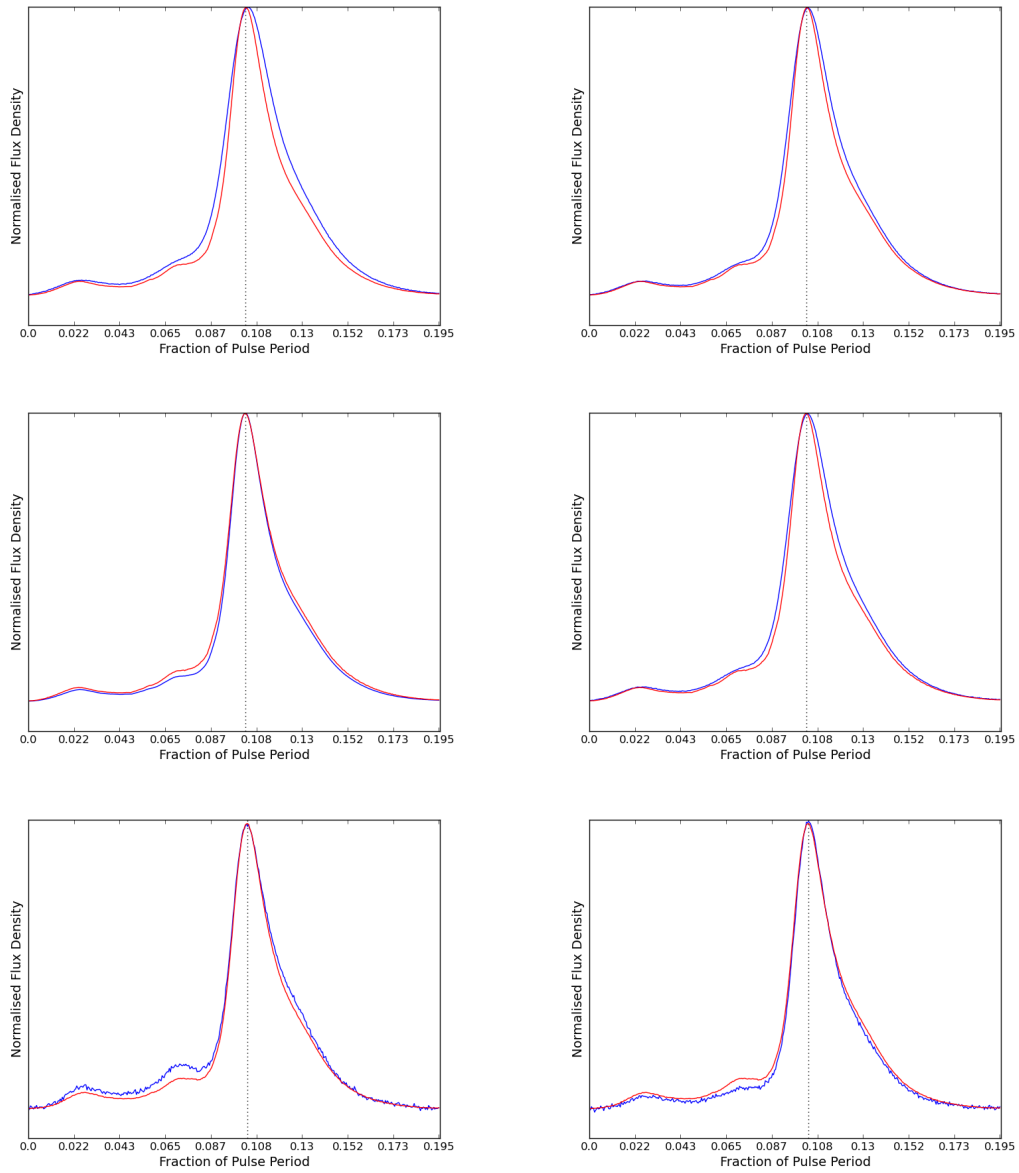


Figure 5.17: From top to bottom, pulse profile changes of PSR J1713+0747 taken at Arecibo at 1400 MHz, Green Bank at 1500 MHz and Arecibo at 2000 MHz. Left column: Observations made on MJD 56360, MJD 56201 and MJD 56319 from top to bottom. Right column: Observations all made on MJD 56598.

is even more clearly seen in the PUPPI and GUPPI data (Figure 5.18 (main pulse) and Figure 5.19 (interpulse). Examples of the kind of profile deviations that produce such variability maps are shown in Figure 5.20.

PSR J1600-3053

PSR J1600-3053 is a binary millisecond pulsar, which has been detected in X-rays and γ -rays (Espinoza et al., 2013), as well as the radio regime. As well as strong and systematic changes in brightness, the pulsars also shows periodic profile shape variations (upper and lower panels of Figure 5.21 respectively). As the profile shape changes are only slight, the lower S/N observations at 800 MHz do not reveal any such shape variability (Figure 5.21). Neither is it seen in any ASP or GASP profiles for the same reason. Figure 5.23 shows the extent of the variability in both the flux calibrated and normalised pulse profiles.

5.4 Discussion

Pulsar timing arrays need to minimise their timing residuals as much as possible in order to directly detect gravitational waves. Unmodelled pulse profile variability is detrimental to this. Through the analysis of the NANOGrav pulsars, we have uncovered profile variability due to instrumental issues and due to astrophysical causes.

Both generations of telescope backend have produced deviant pulse profile, but they are much less common in GUPPI and PUPPI. In any case, the problematic profiles must be flagged.

Additionally, other interesting profile changes have been found that are ostensibly the results of genuine astrophysical processes.

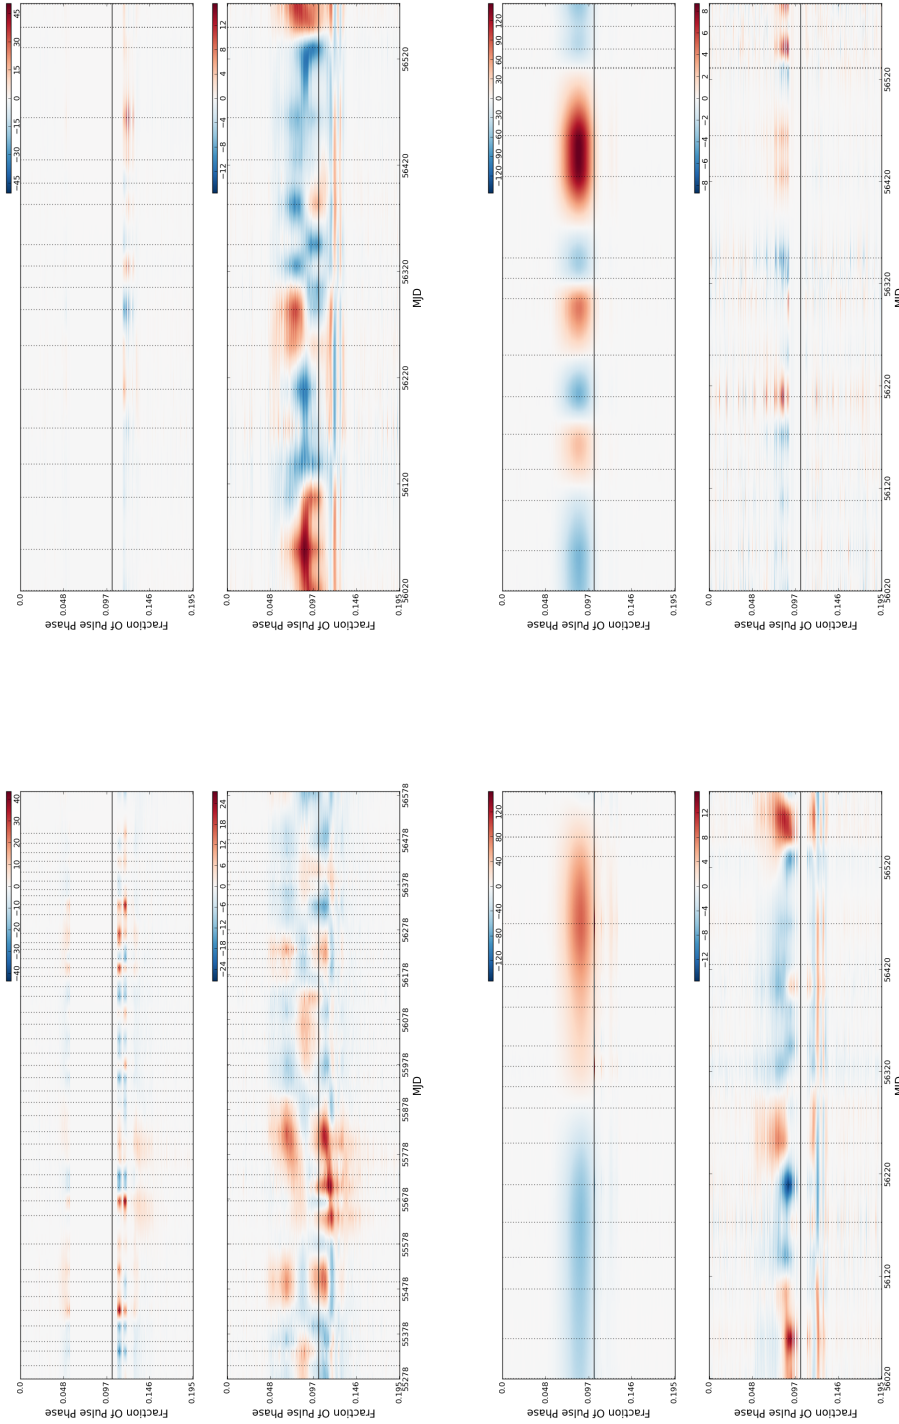


Figure 5.18: Variability maps for the main pulse of B1937+21. Clockwise from top left: Green Bank Telescope observations at 800 MHz and at 1500 MHz, Arecibo observations at 2000 MHz and at 1400 MHz.

