

The highest precision pulsar timing

Stefan Osłowski

Presented in fulfillment of the requirements
for the degree of Doctor of Philosophy

2013

Faculty of Information and Communication Technology
Swinburne University

Abstract

This thesis describes the recent developments in the understanding of limitations of highest precision pulsar timing, focusing on the nearby, bright, binary millisecond pulsar PSR J0437–4715. Pulsar timing is a fruitful astrophysical tool that enables the study of various phenomena, including, but not limited to, testing the general theory of relativity, detecting irregularities in terrestrial time standards and attempting the first direct detection of a stochastic background of gravitational waves. The wealth of information already derived from precision timing of millisecond radio pulsars has led many authors to predict the kind of pulsar timing science possible with the Square Kilometre Array (SKA) by linearly extrapolating current sensitivities.

However, the theoretical limits of pulsar timing have been traditionally based on the radiometer equation in the limit of low signal-to-noise ratio. The timing precision predicted in this way has never been achieved for PSR J0437–4715. The reason for the discrepancy between prediction and reality has long been unknown. Many potential causes have been proposed, such as polarisation calibration, frequency dependence of the pulse profile, hardware and software faults or radio frequency interference. What has been neglected is the self-noise, that is noise arising from the stochastic nature of the pulsar emission itself. At high signal-to-noise ratio, pulsar-intrinsic noise violates the basic premises of the template-matching algorithm commonly used for pulsar timing, leading to statistical bias in arrival time estimates.

In this thesis we describe various statistical properties of the pulsar self-noise, their implications for pulsar timing precision, and statistical methods of removing bias due to self-noise in the estimated times of arrival. Applying these methods reduces the uncertainty-weighted standard deviation of post-fit arrival time residuals by up to 40%. We present the results of studying a large number of single-pulses from PSR J0437–4715, including the first direct detection of orthogonally polarised modes in a millisecond pulsar. We provide evidence for the bright pulses originating from coherent curvature emission from solitons in the pulsar’s magnetospheric plasma.

Acknowledgements

This thesis would not be possible without help, expertise, knowledge, and other contributions from a large number of people. First and foremost, I would like to acknowledge all the hard work and continuing support from my supervisors, Matthew Bailes, Willem van Straten and George Hobbs. They all had their unique contributions to this thesis and shaped me as a scientist. Matthew Bailes always showed incredible enthusiasm and vast knowledge of everything pulsar. His understanding of the big picture has been a guidance throughout this thesis. I appreciate the fact that he allowed me to work mostly independently. I'm humbled by his ability to always learn from all people and situations. Willem van Straten is very knowledgeable and truly an expert in his field. His ability to quickly grasp all the ideas and develop them mathematically has been very useful throughout my candidature. I also thank him for being very supportive, both in professional and personal matters and for his always positive and polite attitude. George Hobbs has been incredibly productive and taught me the importance of focus. His obsession with simulating everything has led me to a number of interesting discoveries and provided supporting arguments for many of the conclusions of this thesis.

I would have never completed this thesis without all the people who have shaped my interests in science from the earliest years. My family has ensured that I received a good education and was supportive throughout my whole life. I would have not achieved anything without them. I must thank Stephen Hawking for his book "A brief history of time" with which my interest in astronomy began. My teachers in primary and high school provided me with good background in mathematics and physics which in turn allowed me to undertake studies in physics at Warsaw University. From that period, people who most influenced my scientific career are Tomasz Bulik, Rafał Moderski, Dorota Gondek-Rosińska, Grzegorz Pietrzyński, Krzysztof Belczynski, and Sylwester Arabas. I would also like to express my gratitude to all the people at Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, Ludwig-Maximilians University, and the European Southern Observatory headquarters in Garching and Munich for expanding my general knowledge of astronomy and social interactions. Special thanks go to Guinevere Kauffmann, Simon White, Rashid Sunyaev, Vivienne Wild, Dimitri Gadotti, Cheng Li, Victor Aguirre Silva, Martin Henze, Katie Dodds-Eden, Francesco Montesano, Anna Raiter, Arek Swat, Till Sawala and many others. I decided to come to Swinburne thanks to Katie Dodds-Eden, Patricia Figueiro Spinelli, and Joris Verbiest, who all encouraged me to do so in one way or another.

Another group of people to whom I owe much are all the students and staff at Swinburne. Special thanks go to Carolyn Cliff and Mandish Webb for ensuring my move to Australia is as smooth as possible, helping me find my way in Melbourne and all the hard work behind the scenes. The pulsar group has provided a very prolific environment. I had many useful discussions with all the group members: Matthew Bailes, Willem van Straten, Andrew Jameson, Ramesh Bhat, Jonathon Kocz, Lina Levin, Sarah Burke-Spolaor, Ben Barsdell, Paul Coster and Emily Petroff. I express my gratitude to radio astronomers in Sydney, especially to Dick Manchester, George Hobbs, Simon Johnston, David Champion, Mike Keith, Ryan Shannon, Jonathan Khoo, Daniel Yardley, Robert Hollow, and Maxim Vornkoff. I had numerous scientific discussions on topics covered in this thesis with people from other institutions as well, in particular with Paul Demorest, Xavi Siemens, Jim Cordes, Willam Coles, Yuri Levin, Joris Verbiest, and Fredrick Jenet. Work in this thesis would not be possible without all the staff at Parkes Radio Observatory. I thank the staff of the Australia Telescope Compact Array as well.

I would like to thank my parents, Swinburne University of Technology, Australia Telescope National Facility (now CSIRO Astronomy and Space Sciences), CAASTRO, and Australian Research Council grants # DP0985272 and DP0878388 for financial support of my thesis, observing, and travels. I am grateful to all people involved in the organisation of my trips and the conferences that I attended. In particular my travels in 2011 to Sydney, Berkeley, and Bonn warrant special thanks to Dan Wertheimer and his wife, Gregory Desvignes, Michael Kramer, Gabi Breuer, Paulo Freire, David Champion, Cherry Ng, Joris Verbiest, Ewan Barr, Masooma Ali, and all other people in Bonn, too numerous to list them here.

Many people have helped me stay sane throughout my PhD and I have spent countless hours discussing with them, relaxing, playing sports and many other activities. Special thanks go to my family, Dominik Sznajder, Tomasz Śliwowski, Łukasz Osuch, Lina Levin, Sarah Burke-Spolaor and her husband Max, Evelyn Caris, Ewa and Stan Walkowiak, Daniel Nugent, Glenn Kacprzak, Victor Silva Aguirre, Dimitri Gadotti, Jan and Martyna Wróbel, Paul Coster, Ben Barsdell, Emily Petroff, Cédric Collet, Jonathan Khoo, Vikram Ravi, all the students at Swinburne, and many others. All the above lists are in mostly random order and by no means complete.

Statement of originality

The work presented in this thesis has been carried out in the Centre for Astrophysics & Supercomputing at the Swinburne University of Technology; Australia Telescope National Facility, CSIRO in Sydney; Radio Astronomy Lab at University of California in Berkeley; and Max Planck Institute for Radioastronomy in Bonn between 2009 and 2012. This thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, this thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis. The content of the chapters listed below have appeared in, or been submitted to, peer-refereed journals. Minor alterations have been made to the published papers in order to maintain argument continuity and consistency of spelling and style.

- Chapters 1 and 2 contain excerpts from “High signal-to-noise ratio observations and the ultimate limits to precision pulsar timing”, 2011, MNRAS, 418, 1258.
- Chapter 3 has been published as “High signal-to-noise ratio observations and the ultimate limits to precision pulsar timing”. The introduction of this publication has been heavily edited to improve the flow of the thesis and the original introduction is partially included in chapters 1 and 2.
- Chapter 4 has been published as “Improving the precision of pulsar timing through polarization statistics”, 2013, MNRAS, 430, 416. The version present in this thesis has been edited to improve the flow of the thesis.

Stefan Osłowski

Melbourne, Australia

2013



Contribution to co-authored published content

Chapter 3 has been published as Osłowski, van Straten, Hobbs, Bailes, Demorest (2011). Chapter 4 has published as Osłowski, van Straten, Demorest, Bailes (2013). The original motivation for this work comes from M. Bailes, G. Hobbs, and W. van Straten. The methodology presented in both these chapters is an extension of the doctoral work of P. Demorest. Dr Demorest aided me in the initial approach I undertook and we discussed the implications of my findings when writing up the work, which warranted his inclusion in the author list. I discussed the progress of the research described in these chapters during regular meetings with my supervisory team. I wrote the complete first drafts of both papers and the co-authors have edited the manuscripts. Their changes have contributed to improving the structure of sentences and usage of correct English expressions. W. van Straten and P. Demorest have made insightful comments on the impact of ionosphere and mixing between Stokes Q and U parameters which enabled me to finalise this work. I performed all of the code development required for me to extend the public domain library PSRCHIVE to create the results for this thesis.

Stefan Osłowski

*Osłowski*²⁰¹³

Matthew Bailes

2013

Paul Demorest

2013

George Hobbs

2013

Willem van Straten

2013

Contents

Abstract	i
Acknowledgements	iii
Declaration	vi
List of Figures	xiii
List of Tables	xv
1 Introduction	1
1.1 From the Big Bang to neutron stars	1
1.1.1 Origin of elements and stars	1
1.1.2 Where no degenerate electron has gone	3
1.1.3 Neutron stars as radio pulsars	4
1.1.4 Neutron stars in binaries and millisecond pulsars	6
1.1.5 Pulsar properties	7
1.1.6 Science with pulsars	11
1.2 General relativity	13
1.2.1 On the equivalence of observers	13
1.2.2 Early tests of general relativity	14
1.2.3 Gravitational waves	14
1.2.4 Pulsar timing and tests of relativistic gravity	16
1.2.5 Pulsars as physics laboratories and tests of gravitation	24
1.2.6 PSR J0437–4715 and implications for the future of pulsar timing	26
1.3 Structure of this thesis	27
2 The fundamentals of radio pulsar observational techniques	29
2.1 Radio waves	29
2.1.1 The symmetry of nature	29
2.1.2 Charges in motion - and a new physical entity	32
2.2 Signal path	36
2.2.1 Antenna	37
2.2.2 Feed horn, radiometer	38
2.2.3 Two stages of downconversion	40
2.3 Hardware and software backends	41

2.3.1	Analogue filterbanks	42
2.3.2	Autocorrelation spectrometers	43
2.3.3	Polyphase digital filterbanks	43
2.3.4	Coherent dedispersion systems	44
2.3.5	Baseband recorders and offline data processing	45
2.3.6	Impact of polarimetric calibration on timing precision	46
2.3.7	Systems used in this thesis	48
3	Stochastic wideband impulse modulated self-noise	49
3.1	Introduction	49
3.2	Stochastic wideband impulse modulated self-noise summary	51
3.3	Observations and data processing	54
3.4	Method	61
3.4.1	Correcting the timing residuals: multiple regression	63
3.4.2	Simulations	65
3.5	Results	69
3.6	Discussion	72
3.6.1	Discussion of the method	72
3.6.2	Application to PSR J0437–4715	74
3.7	Conclusion	80
4	Improvements in dealing with SWIMS	83
4.1	Introduction	83
4.2	Pulsar timing and the stochastic wideband impulse modulated self-noise	86
4.3	Observations and data processing	87
4.4	Method	90
4.5	Results	92
4.6	Discussion	93
4.6.1	Discussion of the method	98
4.6.2	Application to PSR J0437–4715	100
4.7	Conclusion	103
5	PSR J0437–4715 timing and polarimetry - a single-pulse perspective	105
5.1	Introduction	105
5.2	Observations and data processing	106
5.3	Total intensity properties	109

5.3.1	Distribution of instantaneous signal-to-noise ratio	109
5.3.2	Impact of a single pulse on the time of arrival estimate	111
5.3.3	Average pulse profile dependence on sub-pulse S/N	112
5.3.4	Timing properties as a function of sub-pulse S/N	114
5.3.5	Observed electric field intensities and plasma physics	116
5.3.6	Intensity modulation	119
5.4	Polarimetry	123
5.4.1	Polarisation dependence on S/N	123
5.4.2	Phase resolved polarisation histograms	126
5.5	Conclusions	127
6	Conclusions	133
6.1	Pulsar timing before this thesis	133
6.2	Summary	135
6.3	Future directions	137
	Bibliography	139
A	Some details on general relativity	165
A.1	Derivation of the Einstein's field equation	165
A.2	Solutions to the Einstein's field equations	167
A.3	Classical tests of general relativity	169
A.4	Gravitational waves	171

List of Figures

1.1	The $P-\dot{P}$ diagram	10
2.1	Poincaré sphere and the polarisation ellipse	35
3.1	The high S/N (15,000) template for PSR J0437–4715	55
3.2	Timing residuals for 6 days of data	56
3.3	Comparison of measured timing precision with theory	57
3.4	Greyscale image of a single-pulse from PSR J0437–4715	59
3.5	Timing residuals derived from to independent sub-bands	60
3.6	First eigenvector for simplest SWIMS simulation	66
3.7	Distribution of correlation coefficients	70
3.8	First three measured eigenvectors of the noise covariance	71
3.9	The autocorrelation function of noise	77
4.1	High S/N template of PSR J0437–4715	88
4.2	Timing residuals for 16.79 s sub-integrations	89
4.3	Example of $4N_{\text{bin}}$ representation of pulse profile	91
4.4	First, most significant eigenvector in case of full Stokes pca	94
4.5	Second eigenvector in case of full Stokes pca	95
4.6	Third eigenvector in case of full Stokes pca	96
4.7	Distribution of correlation coefficients in case of full Stokes pca	97
4.8	Impact of the ionosphere on measured PA	102
5.1	High S/N template from CASPSR	108
5.2	Observed instantaneous signal to noise ratio	110
5.3	Running ToA estimate for one minute integrations	111
5.4	Comparison of a single pulse with an average pulse profile	113
5.5	Average pulse profile for selective integration	115
5.6	Phase-resolved intensity distribution	118
5.7	2DFS for APSR data	121
5.8	The average pulse profile of the brightest 1% sub-pulses	122
5.9	Polarisation fraction as a function of S/N_n	124
5.10	Examples of bright sub-pulses	125
5.11	Mollweide projection of the Poincaré sphere - OPMs	127
5.12	Mollweide projection of the Poincaré sphere, part 1	128

5.13 Mollweide projection of the Poincaré sphere, part 2	129
5.14 Mollweide projection of the Poincaré sphere, part 3	130

List of Tables

3.1	Parameters of the multi-component simulation.	68
5.1	Rms timing residual for selective integration	116

1

Introduction

In this chapter we lay out the background information to put the rest of this thesis in context and build a science case for this thesis. We start by providing an historical overview of astrophysical research most directly related to the topic of quickly spinning neutron stars. Pulsars are extremely versatile physics laboratories as demonstrated via the broad variety of applications of pulsar timing methodology. While this thesis does not cover all the relevant topics, the introduction provides a broad overview of pulsar related basic physics and astrophysics. Special focus is placed on theories of gravity, their tests, and the search for the gravitational waves predicted by general relativity. These topics currently drive the majority of the currently ongoing research related to pulsars. We provide a broad overview of many topics as this thesis addresses the issue of the ultimate limits of precision timing and provides means of improving the attained precision, especially in the case of high signal-to-noise ratio observations. These will be especially important in the future when such data can be obtained for a larger number of millisecond pulsars. In this manner, this thesis will contribute to all the areas discussed in the introductory materials.

1.1 From the Big Bang to neutron stars

1.1.1 Origin of elements and stars

The Universe started in the Big Bang. Shortly afterwards the primordial nucleosynthesis began, i.e., the collisions of protons which lead to formation of nuclei. In this way elements up to beryllium were synthesised. No heavier elements could be generated in the process of primordial nucleosynthesis, primarily due to beryllium's very short half-life time. The quick decay of beryllium and lack of stable nuclei at masses 5 and 8 is an impassable

barrier for the primordial nucleosynthesis (Boesgaard & Steigman, 1985). The primordial gas formed clouds due to self gravity. Such a gas cloud becomes unstable when its mass exceeds the Jeans mass and the self-gravity is no longer balanced by the gas pressure (Jeans, 1902). During the collapse, the gas cloud fragments and the clumps compress and heat up, creating the first stars.

The Universe is now full of galaxies, each containing millions or billions of stars, which were long thought to be the main constituent of mass balance in the Universe. Nucleosynthesis can resume its progress inside stars which are much denser than primordial plasmas. In these dense conditions, the nuclear stability gap is closed by the much more frequent collisions of three nuclei of helium-4 (Tout, 2006). In this way, many of the heavy elements are created in stars. The heavy elements fused inside stars constitute an important ingredient of the intelligent life, such as evolved on Earth.

With concentrated effort in the last century, it has become understood that the overwhelming majority of stars are giant balls of primarily hydrogen. The nuclear reactions in the hot cores provide the energy necessary to produce the radiation we see and to sustain the structure of the star itself. Without the gas and radiation pressure, the self-gravity would collapse the stars into smaller and denser objects. The early 1930s brought significant theoretical progress in our understanding of what happens when nucleosynthesis stops due to lack of fuel and/or not high enough density and temperature to initiate burning of heavier elements. Eddington (1926) showed that after nucleosynthesis stops at the core of the star, gravity overcomes the internal pressure and the star starts to collapse. Without an additional source of pressure this collapse would never stop. Fowler (1926) suggested that Pauli’s exclusion principle can prevent this and restore the balance between the gravity and pressure. Pauli’s famous principle states that fermions (particles with half-integer spin) cannot occupy the same quantum state simultaneously (Stoner, 1924; Pauli, 1925). Experiments performed by Fermi (1926) and Dirac (1926) confirmed Pauli’s hypothesis soon after he proposed it.

The energy released during stellar collapse ionises the atoms consequently creating electron gas in the interior of the star. The electrons are fermions and are thus bound by the Pauli exclusion principle. The degenerate electron gas generates additional pressure. This pressure is sufficient to balance the gravity in stars whose zero age main sequence (ZAMS) mass is $M_{\text{ZAMS}} < 8.3M_{\odot}$ ¹ (Belczynski et al., 2008). The resulting stellar remnant is called a “white dwarf”. Many such objects, with radii of several thousand kilometres, are observable today as the energy released during the collapse has heated the star enabling

¹ZAMS masses quoted assume solar metallicity. We note that in practice the fate of a star is determined by the mass of its core but the ZAMS mass is a very good indicator overall.

the end product to be observed in optical or infra-red bands until they eventually fade away.