CHAPTER 5. THE EMISSION VARIABILITY OF NANOGRAV MILLISECOND PULSARS

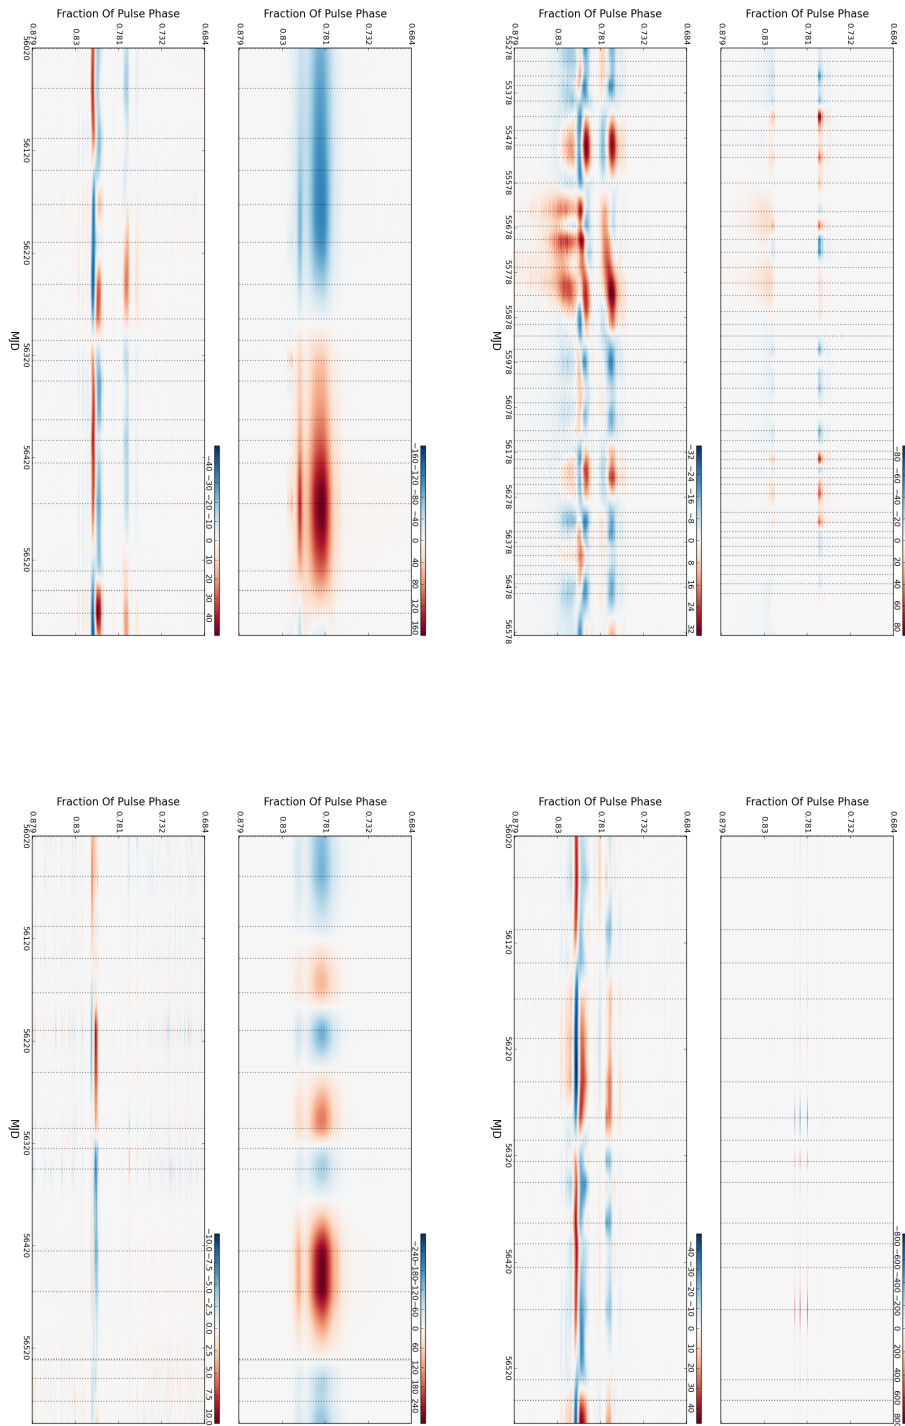


Figure 5.19: Variability maps for the interpulse of B1937+21. As Figure 5.18




Figure 5.20: Observations of the main pulse and interpulse of PSR B1937+21 at 800 MHz with the Green Bank Telescope. The left pair of panels shows the profile on MJD 55430 and the right on MJD 55641. On both days, profile shape deviations can be seen in the trailing edge of the main pulse and interpulse.

PSR B1937+21 shows changes in both the flux calibrated and the normalised pulse profiles, that occur across telescopes, frequencies and backends. The strongest long-term variability with ASP and GASP is seen with the Green Bank Telescope at 800 MHz. With GUPPI and PUPPI, variability is more clearly seen in all observations of PSR B1937+21 and agrees well across the frequency range (Figure 5.18 and Figure 5.19).

In contrast, Shao et al. (2013) describe their non-detection of pulse variation in PSR B1937+21, using data collected by the Effelsberg telescope from 1997 to 2013. The three profile components of PSR B1937+21 were each fitted by parabolas. The shape of, and distance between the parabolas were then tracked with time, only showing variability that was within the measurement error bars. The GP regression technique used in this thesis, however, is sensitive to even very subtle long timescale trends, and may recover some systematic variability in PSR B1937+21 (as we have seen in this chapter), which is missed by other techniques. To elucidate the apparent contradiction, we could isolate the region over which the Shao et al. data overlap with the NANOGrav data, to conduct variability analysis and

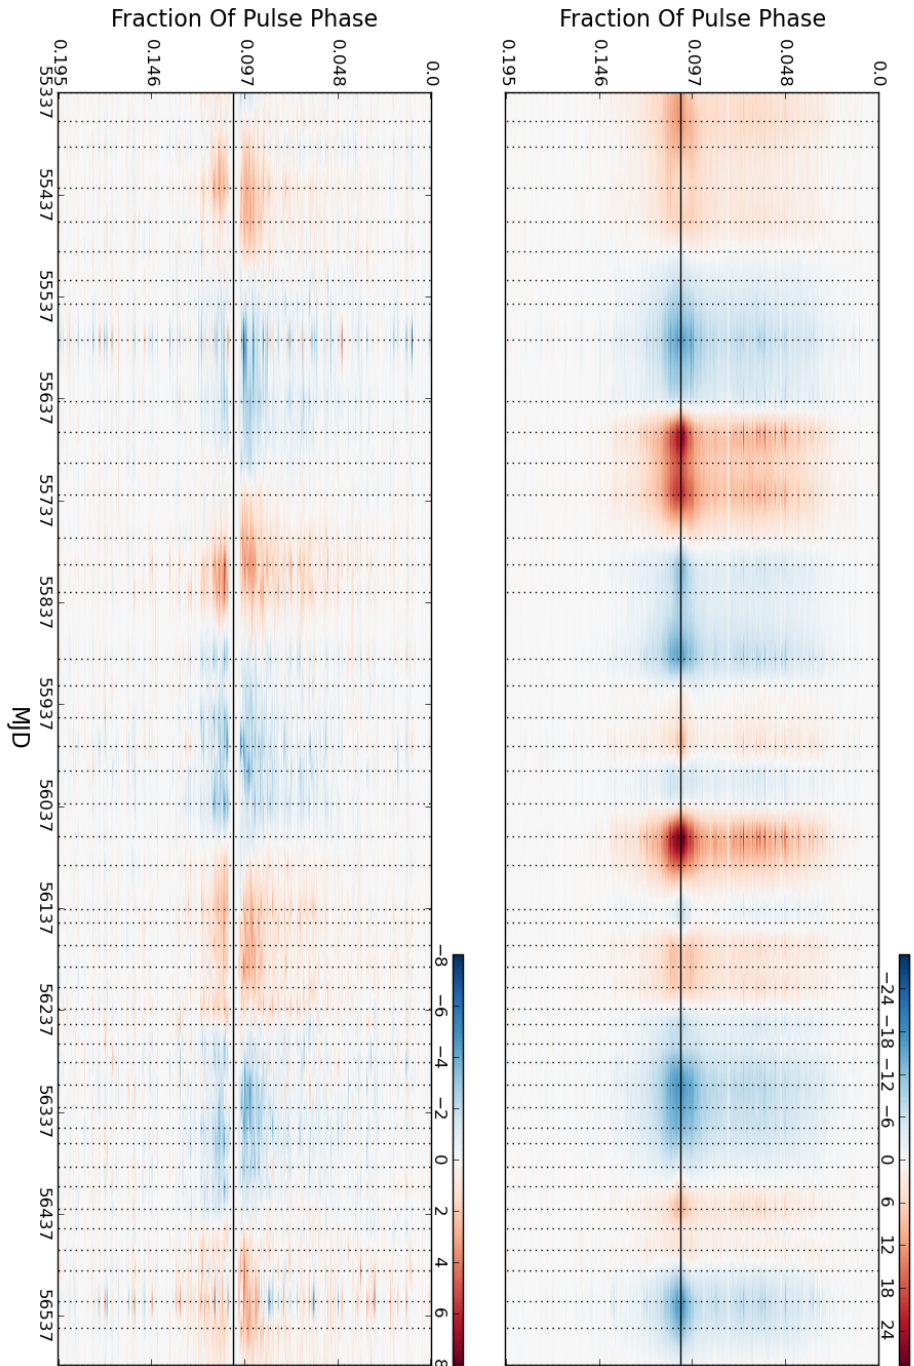


Figure 5.21: Variability maps for PSR J1600-3053, composed of Green Bank Telescope observations at 1500 MHz. Long term variability is seen in the flux calibrated observations, and periodic profile shape variability is seen in the normalised observations. Examples of the degree of variability in both the flux correlated and the normalised profiles are shown in Figure 5.23.

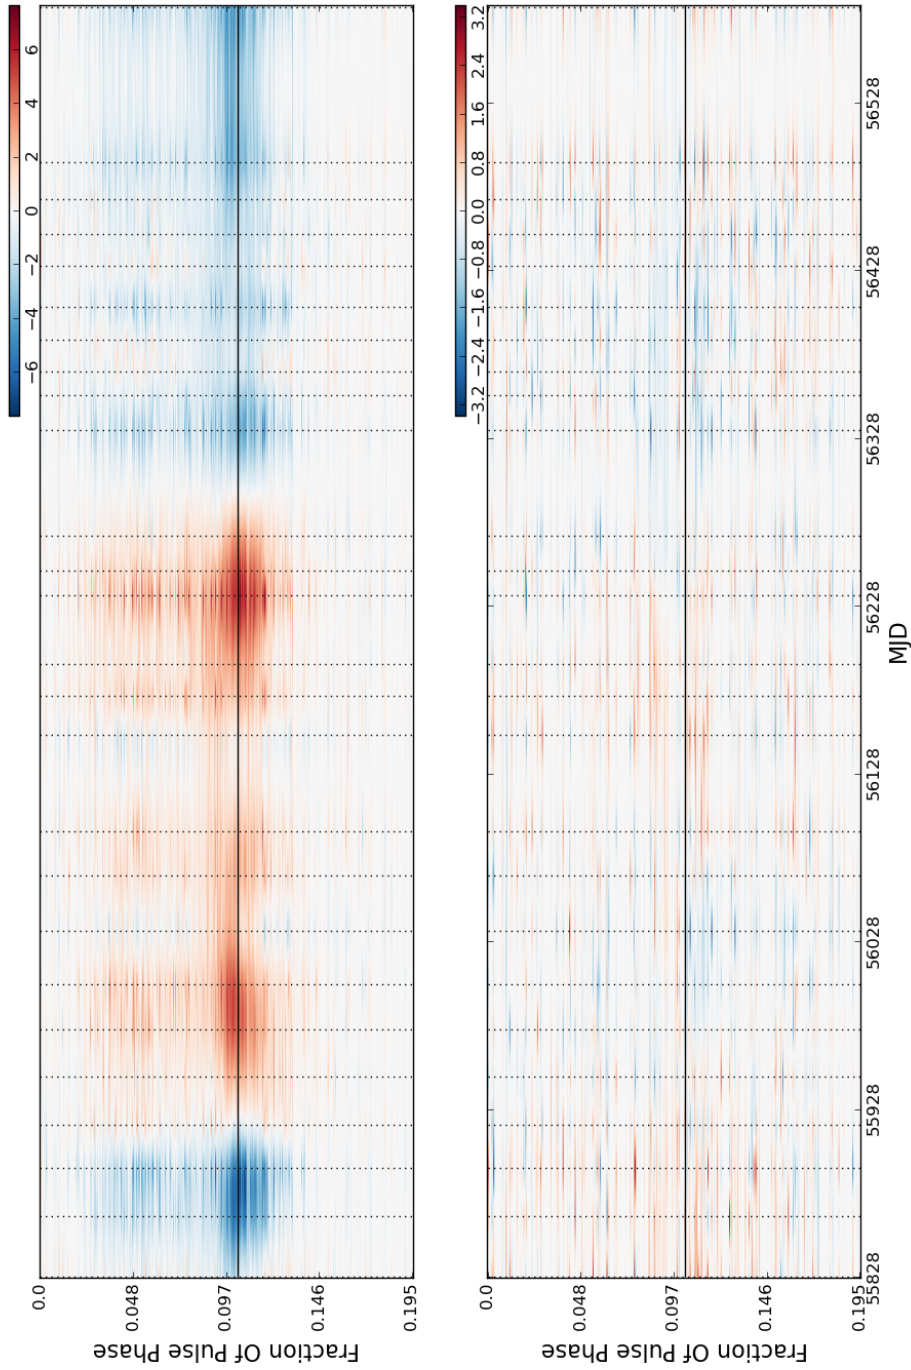


Figure 5.22: Variability maps for PSR J1600-3053, composed of Green Bank Telescope observations at 800 MHz. Long term variability is seen in the flux calibrated observations but not in the normalised profiles.




Figure 5.23: Examples of the pulse profile variability of PSR J1600-3053, as seen in Figure 5.21. The left panel shows a flux calibrated profile from MJD 56066. The right panel shows a normalised profile from MJD 55639. Both observations were made with the Green Bank Telescope at 1500 MHz.

compare the results.

Variability in PSR B1937+21 is not a new discovery. In terms of the timing, Shannon et al. (2013) have shown that the timing variations in PSR B1937+21 are consistent with the signature of an asteroid belt. In terms of the emission variability, it has been demonstrated that PSR B1937+21 produces giant pulses at the near trailing edge of the pulse and interpulse (Zhuravlev et al., 2013); a region where most of the profile variability is seen in our analysis. It stands to reason that if a section of pulse phase is predisposed to produce a relatively low number of high intensity pulses, then the usual observation length will not be long enough for that region of pulse phase to have stabilised under the law of large numbers. We should expect to see profile variability in that region of pulse phase.

Intriguingly, the giant pulsars appear near the trailing edge of both the main pulse and the interpulse. When the individual observations are monitored, the behaviour of the trailing edge of the main pulse is mimicked incredibly well in the interpulse, and is suggestive of a single source. Two examples are seen in

Figure 5.20.

Much profile variability is seen in PSR J1713+0747, but the deviations are mostly seen in single observations only, rather than showing a long-term trend. In one case, there are three observations on PSR J1713+0747 on the same day (MJD 56598). Two were made at Arecibo with the PUPPI backend, one at a frequency of 1400 MHz and one at 2000 MHz. An observation was also made by the Green Bank Telescope at 1500 MHz. The observations made at 1400 MHz and 1500 MHz agree very well, and show a higher flux density than average at either side of the pulse peak. In contrast, the 2000 MHz observation shows a drop in flux density. Is it possible that some radiation has been lost at one frequency, but gained in another, i.e. a spectral change?

We have shown subtle, but remarkably periodic profile changes in PSR J1600-3053. Due to their subtle nature, they can only be seen in the 800 MHz Green Bank Telescope observations with the GUPPI backend. The periodicity seems to be a little under 400 days. We should be wary of any periodicity that has around a yearly cycle, but there is no obvious reason why the profile changes should be affected this way by the Earth's cycle.

PSR J1853+1303 and PSR J1910+1256 have observation days in which the shape of the normalised profile is correlated with the brightness of the pulsar in the flux calibrated observations. The brightest profiles seem to produce narrower normalised pulse profiles. This behaviour is also seen in PSR J1830-1059 (Section 4.4.1).

It is not easy to disentangle the observations in which a pulsar brightens due to intrinsic changes in emission, and ones that *appear* to brighten due to effects of the intervening medium. It is not obvious, however, how the latter could also change the shape of the pulse profile in a consistent way. One conclusion then, is that an

intrinsic pulsar process is able to produce bright pulse profiles that are also more narrow than usual.

In the future, an analysis of the rotational variation of the NANOGrav pulsars would allow us to perform correlation analysis, as was done in the pervious chapter. It would be especially interesting if periodic variability is seen in the timing of PSR J1600-3053 that matches the profile changes.

The profile variability of PSR J1600-3053 is so slight, so as to be noticeable only at one frequency with GUPPI, and not at all in the GASP observations. This makes the case for more sensitive observations in order to observe the true extent of pulse profile variability in millisecond pulsars.