1.1.2 Where no degenerate electron has gone

What happens if the star is initially heavier than $\sim 8.3M_{\odot}$, or if the mass of the white dwarf exceeds the Chandrasekhar limit whose modern value is $1.44M_{\odot}$ (Chandrasekhar, 1931; Landau, 1932)? When the additional pressure of degenerate electrons is not sufficient to counteract the pull of gravity, the collapse of the stellar core continues. Once the density of the collapsing core is sufficiently high, the electrons can cause the inverse beta decay of protons (Wick, 1933; Alvarez, 1937) producing a neutron and electron neutrino. The latter particle easily escapes the core of the collapsing star. This neutrino outflow accelerates the collapse of the star by carrying away the thermal energy. The neutrons on the other hand remain in the core and eventually constitute the neutron superfluid after Bose condensation (Baym et al., 1971). As neutrons also have half-integer spin they follow the Fermi-Dirac statistic and the Pauli exclusion principle. Similarly to electrons, the degenerate neutrons provide additional pressure that can overcome gravity and the stellar remnant reaches a stable configuration, provided the ZAMS mass was smaller than $\sim 21M_{\odot}$. The resulting object is called a neutron star. Stars that are heavier at the beginning of their evolution end up as black holes (e.g., Belczynski et al., 2008), i.e., regions of space from which light cannot escape.

Neutron stars as dense, collapsed end-products of stellar evolution had been predicted as early as 1933 by a Swiss astronomer, Fritz Zwicky (Baade & Zwicky, 1934a,b,c). A few years later, a first quantitative model of such stars was presented by Oppenheimer & Volkoff (1939). After that the scientific community forgot about Zwicky's ideas for years. In 1967 sources of regular radio pulsations were discovered by Jocelyn Bell-Burnell (Hewish et al., 1968), a discovery for which her supervisor, Anthony Hewish, was awarded the Nobel Prize in physics in 1974. The regular pulsations were initially considered to be of artificial origin and hence termed "Little Green Men". It was not until the work of Gold (1968) that the connection between theory and observations was made. He also made a prediction that the pulse period should become longer with time. An observational confirmation followed soon (Boynton et al., 1969) and the quickly spinning, constantly slowing down stars were dubbed "pulsars", short for "pulsating star". Since then many searches for pulsars have been undertaken using steadily improving techniques with the current pulsar population exceeding 2000 objects. Pulsar searching is an extensive topic, beyond the scope of this thesis. For an outline of current techniques and a brief history

of the previous surveys see Keith et al. (2010) and references therein. For the sake of historical completeness, we note that the discovery of pulsars has caused some physicists to finally consider existence of black holes as real objects rather than just a theoretical construct.

Neutron stars are very extreme objects. They are likely highly magnetised, with typical magnetic field strengths of 10^{12}G , but up to 10^{14}G , a quadrillion times more powerful than the Earth's magnetic field. In fact, these are the strongest magnetic fields known and strongly affect the physics of matter. For example, in the presence of such a strong field, the ground energy state of hydrogen is increased 20-fold. The origin of such strong fields is unknown. One of the proposed mechanism for their creation is the freezing of the progenitor star's magnetic field and a dynamo mechanism (Duncan & Thompson, 1992). As we noted previously, the neutron stars are very dense with a non-uniform density profile. The outermost layer of a neutron star contains normal nuclei, mostly iron, and a highly thermally and electrically conductive, degenerate electron gas below the surface (Beskin, 1999). Toward the centre, the pressure increases until it becomes energetically favourable for the neutrons to drip out of atomic nuclei. The strong nuclear forces no longer hold neutrons in the cores of atoms and neutrons form the so called neutron drip layer. At yet slightly higher densities the neutrons become the dominant particle with roughly 8 neutrons for every proton which survived the inverse beta decay. While approaching the nuclear density, the nucleons start forming strands and sheets (so called spaghetti and lasagna; Lattimer & Prakash, 2004). Below that layer our understanding of physics becomes less and less certain and no experimental data are available to study matter at densities exceeding nuclear density. Many common hypotheses for the contents of the very core of neutron stars include strange matter or hyperons (Glendenning, 1985); quark-gluon plasmas (Collins & Perry, 1975); or pion condensates (Kaplan & Nelson, 1986).

1.1.3 Neutron stars as radio pulsars

Once a star reaches the stage of a neutron star, it does not evolve any more as a star, except by cooling down, and can become a radio pulsar. It can also produce visually striking gas distribution heated up from the inside, known as the pulsar wind nebulae, for more details on this topic see Rees & Gunn (1974); Gaensler & Slane (2006). A mechanism for the emission of radio waves is necessary to turn a neutron star into a radio pulsar. The first theoretical works proposed that the radiation originated from charged particles accelerated in the strong magnetic field of neutron stars (Gold, 1968; Pacini, 1968; Goldreich & Julian, 1969). The pulsation is caused by a "lighthouse effect" due to the misalignment of

the emission axis from the spin axis and because the emission is beamed. In other words, the pulsar is thought to emit two beams sweeping across the space-time. When this radio beam crosses the Earth, we see the pulsation with the observed intensity varying over the pulse phase as the beam enters and disappears from our view. The details of the emission mechanism are not yet understood. The extreme observed brightness temperature of the emission, reaching values of $2 \times 10^{41} K$ (Hankins & Eilek, 2007), suggests a coherent plasma process to be responsible for the emission (Melrose, 1992). The first controversial detection of coherency in pulsar emission was claimed by Jenet et al. (2001b) with his results contested by Smits et al. (2003) as inconclusive as the coherence can be interpreted as a result of interstellar scintillation. Many other mechanisms are proposed in the literature. Some of them are linear processes, such as: linear acceleration and maser curvature emission (Luo & Melrose, 1995; Melrose, 1996, 2003); relativistic plasma emission (Melrose, 1996; Asseo, 1996); and a streaming instability into an escaping mode (Gedalin et al., 2002). Other groups have proposed non-linear processes like soliton collapse (Asseo & Porzio, 2006); or the antenna mechanism (Pottelette et al., 1992; Cairns & Robinson, 2000). A variety of less accepted proposals have been brought forward such as the superluminal polarisation currents (not violating causality) observed in laboratories (Bessarab et al., 2006, and references therein) and proposed by Ardavan et al. (2008) as a viable pulsar emission mechanism. Many pulsars are known to emit in the higher energy bands, most notably in X-rays and gamma-rays. The Fermi satellite successfully exploits the latter emission to discover pulsars at a very high detection rate (e.g., Guillemot et al., 2012).

The discovery of pulsars has undoubtedly been one of the major astrophysical discoveries of the previous century. It was only the start of rapid developments in the fascinating field of pulsar research. Many discoveries followed in the years after the initial detection. Six years after the initial detection, Hulse & Taylor (1975) found the first binary system with a neutron star. This feat was rewarded with the Nobel Prize in physics in 1993 for the first measurement of relativistic orbital shrinkage of the order of 1 cm a day, first evidence for gravitational waves carrying energy. Another remarkable discovery occurred in the next decade when Backer et al. (1982) discovered a pulsar spinning at a rate of 642 rotations per second. It was the first member of the millisecond pulsar class and remained the fastest spinning neutron star known for many years. Objects in this class are now central to all high precision timing experiments, in particular pulsar timing arrays (PTAs) pursuing the detection of gravitational waves via their effect on the metric. The first exoplanet was discovered around a pulsar, three years before a discovery using any other methods (Wolszczan & Frail, 1992). Finally, a system with two neutron stars

orbiting each other was found (Hulse & Taylor, 1975), one of ten such systems found so far. Another such system, discovered by Marta Burgay and her collaborators, is special as both stars were visible as radio pulsars at the time of discovery (Burgay et al., 2003; Lyne et al., 2004) which enables many physical experiments, in particular tests of gravitational physics. Many smaller steps in the ongoing quest for knowledge accompanied the listed milestone discoveries every decade. Pulsar astronomy remains a very active field of research.

1.1.4 Neutron stars in binaries and millisecond pulsars

As mentioned at the end of previous section, some pulsars are very rapid rotators. They spin with very short periods of the order of milliseconds and hence are dubbed the millisecond pulsars (MSPs). For example, PSR J1748–2446ad, the fastest pulsar known, rotates at a rate of 716 rotations per second (Hessels et al., 2006). In this section we describe the origin of MSPs. These pulsars are believed to be “recycled” normal pulsars, i.e., neutron stars spun up by the transfer of angular momentum from the matter in an accretion disk. Where is the matter forming the disk coming from? To answer this question we have to take a few steps back and fill in the gaps in our brief story about stellar evolution.

A significant fraction of stars are born in binary or higher order systems where, in case of binaries, two stars are orbiting a common centre of mass. This fraction is highly uncertain with the value 0.5 assumed for many years (Jaschek & Gómez, 1970). More recent findings suggest much smaller values (Lada, 2006) as vast numbers of solitary very low-mass main sequence stars are discovered. However, this result has no direct impact when discussing pulsars in binaries, as the low-mass stars are not progenitors of neutron stars. Thus, the fraction of neutron star progenitors in binaries is still very high.

If two stars are evolving in each other’s neighbourhood, this can have a severe impact on their evolution. Stars change their sizes during their life as the balance between gravity and pressure is affected by changes of the source of energy inside the star. If a star increases its size in a binary system, it may completely fill its Roche lobe, the volume of space bordered by a surface at which the gravitational pull of the two stars in the system balances out. When one of the stars completely fills its Roche lobe, the matter can easily transfer via the Lagrangian point between the two stars and flow towards the companion in a stable manner (Kuiper, 1941). When the companion is a neutron star, the accretion provides energy for the pulsar to emit X-ray photons. Another, much more poorly understood phase in binary evolution is the so called common envelope, related to the aforementioned Roche lobe overflow. The flow of matter can become unstable due

to gravitational perturbations of the outflowing matter (Webbink, 1984). This instability causes the mass transfer rate to increase rapidly, the stars become engulfed by a common gas envelope and the orbit of the binary shrinks on a time-scale of years. For a classical review of binary evolution see Paczyński (1971).

If the mass transfer is stable, an accretion disk will form around a neutron star, producing an X-ray binary and spinning up a neutron star. If the neutron star is spinning faster than the matter at the Alfvén radius, where the ram pressure is balanced by the outwards magnetic pressure from the neutron star, then the pulsar can eject the matter from accretion disk and, possibly slow down. This process is known as the propeller effect (Illarionov & Sunyaev, 1975) but the details of how it happens are still highly uncertain. There is both observational evidence in its support (Cui, 1997) as well as theoretical concerns about its efficiency at lower accretion rates (Romanova et al., 2004; D’Angelo & Spruit, 2010). A neutron star can be either spun up or down during an accretion event, depending on the angular momentum flux (Ghosh et al., 1977; Ghosh & Lamb, 1979a,b). As evidenced by existence of the MSPs, the spin-up mechanism can be very effective and produce MSPs. The magnetic field of the pulsar probably decays during accretion which buries the magnetic field (see Osłowski et al., 2011a, and references therein). This fact is of consequence as the magnetic field is not only responsible for the radio emission and the lighthouse effect, but it also controls how quickly the pulsar loses its rotational energy.

The mass of the heavier star in the binary is determined by an initial mass function, approximately a power law $\sim M^{-2.7}$ (e.g., Scalo, 1986). The mass of the other star is a fraction of the heavier star’s mass assumed to be distributed uniformly (Kuiper, 1935). The main sequence (i.e., hydrogen burning) lifetime of a star is related to its mass and so the two stars in the binary evolve at different speeds. The heavy stars evolve much faster, due to the rate of nuclear reactions increased by the higher density and temperature in their core. With the binary interactions described above, the picture gets more complicated, as the stars can interact and exchange mass, affecting their evolutionary stage and fate as compared to their fate predicted by ZAMS mass. The situation of most interest in the context of this thesis is when one of the two stars explodes in a supernova explosion and, despite the natal kick due to the explosion’s asymmetry (Scheck et al., 2004), the newly born neutron star remains gravitationally bound to its companion.

1.1.5 Pulsar properties

Neutron stars are very simple objects at a first glance but a closer look reveals many more details than we have presented so far. To first order, pulsars are quickly spinning magnets

in a vacuum. As was already predicted by the early works of Gold (1968) and Pacini (1968), pulsars are losing energy as the rotating magnetic field produces a torque acting against the pulsar's rotation. Since we typically assume that a neutron star has a dipolar structure of the magnetic field (Ostriker & Gunn, 1969), the amount of energy radiated by a magnetised neutron star is readily available from classical electrodynamics² (Jackson, 1975):

$$\dot{E} = \frac{32\pi^4}{3c^3} B^2 \frac{R^6}{P^4} \sin^2 \alpha, \quad (1.1)$$

where α is the angle between the rotation and magnetic field axis, B is the magnetic field of a dipole at radius of the neutron star R , and c is the speed of light. By comparison with the rate of loss of rotational energy $\dot{E}_{\text{rot}} = 4\pi^2 I \dot{P} P^{-3}$, we obtain the estimate of the magnetic field strength on the pulsar's surface:

$$B = \sqrt{\frac{3Ic^3 \dot{P} P}{8\pi^2 R^6 \sin^2 \alpha}}, \quad (1.2)$$

where I is the moment of inertia of the neutron star. This formula is often used to calculate the magnetic field of observed radio pulsars but we note that the values obtained from it are only estimates as this equation assumes certain field geometry and a braking mechanism. We can invert this formula and calculate the dependence of spin down angular frequency $\Omega = 2\pi P^{-1}$ on the pulsar properties:

$$\dot{\Omega} = -\frac{2B^2 R^6 \sin^2 \alpha}{3c^3 I} \Omega^3. \quad (1.3)$$

The dependence of the rate of change of spin frequency in equation 1.3 on the spin frequency to the third power is a direct consequence of assuming the dipolar magnetic field. The index of this power law dependence is known as the braking index n and is observed to be anywhere between 2.5 and 3.5 for the small fraction of pulsars for which it has been measured (Manchester et al., 2005). The above spin model determines the characteristic pulsar age T , by assuming an initial spin period P_{ini} , constant value of the magnetic field and of the inclination angle, and that the pulsar's braking index is known and has been constant throughout pulsar's life:

$$T = \frac{P}{(n-1)\dot{P}} \left[1 - \left(\frac{P_{\text{ini}}}{P} \right)^{(n-1)} \right] \quad n=3; \underline{P_{\text{ini}}} \ll P \quad \frac{P}{2\dot{P}}. \quad (1.4)$$

The values obtained from the above equation are typically highly uncertain, due to the

²in this section we use cgs units

difficulties with measuring the braking index and a large uncertainty in P_{ini} . The initial spin periods are believed to be between 10 and 100 milliseconds (Faucher-Giguère & Kaspi, 2006). The initial value of magnetic field is not known very well either but in population synthesis models observations can be reproduced by uniformly distributing the magnetic field in the range $10^{11}\text{G} \leq B \leq 10^{13}\text{G}$ (e.g., Osłowski et al., 2011a).

The issue of magnetic field evolution with time is still highly controversial. Many mechanisms have been introduced to explain an exponential decay of the magnetic field, including the widely accepted Hall-like effect with ohmic dissipation (e.g., Romani, 1990; Geppert & Urpin, 1994). The time-scale of the field decay has been studied in detail by population synthesis and the results suggest values from a few Myr (Gonthier et al., 2002, 2004) up to 2 Gyr (Kiel et al., 2008). Gonthier et al. (2009, 2011) concluded that a successful model used for generating a synthetic population must either correlate the initial magnetic field with P_{ini} or use a very short decay time-scale. In addition, Zhang & Kojima (2006) have shown that there is a lower limit to the strength of the magnetic field of the order of 10^8 G and no pulsars are known with significantly lower values of magnetic field. During the accretion of matter, the magnetic field can be buried (Cumming et al., 2001, e.g.) and the measured strength may be reduced exponentially with the amount of mass accreted (Osłowski et al., 2011a, and references therein).

The magnetic field, spin period and the spin period time derivative are important characteristics of a pulsar. All the other parameters in the equations 1.2 and 1.3 have typically assumed values: radius R of 10 km; moment of inertia $I = 10^{45}$ g cm². An easy way to present the population of pulsars is to plot them on either a B - P or P - \dot{P} diagram. Here we use the latter representation as shown in Fig. 1.1 based on the ATNF pulsar catalogue³ (Manchester et al., 2005). The bulk of the pulsar population in this space is in the upper right part of this diagram with the majority of objects spinning with periods between 0.1 and 1 seconds. These are the so called classical pulsars. They have magnetic fields in the same range as the initial magnetic fields and typically do not have a companion star. The scarcely populated top right corner is the region occupied by magnetars - pulsars with extremely high magnetic fields with strengths of the same order as critical field value of 4.414×10^{13} G. At this value of magnetic field, the cyclotron energy levels become quantised. The main source of energy of magnetars is the magnetic field rather than the rotational kinetic energy. In the opposite, bottom left corner resides a separate population of MSPs with the majority of objects in binary systems. The origin of the solitary MSPs with periods of the order of 30 ms is not fully understood, one

³<http://www.atnf.csiro.au/research/pulsar/psrcat/>

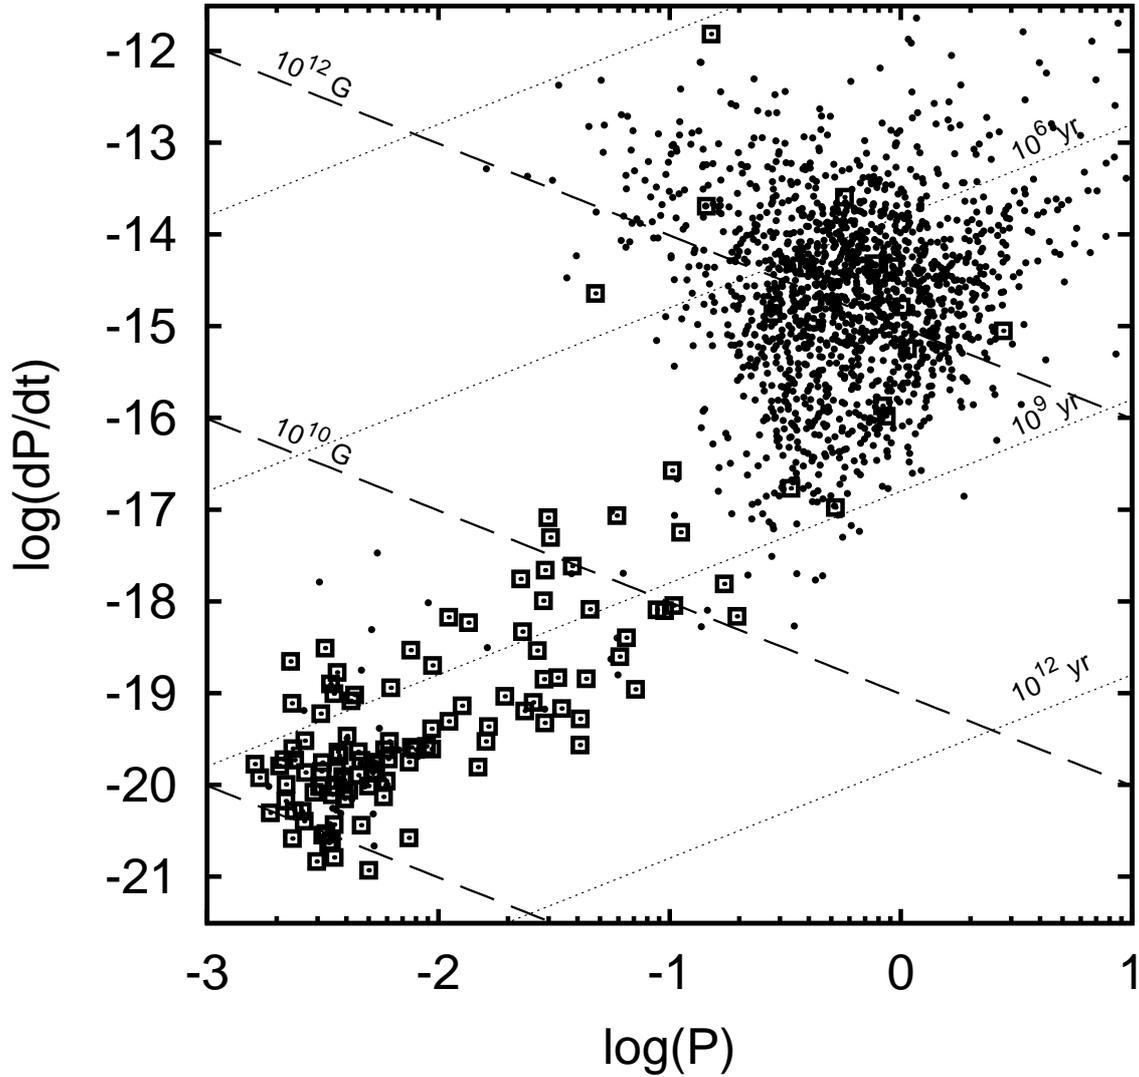


Figure 1.1: The P - \dot{P} diagram demonstrating the distribution of pulsars in the phase space spanned by their spin period and its time derivative. The lines of constant magnetic field and constant age are marked dashed and dotted lines, and based on equations 1.2 and 1.4, respectively. Pulsars in binary systems are marked by squares. The numerical values of period and period derivative were taken in units of s and ss^{-1} , respectively.

potential channel of formation is a binary system that gets disrupted when the second supernova explosion occurs after its progenitor has spun up the already existing pulsar. MSPs have spin down rates much lower than classical pulsars and the derived magnetic field strengths are of the order of 10^8 G. The bottom right region of the figure is empty as the neutron stars which would populate this region cease to emit radio waves and are thus not detectable as radio pulsars. Two theoretical lines, called “death lines” are often calculated, and they correspond to conditions in which the pulsar emission mechanism fails as it is no longer possible to create electron-positron pairs in the magnetosphere, see e.g. Rudak & Ritter (1994). We do not plot these death lines as they are highly uncertain and single cases of pulsars beyond the theoretical lines are known, such as the one described by Young et al. (1999).

As a side note on the naming convention, despite the newly born pulsars’ short spin period we do not call them millisecond pulsars. Only the pulsars in the bottom left corner of the Fig. 1.1 are classified as such. Even in this region, pulsars with periods of tens to hundreds of milliseconds are dubbed mildly recycled pulsars and likely have a different formation history (e.g., van den Heuvel, 1994). For many years the definition of an MSP has been based on arbitrary cuts in this phase space. A recent work by Lee et al. (2012) applies the machine learning algorithm to study the distribution of pulsars in the $P - \dot{P}$ space to obtain the first objective experimental definition of an MSP.

1.1.6 Science with pulsars

Neutron stars are very stable and consistent rotators, spun down at a slow rate by the magnetic torque. As the emission of the pulsar is beamed and the rotation axis is misaligned from the emission axis, the pulsars act like very consistent clocks. While pulsars exhibit a variety of short term instabilities, long-term rotational stability of some MSPs is astounding and challenges that of atomic clocks on time-scales of one year and longer (Matsakis et al., 1997). A wide variety of astrophysical experiments is possible with a number of very stable clocks distributed in the sky and observable via radio telescopes and at other wavelengths. Following Lorimer & Kramer (2004), we list a number of already undertaken and plausible experiments.

As pulsars can be used as clocks, a natural application of pulsar observations is a creation of a pulsar based time-scale as outlined by Petit & Tavella (1996) and Rodin (2008) and implemented by Hobbs et al. (2012). The position and proper motion of the pulsar can be accurately measured from timing of the arriving emission and in the case of PSR J0437–4715 a bright nearby MSP the precision better than $10 \mu\text{as}$ has been

achieved (van Straten et al., 2001). Astrometric measurements derived in this way can be used for constraining physical properties of globular clusters (Phinney, 1992), tying various astronomical reference frames together (Bartel et al., 1996), deriving a pulsar distance scale (Weisberg, 1996), or constraining supernova explosion mechanism (Scheck et al., 2004). First exoplanets were discovered around an MSP by Wolszczan & Frail (1992) and more recently a stellar companion of an MSP transformed into a diamond-like planet has been discovered (Bailes et al., 2011). Pulsar observations allow studies of the interstellar medium and the large scale magnetic fields via their effect on the propagation radio waves (e.g., Freire et al., 2001; Han, 2001; Lyne et al., 2001).

A whole separate class of applications of pulsar observations is related to gravitational physics in the strong field regime. With the dawn of general relativity, a number of predictions different from Newtonian dynamics become available and thus a number of tests were necessary to confirm or refute the new theory. Observations of pulsars, especially of binary systems with two neutron stars, provide an excellent laboratory for gravitational physics, complementary to experiments that can be undertaken in laboratories or by observations within the Solar System. The most striking example is the Nobel prize winning discovery of the orbital shrinkage due to gravitational wave dampening (Weisberg & Taylor, 2005). The deviations of orbits from Kepler's theory enable a number of tests of general relativity (Damour & Taylor, 1992). By additionally considering the spin contribution, effects such as gravitational precession or spin orbit coupling have been predicted and observed (Damour & Ruffini, 1974; Kramer, 1998; Barker & O'Connell, 1975; Wex, 1995). Finally, gravitational waves can also be detected via regular long term observations of a number of MSPs in the so called pulsar timing arrays as implemented by the European Pulsar Timing Array (EPTA, Stappers et al., 2006; Ferdman et al., 2010), North American Nanohertz Observatory for Gravitational Waves (NANOGrav, Jenet et al., 2009), Parkes Pulsar Timing Array (PPTA, Manchester et al., 2013) and a consortium of these groups in the International Pulsar Timing Array (IPTA Hobbs et al., 2010).

The following sections introduce the basics of general relativity on which many of the above applications are based and proceed to a more detailed description of the pulsar timing methodology and implementation of the timing experiments.

1.2 General relativity

1.2.1 On the equivalence of observers

In 1916, Einstein formulated the general theory of relativity (GR), a theory that joined the special theory of relativity with the equivalence principle and removed the necessity for an instantaneous gravitational force. In the special theory of relativity, all frames of reference that are travelling at constant speed are equivalent. By considering the equivalence principle, Einstein extended this in GR to all reference frames that are accelerating, lifting the limitations of special relativity. This means that a laboratory at rest on Earth is fully equivalent to a laboratory in space accelerating at the rate of gravitational acceleration on Earth’s surface. This was a huge breakthrough with the only drawback being the necessity to abolish global inertial frames and only consider local frames to ensure equity of acceleration across the reference frame.

This equivalence of accelerating frames of reference has important consequences. The two famous Einstein’s thought experiments yielded important insights into the new physics, without considering yet the mathematical background. These insights included the fact that the space-time is curved by matter and gravitational redshift, i.e., the gravitational time dilation. To provide a complete theory, more thorough and mathematical approach was warranted. Building on Gauss’ “remarkable theorem” from 1828 and its extension to a higher number of dimensions by Riemann, Einstein provided a complete theory of gravity (Einstein, 1916).

The curvature of space-time, a central premise of general relativity, can be calculated from Einstein’s field equations, which relate the metric tensor, via the Einstein’s tensor $G_{\mu\nu}$, to the distribution of matter and energy as described by the stress-energy tensor $T_{\mu\nu}$:

$$G_{\mu\nu} = 8\pi T_{\mu\nu} . \tag{1.5}$$

Due to symmetry, only 10 of these equations are independent. In this view gravity is no longer a force but the curvature of space-time caused by the energy-momentum distribution. See also Appendix A.1 for a more complete derivation of these field equations.

Other scientists have derived alternative theories of gravity in which the gravitational field is transmitted by a scalar field, rather than a tensor field. The most famous example of such a theory is Newtonian gravitation. Many other variants have been proposed throughout the years. Nordström’s theory had some problems, which motivated Einstein to try his own scalar theory and move on to a tensor theory (Nordström, 1913; Einstein

& Fokker, 1914). A five dimensional extension of GR attempting to unify gravitation and electromagnetism was introduced by Kaluza (1921). A mix between a scalar and tensor theory was first proposed by Jordan and is equivalent to the theory of Brans and Dicke (Jordan, 1959; Brans & Dicke, 1961). Another idea for a scalar-tensor theory came from Zee who used the Higgs field as the scalar part of his theory (Zee, 1979). So far GR remains the most successful theory of gravitation as measured by the degree of agreement with observations (Will, 2001).

1.2.2 Early tests of general relativity

General relativity was tested soon after it was proposed by Einstein. The first classical test was the measurement of the precession of Mercury's perihelion (Einstein, 1916; Clemence, 1943). The second test was completed soon after by Sir Arthur Eddington who measured the deflection of light by the Sun during an eclipse (Dyson et al., 1920). Gravitational redshift was observed both with astronomical observations (Adams, 1915, but somewhat controversial) and on Earth (Pound & Rebka, 1960; Pound & Snider, 1964). More details on these classical tests are presented in the Appendix A.3.

One more test, directly relevant to this thesis, was proposed by Irwin Shapiro (Shapiro, 1964). He assumed that the metric of the Solar System is well represented by the Schwarzschild's metric (Schwarzschild, 1916) and calculated the difference between the light travel time in a flat and curved space-time. This effect is parametrised by two quantities: the range, corresponding to the amplitude of the delay and the mass of object affecting space-time; and the shape, determined by how close to the mass curving the space-time the light rays pass. Shapiro noticed that radar technology had improved sufficiently to measure these relativistic delays of radar pulses bounced off one of the inner planets of the Solar System. Four years later, the results of the experiments carried out with the Haystack Observatory were published (Shapiro et al., 1968), confirming the predictions. Many further tests have been carried out and ranging test by the Cassini spacecraft on its orbit near Saturn yield measurements consistent with GR to 20 parts in a million (Bertotti et al., 2003), the most stringent test of GR so far.

1.2.3 Gravitational waves

The equation 1.5 implies the existence of gravitational waves. Consider a small perturbation to the metric so that $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$ where $\eta_{\mu\nu}$ is the unperturbed metric far away from any mass and $h_{\mu\nu} \ll 1$ is the small perturbation. After some algebra and exploiting

some conservation laws, we can show that:

$$\square^2 h_{\mu\nu} = 0, \quad (1.6)$$

where \square^2 is the d'Alembert operator for flat space. This homogeneous equation has known solutions, i.e., plain waves, and represents the gravitational wave in the far field, away from the source. For details of the derivation see Appendix A.4

General relativity is a geometric theory with physical interpretations being made on top of the mathematical theory. One of the physical concepts difficult to introduce in GR is the concept of energy and it is beyond the scope of this thesis. Careful analysis reveals that gravitational waves are emitted from physical systems with non-zero second time derivative of quadrupole mass moment. Peters (1964) calculated the amount of energy emitted in a binary system in the approximation of two point masses. While the trajectories of masses are not Keplerian in GR, given that most systems are not highly relativistic, it is convenient to model some of the effects of GR as a time variability of Keplerian parameters. In the case of energy carried away by gravitational waves to first order, the calculated change in the Keplerian semi-major axis a is:

$$\frac{da}{dt} = -\frac{64(M_1 + M_2)^2}{5a^3(1 - e^2)^{3.5}} \left[1 + e^2 \left(\frac{73}{24} + \frac{37}{96}e^2 \right) \right], \quad (1.7)$$

and the change of Keplerian eccentricity is given by:

$$\frac{de}{dt} = -e \frac{304(M_1 + M_2)^2}{5a^4(1 - e^2)^{2.5}} \left(1 + \frac{121}{304}e^2 \right). \quad (1.8)$$

Note that both time derivatives are negative implying that the orbits are tightening and becoming more circular with time. These changes are typically very small.

The frequency of the emitted gravitational waves is related to the typical time-scale of variability of the emitting system. Examples of astrophysical sources of gravitational waves are: binary stars, especially degenerate, such as white dwarf, neutron star and / or black hole binaries (objects with compact degenerate remnants are more likely to be in tight, relativistic orbits owing to the common envelope phase); asymmetric supernova explosions; spinning non-axisymmetric bodies such as a non-spherical neutron star; binary supermassive black holes; quantum fluctuations or phase transitions in the early Universe, in particular cosmic strings (Abbott & the LIGO collaboration, 2007). The gravitational

wave background is typically described by the gravitational wave spectrum:

$$\Omega_{GW}(f) = \rho_c^{-1} \frac{d\rho_{GW}}{d\ln f}, \quad (1.9)$$

where f is the frequency of gravitational wave, ρ_c is the critical energy density of the Universe and ρ_{GW} is the energy density of gravitational waves in an infinitesimal logarithmic frequency interval (Allen & Romano, 1999). Detection of the gravitational wave background is what pulsar timing experiments like EPTA, NANOGrav and PPTA are pursuing.

1.2.4 Pulsar timing and tests of relativistic gravity

When Albert Einstein was thinking about his special relativity, which is based on equivalence of observers travelling at constant speed, he thought of space filled with a regular grid of clocks for his thought experiments. It turns out that, while not really distributed on a regular grid, pulsars are very good clocks in space creating a possibility of conducting exciting experiments. In this section we describe pulsar timing, a methodology from which the majority of pulsar related astrophysics originates.

Time of arrival measurement and clock transfer

As the first step of pulsar timing, the times of arrival (ToAs) of the pulsations needs to be determined. The telescope observes the discrete radio pulsations from the pulsar and records a phase-resolved (average) light curve $O(t)$. This light curve is commonly referred to as the “(average) pulse profile”. While the observed light curve is different during every rotation of the pulsar, summation (known as “folding”) of hundreds of pulses yields the $O(t)$ curve that is unique for every pulsar (Helfand et al., 1975). A high signal-to-noise ratio (S/N) template (or “standard”) for this pulse profile is created from many hours of observations or from an analytical model describing the pulse profile comprised of one or more von Mises components, i.e., cyclic Gaussian distributions. The observed data are then cross correlated with this template by assuming the pulse profile is well described as a shifted by some phase ϕ , scaled by a factor A template S with a constant offset and an additive white radiometer noise N (Downs & Reichley, 1983; Taylor, 1992; Bailes, 2010).

$$O(t) = AS(t - \phi) + B + N(t). \quad (1.10)$$

The obtained ToA is typically calculated in the observatory’s local time as maintained by a hydrogen maser frequency standard which is very stable on short time-scales. These

time tags need to be transferred to the Coordinated Universal Time (UTC) via the Global Positioning System. UTC is then converted to International Atomic Time by adding an integral number of leap seconds. The final time standard is the Terrestrial Time (TT). This is an atomic time that reflects an idealised atomic clock at a mean sea level on the Earth’s surface (Lorimer & Kramer, 2004).

We note that the presence and importance of pulse profile variability have been discussed in many different contexts. Some “classical” pulsars have been observed to change between two or more stable profiles – a phenomenon known as mode changing – on time-scales of minutes to hours (e.g. Backer, 1970c; Bartel et al., 1982). On longer time-scales, pulsars have been discovered whose emission completely switches off for many days, weeks or even months (e.g. Durdin et al., 1979; Kramer et al., 2006a). Recently, Lyne et al. (2010) have shown that the pulse profiles for many pulsars switch between two unique states on time-scales of months to years. Karastergiou et al. (2011) recently detected a transient component in PSR J0738–4042, varying on time-scale of years or decades.

Spin down and astrometric parameters

If the rotation of a pulsar was fully described by equation 1.3, the situation would be simple and modelling pulsars’ rotation would be straightforward. We could observe pulsar’s pulsations at Earth to determine the pulsar spin period and its time derivative. With this information we could form a simple model to predict when the next pulse will arrive and compare it to the observed ToA. The difference between the modelled and observed ToA is called the timing residual. In reality, the situation is quite a bit more complicated as many other astrometric and orbital parameters are involved. Our observations are conducted from Earth which is not at rest and this has to be considered. General relativistic effects are affecting the travel of the signal as the light has to travel on geodesics through space-time between the pulsar and the telescope. Most of the important effects are well understood and here we describe them in more detail based on the description of the timing model used by the standard pulsar timing software TEMPO2 (Edwards et al., 2006).

The pulsar is in motion and to get a precise model of when the next pulse will arrive, the position and motion of the pulsar need to be determined. These parameters affect the pulsar ToAs by purely geometric means. This geometric delay, to the second order, can be derived as:

$$\Delta_{\text{geometric}} = \frac{v_{\parallel}}{c}t + \frac{r_{\parallel}}{c} + \frac{v_{\perp}^2}{2cd}t^2 + \frac{\vec{\mu} \cdot \vec{r}_{\perp}}{c}t + \frac{r_{\perp}^2}{2cd}, \quad (1.11)$$

where c is the speed of light, \vec{v} is the pulsar’s velocity with radial and transverse (with respect to the line of sight) magnitude components denoted respectively by indices \perp and

$\|\cdot\|$, \vec{r} is the vector from the observatory to the Solar System barycentre (SSB), \vec{d} is the distance from the SSB to the pulsar, and μ is the proper motion of the pulsar.

The terms in the equation 1.11 are readily interpreted in geometric terms. The first term is the secular increase of radial distance to the pulsar and is highly covariant with pulsar's spin period. As the Earth moves along the geodesic around the Sun, the distance between the telescope and pulsar changes, as measured by the second term known as the Römer delay. The third term is the apparent transverse quadratic Doppler shift (or Shklovskii effect or secular acceleration, Shklovskii, 1970). This delay arises from the increase of the projected distance to the pulsar due to the transverse motion. The fourth term is the proper motion of the pulsar. The last term is the timing parallax caused by the apparent shift of the pulsar against the background (e.g., distant quasars) and corresponds to curvature of the incident radio wave.

If the pulsar's position or proper motion μ are not well determined, the timing residual will exhibit a sinusoidal variation in time with constant or increasing amplitude, respectively, as the transfer of time tags from the SSB (known as the barycentric ToAs, BATs) to the observatory time tags (known as the site arrival times, SATs) will be inaccurate. The increasing amplitude of the sinusoidal variation in the case of an incorrect proper motion stems from the temporal increase of the error in position.

Relativistic corrections and binary pulsars

Since pulsar astronomy deals with very high precision experiments, pulsar timing models need to take relativistic effects into account as well. There are two effects that are typically taken into account simultaneously and dubbed the Einstein delay and can be calculated to a high precision based, e.g., on the work of Blandford & Teukolsky (1976). These two effects are the special relativistic time dilation and general relativistic time dilation, i.e., the gravitational redshift. The first one is caused by the Earth's motion relative to the Sun. The general relativistic dilation is caused by the space-time curvature at Earth due to all the bodies in the Solar System with main contributions from the Sun and Jupiter. Another relevant relativistic effect, as described in section 1.2.1, is the Shapiro delay with the same two bodies having highest impact on the ToA.