Chapter 6

Detecting variable spindown rates in mode-changing and nulling pulsars

This chapter describes an approach designed to discover whether the different emission modes of nulling and mode-changing pulsars are each associated with different rates of spindown.

6.1 Motivation

Previous chapters have provided much evidence of the link between a pulsar's radio emission and its rate of spindown, yet so far we have only been able to observe this relationship on timescales of months and years. The reason for this, is that the rate of spindown is generally so slight, that it can only be detected over the span of weeks or longer. Therefore, when emission changes occur on short timescales, e.g. in mode-changing and nulling pulsars, we are unable to measure

the behaviour of the spindown rate to know if it is changing in a similar way.

Because similarities in emission behaviour are observed over a wide variety of timescales (e.g. nulling, intermittent pulsars and RRATS), it is often hypothesised that the variability timescale is a continuum. This paradigm leads us to expect that emission changes in nulling and mode-changing pulsars may also produce correlated rotational behaviour, as their longer-term counterparts do.

Despite the difficulties, there is a scenario in which it would be possible to infer whether mode-changing and nulling are behaviours that are accompanied by a change in spindown rate. We first conceptualise a pulsar that has two distinct emission states, each having a different rate of spindown. The pulsar switches between states on timescales of minutes and hours. If we observe the pulsar continuously for a span of time, we can say what fraction it spent in state A and what fraction it spent in state B. Depending on the length of the observation span and the nature of the pulsar, these *state-fractions* may be different (to some degree) for each span. If the observations are long enough, so that an average spindown rate could be measured, then it follows that two separate monitoring spans in which the state-fractions are different, would have a different average spindown rate.

The relationship between short-term emission changes and pulsar rotation could be elucidated by continuous monitoring of a sufficiently bright mode-changing or nulling pulsar. An analysis of its emission will reveal the fraction of time spent in each state over a span of time. If we begin to see a correlation between the fraction of time spent in an emission state and the measured spindown rate, then we can infer that each emission state also has a distinct spindown rate associated with it.

We have produced a simulation in order to explore the range of parameters over

which such an experiment would be possible. For example, if the state-fraction of the pulsar has only small variance between each observation span, then the average spindown rate will be similar each time. A link between the state-fraction and the spindown rate would be difficult to detect under these circumstances. A large variance in the state-fraction would be much more conducive to this type of analysis.

We attribute typical properties to a simulated pulsar and allow its state-fraction to change in various ways. This varying fraction is then compared to the measured spindown rates, and the results are presented below.

6.2 Simulation and results

We have written a Python code that is able to simulate the behaviour of a mode-changing or nulling pulsar, and produce artificial TOAs.

Expressed as a Taylor expansion, a pulsar's rotation frequency is given by

$$\nu(t) = \nu_0 + \dot{\nu}_0(t - t_0) + \frac{1}{2}\ddot{\nu}_0(t - t_0)^2 + \dots, \quad (6.1)$$

where the subscript 0 denotes the value of a variable at some reference epoch t_0 . Rotation frequency is \dot{N} , where N is the pulse number. We can integrate equation 6.1 to show that

$$N = N_0 + \nu(t - t_0) + \frac{1}{2}\dot{\nu}(t - t_0)^2 + \frac{1}{6}\ddot{\nu}(t - t_0)^3 + \dots, \quad (6.2)$$

where N_0 is the pulse number at t_0 .

With the exception of very young pulsars, the $\ddot{\nu}$ term is too small to be measured and N can be accurately approximated by the first three terms on the right-hand

side of equation 6.2.

The following variables can be set: the initial rotation frequency ν_0 ; the state-fraction; the $\dot{\nu}$ value of each state; the total observation span; the period over which $\dot{\nu}$ and the state-fraction are measured; the timescale of state changes and the intervals in the code iterations at which N is evaluated and ν is updated. The more often ν is updated, the more precise the calculation of the pulsar TOAs will be.

As frequently as desired, the pulse number N of the pulsar can be evaluated from equation 6.2. The value of ν is then updated so that the next calculation of N will maintain the high level of precision needed for this analysis.

We choose the effective TOA for the simulated pulsar to be taken when N has an integer value, i.e. when the pulsar beam is pointing towards Earth. The simulation calculates N whenever the evaluation takes place. The time for the pulsar to rotate so that N has the next integer value is easily calculated and so this is added to the evaluation time to produce a simulated TOA.

To add noise to the simulated TOAs, a sample is drawn from a Gaussian distribution, with a zero mean and a standard deviation of 1×10^{-4} seconds, which corresponds to a 10 ms period pulsar observed at a S/N ratio of 100 (Equation 1.10).

At an interval determined by the user, a crossroad is reached, at which point the pulsar can remain in its current state or switch to the other. This is decided by the generation of a random number, weighted by the value of the underlying state-fraction at that point in the simulation.

Once the whole observation length has been simulated, the TOAs are then saved in a TOA file format that is recognised by TEMPO2.

With these TOAs as an input file, TEMPO2 then produces timing residuals with respect to a simple timing model consisting of constant values of ν and $\dot{\nu}$. The variability of $\dot{\nu}$ can then be calculated from the timing residuals as described in Section 2.2.4.

The length of the continuous observation that we simulated was 800 days. In 15 day blocks, the state-fraction is recorded as the fraction of time spent in one state, and the spindown rate is evaluated. Within these blocks, at hourly intervals, the rotation rate of the pulsar is updated and the spindown rate is permitted to switch between two values. The values chosen were $\dot{\nu}_1 = -1.05 \times 10^{-13} \text{ s}^{-2}$ and $\dot{\nu}_2 = -1.1 \times \dot{\nu}_1$.

In each run of the simulation, the state-fraction is initially set to 0.7 (i.e. the pulsar spends $\sim 70\%$ of the time in $\dot{\nu}_1$), but it is then permitted to change in various ways:

1. The state-fraction is drawn from a Gaussian distribution with a standard deviation of 0.05 and a mean of 0.7.
2. The state-fraction is drawn from a Gaussian distribution with a standard deviation of 0.5 and a mean of 0.7.
3. The state-fraction drops by 0.1 every 15 days.
4. The state-fraction drops by 0.01 every 15 days.
5. The state-fraction drops by 0.005 every 15 days.
6. The state-fraction drops by 0.002 every 15 days.
7. The state-fraction varies sinusoidally.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.1: Correlation between spindown rate and a state-fraction varying randomly with time. Top: The state-fraction is randomly drawn from a Gaussian distribution about a mean of 0.7, with a standard deviation of 0.05. Middle: The spindown rate as measured from the simulated TOAs. Bottom: The measured spindown rate as a function of the state-fraction. No trend is obvious if the state-fraction behaves in this way.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.2: Correlation between spindown rate and a state-fraction varying randomly with time. The state-fraction is drawn from a Gaussian distribution with a standard deviation of 0.5. Otherwise as Figure 6.1.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.3: Correlation between spindown rate and a state-fraction varying linearly with time. The state-fraction drops by 0.1 every 15 days. Otherwise as Figure 6.1.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.4: Correlation between spindown rate and a state-fraction varying linearly with time. The state-fraction drops by 0.01 every 15 days. Otherwise as Figure 6.1.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.5: Correlation between spindown rate and a state-fraction varying linearly with time. The state-fraction drops by 0.005 every 15 days. Otherwise as Figure 6.1.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.6: Correlation between spindown rate and a state-fraction varying linearly with time. The state-fraction drops by 0.002 every 15 days. Otherwise as Figure 6.1.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS




Figure 6.7: Correlation between spindown rate and a state-fraction varying sinusoidally with time. Otherwise as Figure 6.1.

In Figure 6.1 and Figure 6.2 the state-fraction is randomly drawn from a Gaussian distribution about a state-fraction of 0.7 and with standard deviations of 0.05 and 0.5 respectively. When the standard deviation of the state-fraction is only 0.05, the variation seen in $\dot{\nu}$ is also small in comparison to its uncertainty. There is no correlation between $\dot{\nu}$ and state-fraction as a result. When the standard deviation of the state-fraction is 0.5, however, $\dot{\nu}$ is seen to vary over a wider range and a weak correlation can be seen.

In figures 6.3 to 6.6 the state-fraction changes linearly with time. In Figure 6.3 it drops by 0.1 every 15 days until it is entirely in one emission mode. In this case, $\dot{\nu}$ a correlation between state-fraction and spindown rate becomes clear. As the nature of this change of state-fraction may be unrealistic, figures 6.3, 6.4, 6.5 and 6.6 show how the correlation changes as the state-fraction drops more slowly over time. Even in the case when the state-fraction is systematically dropping by only 0.02 every 15 days, the correlation with spindown rate remains clear.

In Figure 6.7 the state-fraction is simulated to vary sinusoidally. The systematic nature of the variation again results in a strong correlation with the measured spindown rate.

In general, a correlation is most clearly visible when the range over which $\dot{\nu}$ varies, is large with respect to its uncertainty. This occurs when the two $\dot{\nu}$ values are very different, when the spread in state-fraction is wide, when the state-fractions vary in a systematic way, or any combination of the three.