As explained in section 1.1.4, some of the pulsars, and most of MSPs, have a companion star. In this case, even when the orbit of the pulsar is very well approximated by Newtonian dynamics, the orbital motion of the pulsar must be modelled (Roy, 1988). Care needs to be taken while calculating the arrival times to include the acceleration in the binary to compensate for varying Doppler shifting of the measured orbital period. The Römer delay,

analogous to the Solar System's case, is important for the pulsar as the line of sight is changing in time due to the periodic change of the pulsar's position (for details of binary models see Edwards et al., 2006). If the orbit is relativistic (i.e., the curvature of space-time cannot be neglected) the Einstein and Shapiro delays in the binary system must be modelled. In addition, as described in A.3, the trajectories in the curved space-time are different from the conic sections which warrants updating the binary model used for correcting the line of sight variation to include the precession of periastron and change in the orbital period (e.g., Damour & Deruelle, 1985, 1986). In total, up to five relativistic parameters are included in the models and measurement of at least two of them allows the masses of the pulsar and its companion to be uniquely determined. This enables studies of the equation of state of matter at extreme densities as it constrains the models by detection of heavy neutron stars inconsistent with some models. The proper motion of the pulsar binary, as well as the Earth's motion around the SSB, cause an apparent change in the semi-major axis projected on the sky and the longitude of periastron (Kopeikin, 1995) and has been measured for PSR J0437–4715 by van Straten et al. (2001). The changes caused by the proper motion of the binary are secular (Kopeikin, 1996), which allows us to distinguish them from periodic changes originating from the Earth's motion. A number of other effects can cause secular changes of the observed projected semi-major such as presence of a planet (Rasio, 1994), spin-orbit coupling (Smarr & Blandford, 1976), or geodetic precession of the pulsar spin axis (Damour & Ruffini, 1974).

The effects of interstellar medium

Another set of delays to the ToA are due to the interstellar medium (ISM) along the line of sight to the pulsar. Interstellar space is filled with a tenuous plasma, i.e., ions and free electrons⁴ which has several effects on the radio waves propagating from the pulsar towards the Earth. From classical electrodynamics (e.g., Jackson, 1975), and noting that the ISM is a cold plasma, i.e., by neglecting the collisions, we can write the dielectric permittivity of ISM as:

$$\epsilon = 1 - \frac{\omega_p^2}{\omega^2} \mp \frac{\omega_p^2 \omega_B}{\omega^3}, \quad (1.12)$$

where ω is the frequency of a radio wave propagating through the plasma, $\omega_p = \sqrt{n_e e^2 (m_e \epsilon_0)^{-1}}$ is the characteristic (plasma) frequency, n_e is the number density of the free electrons, e and m_e are the charge and mass of an electron, respectively and ϵ_0 is the dielectric per-

⁴Note that ions are much heavier than electrons and therefore their motion is at least three orders of magnitudes less than that of electrons. All the considerations in this thesis ignore the presence of ions in the tenuous plasma.

mittivity of vacuum, and ω_B is the cyclotron frequency, related to the Galactic magnetic field along the line of sight. We note that ϵ is very close to unity at frequencies of interest in radio astronomy with typical values of n_e in the ISM (e.g., as measured by Ables & Manchester, 1976). The phase velocity of a plane wave is given by $v_p = \omega k^{-1}$, where k is the wave number. Since the phase velocity of the wave, as it travels through the ISM, has to equal $c\epsilon^{-0.5}$, we arrive at the expression for the wave frequency $\omega = kc\epsilon^{-0.5}$. By combining this result with equation 1.12 (where we ignore the term with cyclotron frequency as it is much smaller than the second term) we obtain the dispersion relation for the ISM:

$$\omega = \sqrt{k^2 c^2 + \omega_p^2}, \quad (1.13)$$

and as a consequence:

$$v_p = \frac{c}{\sqrt{1 - \omega_p^2 \omega^{-2}}}. \quad (1.14)$$

If we now consider a discrete pulse of radiation, as expected from a pulsar, it will consist of a superposition of plane waves with a wide range of wave numbers. By considering the peak of such a pulse, we conclude that the group velocity of this wave pulse is given by $v_g = \frac{d\omega}{dk}$. By using the dispersion relation 1.13 the peak of a discrete pulse of radiation travels through the ISM at a velocity given by:

$$v_g = c \sqrt{1 - \frac{\omega_p^2}{\omega^2}}. \quad (1.15)$$

Following Lorimer & Kramer (2004), we use this velocity to consider what is the difference between the travel time of a radiation pulse of infinite frequency travelling over the distance d as compared to a wave of frequency ω , i.e., $\Delta t = \int_0^d v_g^{-1} dl - d/c$, we obtain the following expression:

$$\Delta t = \frac{1}{c} \left(\int_0^d \left[1 + \frac{\omega_p^2}{2\omega^2} \right] dl - d \right) = \mathbb{D} \frac{DM}{\omega^2}, \quad (1.16)$$

where we used the Taylor expansion of v_g for $\omega_p \ll \omega$, $\mathbb{D} = e^2 (2\pi m_e c)^{-1}$ is the dispersion constant and $DM = \int_0^d n_e dl$ is the dispersion measure. The radio waves propagating through the ISM are delayed with respect to a radiation of infinite frequency by time inversely proportional to their frequency squared, i.e., the radio waves at higher frequency arrive earlier than the radio waves at lower frequency. We note that the dispersion measure is not a constant quantity and varies in time due to the relative motion of the Solar System, observed pulsar and the intervening ISM (see e.g., You et al., 2007).

The ISM can also cause other effects such as scintillation or scattering. These effects are more complicated and not fully understood yet and we refer to an excellent review by Rickett (1990). In short, both these effects are caused by multi-path propagation through the varying distribution of the ISM plasma. The scintillation is related to long term variability and demonstrates itself by flux variation of the pulsar, while scattering is related to short term variability and causes smearing of the pulsar signal. Finally we note that the plasma need not be in interstellar space. The effects of the solar plasma can be seen in the pulsar timing data when the line of sight towards a pulsar passes near the Sun (e.g., Ord et al., 2007; You et al., 2007).

Timing noise and profile variability

When all the effects listed in previous paragraphs, including corrections for DM variations, are taken into account, we expect that in an ideal case the post-fit timing residuals, i.e., the difference between the measured ToAs and the observed ToAs after fitting for all the model parameters, will be white. In practice, there is a plethora of physical processes which can contribute additional noise to the timing residuals, for an excellent review of such processes see Cordes & Shannon (2010). Briefly speaking, such processes include pulsar spin rate variations, variation of pulsars' magnetosphere, effects related to the ISM, irregularities of Earth's rotation, and more. If no explanation can be found for the source of additional variations in ToAs, or we are unable to model and remove them, we say that the pulsar exhibits timing noise⁵.

The connection between the pulse shape changes and timing noise was made soon after the discovery of timing noise in the pulsar observations by Boynton et al. (1972). They studied optical timing observations of the Crab pulsar and discovered a noise component in the timing residuals which was well modelled as a random walk in the pulsar spin frequency. The authors also considered a random walk in the pulse phase and spin frequency derivative, but found no evidence of such noise in their data. This analysis was extended in a series of papers (Groth, 1975a,b,c). The author presented an analysis method suitable for studies of data with inherent timing noise. This improved methodology led to the conclusion that noise in the Crab pulsar timing is dominated by a random walk in the spin frequency but a random walk in pulsar phase might also be present. In the meantime, Manchester & Taylor (1974) described timing noise for two slow pulsars from radio observations. A few years later, another series of papers (Helfand et al., 1980; Cordes,

⁵We note that there is no widely accepted definition of timing noise and some astronomers would only include red noise processes under this term.

1980; Cordes & Helfand, 1980) presented statistics of timing noise for 37 bright pulsars and concluded that it is a ubiquitous phenomenon. These authors presented a careful analytical description of random walks in pulsar phase, spin frequency and its derivative and the uncertainties in the estimation of their parameters. The last paper in the series pointed out that the random walk in the pulsar phase can be due to the random pulse shape changes but concluded that it was unlikely to be the dominant source of timing noise in the available data. A different dataset was analysed in a similar manner by Cordes & Downs (1985) who stated that either excessive jitter or pulse shape changes are important for a significant fraction of their sample. They also pointed out that the pulse shape changes are likely to be universal but their importance varies from object to object. Later, Cordes (1993) detected pulse shape variability in 11 out of 14 studied objects. These variations were consistent with being caused by summation of a finite number of pulses. Similar studies with same conclusions were performed also by D’Alessandro et al. (1993). A year later Kaspi et al. (1994) studied two millisecond radio pulsars and discovered timing noise in one of them. The general continuity of properties between classical and MSPs suggests that pulse shape changes may be common in MSPs as well.

The profile variability of MSPs has been studied in relatively few cases. The single pulses from PSR J1939+2134 show no sub-pulse structure over selected ranges of pulse longitude (Jenet et al., 2001a) but emit giant pulses as much as 300 times brighter than the average pulse, that are narrower and systematically delayed with respect to the main and interpulse components (Cognard et al., 1996; Kinkhabwala & Thorsett, 2000). Several other groups have argued that mildly recycled pulsars and MSPs exhibit profile shape changes. Some are associated with different viewing geometries due to gravitational spin-orbit coupling; e.g. PSR B1913+16 (Weisberg et al., 1989; Kramer, 1998) and PSR B1534+12 (Arzoumanian, 1995; Stairs et al., 2000b). Backer & Sallmen (1997) claimed erratic emission modes from PSR B1821–24 but at only one observing frequency. In another work, Kramer et al. (1999) studied PSRs J1022+1001 and J1730–2304. In both cases, they detected profile variations over time-scales of the order of 10 to 15 minutes; however the data quality for the latter did not allow a rigorous statistical analysis. On the other hand, Hotan et al. (2004a) detected no significant variations in the pulse profile of PSR J1022+1001 and demonstrated that the reported profile shape variations could be explained by polarisation calibration errors.

Small profile changes in PSR J0437–4715 were described by Vivekanand et al. (1998) using observations performed at a very low frequency with only a single polarisation. This result was contested by Sandhu et al. (1997), who argued that calibration errors were

the true origin. Vivekanand (2001) later argued that the variations are intrinsic to the pulsar and correlated with spiky emission in the varying component. Variations in the central region of the profile were also reported by Navarro et al. (1997) with 24 minute integrations at 428 MHz but they were not investigated in detail.

Attainable timing precision

We have described what affects the ToAs and can now consider what accuracy is expected and achieved in pulsar timing experiments. The most precise experiments are conducted with MSPs as their spinning is so regular that it challenges the stability of atomic clocks (Matsakis et al., 1997). While classical pulsars display significant amounts of red timing noise (e.g., Shannon & Cordes, 2010), connected to magnetospheric changes in the pulsar in at least some cases (Kramer et al., 2006a; Lyne et al., 2010), MSPs show much lower levels of red noise (Verbiest et al., 2009). It remains to be understood yet how steep the red noise spectrum is in MSP timing.

The limit of attainable timing precision is the uncertainty related to the template matching algorithm (equation 1.10). In the case of a pulsar with no timing noise, a timing model including all the effects, and conforming to all the assumptions of this ToA estimation algorithm, the root mean square (rms) timing residual should be equal to the uncertainty of a ToA measurement. This uncertainty can be shown to be well approximated with the following expression (Downs & Reichley, 1983; van Straten, 2006; Cordes & Shannon, 2010):

$$\sigma_{\text{ToA}} = \frac{W_{\text{eff}}}{2S/N}, \quad (1.17)$$

where W_{eff} is the effective pulse profile width. Any additional effects will worsen the timing precision and we can hope to at best achieve rms timing residual equal to the uncertainty of a single ToA as measured by the equation above. Cordes & Shannon (2010) present a wide discussion of noise balance in pulsars in the context of precision timing and describe a number of phenomena, such as pulse jitter and spin-noise, which can increase the rms timing residual above this limit.

The current best MSP timing experiments yield rms timing residuals of the order of a few hundred nanoseconds over time-scales of the order of years. Hotan et al. (2006) achieved ~ 300 ns timing rms for one object and ~ 200 ns for two other pulsars over 2 years. PSR J1713+0747, another bright MSP has been timed with rms timing residual below 200 ns over a time span of 6 years (Verbiest et al., 2008). Hobbs et al. (2009) and Verbiest et al. (2008) analysed 5 and 10 years of observations of PSR J0437–4715, respectively, and both demonstrated precision to within 200 ns. van Straten et al. (2001) timed this pulsar

with the rms timing residual of 130 ns over 3.4 years. Sub 50 ns rms timing residuals were reported by Demorest et al. (2012), however these values are not directly comparable to the other quoted results due to inherently different analysis techniques. The timing precision of PSR J0437–4715 has never approached the limit of its timing precision predicted by equation 1.17. This thesis scrutinises the reasons for this discrepancy and introduces methodology which brings the short term timing precision closer to this theoretical limit.

1.2.5 Pulsars as physics laboratories and tests of gravitation

While the more complete picture of pulsar timing presented above is much more complicated than the initial simple description of a spinning neutron star, all the complications are not a nuisance but in fact provide an excellent opportunity to study all the described physics relevant to these delays. Given how broad the range of areas are relevant for timing, the pulsars are extremely versatile laboratories of physics on scales from intra-nuclear to cosmological. Some highlights of the research completed by means of pulsar timing are: most precise published mass measurement of planets in the Solar System (Champion et al., 2010) and exoplanets (Wolszczan & Frail, 1992); characterisation of the ISM (You et al., 2007); development of an independent terrestrial time-scale (Petit & Tavella, 1996; Rodin, 2008; Hobbs et al., 2012); studies of highly magnetized plasmas (Cairns et al., 2003a); spacecraft navigation (Tartaglia et al., 2011); and studies of the pulsars themselves (Lyne et al., 2010). This thesis is concerned with improving the precision of pulsar timing experiments to make new and even more accurate experiments plausible.

One class of application of pulsar timing deserves a special interest as they are strongly related to what is driving current international efforts in pulsar astronomy research. This application is testing the theories of gravity with pulsar timing. The most famous test is related to the discovery of a highly relativistic binary by Hulse & Taylor (1975). The change of orbital period in this system has been measured by Taylor & Weisberg (1982) and demonstrated to match the predictions of general relativistic decay due to the emission of gravitational waves. Since then this effect has been measured for a few more pulsars (Stairs et al., 1998; Kramer et al., 2006b; Jacoby et al., 2006; Bhat et al., 2008; Freire et al., 2012) (Stairs et al., 1998), again in agreement with GR. Timing of other pulsars have yielded the measurement of the Shapiro delay. A notable case is the work by Demorest et al. (2010) who used the Shapiro delay to estimate the mass of the pulsar PSR J1614–2230 to be $1.97 \pm 0.04 M_{\odot}$, the highest known so far, thus constraining the equation of state of dense matter. This effect has been also measured for PSR J0437–4715. Not only the delay of light propagating through the curved space-time has been measured (van Straten et al.,

2001), but the same authors independently determined the full three dimensional orbit of the pulsar by measuring the apparent variation in the semi-major axis of the binary. The inclination of the orbit measured from the Shapiro delay agreed with the geometric value, thus providing a strong confirmation of GR predictions.

A relativistic system discovered by Burgay et al. (2003) was shortly after confirmed as a double pulsar system by Lyne et al. (2004). This created the opportunity to perform even more stringent tests of GR. While all the tests in the Solar System have tested GR in the weak field approximation, this system enabled the strong field regime to be explored. All five post-Keplerian parameters were measured by Kramer et al. (2004). As the double pulsar system provides two lines of sight, the projected semi-major axis for both objects were derived from observations and thus the mass ratio was estimated independently from relativistic parameters. In this way the authors over-constrained the problem of estimating the masses of the two pulsars with six equations. Such over-constraints enable stringent tests of non-GR theories of gravity. The measured values were all yielding consistent values of the two masses by assuming GR holds true in the strong field regime. More stringent tests are possible in this system by measurements of higher order corrections to the Keplerian orbits and possibly the first direct measurement of the pulsar's moment of inertia.

Even with such stringent tests, regions of phase space of GR predictions remain unexplored (Kramer et al., 2006b). When a neutron star - black hole binary is discovered, the black hole mass, spin and spin orbit coupling can be measured (Wex et al., 1998) in turn enabling tests of the cosmic censorship conjecture (no singularities exists without an event horizon; Hawking & Penrose, 1970) and the no-hair theorem (black holes are fully characterised by three parameters: mass, electric charge, and angular momentum; Thorne, 1980). Some of the gaps have been closed, e.g., very accurate tests of tensor-scalar theories with timing of neutron-star white dwarf binaries (Lange et al., 2001) with the recently published most stringent test so far (Freire et al., 2012). Existence of preferred reference frames violate GR and can be tested by pulsar timing (Stairs, 2003; Stairs et al., 2005).

The stochastic gravitational wave background, described by equation 1.9, provides another way of testing theories of gravity. In the case of gravitational waves originating from supermassive black hole binaries, the frequency region of interest is around a few nano-Hertz (Sazhin, 1978). The amount of energy emitted in gravitational waves can be extremely large, reaching expected fluxes at Earth of the order of 10^{20} Jy (Sathyaprakash & Schutz, 2009). The gravitational wave radiation is very weakly coupled with matter and therefore the contribution of this low frequency stochastic background is expected to

introduce variations in the timing residuals as small as a few to tens of nanoseconds in the optimistic cases of current models (Detweiler, 1979; Rajagopal & Romani, 1995; Sesana, 2010). We note that this is well below the currently attainable precision as described in section 1.2.4. This thesis contributes to the goal of gravitational wave detection and other timing experiments by increasing the precision of pulsar timing.

Kaspi et al. (1994) measured an excess of power at the low frequency end of the timing residuals spectrum. By ignoring potential timing noise and assuming all these unexplained noise was due to the stochastic background she obtained an upper limit of the amplitude of this background. Since then this result was superseded by Lommen & Backer (2001) with better data. Much stronger constraints on the gravitational wave background are placed by observing multiple MSPs within a so called pulsar timing array (Foster & Backer, 1990). In principle, the predicted correlations between the gravitational wave signal between pairs of pulsars distinguish the stochastic background from other effects on the pulsar timing data and allow to overcome problems with detecting the very small expected signal. The quadrupolar character of gravitational wave emission produces a characteristic angular distribution of these correlations (Hellings & Downs, 1983). So far, no PTA has detected gravitational waves but the current limits as presented by Jenet et al. (2006); van Haasteren et al. (2011); Demorest et al. (2012) are on the edge of expected detection of the stochastic background or single sources (e.g., Sesana, 2010). The three major PTAs have now joined together to enable faster detection of gravitational waves (Hobbs et al., 2010). We note that the detection via PTA is qualitatively different from the detected changes in the orbital period as it does not depend on the gravitational waves carrying energy away but directly measures the variations in the space-time curvature. Similar searches for gravitational wave backgrounds at different frequencies are undertaken with several ground based interferometers such as LIGO (Abramovici et al., 1992) or VIRGO (Bradaschia et al., 1991).