The vital parameter, as shown in Table 6.1, is

$$M = \frac{\dot{\nu}_{max} - \dot{\nu}_{min}}{\delta\dot{\nu}}, \quad (6.3)$$

Figure	M Value	Correlation Coefficient
6.1	-0.5	0.31
6.2	-3.4	0.64
6.3	-29.4	0.99
6.4	-13.7	0.99
6.5	-14.0	0.95
6.6	-3.1	0.82
6.7	-3.5	0.95

Table 6.1: M value and correlation coefficient for the state-fraction and spindown rate of figures 6.1 to 6.7.

where $\dot{\nu}_{max}$ and $\dot{\nu}_{min}$ are the maximum and minimum values of spindown rate, and $\delta\dot{\nu}$ is the uncertainty in the spindown measurements.

6.3 Proposal 1: Do state-fractions change over time?

The best opportunity to observe whether short timescale emission changes are indeed linked to rotation, would be provided by a mode-changing or nulling pulsar with two significantly different rates of spindown, and a state-fraction that varies systematically over a wide range. The long term behaviour of the state-fraction of mode-changing and nulling pulsars in this context is currently unknown. Published mode-changing and nulling fractions, i.e. the fraction of the observation duration in which a pulsar is in a null state or an alternative mode, have typically been obtained through single, long-duration observations (e.g. two hours for Wang et al. (2007)). In order to learn more about the nature of the state-fraction variance, we propose observations of pulsars which change their states on short timescales, to reveal whether the mode-changing/nulling fractions vary systematically over a span of observations, spaced weeks apart.

When selecting target pulsars, we consider the evidence seen in pulsars showing

long-term variability (e.g. Kramer et al., 2006; Lyne et al., 2010), and hypothesise that greater differences in the emission properties of state-switching pulsars will be related to greater changes in the rotational properties, as both phenomena are associated with streaming energetic particles. We will, therefore, target pulsars with pronounced differences between the two emission states, nulling pulsars being the extreme case and mode-changing pulsars with significantly different mode profiles.

6.4 Proposal 2: Continuous monitoring

If the above proposal shows evidence that the mode-changing and/or nulling fractions vary systematically, such as to allow the determination of a correlation with the spindown rate, then we make a further proposal: an experiment that would continuously observe a bright nulling and/or mode-changing pulsar in order to monitor its state-fraction. The target pulsar should be circumpolar and bright enough so as to be observable by a small radio telescope that could be commandeered for such an experiment; given the amount of time needed to evaluate $\dot{\nu}$ with a small uncertainty, at least a year of continuous observations would be needed, giving us 24×15 day samples. Furthermore, if the monitoring were to be conducted using an aperture array, then we could potentially point at a small number of mode-changing/nulling pulsars and obtain a unique dataset.

The mode-changing/nulling fraction will be determined in quasi real-time, using a process of template matching and the well known pulse profile of each mode of each mode-changing pulsar. The amount of storage space needed to store a year's worth of single pulses from a ~ 0.5 second period pulsar is of the order of hundreds of gigabytes.

CHAPTER 6. DETECTING VARIABLE SPINDOWN RATES IN MODE-CHANGING AND NULLING PULSARS

Continuous monitoring of a pulsar with properties described above, would enable us to discover if the emission states of mode-changing and nulling pulsars do have distinct spindown rates and would be a step toward the unification of pulsars that show long and short-term emission variability.

Chapter 7

Conclusions

Throughout this work, new techniques have been developed which have yielded many new findings. In this final chapter we recap both, before offering a framework for future studies of pulsar variability.

7.1 New techniques

Pulse profile monitoring technique

We have developed a new technique for monitoring long-term emission variability in pulse profiles. In the latest version, the emission variability in each pulse profile phase bin is modelled by GP regression to produce a smooth continuous variability map, despite sparse and irregularly-spaced observations.

What could be considered a limitation of the GP regression technique, is that it is predominantly sensitive to systematic trends, and so profile variability that occurs in individual observations only, often do not appear in the variability maps. An extreme outlier, however, may influence the GP regression enough to appear in the variability map and serve as a flag to the presence of potentially problematic data.

The pulse profile monitoring technique has been used to monitor the absolute brightness of the pulse profiles, as well as the shape changes in profiles that have been normalised to their peak.

Spindown monitoring technique

A new technique for modelling the spindown rate of a pulsar has also been developed. The basis is again GP regression, which is used to produce an analytical model that describes the timing residuals. The second time derivative of this model can be taken, to find the spindown rate at any point. Prior to this work, only numerical techniques have been used to obtain rates of spindown.

A comparison of the spindown rate and the pulse profile variability reveals that correlations exist at certain pulse phases, sometimes with time lags. This information could hold vital clues regarding the nature of pulsar variability.

The identification of any correlation between spindown rate and profile variability, hinges on our ability to accurately measure both. These measurements come from pulsar timing campaigns in which integration times used are based on the S/N required to obtain high quality TOAs. This S/N however is not always sufficient to track the variability in individual pulse phase bins. Except in rare cases, it is only in the brightest radio pulsars that we are currently able to uncover profile variability and make comparisons with spindown rate.

We do not currently have a physical model for the processes that are driving the rotational variability, seen as systematic noise in the timing residuals. It is necessary, therefore, to first examine the data, in order to inform the nature of the covariance function that will provide the best-fitting model.

Using both the old and new generations of variability monitoring techniques with large pulsar datasets has yielded many new findings, presented in the next section.

7.2 New findings

Simultaneous emission and rotational changes in PSR J0738-4042

We have shown that in 2005, both the pulse profile and spindown rate of PSR J0738-4042 underwent dramatic and sudden changes. We have demonstrated that the spindown changes can be associated with a mass that falls within the range of typical asteroid masses, offering one possible explanation of the event.

As was discussed in Chapter 3, it is not obvious how other, less pronounced profile changes throughout the history of PSR J0738-4042 are related to the 2005 event. It is also unclear how the 2005 event is related to other forms of emission and rotational variability seen in various pulsars over a range of timescales, however the event does share traits with other forms of pulsar variability and so may be considered to have a common cause. Explanations of variability beyond the asteroid hypothesis are laid out in Section 7.3.

Observationally, PSR J0738-4042 is bright and we have the benefit of a quarter of a century of observations of it. It is possible that such variability is not uncommon when a bright pulsar is observed for long enough.

Throughout the 2005 event, and in subsequent observations with the Parkes telescope, hints of communication between regions of the pulse profile are often seen. Some phase regions show correlation, while some are anti-correlated. This is seen clearly in the correlation map of Figure 4.8. The anti-correlation in some adjacent phase regions is entirely expected when one also considers that changes in flux density appear to drift from one region of pulse phase to another, seen most notably in Figure 3.7.

Simultaneous emission and rotational changes in PSR J1602-5100

In Chapter 4, the analysis of nine out of a total of 168 young, energetic pulsars observed as part of the Fermi Timing Programme provided us with a handful of examples of correlated profile and spindown variability. The outstanding case was seen in PSR J1602-5100, and shows a huge profile change and simultaneous drop in spindown rate, with a duration of around two years.

Following the calculations introduced by Kramer et al. (2006) to relate magnetospheric currents to $\dot{\nu}$ for the intermittent pulsar B1931+24, the injected mass calculations associated with the $\dot{\nu}$ changes of PSR J1602-5100 give similar results to PSR J0738-4042.

Pulse profile variability seen in millisecond pulsars

Analysis of the NANOGrav observations shows that pulse profiles displaying indications of instrumental problems are sometimes included in the pulsar timing array analysis, in the data processed by both the ASP/GASP and PUPPI/GUPPI backends. These profiles need to be accounted for, in order to increase the timing precision of the millisecond pulsars involved.

We have also shown for the first time, that long-term and coherent pulse profile variations exist in millisecond pulsars, due to astrophysical and not instrumental effects.

One pulsar that consistently seems to show small profile changes is B1937+21. The variability seems to mainly appear at the trailing edge of both the main pulse and the interpulse, possibly due to the appearance of giant pulses at those profile regions (Section 5.4). The similarity of the profile changes that occur, suggests prompt communication between the main pulse and interpulse and is

indicative of a common source. Such rapid communication between main pulse and interpulse has been witnessed before in other interpulse pulsars (Gil et al., 1994; Weltevrede et al., 2007).

A link between changes in intrinsic pulsar brightness and profile shape

In both the Fermi Timing Programme and NANOGrav datasets, there are some normalised pulse profiles that show subtle drops in flux density either side of the main pulse, i.e a slight narrowing is observed (e.g. Figure 5.14). When the corresponding calibrated flux profiles are reviewed, they show much brighter profiles than the average, i.e. profile components that are more narrow than average seem to be substantially brighter than average. This is not a reciprocal relationship, because observations which appear bright can be caused both intrinsically and by propagation effects. Only the former would be expected to simultaneously affect the brightness and the shape of a pulse profile.

For illustrative purposes, Figure 7.1 shows how the shape of a normalised Gaussian function changes when its amplitude and its standard deviation are modified in turn. As expected, only changing the latter produces a shape change in the normalised function. This demonstrates that an observation which is amplified by effects of propagation would not show shape modulation when normalised.

This effect is striking in PSR J1830-1059 (Section 4.4.1); Figure 4.4 shows that the absolute flux density changes in the top panel occur in synchronicity with the profile shape changes seen in the middle panel (which are both related to the change in $\dot{\nu}$). The changes in brightness of this pulsar, therefore, appear to be intrinsic, as opposed to being due to refractive scintillation. With its timescale of variability, intrinsic and periodic jumps in brightness and their relationship with

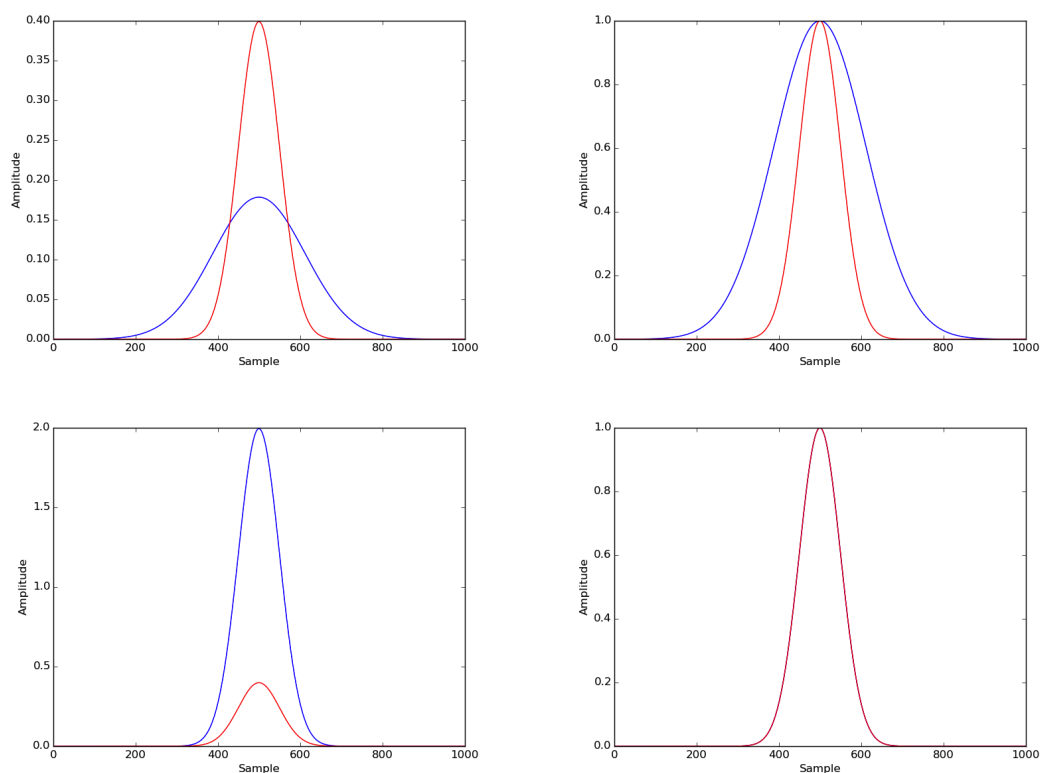


Figure 7.1: Effects of normalisation on Gaussian functions. In the top left panel, the blue Gaussian function has a standard deviation that is a factor of 5 larger than the red Gaussian. The top right panel demonstrates how difference in shape of these two functions when they are each normalised to their peak. In the bottom left panel, the blue Gaussian function has an amplitude that is a factor of 5 larger than the red Gaussian, but they have the same standard deviation. The bottom right panel shows that both the blue and red traces share an identical shape when normalised to their peak.

spindown rate, PSR J1830-1059 shares much in common with the intermittent pulsars and the transient component of PSR J0738-4042 seen in Chapter 3.

Observations showing the same effect are seen in another of the Fermi Timing Programme pulsars, PSR J1359-6038 (Section 4.4.8), as well as the NANOGrav pulsars PSR J1853+1303 and PSR J1910+1256 (Section 5.2.2).

Throughout our analysis, the best examples of pulse profile variability are primarily the bright, high S/N pulsars. This is unsurprising, as subtle profile changes are easily masked by noisy profiles. The possibility also exists, however, that intrinsically bright radio emitting pulsars have a propensity to display profile fluctuations.

How to detect variable spindown rates from mode-changing and nulling pulsars

We have shown that it is possible in principle to determine if mode-changing and nulling are associated with spindown changes, via continuous pulsar monitoring, and we propose a method to do so. Evidence of a correlation would provide a valuable link between short- and long-term variability.

7.3 Proposed framework for the interpretation of pulsar variability

With each new finding in this work, new areas for study open up in the field of pulsar variability, that can be addressed in the future. I propose the following framework within which further investigation can be conducted.

Unification

In an attempt to unify the various types of pulsar variability, it is natural to ask two questions:

1. Are pulsars that switch *between* emission states and those which appear to switch off completely, controlled by the same underlying mechanisms?
2. Are common processes responsible for the very different timescales of pulsar variability that we observe?

In answer to question (1), we know that nulling seems to just be low emission in some cases (e.g. Young et al., 2015). This connection could be fortified by a study which finds that the frequency of each phenomenon and the distribution of their state-fractions to be comparable.

The transient component seen in PSR J0738-4042 is considered to be the result of a state-switch, but the behaviour of the component in isolation is comparable to that of an intermittent pulsar, i.e. it appears and disappears on timescales of many years. Similarly, other pulsars that are considered to have variable emission have components that have RRAT-like behaviour.

Question (2) asks if nulling, intermittency and RRAT-like behaviour are members of the same family and if the same can be said of mode-changing and state-switching. As mentioned at the end of the last section, a discovery that nulling or mode-changing pulsars have different spindown rates associated with their emission states, would strengthen their ties with long timescale emission variability, which has conclusively been shown to be linked to rotation.

We have also seen in the previous chapter, that if the state-fraction of a nulling/mode-changing pulsar is allowed to wander, timing noise is naturally produced when

emission and rotation are linked. Similarly, the behaviour of state-switching pulsars is reproducible by the a mode-changing pulsar with a state-fraction which wanders over long timescales.

Causes

There are four possible ways in which emission from a pulsar can appear to cease. If we make the hypothesis that the various timescales and types of emission variability are all different facets of an underlying process, then we can attribute all variability to the same cause.

1. Movement of the pulsar emission beam direction:

Radio emission can disappear if the emission beam moves away from our line of sight. Do the last (or first) few profiles of an intermittent pulsar show any change in their polarisation position angle before they enter (or after they exit) their inactive phase? Movement of the magnetic axis would also affect the rotation of the pulsar.

2. A failure of the emission mechanism:

The currently accepted explanation for the behaviour of intermittent pulsars is that changing magnetospheric currents are responsible for correlated changes in emission and rotation (Kramer et al., 2006). Under the assumption that this hypothesis is correct, and that there are different $\dot{\nu}$ values for the active and inactive states, we should be able to see a correlation between the lengths of the inactive phases and the discrepancy between the observed ν and the ν predicted by an unchanging spindown rate when the next active phase begins. In other words,

the longer the inactive phase lasts, the larger the difference between the observed and expected frequency should be. It seems, however, that no such correlation currently exists. This is due in part to the fact the errors in the lengths of the active and inactive phases, and the fitted values for $\dot{\nu}$ while the pulsar is active, are likely to corrupt the correlation analysis. In principle, however, this method provides a test of the current intermittent pulsar paradigm.

An alternate possibility, is that an intermittent pulsar maintains the same spindown rate when both active and inactive, and that glitches are responsible for the switch from an inactive to an active state. This jump in ν could also produce the timing signatures seen in the intermittent pulsars (Figure 1.19).

A method of testing whether an intermittent pulsar's radio emission does indeed cease due to a global change in magnetospheric current, is if the pulsar were embedded within a pulsar wind nebula. Variability in the nebula's emission on a timescale similar to that of the intermittent pulsar would strongly support the hypothesis that a cessation of outflowing currents is responsible for the quiescent phases.

Correlated emission and spindown variability has also been observed in a γ -ray pulsar (Allafort et al., 2013). In this object, the spindown rate *drops* when the γ -ray flux increases. This is at odds with the idea that an increase in global magnetospheric currents is responsible for increased emission along with an increase in braking torque.

Finally, it is not understood why pulsars that are far from the $P - \dot{P}$ diagram deathline should temporarily become inactive.

3. Obscuration of the emission beam:

This potential explanation is included for completeness, but it is difficult to see

how beam obscuration could account for either the emission or rotation variability signatures that we see in pulsars.

4. Shifts of emission frequency:

PSR B0943+10 is thought to undergo rapid and global magnetospheric changes that induce a switch between radio and X-ray emission modes (Hermsen et al., 2013). When the pulsar is in the radio-bright mode, only non-thermal unpulsed X-rays are seen. In the radio-quiet mode, the X-ray luminosity increases by more than a factor of two, and a pulsed thermal component is observed.

The binary millisecond pulsar system PSR J1023+0038 has been seen to enter a radio-quiet mode in which the γ -ray emission increases by a factor of five. Stappers et al. (2014) conclude that even when the radio signal is quiescent, the pulsar mechanism is still active and creating the observed γ -ray emission. PSR J1023+0038 is thought to be experiencing periods in which it is accreting material from a companion star. In the accretion phases, the radio emission is suppressed and X-ray and γ -ray emission increase.