1.2.6 PSR J0437–4715 and implications for the future of pulsar timing

PSR J0437–4715 is the brightest millisecond pulsar and closest pulsar known. This pulsar, spinning 173 times a second, was discovered in the Parkes 70cm survey (Johnston et al., 1993) and has flux at wavelength of 20cm of 150 mJy, approximately 60 times brighter than the median MSP flux. It is not only a very bright radio source but also lies off the Galactic plane near the anti-centre longitude, i.e., it is visible when the majority of Galactic sources are set and the competition for telescope time is not severe. This allows for frequent and extensive observing sessions and detailed studies of this pulsar. By examining this pulsar

we are able to predict properties of other MSPs when observed in future with the next generation of telescopes such as the Square Kilometre Array (SKA, Kramer et al., 2004) with which other pulsars could be observed at the same level of detail as PSR J0437–4715 is currently studied using facilities like the Parkes Observatory.

PSR J0437–4715 is one of the best celestial clocks known and timing precision of 130 ns has been achieved over the timespan of 3.4 years allowing full three dimensional determination of the orbit’s geometry providing stringent tests of general relativity and very accurate measurement of the pulsar’s position on the sky and distance to it (van Straten et al., 2001). It also is a key object for the PPTA in attempts to detect gravitational waves and create pulsar based time-scale.

This pulsar lies nearby, at a distance of 0.16 kpc which is also advantageous as it implies a relatively small integrated density of the tenuous plasma along the line of sight between the pulsar and the Earth. This makes drawing conclusions about the intrinsic behaviour of the pulsar easier than for other objects further away in the Galactic plane.

This thesis is concerned with detailed studies of PSR J0437–4715 concentrated around improving the attainable timing precision. By understanding this pulsar we lay ground for future extremely precise timing experiments and enable more stringent tests and experiments in all the areas listed in this introductory chapter. While many of the listed topics will not be further related to in the remainder of this work, they can all benefit from the significant progress in understanding the limitations and pushing the limits of precision timing achieved with the subsequently presented theoretical understanding and practically applicable techniques made available to the whole pulsar community.

1.3 Structure of this thesis

This thesis is structured as follows. The second chapter discusses the signal path from the pulsar through the telescope systems. It includes the basics of radio astronomy and a description of the hardware used in observations of pulsars. The third presents the results of the research conducted during the author’s candidature on the limitations of pulsar timing precision. It discusses the statistical properties of noise intrinsic to the pulsar and how to exploit them to achieve the best possible timing precision. The fourth chapter unveils further improvements to the exploitation of noise statistics by using the full polarisation signal. Chapters three and four together provide means of achieving higher timing precision in all types of experiments, relevant to all the research areas outlined above. An analysis of vast amounts of data from single pulse observations of PSR J0437–4715 is described in detail in the fifth chapter. The topics include profile evolution with pulse

amplitudes relative to the mean; precision timing with high time resolution observations; physics of the magnetosphere; and evidence of orthogonally polarised modes in single pulse observations of PSR J0437–4715. Finally, the sixth chapter concludes by summarising all the findings presented in previous chapters and discussing the future of pulsar timing.

2

The fundamentals of radio pulsar observational techniques

Chapter 1 outlined the formation of neutron stars, their connection to the general theory of relativity and a variety of possible applications of pulsar studies. In this chapter one missing piece of information is provided, the means of observing the electromagnetic radiation that originates from the pulsars we want to observe. Section 2.1.1 briefly summarises the basic physics of fundamental interactions. We proceed to explain the significance of photons, that we use in astronomy to receive information from distant stars. After that, section 2.1.2 presents a more classical approach to electromagnetism that is traditionally used in astronomy. At the end of this section properties of electromagnetic radiation are presented along with information about the polarisation nomenclature and its description. This provides the necessary background for the parts of the thesis which consider the polarisation properties of pulsars. Section 2.2 provides a description of the signal path from the antenna, via the feed and receiver to the intermediate frequency, baseband, and detected signal devices. Finally, section 2.3 introduces a variety of radio astronomy pulsar “backends”, i.e., data recording and initial processing systems.

2.1 Radio waves

2.1.1 The symmetry of nature

Traditionally, physicists consider four separate fundamental forces: electromagnetism; the strong and weak nuclear forces; and gravitation. This view has some caveats: as discussed in section 1.2.1, gravitation is considered in modern physics to be a manifestation of space-time curvature rather than a force; electromagnetic force is a unification of the electric and magnetic forces which were treated separately in the past; since the work of Glashow

(1961) weak and electromagnetic forces are unified in the electroweak force; it is envisioned that all fundamental forces can be unified into one universal force at higher energies in the so called Grand Unification Theory as first proposed by Georgi & Glashow (1974). We note that some physicists disagree on the possibility of unifying gravity with any of the other fundamental interactions. In the discussion of fundamental interactions below we exclude gravitation.

All fundamental interactions are described in the Standard Model (for a review see Beringer J. on behalf of Particle Data Group, 2012) which explains the structure of matter as consisting of elementary particles and the interactions between them. The structure of the model is exactly the same as certain groups from the Lie algebra (e.g., Erdmann & Wildon, 2007). In this way modern physics owes much of the current understanding of the Universe to the mathematical studies of differentiable manifolds (see Castillo, 2011, for a review with focus on application in physics). In the Standard Model, all the elementary particles constituting matter are known as fermions which we already described in chapter 1. The fundamental interactions listed in the paragraph above occur as an exchange of force carrying particles known as gauge bosons and correspond to the generators¹ of the aforementioned groups. The three fundamental interactions described here can be seen as a gauge field, i.e., a field arising from the requirement that gauge symmetry² holds everywhere in space-time. All these gauge symmetries correspond to a mathematical group of symmetries. For example, the electromagnetic field corresponds to the Lie algebra of dimension one (or the unitary group of degree 1) $U(1)$, a group of symmetries of a complex number with respect to its phase. Just as expressed by the Noether's theorem (Noether, 1971), these symmetries correspond to laws of conservation. The other two fundamental interactions correspond to the symmetry groups $SU(2)$ and $SU(3)$ for the weak and strong forces respectively.

The predicted gauge bosons for the electromagnetic, weak and strong forces are: photons³; intermediate vector bosons $W+$, $W-$ and Z_0 ; and eight gluons, respectively. The experimental discovery of the intermediate vector bosons, first indirectly via the Nobel Prize winning (1979) measurement of neutral currents (Hasert et al., 1973), for which the Z_0 bosons are responsible, and later directly by Arnison G. et al. (UA1 collaboration) (1983a,b); Banner M. et al. (UA2 collaboration) (1983); Bagnaia P. et al. (UA2 collaboration) (1983), was a grand success of the Standard Model. In fact, the latter discovery was

¹i.e., elements of the group which can yield all the elements of the group under group operation

²Gauge symmetry states that the physics laws are the same in different systems of coordinates

³first invoked by Einstein (1905, (Nobel Prize in 1921)) to explain the photoelectric effect and based on Max Planck's theory of black-body radiation (Planck, 1900a,b, Nobel Prize 1918)

awarded a Nobel Prize only one year later, unusually quickly after the discovery. It turns out that the correspondence between the weak force and the symmetry group $SU(2)$ is not exact as the intermediate vector bosons are very massive. This is contrary to the other two types of gauge bosons, photons and gluons. Work by Weinberg (1967) and Salam & Strathdee (1972) led to our understanding that this is because the electromagnetic and weak interactions are actually one and the same interaction with the intermediate vector bosons having mass as a result of symmetry breaking as caused by the Higgs field (Englert & Brout, 1964; Guralnik et al., 1964; Higgs, 1964), winning yet another Nobel Prize in Physics for the particle physicists in 1999. In modern physics, the strong force is described by quantum chromodynamics (e.g., Greiner et al., 2007) with the electroweak force describing the combined electromagnetic and weak forces.

The understanding of fundamental interactions is crucial for creating models of the interior of neutron stars. Pulsars provide a unique laboratory where all the four forces are important due to the extreme densities in the neutron star interiors. The conditions inside a neutron star cannot be recreated in a laboratory and thus constraining the equation of state of a neutron star via observations of radio or X-ray pulsars may be the only insight we get into the fundamental forces except the particle accelerators (see, e.g., Kurkela et al., 2010). However, the latter instruments are typically studying only three of the fundamental forces excluding gravity while neutron stars are also highly relativistic objects.

Astronomy is different from other areas of physics in the fact that it is mostly constrained to observing the photons from distant objects without any possibility to adjust the experimental object. All the objects of interest in astronomy lie at vast distances. Even the Moon is 384,400 km away from the Earth on average⁴. Due to the symmetry breaking and intermediate vector bosons being massive, the weak aspect of the electroweak force is only detectable on very short distances, shorter than the size of a hydrogen nucleus. The main relevant phenomenon on the small scale is the inverse beta decay we invoked to describe formation of neutrons inside a neutron star in section 1.1.2. To study gauge bosons from distant astronomical objects, we can consider only photons as described by quantum electrodynamics (QED, Dirac, 1927; Fermi, 1932), i.e., a relativistic version of classical electrodynamics that is consistent with quantum mechanics. In radio astronomy in particular, we can conveniently make use of the particle-wave duality as described by de Broglie (1924) and limit ourselves to a classical description of the electromagnetic waves. This is justifiable and convenient as we are dealing with a very large number of photons of very low energy. In these conditions, the detection of a single photon is impractical and

⁴as provided by the Jet Propulsion Laboratory <http://ssd.jpl.nasa.gov/?horizons>

instead of counting single photons, like in many other areas of astronomy, radio astronomy is concerned with the wave description of the electromagnetic field. In this way we resign from the beauty of symmetries and the current best understanding of nature and we opt for the mathematically simpler formulas which are still a very good approximation in the context of radio astronomy. The next section outlines the mathematical basis of the relevant parts of classical electrodynamics.

We note that with progress of technology other information carriers are being used in astronomy. In particular, neutrinos are detected in a number of facilities (e.g., IceCube, Halzen & Klein, 2010), cosmic rays can be observed via the cascades of particles they produce (e.g., Aharonian et al., 1997), and in near future gravitational wave detectors are expected to reach sensitivities necessary to open this window on the Universe (e.g., Ju et al., 2000).

2.1.2 Charges in motion - and a new physical entity

The existence of an oscillating electromagnetic field was shown by Maxwell (1861) who also identified light as an electromagnetic wave (Maxwell, 1865). His equations, in the classical case⁵, are:

$$\vec{E} = -\vec{\nabla}\phi - \frac{\partial\vec{A}}{\partial t}, \quad (2.1)$$

$$\vec{B} = \vec{\nabla} \times \vec{A}, \quad (2.2)$$

$$\vec{\nabla} \cdot \vec{D} = \rho, \quad (2.3)$$

$$\vec{\nabla} \times \vec{H} = \vec{J} + \frac{\partial\vec{D}}{\partial t}, \quad (2.4)$$

where \vec{E} and \vec{B} are the vectors of electric and magnetic field, respectively; ϕ and \vec{A} are the electric and magnetic vector potential, respectively; ρ is the density of free charges; \vec{J} is the total electric current; \vec{D} is the electric displacement field; and \vec{H} is the demagnetising field. In the case of isotropic non-magnetic and non-dielectric non-polarized and non-magnetized materials, the last two quantities are simply scaled versions of the electric and magnetic field: $\vec{D} = \epsilon\vec{E}$ and $\vec{H} = \mu^{-1}\vec{B}$, where ϵ and μ are the dielectric permittivity and magnetic permeability, respectively. The field equations were mostly known before Maxwell and, in order appearing above, they are known as Gauss's law, Gauss's law for magnetism, Faraday's law, and Ampère's law. One of Maxwell's major contributions was the inclusion of the displacement term in the equation 2.4. While seemingly a small change, it has

⁵for a review of classical electrodynamics see Jackson (1975) and Johnson (1998)

dramatic implications and Einstein considered it the most important progress in physics since Newton's time. Maxwell was the first person to unify two fields within one self-consistent theory, laying the ground work for modern field theories. In his equations, the electromagnetic field is an independent physical entity that does not have to be associated at all times to a source, which was another conceptual breakthrough. The full set of Maxwell's equations governs the creation and evolution of the electromagnetic field. The electric and magnetic fields are created by free charges and currents, but also by the variation of the other field, i.e., these fields are coupled. Therefore, an oscillating charge will create an electromagnetic field.

To study how this field will evolve some distance away from the charge, we can look at Maxwell's equation in the vacuum by setting $\rho = 0$ and $\vec{J} = 0$ in the equations above. After this operation, by calculating the curl of the equation 2.3 and joining it with equation 2.4 we obtain a wave equation for the electric field. Similar algebra yields the wave equation for the magnetic field. Just as it was the case for equation 1.6, the solutions can be represented by linear combinations of plane waves:

$$\vec{E} = E_0 e^{-i\vec{k}\cdot\vec{r}}, \quad (2.5)$$

$$\vec{B} = B_0 e^{-i\vec{k}\cdot\vec{r}}. \quad (2.6)$$

Other important parametrisations of the electromagnetic waves are the spectral decompositions as linear combinations of sine and cosine waves. This is possible owing to the linearity of Maxwell's equations and the Fourier transform. Another commonly used representation of electromagnetic waves is the multipole expansion which provides a convenient way of parametrising the dependence of the waves on angles in a chosen coordinate system. The radial component is chosen so that it diminishes strongly with distance to ensure that the leading non-vanishing terms in the multipole expansion are the strongest.

Polarisation

The electromagnetic waves far from the source are propagating with the electric and magnetic field oscillating in phase with each other, with a constant ratio of their maximum amplitudes, and perpendicular to each other. The oscillations of the electric and magnetic fields occur in the plane perpendicular to the direction of propagation of the wave, i.e., the electromagnetic waves are transverse waves. As Maxwell's equations allow two independent solutions, we can consider one more important property of the electromagnetic waves, i.e., the correlation between these two solutions. When the statistical independence

of the two solutions is lost, we say that an electromagnetic wave is polarised.

A polarised wave exhibits organised oscillations of the electric and magnetic fields with a preferred orientation or ellipticity. To study the polarisation of electromagnetic waves, we typically analyse the curve traced by the vector of electric field projected onto an arbitrary plane perpendicular to the propagation direction. If the wave is polarised we classify the polarisation state as one of the three types: linear polarisation, when the Lissajous figure drawn by the electric field is a straight line; circular, if it is a circle; and elliptical if it is an ellipse. If the electromagnetic radiation is coming from a large number of sources whose motions are not correlated, then there is no preferred direction of oscillations and the electromagnetic wave will be unpolarised. The electromagnetic radiation can be only partially polarised if the correlation between the two solutions of Maxwell's equations is partial. A typical example of an astrophysical coherent source producing highly polarised radiation are molecular masers based on the stimulated emission of microwave and radio waves. Pulsar emission often shows a high degree of polarisation as well.

It is convenient to represent electromagnetic radiation by the means of complex-valued quantities, known as the Gabor, or analytic signal (Gabor, 1946). For simplicity, all the examples are given for electric field only without any loss to generality of the following considerations. Furthermore, we can assume that the electromagnetic wave is propagating in the x_3 direction of a Cartesian coordinate system. Then, the non-vanishing components of the electric field can be described by only the two components in the plane spanned by x_1 and x_2 axis, known as the Jones vector \vec{e} . By calculating the outer product of the Jones vector with its Hermitian transpose and taking the ensemble average we can calculate the loss of independence between the two solutions of the Maxwell's equations. This coherency matrix (e.g., Wolf, 1959; Born & Wolf, 1999) of the electric field can be written as:

$$\hat{\rho} = \langle \vec{e} \vec{e}^\dagger \rangle = \begin{pmatrix} \langle a_1^2 \rangle & \langle a_1 a_2 e^{i\delta} \rangle \\ \langle a_1 a_2 e^{-i\delta} \rangle & \langle a_2^2 \rangle \end{pmatrix}, \quad (2.7)$$

where a_i are the amplitudes of the i -th component of the Jones vector, and δ is the phase difference between the two components of the Jones vector. Note that we choose the amplitudes a_i as real-valued as the Jones vector is not unique with respect to complex phase. If the electromagnetic wave in question is monochromatic, the averages on the right hand side of last equality can be removed as a_i and δ are constant in time for such a wave. We note that such a wave is fully polarised.

One possible choice of an orthogonal basis for the space of two dimensional Hermitian

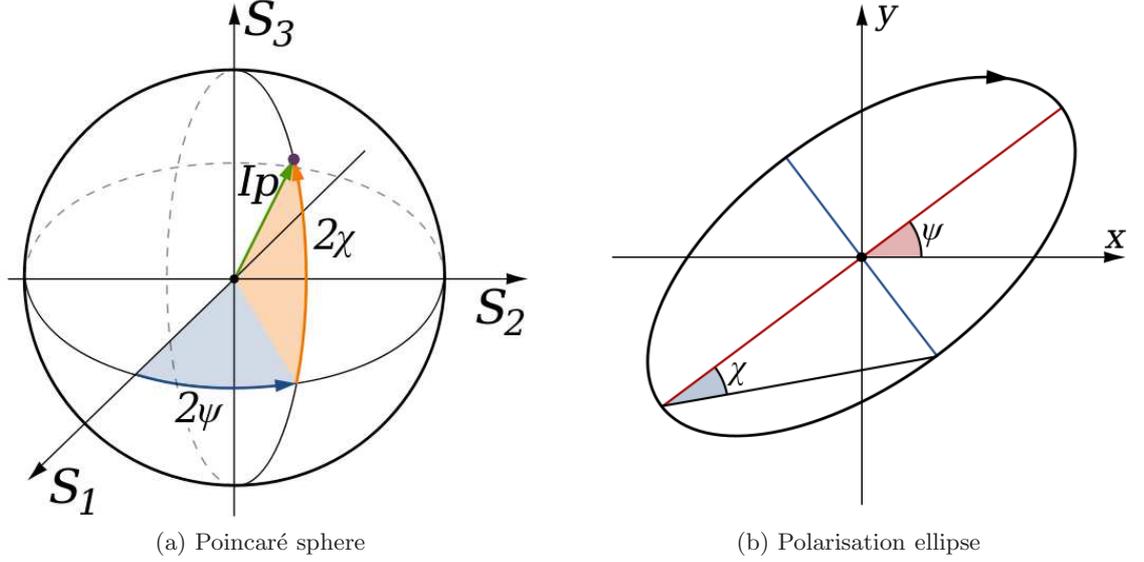


Figure 2.1: Polarisation can be characterised by: a) the Poincaré sphere b) polarisation ellipse. Both images are taken from the public domain.

matrices are the Pauli matrices, i.e., the generators of the $SU(2)/Z_2$ ⁶ or $SO(3)$ ⁷ groups. In the latter view, the i -th Pauli matrix can be seen as a rotation around the i -th axis. This parametrisation provides a convenient way to study the polarisation (e.g., McMaster, 1961). The coherency matrix ρ in this basis is $\rho = 0.5 \sum_{\mu} S_{\mu} \sigma_{\mu}$, where S_{μ} are known as the Stokes parameters⁸. The S_0 is the total intensity of the radiation, often denoted by I ; and S_k represent the polarisation state of the radiation. The degree of polarisation p is defined as $p^2 = I^{-2} \sum_k S_k^2$. The polarisation state can be represented on the so called Poincaré sphere as shown in Fig. 2.1a. The two angles visible on this plot are the position (or orientation) angle ψ , and the ellipticity angle χ . The Fig. 2.1b shows these two angles in the plane of the electric field oscillation. We can now provide two often used parametrisation of the Stokes parameters, without expressing them via the Pauli matrices:

$$\begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix} = \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} I \\ Ip \cos 2\psi \cos 2\chi \\ Ip \sin 2\psi \cos 2\chi \\ Ip \sin 2\chi \end{pmatrix} = \begin{pmatrix} a_1^2 + a_2^2 \\ a_1^2 - a_2^2 \\ 2\text{Re} \left(\langle e_{x_1}^{\dagger} e_{x_2} \rangle \right) \\ 2\text{Im} \left(\langle e_{x_1}^{\dagger} e_{x_2} \rangle \right) \end{pmatrix}, \quad (2.8)$$

where often used naming convention for the Stokes parameters is introduced. The sign of

⁶i.e., the quotient group of 2 by 2 unitary matrices with determinant one and cyclic group of order 2

⁷i.e., the group of rotations in three dimensions

⁸In this chapter Greek indices include values from 0 to 3 while Latin span values from 1 to 3.

V determines the sense in which the field vector is rotating and we say that the radiation has right hand sense of circular polarisation for negative values of V and left hand sense for the opposite sign of V . We note that the opposite convention is also widely used, see van Straten et al. (2010) for details. The linear polarisation is defined as $L = \sqrt{Q^2 + U^2}$. This quadrature summation biases the linear polarisation L towards positive values and results in a non-Gaussian probability distribution of its values. The latter is true also for the distribution of the position angle ψ . Everett & Weisberg (2001) and Naghizadeh-Khouei & Clarke (1993) provide schemes to retrieve unbiased values for L and ψ , respectively.

The polarisation angle of the electromagnetic radiation propagating through vacuum can be affected by an external magnetic field. Polarised radiation from pulsars propagating through the ISM can be affected by the Galactic magnetic field. This is known as the Faraday effect (Faraday, 1845) and is caused by circular birefringence, i.e., a difference in the propagation speed of the left and right hand circularly polarised components, as motivated by the third term in equation 1.12. This causes a change of the polarisation angle by the amount given by $RM\lambda^2$, where λ is the radiation wavelength and RM is the rotation measure:

$$RM = \frac{e^3}{2\pi m_e^2 c^4} \int_0^d n_e B_{\parallel} dl, \quad (2.9)$$

where B_{\parallel} is the component of the Galactic magnetic field along the line of propagation of the polarised radiation and the integration is along the line of sight of length d . Pulsar observations provide a good way to study the Galactic magnetic field by observing the Faraday rotation and dispersion (e.g., Mitra et al., 2003; Noutsos, 2012). Other effects of external fields on propagation of radiation, such as the Kerr effect (Weinberger, 2008), are typically not considered in astrophysical contexts.

2.2 Signal path

In this section we describe what happens to the electromagnetic radiation when it arrives at a radio telescope, starting with the antenna, via the receiver and mixers to the backends. All the statements without references provided are based on the standard textbooks Rohlfs & Wilson (2004) and Lorimer & Kramer (2004). The topics covered in this thesis are a small part of all the relevant considerations which are discussed in detail in the first of these books.

2.2.1 Antenna

The first experiments with radio waves, conducted by Heinrich Hertz (Hertz, 1889) were the first spark to ignite interest in radio. Various people, including Oliver Lodge and Thomas A. Edison, attempted detection of radio waves from the Sun, but to no avail. Radio astronomy had to wait until the 1930s, for Karl Jansky, an engineer working for the Bell Labs, to detect the first radio emission coming from the Milky Way with a directional antenna used for communications (Jansky, 1933). Bell Labs reassigned him to different tasks and he couldn't pursue his idea of building a parabolic antenna to follow up on his discovery. His idea came to life with Grote Reber's backyard telescope. His first radio map of the Galaxy marks the beginning of modern radio astronomy (Reber, 1940). The Second World War slowed down the progress in some sciences, but radar technology was actively developed for military purposes with most of antennas built in paraboloid shapes.

Paraboloid antennas owe their popularity to the geometric properties of this type of a reflector. It is practical in astronomy to use geometric optics as the antenna sizes are much larger than the radiation wavelengths and thus the wave effects are negligible. However, the shape of the surface needs to be accurate to within a fraction of the wavelength. Any deviations from the parabolic shape will cause antennas to lose sensitivity as described by Ruze's equation (Ruze, 1952). Assuming the antenna has perfect shape, all the incident waves coming from a distant source will form a spherical wave front after the reflection and will arrive at the focus at the same phase. To avoid aberrations, in case of a Cassegrain system, the focus of a secondary hyperboloidal reflector must coincide with the focus of the primary reflector. From there the incoming waves are reflected into the feed. Another possibility is placing the feed directly in the primary focus. The disadvantages of the latter solution are: a larger fraction of the antenna's surface will be blocked by the feed and other instruments in the primary focus than by a secondary reflector; and to minimise the spillover radiation from the ground T_{spill} the antenna must be under-illuminated, a problem not present in Cassegrain telescopes. The antenna measures an average of the incident electric field with weighting defined by antenna's Fraunhofer diffraction pattern referred to as the beam shape. The origin of the beam shape is readily understood by considering a transmitting antenna and using the reciprocity theorem. This diffraction pattern defines the antenna's field of view and is described by a Bessel function. The full width at half maximum and the size of primary beam, i.e., distance to the first null are proportional to the ratio of radiation wavelength to the antenna's diameter by multiplicative factors of 1.22 and 2.44, respectively. All other non-vanishing regions of the diffraction pattern are known as the side lobes.

The ratio of the effective area to the geometric area of the antenna characterises the efficiency η of any telescope. This efficiency depends on a number of properties such as: the wavelength of the incident radiation λ ; the illumination efficiency, related to the radiation pattern of the feed, often chosen to minimise the side lobes (e.g. Krauss, 1966; Baars & Swenson, 2008); spillover, i.e., the fraction of incoming power that originates from the ground, e.g., via side lobes; ohmic losses, which can be important at high frequencies or for antennas with a non-metallic layer on the reflecting surface; mechanical displacement of the feed; and blocking of a fraction of antenna’s surface by the support structure of the feed and other instruments near the focus. The description of antennas is necessarily limited here; for an excellent review of parabolic antennas covering many more details see Baars & Swenson (2008).

2.2.2 Feed horn, radiometer

After the antenna has focused the radiation it arrives at the focus, coherent in phase as described in section 2.2.1. The feed is often shaped like a horn guiding the radio waves focused by the antenna along its length. Two dipoles are located at the end of the feed horn. These dipoles can have either a linear or a circular structure. For clarity we will focus on linear dipoles without losing any generality. The linear dipole consists of two perpendicular wires that sample the incident electric field in a plane perpendicular to its propagation. This allows a reconstruction of the full polarisation information by calculating the cross-products between the two input signals to produce all four Stokes parameters, defined in the equation 2.8. The coherent electric field forces the electrons in one or both wires to move in an organised manner, thus producing an electric current in the dipole. Cryogenically cooled low noise amplifiers (LNAs)⁹ are used for the amplification of this current in the receiver. The receiver is placed near the feed horn to minimise losses and any additional noise that could occur if a feed line was used. At this stage, the signal is referred to as radio frequency (RF) as it is at the same frequency as the incident radiation. The RF signal is typically band limited by application of low and/or high pass filters to reject any external signals such as radio frequency interference (RFI) and prevent aliasing, i.e., the spectral reflection of the signal. Some coaxial feed designs allow two separate frequency bands to be received simultaneously like the “10cm/50cm” receiver at the Parkes Radio Telescope (Granet et al., 2005).

A convenient way to measure the power of the incident radiation is by comparing it

⁹Note that as there are several stages of amplification in the signal path, minimisation of noise requires that the lowest noise amplifiers are used as the first stage.

to the temperature of a resistor dissipating power: $P_D = kT_r$, where P_D is the dissipated power, k is the Boltzmann constant, and T_r is the temperature of the resistor. This can be equated to the observed power $P_O = 0.5\eta AS$, where S is the flux density of the observed source, η is the antenna's efficiency, and A is its area. By balancing these two expressions, we can associate temperature T_s with the observed source. From this we obtain:

$$S = \frac{2kT_s}{\eta A} = \frac{T_r}{G}, \quad (2.10)$$

where G is the antenna's gain $G = 0.5\eta Ak^{-1}$. We can now define the system temperature T_{sys} as:

$$T_{\text{sys}} = T_{\text{rec}} + T_{\text{spill}} + T_{\text{atm}} + T_{\text{sky}} + T_s, \quad (2.11)$$

where the various contributions are, respectively: Johnson noise associated with thermal fluctuations in the receiver (typical values ~ 30 K for cryogenically cooled receivers), in pulsar astronomy known as the receiver noise; spillover noise (~ 10 K); atmospheric emission (only important at high frequencies, above 5 GHz); and the sky temperature, which includes the synchrotron emission from the tenuous plasma in our Galaxy (~ 10 K for observations at 1400 MHz but strongly dependent on frequency and direction, with values reaching ~ 800 K for low frequencies in the Galactic plane) and the cosmic microwave background radiation. The contribution of typical pulsar flux density T_s is usually much smaller than 1 K and traditionally neglected. The research described in the remainder of this thesis focuses on the case of a pulsar for which T_s is no longer negligible and considers its effect on the precision of pulsar timing experiments.

By using the antenna's gain G we can define the system equivalent flux density (SEFD), i.e., the flux density corresponding to the system temperature: $\text{SEFD} = T_{\text{sys}}G^{-1}$. At the Parkes Radio Telescope, the receiver typically used for pulsar timing observations is the multibeam receiver (Manchester et al., 2001) which has $\text{SEFD} \approx 27$ Jy in the centre beam. By observing a wide range of frequencies (a wide bandwidth B), using the signal from both perpendicular dipoles probing independent parts of the incident radiation, and integrating over time τ , we can hope to detect a weak pulsar as the variance of noise contributions are going to be reduced by averaging over all the available independent samples of voltage in the receiver. This leads to an expression for a minimum detectable flux (Vivekanand et al., 1982; Dewey et al., 1985; Bhattacharya, 1998):

$$S_{\text{min}} = \beta \frac{S/N_{\text{min}} T_{\text{sys}}}{G \sqrt{n_p \tau B}} \sqrt{\frac{W}{P - W}}, \quad (2.12)$$

where S/N_{\min} is the minimum signal to noise ratio that satisfies our detection threshold; n_p is the number of independent dipoles (one or two) and β is a correction factor due to other losses of signal in the system, such as digitisation errors. The last factor corresponds to the fact that the pulsar's emission is periodic with duty cycle dependent on the width of emitted pulses W and the pulsation period P . The relation 2.12 is often reversed and used to calculate the S/N for a pulsar with known flux S . For PSR J0437–4715 observed at Parkes at wavelength of 20 cm with bandwidth of 256 MHz we expect S/N of the order 10,000 after an hour of integration based on measurement of noise in the off-pulse phase range.

We note that the equation 2.12 is only valid for pulsars with the observed width less than the pulse period. In case of pulsar where the pulsed flux extends over more than a period, for example due to interstellar scattering, a different expression based on considerations of the pulsar signal in the Fourier domain needs to be used (Cordes & Chernoff, 1997, equation A4). For the purposes of this thesis, which only deals with observations of PSR J0437–4715 at a frequency of the order of 1.4 GHz the equation 2.12 provides a valid approximation.

2.2.3 Two stages of downconversion

After the incident radiation is focused by the antenna, fed into the receiver and amplified by the LNAs, the RF signal is mixed (i.e., multiplied) with a monochromatic signal from a local oscillator (LO). The resulting signal is referred to as the intermediate frequency (IF) signal. By combining the signals at two frequencies, RF and the frequency of LO, two beat signals are produced. The beat frequencies are equal to the sum and difference of the RF and LO frequencies. The goal of this mixer is to change the frequency of signal to a range in which the electronics have good impedance characteristic, wave effects can be ignored, and to avoid the feedback effects of the amplifying stages. In practice this means that we are interested in reducing the frequency of the signal as equipment that can handle high frequency signals is very expensive. Therefore, we take this new intermediate frequency (IF) signal to be at the frequency lower than RF and refer to the mixing as downconversion. If the RF is higher than LO frequency, an increase of frequency in the IF band corresponds to an increase of frequency in the RF band, known as the upper sideband. On the contrary, if the LO generates a signal of higher frequency than RF we have the case of a lower sideband with frequency inversion. The IF signal is split and one copy is sent to a second stage mixer for further downconversion and the other copy is sent to data acquisition hardware, such as filterbanks and correlators, i.e., the IF devices

described in section 2.3. Typically, the IF signal is amplified and further filtered to remove artefacts and unwanted signals. Application of low- and high-pass filters generates the desired bandpass thus avoiding frequency aliasing, i.e., the reflection of higher frequencies back into the band of interest. For some modern systems the first stage conversion is an up-conversion, which makes image rejection, i.e., the rejection of the other beat frequency, much easier.

The second stage of downconversion works on the same principles as the first stage of mixing. After the first stage of mixing, the amplified and band-limited IF signal is mixed with a signal from another LO to bring the signal down to baseband, i.e., frequencies ranging from zero to some cut-off frequency. Note that sometimes the baseband signal is referred to as the lowpass signal. The baseband signal is used for detection and integration or for baseband recorders described in section 2.3. Reducing the frequency to the minimum possible allows the lowest achievable sampling rates that still recover full information present in the signal while keeping the sampling clock errors at a minimum level.

The downconversions at both stages are typically performed separately for the two RF signals corresponding to two orthogonal senses of polarisation. There are two types of downconversion: single- and dual-band downconversion; the latter is also known as quadrature downconversion. In the first case, the signal is first band-limited and then mixed with the signal from LO, while in the quadrature downconversion, the RF signal is split into two separate paths. Both these paths are mixed before band limiting, with one path mixed with the LO signal with a phase that has been offset by 90° with respect to the LO signal mixed with the other path. Recombination of the two signals, known as the in-phase and quadrature signal, after digitisation yields the complex-valued representation of the real-valued voltage.

2.3 Hardware and software backends

Both IF and baseband signals are output to data acquisition hardware, referred to as backends. These devices can be analogue or digital and can be used to either record the data or process the received signal in real time. The primary goal of many of the various backends is to record the data and remove the effects of the propagation of the electromagnetic wave through the ISM, especially phase dispersion as described in section 1.2.4. The correction of the dispersion on detected signal is referred to as the incoherent dedispersion as the phase information of the incident radiation is unavailable at the time of dedispersion. Systems which perform dedispersion while the phase information is still available are known as the coherent dedispersion backends.

Recording and processing of pulsar observations requires different hardware and algorithms than those used for most other astronomical data analysis due to rapid time variability of the incident radiation. To resolve the pulsar’s light curve, short time sampling, typically of the order of $10^{-3}P$, is used to record the data in any pulsar backend. At this time resolution the minimum detectable flux density as given by equation 2.12 is very high. To successfully detect the incident radiation from radio pulsars, which have mean flux densities of order of milli-Janskys (Yan et al., 2011), the periodicity of the pulsar is typically employed. Pulsar backends perform “folding”, i.e., a phase resolved summation of the signal over an arbitrarily large number of pulse periods, thus decreasing the minimum detectable flux density. Pulsars can be also detected from single pulse searches or imaging (e.g., Burke-Spolaor et al., 2011; Cordes & Lazio, 1997).

The following sections briefly describe the most commonly used types of devices. Much of the processing in astronomy and many other areas of science and engineering is done in the frequency domain. The observed time domain signals are converted into the frequency domain using the Fourier transform (Fourier, 1822), for an extensive review of properties and applications see Bracewell (2000).

2.3.1 Analogue filterbanks

The delays caused by interstellar dispersion, as calculated in equation 1.16, can be removed in post-processing. However, to minimise the adverse effects of interstellar dispersion within one frequency channel, a sufficient frequency resolution must be generated. As the correction of the frequency dependent delay can be only applied at a post-detection stage in this type of instrument, the desired width of frequency channels depends upon the DM in the direction of a given pulsar. In an analogue filterbank (AFB), the frequency resolution is achieved by splitting the IF signal and feeding it into multiple frequency filters. Each of these frequency channels undergoes digitisation and detection separately. In order to record the output, the data rate can be maintained at a small level by using low bit digitisation schemes. For example, 1-bit digitisation leads to a small data rate, is very robust in the presence of RFI, and effectively provides a low pass filter preventing aliasing effects. However, low bit digitisation has also deleterious effects on the quality of data (Jenet & Anderson, 1998).

In practice the number of channels produced by an AFB is fixed as it depends on the hardware and cannot be adjusted. AFBs are now not normally used for polarimetry as better systems are now relatively inexpensive. An example of an application of the analogue filterbanks is the Parkes Multibeam Pulsar Survey (Manchester et al., 2001), the

most successful survey of the Galactic plane which originally found over 600 pulsars.

2.3.2 Autocorrelation spectrometers

Another type of backend that forms a filterbank but provides much more flexibility in the number of frequency channels than AFBs is an autocorrelation spectrometer (Weinreb, 1963). Here the amplified IF signal is split into two copies. One copy of the signal is delayed by some amount τ and multiplied by the original signal producing the first output signal, or “zero-lag”. The delayed signal is then again delayed and again multiplied by the original signal, producing a second output signal. This procedure can be repeated an arbitrary number of times. In this way, the autocorrelation function of the IF signal is calculated.

The Wiener-Khinchin theorem states that the Fourier transform of the autocorrelation function is the power spectrum, i.e., by transforming the output of the autocorrelation spectrometer we arrive at the same product as from the analogue filterbank. The optimal number of delayed copies of the signal is determined by the Nyquist’s sampling theorem and desired spectral resolution. The digitisation of the signal biases the autocorrelation function with corrections available for one and two bit sampling schemes (van Vleck & Middleton, 1966; Kulkarni & Heiles, 1980).

Polarimetry is much easier in the case of an autocorrelation spectrometer than it is for AFBs. To produce all four Stokes parameters, the two signals corresponding to the orthogonal polarisations of the incident radiation can be processed independently in the way described above. In addition, cross-products are formed by multiplying the two polarisations yielding all the quantities necessary to form four Stokes parameters.