Is it possible that intermittent pulsars are behaving in a similar way? X-ray observations were made with the Chandra X-ray Observatory for two of the three known intermittent pulsars (PSR J1841-0500 and PSR J1832+0029), but no emission was detected.

Do we also see evidence for less extreme cases of radiation frequency transfer, in which the shift is only from one radio frequency to another? In Chapter 5, the millisecond pulsar PSR J1713+0747 was observed to show a flux density increase at 1400 MHz and 1500 MHz, but a *decrease* at 2000 MHz (Figure 5.17).

What could cause such frequency shifts? The RFM model discussed under the heading of *Integrated pulse profiles* in Section 1.2.2 would suggest that an observed

spectra change in the flux density could be attributed to a change in emission height. These emission changes must also be related to the rotation rate of the pulsar if they are able to explain the behaviour of intermittent pulsars. Multi-frequency observations of pulsars that show profile variability would be useful to investigate this idea further.

Final thoughts

Although the exotic nature of neutron stars does not lend itself to simple physical interpretation, careful observations and inventive analysis techniques can allow us to learn more. The braking index, for example, contains information about a pulsar's winds and magnetic fields. With the exception of very young pulsars, however, the braking index can rarely be found as the $\dot{\nu}$ term is contaminated by timing noise. Furthermore, in the cases where it can be determined, it is often far from the value of 3 that is predicted by a simple rotating magnetic dipole model (Kaspi & Helfand, 2002).

Our GP regression method of measuring the $\dot{\nu}$ could be taken one step further to find $\ddot{\nu}$ and, hence, obtain braking index variability. Monitoring the braking index in variable pulsars could offer vital new insights into neutron star physics, especially when considering the causes of the dramatic $\dot{\nu}$ variability that we have witnessed already.

Immediate future plans involve the analysis of further pulsar datasets with the techniques that have been developed throughout this work. Improving the S/N of our observations, either with more sensitive instruments or longer integration times, would permit us to detect cases of pulse profile variability that currently remain hidden in noisy observations. As well as the potential to see new cases of variability, working with different instruments at various frequencies would help

to distinguish between instrumental and astrophysical effects by cross-referencing results. One longer-term goal is to be able to predict the timing noise of a variable pulsar by analysis of the pulse profile shape only. The ability to successfully mitigate timing noise in this way would be extremely valuable to the field of precision pulsar timing and would bring us a step closer to understanding pulsars and their environments.

References

- Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *Classical and Quantum Gravity*, 32, 024001
- Allafort, A., Baldini, L., Ballet, J., et al. 2013, *Astrophysical Journal, Letters*, 777, L2
- Andersson, N., Glampedakis, K., Ho, W. C. G., & Espinoza, C. M. 2012, *Physical Review Letters*, 109, 241103
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Science*, 340, 448
- Baade, W. & Zwicky, F. 1934, *Proceedings of the National Academy of Science*, 20, 254
- Backer, D. C. 1970a, *Nature*, 228, 1297
- Backer, D. C. 1970b, *Nature*, 228, 42
- Backer, D. C. 1976, *Astrophysical Journal*, 209, 895
- Bates, S. D., Lorimer, D. R., & Verbiest, J. P. W. 2013, *Monthly Notices of the RAS*, 431, 1352
- Biggs, J. D. 1992, *Astrophysical Journal*, 394, 574
- Brook, P. R., Karastergiou, A., Buchner, S., et al. 2014, *Astrophysical Journal, Letters*, 780, L31
- Burgay, M., D'Amico, N., Possenti, A., et al. 2003, *Nature*, 426, 531
- Burke-Spolaor, S., Bailes, M., Johnston, S., et al. 2011, *Monthly Notices of the RAS*, 416, 2465
- Camilo, F., Cognard, I., Ransom, S. M., et al. 2007, *ApJ*, 663, 497
- Camilo, F., Ransom, S. M., Chatterjee, S., Johnston, S., & Demorest, P. 2012, *Astrophysical Journal*, 746, 63

- Campana, S., Lodato, G., D'Avanzo, P., et al. 2011, *Nature*, 480, 69
- Chandrasekhar, S. 1931, *Astrophysical Journal*, 74, 81
- Chukwude, A. E. 2002, *African Skies*, 7, 41
- Cognard, I., Shrauner, J. A., Taylor, J. H., & Thorsett, S. E. 1996, *Astrophysical Journal, Letters*, 457, L81
- Coles, W., Hobbs, G., Champion, D. J., Manchester, R. N., & Verbiest, J. P. W. 2011, *Monthly Notices of the RAS*, 418, 561
- Cordes, J. M. & Shannon, R. M. 2008, *Astrophysical Journal*, 682, 1152
- Debes, J. H. & Sigurdsson, S. 2002, *Astrophysical Journal*, 572, 556
- Demorest, P. B., Ferdman, R. D., Gonzalez, M. E., et al. 2013, *Astrophysical Journal*, 762, 94
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, 467, 1081
- Drake, F. D. & Craft, H. D. 1968, *Nature*, 220, 231
- DuPlain, R., Ransom, S., Demorest, P., et al. 2008, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7019, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 1
- Espinoza, C. M., Guillemot, L., Çelik, Ö., et al. 2013, *Monthly Notices of the RAS*, 430, 571
- Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, *Monthly Notices of the RAS*, 414, 1679
- Farihi, J., Brinkworth, C. S., Gänsicke, B. T., et al. 2011, *Astrophysical Journal, Letters*, 728, L8
- Faucher-Giguère, C.-A. & Kaspi, V. M. 2006, *Astrophysical Journal*, 643, 332
- Faulkner, A. J., Stairs, I. H., Kramer, M., et al. 2004, *Monthly Notices of the RAS*, 355, 147
- Foster, R. S., Wolszczan, A., & Camilo, F. 1993, *Astrophysical Journal, Letters*, 410, L91
- Gaensler, B. M., Stappers, B. W., Frail, D. A., & Johnston, S. 1998, *Astrophysical Journal, Letters*, 499, L69
- Gavriil, F. P. & Kaspi, V. M. 2002, *Astrophysical Journal*, 567, 1067

REFERENCES

- Gavriil, F. P. & Kaspi, V. M. 2004, *Astrophysical Journal, Letters*, 609, L67
- Gavriil, F. P., Kaspi, V. M., & Woods, P. M. 2002, *Nature*, 419, 142
- Gil, J., Gronkowski, P., & Rudnicki, W. 1984, *Astronomy and Astrophysics*, 132, 312
- Gil, J. A., Jessner, A., Kijak, J., et al. 1994, *Astronomy and Astrophysics*, 282, 45
- Gil, J. A., Kijak, J., & Seiradakis, J. H. 1993, *Astronomy and Astrophysics*, 272, 268
- Goldreich, P. & Julian, W. H. 1969, *Astrophysical Journal*, 157, 869
- Gotthelf, E. V., Gavriil, F. P., Kaspi, V. M., Vasisht, G., & Chakrabarty, D. 2002, *Astrophysical Journal, Letters*, 564, L31
- Hankins, T. H. & Eilek, J. A. 2007, *Astrophysical Journal*, 670, 693
- Harry, G. M. & LIGO Scientific Collaboration. 2010, *Classical and Quantum Gravity*, 27, 084006
- Hermsen, W., Hessels, J. W. T., Kuiper, L., et al. 2013, *Science*, 339, 436
- Hobbs, G., Jenet, F., Lee, K. J., et al. 2009, *Monthly Notices of the RAS*, 394, 1945
- Hobbs, G., Lyne, A. G., & Kramer, M. 2010, *Monthly Notices of the RAS*, 402, 1027
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *Monthly Notices of the RAS*, 369, 655
- Holsclaw, T., Sans, B., Lee, H. K. H., et al. 2013, *Technometrics*, 55, 57
- Hulse, R. A. & Taylor, J. H. 1975, *Astrophysical Journal, Letters*, 195, L51
- Janssen, G., Hobbs, G., McLaughlin, M., et al. 2015, in *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 37
- Johnston, S., Karastergiou, A., & Willett, K. 2006, *MNRAS*, 369, 1916
- Johnston, S., Kramer, M., Karastergiou, A., et al. 2007, *MNRAS*, 381, 1625
- Johnston, S., Manchester, R. N., Lyne, A. G., Nicastro, L., & Spyromilio, J. 1994, *Monthly Notices of the RAS*, 268, 430
- Johnston, S. & Romani, R. W. 2003, *Astrophysical Journal, Letters*, 590, L95