Examples of autocorrelation spectrometers used for pulsar astronomy are the Caltech Fast Pulsar Timing Machine (Navarro et al., 1992) and Wideband Arecibo Pulsar Processor (Dowd et al., 2000) backends.

2.3.3 Polyphase digital filterbanks

With the progress of digital technology, all of the signal processing can be improved by shifting the implementation to purely digital designs by employing Field-Programmable Gate Arrays (FPGAs), specialised digital signal processors. FPGA contains a large number of programmable logical gates, allowing an engineer to design custom data paths. In particular, an implementation of the discrete Fourier transform (DFT) is easily achieved with the FPGA chip. DFT suffers from two significant problems, i.e., spectral leakage and scalloping loss. A monochromatic input signal can appear in multiple frequency bins

due to the finite length of the input signal, producing a sinc function in the frequency response, inducing the first effect. The latter effect is the loss of signal between frequency bin centres due to the shape of the single-bin response. Both problems can be alleviated by implementing a polyphase filterbank¹⁰, i.e., not performing the Fourier transform of the input signal in one step. Instead, the signal is weighted by a Fourier transform of the desired frequency channel response, and split into a chosen number of shorter time series. The latter number is known by the name of “taps” of the polyphase filterbank. The taps are then summed point-to-point before undergoing a Fourier transform. While a polyphase filterbank is more computationally costly to implement than a plain DFT filterbank, the increased quality of the resulting output justifies the additional cost.

Polyphase digital filterbanks are flexible devices with the user providing required frequency and time resolution at the beginning of observing with reprogramming of the FPGA lasting a negligible fraction of time of observing. Higher quality of data and improved properties allow higher precision in pulsar timing experiments. All four Stokes parameters can be easily derived and subsequently calibrated. For more details on polyphase filterbanks see the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) website¹¹ and references therein.

2.3.4 Coherent dedispersion systems

Another class of commonly used devices contains the baseband recording and real-time processing systems. All of the backends described so far are incoherent dedispersion systems which do not remove the dispersion of radiation within each frequency channel. It is possible to use the phase information of the incident radiation and recover the original signal using coherent dedispersion (Hankins, 1971; Hankins & Rickett, 1975). As the phase information of the incident radiation is necessary for successful removal of dispersion, coherent dedispersion requires pre-detection baseband data. Many such systems, progressively improved, existed at the Parkes Observatory, including two generations of the Caltech Parkes Swinburne Recorders.

Hankins & Rickett (1975) demonstrated that the effects of the ISM on the radiation propagating through it can be described as a multiplicative transfer function H in the frequency domain. By multiplying the Fourier transform of the observed signal by the inverse of the transfer function, the original signal is recovered to within an arbitrary phase shift and a term linear in frequency of the observed radiation. The phase shift is unrecoverable

¹⁰Also known as the weighted overlap-add or window pre-sum-FFT method

¹¹https://casper.berkeley.edu/wiki/The_Polyphase_Filter_Bank_Technique

as it is lost during square-law detection, while the latter term corresponds to a time shift as per the shift-theorem (Bracewell, 2000). Typically, in any timing considerations, we consider ToA of signals at infinite frequency, unaffected by the dispersion thus determining this time shift. The only term left in the transfer function is a term quadratic in frequency and dependent on the measured dispersion measure. Often a low-pass filter, or “taper” is applied at the same time as the inverse of the transfer function to limit aliasing. The combination of taper and H^{-1} is known as the chirp function.

The advantage of coherent dedispersion is the recovery of instrumental time resolution by increasing the practical maximum width of the frequency channels. In practice, to increase the S/N, time resolution is lowered but coherent dedispersion produces many sharper features in the observed pulse profiles. As the precision of timing is related to the width of the profile, coherent dedispersion enables higher precision of timing experiments. Coherent dedispersion requires more computational resources but with the progress of technology near real-time coherent dedispersion systems became possible in the 2000s. Many of the coherent dedispersion backends are implemented in software rather than hardware.

For a long time the main limitation of the coherent dedispersion backends was the lack of technology able to sustain processing of high data rates. This has limited the bandwidth, sampling rate and duration of observations with these systems. With the progress of electronics and doubling of the processing capabilities every two years (so called Moore’s law Moore, 1965), many real time coherent dedispersion systems were developed during the early 1990s and early 2000s. Hotan et al. (2006) used such a backend implemented in software, the Caltech Parkes Swinburne Recorder mark 2 (CPSR2 Bailes, 2003), to monitor 15 MSPs over a 3 year period. They measured proper motions and parallaxes for many of the observed pulsars, constrained models of plasma distribution in the Galaxy, and Shapiro delays. They concluded that the high fidelity of observations using coherent dedispersion has been the single most important factor in achieving all the results and better timing precision than observations with other more sensitive telescopes without the implementation of coherent dedispersion. With the progress of computational resources, coherent dedispersion systems can now process as wide bands as many of the incoherent systems and are a better type of instrument.

2.3.5 Baseband recorders and offline data processing

During the same time when the increased processing capabilities allowed real time coherent dedispersion, fast hard drives became available at low cost. In this way, instead of

processing the data in real time, the raw voltages can be recorded to the hard drive with high time resolution over a wide frequency range. Storing baseband data is advantageous as the stored data can be analysed using processing methods that cannot be used in real time. As new data analysis techniques or mistakes in previously applied analysis are discovered, the data can be loaded and processed again. An example of such an application is presented in Chapter 5 which discusses an analysis of single pulses from PSR J0437–4715, originally recorded as baseband data with the CASPER Parkes Swinburne Recorder.

Another technique of signal processing recently introduced to pulsar analysis, and typically performed offline, is the cyclostationary spectroscopy (Gardner & Cadzow, 1990). This technique is based on the fact that the statistics of radiation originating from pulsars are periodic in time, rather than stationary as assumed in all the previously described methods. The initial application in pulsar astronomy was by Feliachi (2010), who employed it to identify and remove radio frequency interference. Demorest (2011) was the first to apply it directly to the pulsar signal based on the model of pulsar radio signal as an amplitude modulated noise (Rickett, 1975).

Cyclic spectroscopy removes the dependence between frequency and time resolution of pulsar observations. This is achieved by incorporating the non-stationarity of the baseband signal into the analysis. In this case, the frequency resolution is not limited by the time resolution but by the total integration time. Cyclic spectroscopy is an ideal tool to study the ISM, separate its effects from pulsar intrinsic phenomena and remove their detrimental impact on timing precision. The cyclic spectrum contains the information about the signal's phase. This allows one to recover the original emitted signal corrected for the effects of multipath propagation, i.e., the scattering as opposed to coherent dedispersion which only allowed correction of the interstellar dispersion delays. Currently no backends exist that provide cyclic spectra in real time and the baseband data need to be processed off-line.

2.3.6 Impact of polarimetric calibration on timing precision

Most modern precision pulsar timing systems provide all four polarisation cross-products and subsequently all four Stokes parameters can be formed. However, the polarisation information needs to be calibrated before being used for most scientific purposes. The necessity of calibration arises from a variety of reasons that introduce differences between the measured and intrinsic Stokes parameters. Hans Müller noticed that all the distortions

to the polarised signal can be contained in a matrix transfer function:

$$S_m = \mathbb{M} \times S_i, \quad (2.13)$$

where \mathbb{M} is known as the Müller matrix acting on the intrinsic Stokes parameters S_i to produce the measured Stokes parameters S_m (e.g., Hecht, 2001). The Müller matrix can be decomposed into several components, each corresponding to a different effect on S_i . Firstly, many telescopes have an “alt-azimuth” mount for which the feed rotates with respect to the sky during the observation. This causes mixing between the Stokes Q and U parameters. Furthermore, the two dipoles in the feed are not perfectly orthogonal, causing mixing between the Stokes U and V parameters and/or Lorentz boosts of the Stokes four-vector, depending on the choice of feed type. The signals from the two orthogonal polarisations are then processed by a different set of amplifiers along different cables, introducing differential phase and gain into the Stokes parameters. The first attempts at measuring the matrix \mathbb{M} and correcting the Stokes parameters were necessarily simplified by the limited amount of computational power available at the time (e.g., Stinebring et al., 1984; Turlo et al., 1985). More recently, Johnston (2002) solved equation 2.13 for one pulsar without any additional assumptions about the orthogonality of the receptors.

Since then, significant progress has been achieved by van Straten (2002, 2004, 2006) and van Straten (2013). In this series of papers, the author presents an equivalent formulation of the problem using the Jones calculus following on ideas first developed by Britton (2000) and Hamaker (2000). He also derives formulae for the expected improvement of timing precision for full polarisation data. These papers also present progressively improved methodologies for determining the effects of the antenna and instrumentation on the polarisation. The most up-to-date method employs one pulsar as a polarised reference source to calibrate other pulsars observed with the same system, significantly reducing the amount of effort required to obtain well-calibrated data.

PSR J0437–4715 has been studied by numerous authors with steadily improving instrumentation and progressively employing more sophisticated calibration techniques. The pulse profile of this pulsar exhibits a sudden change in the position angle as noted by (Navarro et al., 1997) which warrants careful polarimetric calibration if high timing precision is to be achieved (van Straten, 2006). Sandhu et al. (1997) observed the pulsar using an auto-correlation spectrometer with 128 MHz of bandwidth, and could model pulse arrival times over two years with a post-fit residual standard deviation of 500 ns. Noting that the formal uncertainty of arrival time estimates was typically around 50 ns, the authors concluded that their results were limited by polarimetric calibration errors. Britton (2000)

first proposed the use of the Stokes invariant interval to mitigate the problems caused by polarisation calibration. This was later implemented by van Straten et al. (2001), who used a combination of 16 MHz and 20 MHz baseband recording systems, typical integrations of 1 hour duration, and coherent dedispersion to obtain a root-mean-square (rms) timing residual of 130 ns over 3.4 yr. Using new and improved calibration methods developed by van Straten (2004, 2006) and a new baseband recording and processing system (CPSR2; Bailes, 2003; Hotan, 2006) with 128 MHz of bandwidth, Verbiest et al. (2008) achieved 199 ns over 10 years.

2.3.7 Systems used in this thesis

All of the data presented in this thesis were recorded at the Parkes Observatory, part of the Australia Telescope National Facility, funded by the Commonwealth of Australia for operation as a national Facility and managed by CSIRO. Several observing backends were used in the presented studies of PSR J0437–4715. The findings reported in chapter 3 are based mostly on a third generation of the polyphase digital filter bank (PDFB3). CPSR2 and the ATNF Parkes Swinburne Recorder (APSR), a software backend allowing coherent dedispersion and baseband recording, were used to demonstrate some arguments in this chapter. Chapter 4 is based solely on data recorded with CPSR2. The fifth chapter is based on a new generation of baseband recorder with coherent dedispersion capability, CASPER Parkes Swinburne Recorder which contains several improvements over APSR. These include wider maximum bandwidth, use of graphics processing units (GPUs), real time RFI mitigation, and time-division multiplexing instead of frequency-division multiplexing using a two-stage analysis-synthesis polyphase filterbank, that introduces significant spectral leakage.

3

Stochastic wideband impulse modulated self-noise

The first two chapters of this thesis have briefly outlined many of the areas of physics and astrophysics that are relevant for neutron stars. The observations of pulsars provide a unique opportunity to study all these topics, otherwise studied disjointly or not possible to explore experimentally whatsoever. This and the next chapter both discuss the limitations of the pulsar timing experiments and how to attain the best possible timing precision. The results presented are currently directly relevant for a handful of objects, but as the next generations of radio telescopes are built, they will become applicable to the majority of pulsar population and thus to most of the pulsar timing experiments. In this way, by presenting means of achieving the highest timing precision, this thesis contributes to all the areas of physics discussed in the introductory chapters.

3.1 Introduction

Chapter 1, among other information, states that the most fundamental property of radio pulsars was their periodic series of radio pulses that enable both their discovery and a myriad of timing applications. A sub-class of pulsars, namely the MSPs, with their short periods and low braking torques are especially good clocks. These pulsars exhibit the highest rotational stability (Matsakis et al., 1997). The wealth of information already derived from the precision timing of radio MSPs has led many authors to predict the kind of pulsar timing science possible with the Square Kilometre Array (SKA) by linearly extrapolating timing precision attainable with current telescopes to that of the SKA.

Section 1.2.4 listed examples of the rms timing residuals achieved for a number of MSPs. We also pointed out that none of the studies of PSR J0437–4715 achieved the timing precision predicted by the formal uncertainty in arrival time estimates. When observing PSR J0437–4715 in the 20 cm band at the Parkes 64m observatory, the expected

rms timing residual from a 256 MHz bandwidth with 21 K system temperature is about 10 and 80 ns in one hour and one minute of integration, respectively. We remind that these uncertainties are derived from the template-matching method used for pulsar timing (Downs & Reichley, 1983; Taylor, 1992; Bailes, 2010) and summarised in equation 1.10.

It is generally assumed that the summation of many hundreds or thousands of pulses approaches a stable pulse profile that is characteristic of the pulsar (Helfand et al., 1975). Consideration of only additional white noise, $N(t)$, in the aforementioned equation is equivalent to assuming that the system equivalent flux density (SEFD) is the only significant source of noise. However for bright sources and/or high gain antenna, this assumption is no longer tenable in at least two circumstances. Firstly, when the flux density of the pulsar approaches the SEFD of the receiver additional noise proportional to the pulsar’s flux density becomes significant (e.g. Kulkarni, 1989; Gwinn, 2001, 2004, 2006; van Straten, 2009; Gwinn & Johnson, 2011; Gwinn et al., 2011). Secondly, it is known that each single pulse can have very different morphology and can occur at different pulse phases (Drake & Craft, 1968b; Helfand et al., 1975; Jenet et al., 1998). Even after integrating over many pulse periods, this single-pulse modulation¹ can introduce detectable fluctuations in the average profile shape and thereby contribute additional noise to timing data. We discuss the noise balance in more detail in section 3.2, and argue that these two contributions should be considered together as they are related and are described by the same statistical model.

As previously noted in section 1.2.4 and 2.3.6, pulse shape changes are likely to be as prevalent in MSPs as they are in classical pulsars and have been detected previously in PSR J0437–4715, although not studied in detail. With the progress of polarimetric calibration methods, we will demonstrate that these variations are not likely to be caused by calibration errors as was the case in some of the previous studies.

As described in more detail in section 3.2, this chapter focuses on the stochastic fluctuations in total intensity that arise from the sub-pulse structure observed in single pulses and their effect on the timing precision attainable for PSR J0437–4715. In section 3.3 our observations are described along with the applied data processing followed by results of timing our observations. In section 3.4 we describe a statistical method useful for detecting profile shape variations, then apply it to simulated data as a demonstration of how it can be used to correct ToA residuals. The results of the statistical analysis are presented in section 3.5. We summarise our findings and discuss their consequences in section 3.6, which also contains a discussion of other possible problems that prevent us from reaching

¹The term “phase jitter” (e.g. Cordes & Downs, 1985; Cordes, 1993; Cordes & Shannon, 2010) is also used to describe this phenomenon.

the theoretical timing accuracy. Finally we draw our conclusions in section 3.7.

3.2 Stochastic wideband impulse modulated self-noise summary

The noise $N(t)$ in equation 1.10 is normally assumed to be dominated by the white radiometer noise. In practice, for bright sources and/or high gain antennae, two additional effects may contribute significantly. In this section we discuss the noise balance for pulsars and demonstrate that the different contributions are well described by a single statistical model.

Firstly, the noise balance has to include the source itself when the flux density of the pulsar approaches the SEFD. This noise is intrinsic to the source and is accordingly called “self-noise”. In the case of PSR J0437–4715, the mean flux at the peak of the pulse profile is of the order of 5 Jy; this contributes only $\sim 2\%$ to the standard deviation of the total intensity, which is dominated by the SEFD ~ 27 Jy of the 20 cm multibeam receiver (Manchester et al., 2001) commonly used for pulsar timing observations at Parkes. Any source that can be described as noise (e.g. thermal emission) will contribute to the variance of the observed total intensity of the source. When the signal-to-noise ratio (S/N) is low, this contribution is negligible. We note that throughout the thesis we use S/N values calculated using the noise measured in the off-pulse baseline.

Secondly, dramatic single-pulse amplitude modulation is a ubiquitous feature of radio pulsar emission (Rickett, 1975) that spans orders of magnitude in intensity and duration and is a broadband phenomenon (e.g. Staelin & Reifenstein, 1968; Taylor et al., 1975; Manchester et al., 1975; Hankins et al., 1993; Hankins & Eilek, 2007; Wang et al., 2007). The single-pulse emission from PSR J0437–4715 is well studied; Jenet et al. (1998) observed an exponential distribution of peak single-pulse intensities with a mean flux density of 16.6 Jy, which is comparable to the SEFD. More importantly, the mean sub-pulse width of $65 \mu\text{s}$ (Jenet et al., 1998) is about an order of magnitude larger than the time resolution used to sample the mean pulse profile of PSR J0437–4715. Consequently, single-pulse intensity fluctuations introduce detectable variations in the average pulse profile. Given the stochastic nature of single-pulse structure, these fluctuations in mean pulse profile shape can be expected to introduce significant additional noise in derived arrival time estimates.

In one minute, PSR J0437–4715 turns $\sim 10^4$ times and emits at least $\sim 2 \times 10^3$ sub-pulses (Jenet et al., 1998). After integrating over such a large number of emission events, it is no longer practical to consider the impact of an individual single-pulse. Rather, it becomes necessary to describe the effects of single-pulse modulation in purely statistical terms using the fourth moments of the electric field (Rickett, 1975). From this perspective,

the single-pulse modulation phenomenon is a noise process that contributes to the auto-correlation of the total intensity (Rickett, 1975). For a given source flux density, amplitude modulation increases the variance of the total intensity and, depending on the time-scale of the modulations and the sampling interval of the instrument, introduces power at non-zero delays in the auto-correlation of the total intensity.

Measured statistical distributions of single-pulse intensities vary between sources and as a function of pulse longitude (e.g. power law, log normal, etc.; for an excellent review, see Cairns, 2004). Regardless of the original distribution, after a large number of pulses have been integrated, the central limit theorem applies and profile shape variations are well described by a multivariate normal distribution. The covariance matrix that quantifies this distribution contains phase-resolved information about the mean auto-correlation of the total intensity.