- Johnston, S. & Romani, R. W. 2004, in IAU Symposium, Vol. 218, Young Neutron Stars and Their Environments, ed. F. Camilo & B. M. Gaensler, 315
- Johnston, S. & Weisberg, J. M. 2006, Monthly Notices of the RAS, 368, 1856
- Karastergiou, A., Hotan, A. W., van Straten, W., McLaughlin, M. A., & Ord, S. M. 2009, Monthly Notices of the RAS, 396, L95
- Karastergiou, A. & Johnston, S. 2006, MNRAS, 365, 353
- Karastergiou, A. & Johnston, S. 2007, Monthly Notices of the RAS, 380, 1678
- Karastergiou, A., Roberts, S. J., Johnston, S., et al. 2011, Monthly Notices of the RAS, 415, 251
- Kaspi, V. M., Bailes, M., Manchester, R. N., et al. 1997, Astrophysical Journal, 485, 820
- Kaspi, V. M. & Gavriil, F. P. 2003, Astrophysical Journal, Letters, 596, L71
- Kaspi, V. M., Gavriil, F. P., Chakrabarty, D., Lackey, J. R., & Muno, M. P. 2001, Astrophysical Journal, 558, 253
- Kaspi, V. M. & Helfand, D. J. 2002, in Astronomical Society of the Pacific Conference Series, Vol. 271, Neutron Stars in Supernova Remnants, ed. P. O. Slane & B. M. Gaensler, 3
- Keith, M. J., Johnston, S., Weltevrede, P., & Kramer, M. 2010, Monthly Notices of the RAS, 402, 745
- Keith, M. J., Shannon, R. M., & Johnston, S. 2013, Monthly Notices of the RAS, 432, 3080
- Koester, D. & Wilken, D. 2006, Astronomy and Astrophysics, 453, 1051
- Komesaroff, M. M., Morris, D., & Cooke, D. J. 1970, Astrophysics Letters, 5, 37
- Kouveliotou, C., Dieters, S., Strohmayer, T., et al. 1998, Nature, 393, 235
- Kouveliotou, C., Strohmayer, T., Hurley, K., et al. 1999, Astrophysical Journal, Letters, 510, L115
- Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006, Science, 312, 549
- Kramer, M., Wielebinski, R., Jessner, A., Gil, J. A., & Seiradakis, J. H. 1994, Astronomy and Astrophysics, Supplement, 107, 515
- Kramer, M., Xilouris, K. M., Lorimer, D. R., et al. 1998, Astrophysical Journal, 501, 270

REFERENCES

- Liu, K., Yue, Y. L., & Xu, R. X. 2007, *Monthly Notices of the RAS*, 381, L1
- Lorimer, D. & Kramer, M. 2005, *Cambridge observing handbooks for research astronomers*; 4, Vol. 4, *Handbook of pulsar astronomy* (Cambridge [ua]: Cambridge Univ. Press), IX, 301 S.: Ill., graph. Darst.
- Lorimer, D. R., Lyne, A. G., McLaughlin, M. A., et al. 2012, *Astrophysical Journal*, 758, 141
- Lyne, A., Hobbs, G., Kramer, M., Stairs, I., & Stappers, B. 2010, *Science*, 329, 408
- Lyne, A. G., Burgay, M., Kramer, M., et al. 2004, *Science*, 303, 1153
- Lyne, A. G. & Graham-Smith, F. 2012, *Pulsar astronomy*; 4th ed., *Cambridge monographs on astrophysics* (Cambridge: Cambridge Univ. Press)
- Lyne, A. G. & Manchester, R. N. 1988, *Monthly Notices of the RAS*, 234, 477
- Lyne, A. G., McLaughlin, M. A., Keane, E. F., et al. 2009, *Monthly Notices of the RAS*, 400, 1439
- Lyne, A. G., Pritchard, R. S., Graham-Smith, F., & Camilo, F. 1996, *Nature*, 381, 497
- Lyne, A. G., Pritchard, R. S., & Smith, F. G. 1988, *Monthly Notices of the RAS*, 233, 667
- Manchester, R. N., Hamilton, P. A., & McCulloch, P. M. 1980, *MNRAS*, 192, 153
- Manchester, R. N. & IPTA. 2013, *Classical and Quantum Gravity*, 30, 224010
- Maron, O., Kijak, J., Kramer, M., & Wielebinski, R. 2000, *Astronomy and Astrophysics, Supplement*, 147, 195
- McCulloch, P. M., Hamilton, P. A., Manchester, R. N., & Ables, J. G. 1978, *Monthly Notices of the RAS*, 183, 645
- McLaughlin, M. A. 2013, *Classical and Quantum Gravity*, 30, 224008
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, *Nature*, 439, 817
- Noutsos, A., Karastergiou, A., Kramer, M., Johnston, S., & Stappers, B. W. 2009, *MNRAS*, 396, 1559
- Oppenheimer, J. R. & Volkoff, G. M. 1939, *Physical Review*, 55, 374
- Radhakrishnan, V. & Cooke, D. J. 1969, *Astrophysics Letters*, 3, 225

- Radhakrishnan, V. & Shukre, C. S. 1985, in *Supernovae, their Progenitors and Remnants*, ed. G. Srinivasan & V. Radhakrishnan, 155
- Rankin, J. M. 1993, *Astrophysical Journal*, 405, 285
- Rankin, J. M., Ramachandran, R., & Suleymanova, S. A. 2005, *Astronomy and Astrophysics*, 429, 999
- Rasmussen, C. E. & Williams, C. K. I. 2005, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)* (The MIT Press)
- Redman, S. L., Wright, G. A. E., & Rankin, J. M. 2005, *Monthly Notices of the RAS*, 357, 859
- Ritchings, R. T. 1976, *Monthly Notices of the RAS*, 176, 249
- Roberts, S., Osborne, M., Ebdon, M., et al. 2012, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371
- Romani, R. W. & Johnston, S. 2001, *Astrophysical Journal, Letters*, 557, L93
- Sathyaprakash, B. S. & Schutz, B. F. 2009, *Living Reviews in Relativity*, 12, 2
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *Monthly Notices of the RAS*, 390, 192
- Shannon, R. M., Cordes, J. M., Metcalfe, T. S., et al. 2013, *Astrophysical Journal*, 766, 5
- Shao, L., Caballero, R. N., Kramer, M., et al. 2013, *Classical and Quantum Gravity*, 30, 165019
- Smith, F. G. 1973, *Monthly Notices of the RAS*, 161, 9P
- Somiya, K. 2012, *Classical and Quantum Gravity*, 29, 124007
- Stairs, I. H., Lyne, A. G., & Shemar, S. L. 2000, *Nature*, 406, 484
- Stappers, B. W., Archibald, A. M., Hessels, J. W. T., et al. 2014, *Astrophysical Journal*, 790, 39
- Stinebring, D. R., Smirnova, T. V., Hankins, T. H., et al. 2000, *Astrophysical Journal*, 539, 300
- Taylor, J. H. 1992, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 341, 117
- Taylor, J. H. & Huguenin, G. R. 1971, *Astrophysical Journal*, 167, 273

REFERENCES

- Taylor, J. H. & Weisberg, J. M. 1982, *Astrophysical Journal*, 253, 908
- Thorsett, S. E. 1992, in *IAU Colloq. 128: Magnetospheric Structure and Emission Mechanics of Radio Pulsars*, ed. T. H. Hankins, J. M. Rankin, & J. A. Gil, 143
- Thorsett, S. E., Arzoumanian, Z., Camilo, F., & Lyne, A. G. 1999, *Astrophysical Journal*, 523, 763
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, 556, A2
- van Ommen, T. D., D'Alessandro, F., Hamilton, P. A., & McCulloch, P. M. 1997, *MNRAS*, 287, 307
- van Straten, W. 2006, *Astrophysical Journal*, 642, 1004
- Verbiest, J. P. W., Bailes, M., Coles, W. A., et al. 2009, *Monthly Notices of the RAS*, 400, 951
- Wang, N., Manchester, R. N., & Johnston, S. 2007, *Monthly Notices of the RAS*, 377, 1383
- Wang, N., Manchester, R. N., Zhang, J., et al. 2001, *Monthly Notices of the RAS*, 328, 855
- Wang, Z., Chakrabarty, D., & Kaplan, D. L. 2006, *Nature*, 440, 772
- Weltevrede, P., Edwards, R. T., & Stappers, B. W. 2006, *Astronomy and Astrophysics*, 445, 243
- Weltevrede, P. & Johnston, S. 2008, *Monthly Notices of the RAS*, 391, 1210
- Weltevrede, P., Johnston, S., & Espinoza, C. M. 2011, *Monthly Notices of the RAS*, 411, 1917
- Weltevrede, P., Johnston, S., Manchester, R. N., et al. 2010, *Publications of the Astron. Soc. of Australia*, 27, 64
- Weltevrede, P., Wright, G. A. E., & Stappers, B. W. 2007, *Astronomy and Astrophysics*, 467, 1163
- Wiringa, R. B., Fiks, V., & Fabrocini, A. 1988, *Physical Review C*, 38, 1010
- Wolszczan, A. & Frail, D. A. 1992, *Nature*, 355, 145
- Woods, P. M., Kouveliotou, C., Göğüş, E., et al. 2002, *Astrophysical Journal*, 576, 381
- Xilouris, K. M., Kramer, M., Jessner, A., Wielebinski, R., & Timofeev, M. 1996, *Astronomy and Astrophysics*, 309, 481

REFERENCES

- You, X. P., Hobbs, G., Coles, W. A., et al. 2007, Monthly Notices of the RAS, 378, 493
- Young, N. J., Weltevrede, P., Stappers, B. W., Lyne, A. G., & Kramer, M. 2015, ArXiv e-prints
- Zhuravlev, V. I., Popov, M. V., Soglasnov, V. A., et al. 2013, Monthly Notices of the RAS, 430, 2815