To summarise, depending on the pulsar’s flux density, its emission properties and the used instrument, we can distinguish three noise regimes:

1. *First regime:* The pulsar’s flux density is much smaller than the SEFD of the instrument. This is the classic regime, in which the noise is temporally uncorrelated between the phase bins and homoscedastic (i.e., the variance of noise is the same in each phase bin). In this regime the covariance matrix of the pulse profiles is well approximated by a diagonal matrix with all the elements on the diagonal equal to the variance of SEFD.
2. *Second regime:* The pulsar’s flux density approaches or exceeds the SEFD of the receiver used. In this regime the self-noise cannot be neglected. The noise is still temporally uncorrelated between the phase bins, but it is heteroscedastic; that is the variance of the noise is different in each bin and proportional to the sum of squares of the pulsar’s flux density and the SEFD. The covariance matrix of the data is still diagonal, but the non-zero elements are no longer equal. In this regime, the on-pulse noise is no longer measured by the off-pulse noise and using the latter to calculate the S/N can lead to overestimating the achievable timing precision.
3. *Third regime:* The pulsar is heavily amplitude modulated with the modulated flux approaching or exceeding the SEFD of the instrument. Even though the self-noise contribution may be negligible, the modulated single pulses can approach the SEFD of the receiver thus contributing significant “noise” to the averaged pulse profile. If the sampling rate is high enough to resolve the single-pulse structure, the noise in different phase bins will be heteroscedastic and temporally correlated. The off-

diagonal elements of the covariance matrix will be non-zero in this regime. If the single pulses are not resolved, amplitude modulation may still be evident in the variation of the modulation index as a function of pulse phase. The broadband nature of the impulses will also lead to spectral correlation of the noise, which can be detected by measuring the covariance of intensity fluctuations in different frequency channels. Therefore, when analysing only the covariance matrix stochastic wideband impulse modulated self-noise (SWIMS) might be confused with self-noise.

In this chapter, we investigate the effects of the third regime, which we call SWIMS, on pulse arrival time estimation. This pulsar-intrinsic noise has also been called pulse-phase jitter (or jitter noise), “intermittent emission” (Gwinn et al., 2011) or simply self-noise. We will demonstrate that the timing precision of PSR J0437–4715 is currently limited by SWIMS and that its effect on the mean pulse profile is readily detectable.

As a function of integration length T , the covariance matrix scales as T^{-1} ; that is, the effects of SWIMS are reduced by integration, regardless of the dominating source of noise. The fact that the covariance matrix scales as T^{-1} allows us to study the statistics of single pulses with longer integrations. If any pulse-to-pulse correlation were present in data, such as arising from drifting sub-pulses, nulling, mode changing, scattering or polarisation calibration errors, the scaling of the covariance matrix with integration length would deviate from the above proportionality. We note that the relative contribution of source-intrinsic noise to the covariance matrix will vary as the flux density of the pulsar varies, primarily owing to interstellar scintillation. However, after averaging over many scintillation time-scales, the relative contribution of SWIMS compared to the SEFD is constant (Kulkarni, 1989); therefore, the relative importance of SWIMS is independent of integration length. We note that in the first or second regime, noise can be reduced by increasing the bandwidth. However, because the intensity fluctuations are typically correlated over wide bandwidths, the noise due to single-pulse modulation is not reduced by increasing the bandwidth and only longer integration times and active mitigation can improve timing precision in the third regime.

Gwinn et al. (2011) recently performed a detailed analysis of impulse-modulated self-noise in the context of interstellar scintillation observations and concluded that self-noise may limit pulsar timing precision. In this chapter, we explore the impact of both temporal and spectral correlations of intensity fluctuations on pulsar timing and consider active mitigation of SWIMS.

3.3 Observations and data processing

Observations of PSR J0437–4715 were recorded during one week of February 2010 using the Parkes 64m radio telescope and the central beam of the 20 cm multibeam receiver (Staveley-Smith et al., 1996). The third generation of the Pulsar Digital Filterbank (PDFB3) digitised the voltage data from two orthogonally polarised 256 MHz bands and formed 1024 frequency channels using a polyphase filterbank. After full polarisation detection (following the definitions described by van Straten et al., 2010), the data were folded at the topocentric period of the pulsar into 1024 phase bins. The mean polarisation profile was output every minute and a total of 25 hours of data were recorded. The multibeam receiver is equipped with a noise diode that is coupled to the receptors and driven with a square wave to inject a pulsed polarised reference signal into the feed horn. This signal was recorded for three minutes before and after every 64 minute observation of the pulsar. These data are in the Parkes Pulsar Data Archive² (Hobbs et al., 2011a).

The data are stored in the PSRFITS format and all processing was done with the PSRCHIVE data processing software suite (Hotan et al., 2004b). First we ensured that a recent model for the pulsar spin, astrometric and orbital parameters (Verbiest et al., 2008) was used throughout for our data processing. To remove narrow band radio frequency interference (RFI), median filtering was applied by comparing the total flux in each frequency channel with that of its 49 neighbouring channels. To avoid distortions at the edge of the observing band, we rejected five per cent of the frequency channels on each side of the band. A search for impulsive RFI was performed with the “lawn mower” method³. Pulsar observations were calibrated for polarisation as in van Straten (2004). The flux density was calibrated by observing the Hydra A radio galaxy which is assumed to have a flux of 43.1 Jy at 1400 MHz and a spectral index of -0.91 (Baars et al., 1977).

A high S/N template for each frequency channel was created by integrating the observations obtained during the first day of data and then used to identify and remove data affected by broadband impulsive RFI as follows. First, for each frequency channel, the best-fit scale, baseline offset and phase shift (Taylor, 1992) between the profile and template were applied to compute the difference between the template and the data. Second, the rms flux of this difference was computed after integrating over frequency at the DM value of the pulsar and at zero DM. RFI will induce a high rms flux at zero DM while, at the pulsar’s DM, impulsive interference will be smeared across multiple phase bins. Hence, if the difference has an rms flux value at zero DM that is higher than at the pulsar’s DM

²<http://datanet.csiro.au/dap/>

³<http://psrchive.sourceforge.net/manuals/paz>

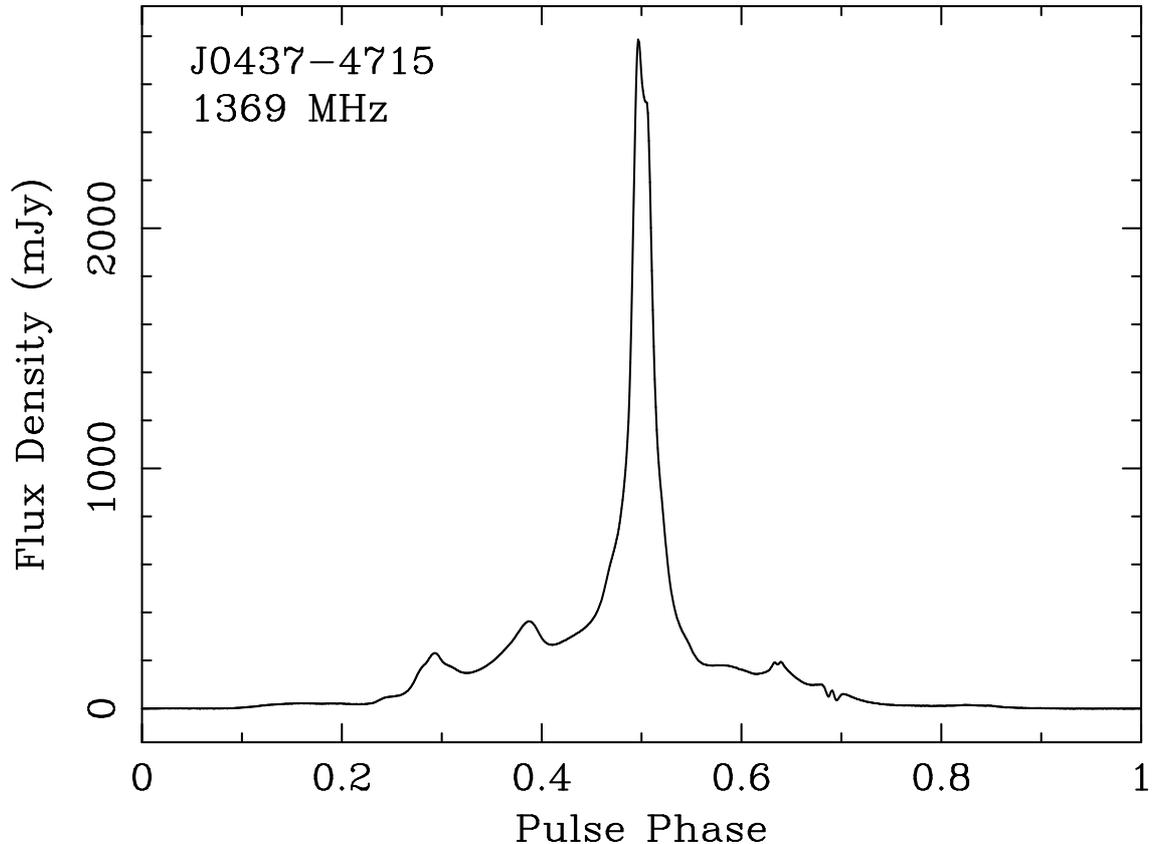


Figure 3.1: The high S/N (15,000) template for PSR J0437–4715, created from 5.5 hours of observations. The solid black line represents the total intensity.

the profile is potentially polluted by RFI. A few hundred profiles were examined by eye. The rms ratio at both DMs for the RFI polluted difference profiles allowed the determination of a threshold ratio above which the remaining profiles were automatically tagged as being affected by RFI and rejected from further analysis. After all the RFI removal stages we were left with 1145 one minute integrations, or about 70 per cent, of the complete dataset, that are considered RFI free. From the 5.5 hours of observations taken during the first day of observing, we created a final, frequency integrated total intensity template shown in Fig. 3.1 with a S/N of 15,000.

The ToA of each observation not affected by RFI in the remaining six days was determined by cross-correlation with the template (Taylor, 1992). Timing residuals were formed from these ToAs and the pulsar model using the TEMPO2 software package (Hobbs et al., 2006). The ToA residuals for these data are shown in Fig. 3.2 as a function of ToA number. Note that they are not evenly spaced throughout each day. The mean ToA estimation error is only 72 ns and the mean S/N is 770. The weighted rms of the timing

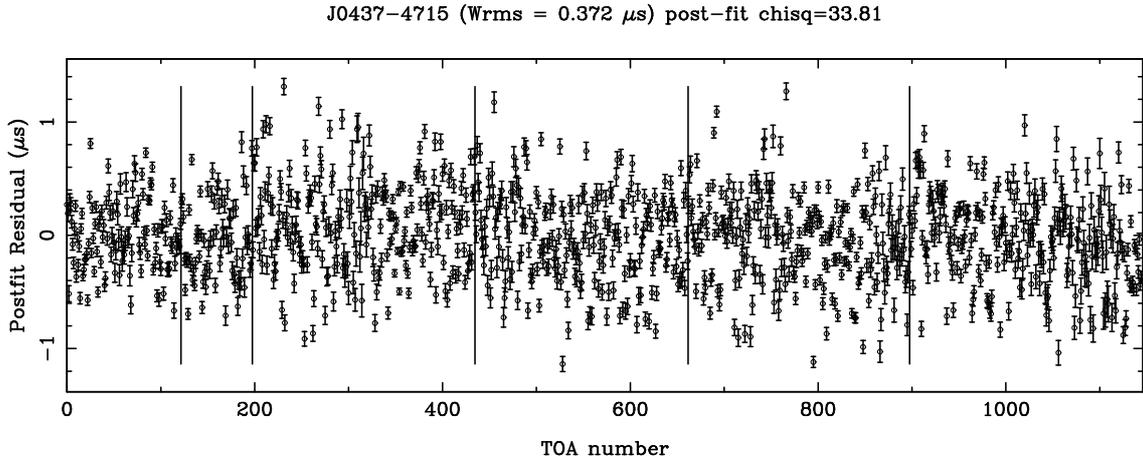


Figure 3.2: Timing residuals for 6 days of data timed against the standard from Fig. 3.1. The mean ToA estimation error is of the order of 72 ns, whereas the weighted rms of the residuals σ_{ToA} is 372 ns. The fit has $\chi^2/\text{d.o.f.}^2$ of 33.8. For clarity we have plotted the residuals as a function of ToA number. The vertical lines are plotted between observations taken on different days.

residuals however is $\sigma = 372$ ns and the reduced chi squared of the fit, $\chi^2/\text{d.o.f.}$, where d.o.f. denotes the number of degrees of freedom, is 33.8. The unweighted rms timing residual is similar: 389 ns.

The high $\chi^2/\text{d.o.f.}$ value could be caused by underestimation of the ToA uncertainty or because the pulsar model does not accurately predict the observed ToAs. To verify the estimated arrival time uncertainties, we carried out a Monte-Carlo simulation in which each observed profile is replaced by an exact copy of the template with a suitable amount of white noise added to yield a S/N equal to that of each observation when averaged over many realisations. The $\chi^2/\text{d.o.f.}$ of the timing residuals is always very close to unity in these simulations, implying that the ToA uncertainties are calculated correctly under the assumption of equation 1.10.

In the above simulation, the white noise added to each simulated pulse profile is statistically independent of the noise in every other profile; in this case, as profiles of roughly equal S/N are integrated together, the rms of arrival times derived from the integrated totals will be roughly proportional to $T^{-1/2}$, where T is the integration length. Consequently, the statistical independence of errors in arrival time estimates is commonly verified using a plot of residual rms as a function of integration length, as shown in Fig. 3.3. Here, the thin line indicates the mean theoretical expectation based on 10^5 simulations of 64 ToAs derived from template with white noise added. The shaded region shows the 95% confidence levels derived from the same simulations. The thick line represents the

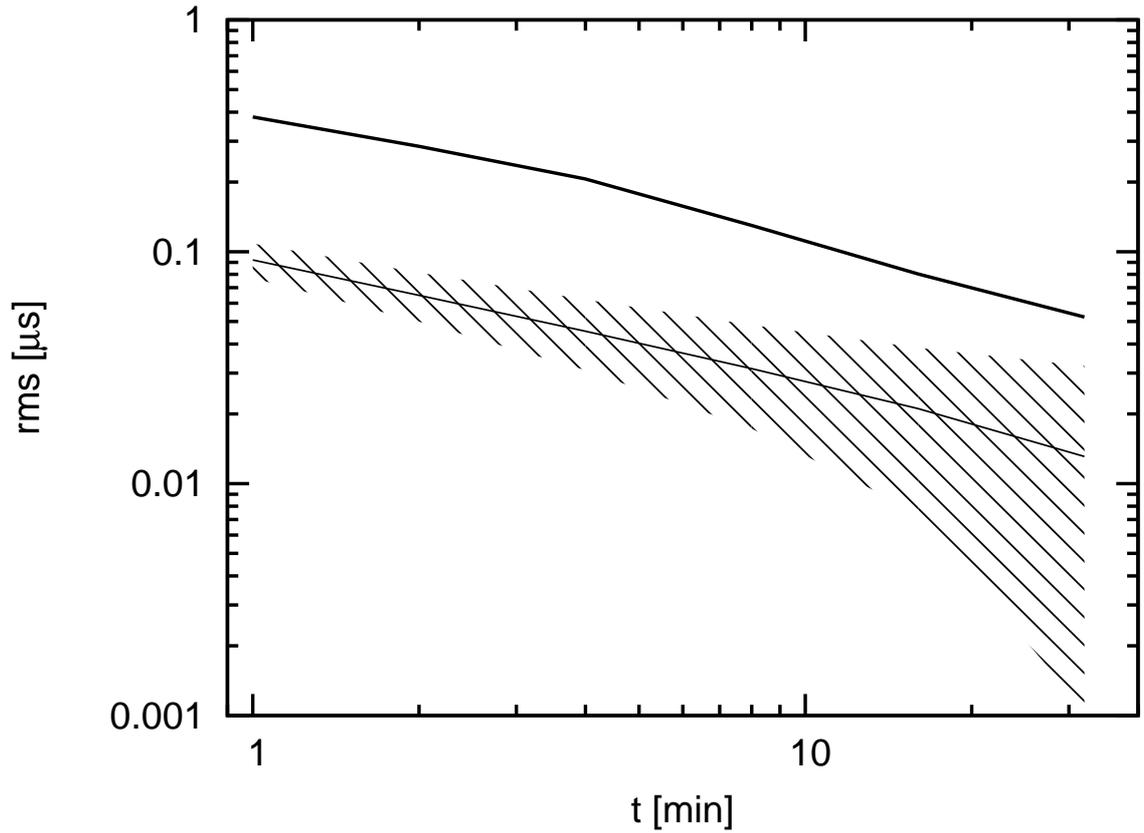


Figure 3.3: Comparison of achieved timing precision with the theoretically attainable precision as a function of integration time. The line in the dashed region plots the theoretical rms, equal to the mean rms obtained from 10^5 different realisations of white, homoscedastic noise. This values agree with expectations based on the width of the mean pulse profile and the S/N . In 95% of the simulations, the rms falls within the dashed region. The confidence interval is much broader at long integration lengths because a fixed number of initial pulse profiles is used; hence, at large t , fewer independent instances remain, thus biasing the estimate in a random manner.

measured rms timing precision. The deviation of this line from $1/\sqrt{t}$ behaviour is solely due to the small number of points at longer integrations, e.g. only 2 points are available at 32 minute integration length. The fact that the observed rms follows the expected behaviour suggests that no pulse-to-pulse correlations or anti-correlations are important in our dataset. We note that the precision measured is far worse than the theoretical expectation. For example, with 32 minute integrations, we expect an rms timing residual of ~ 13 ns but we observe a typical value of 52 ns. This factor of ~ 4 worse than the theoretical prediction implies that, if this problem was understood and fully corrected, then the same observing precision could be achieved with integrations 16 times shorter than currently required.

The above simulation does not include any self-noise or single-pulse modulation; therefore, the predictions in Fig. 3.3 are those expected from the radiometer equation, see equation 2.12. However, in the real data the variance of the noise in the off-pulse region underestimates the variance and completely neglects the temporal correlation of the noise in the on-pulse region; that is, the actual noise is heteroscedastic and correlated. Increasing the gain of the antenna will amplify SWIMS and while the S/N calculated using only the off-pulse noise will increase the rms timing residual will not decrease. Therefore, longer integrations will be necessary to achieve a lower rms timing residual. More importantly, the SWIMS in the total intensity from single-pulse modulation is spectrally correlated. Fig. 3.4 shows a greyscale image of a single pulse from PSR J0437–4715 as function of pulse phase and radio frequency taken with the ATNF Parkes Swinburne Recorder (APSR; van Straten & Bailes, 2011). The emission clearly extends across the entire observed band, producing a high degree of spectral correlation of single-pulse intensity fluctuations. To see if ToAs from independent bands are correlated we divided our template and each one minute observation into two independent frequency channels and determined the timing residuals for each channel separately. Fig. 3.5 shows the timing residuals plotted against each other for one hour of data processed in this way and shows that the ToAs are highly correlated between the two channels (the average Pearson product moment correlation coefficient between the two sub-bands is 0.91). If the single-pulse modulation contribution to SWIMS had no impact on the timing residual, no such correlation would be present. In addition, the rms timing residual of each of the two sub-bands, as well as the combined rms timing residual is similar to the value obtained for the combined data. Under the assumption that the broadband single-pulse properties of PSR J0437–4715 are responsible for the scatter in the timing residuals, increasing the observing bandwidth will not reduce the SWIMS component due to single-pulse modulation.

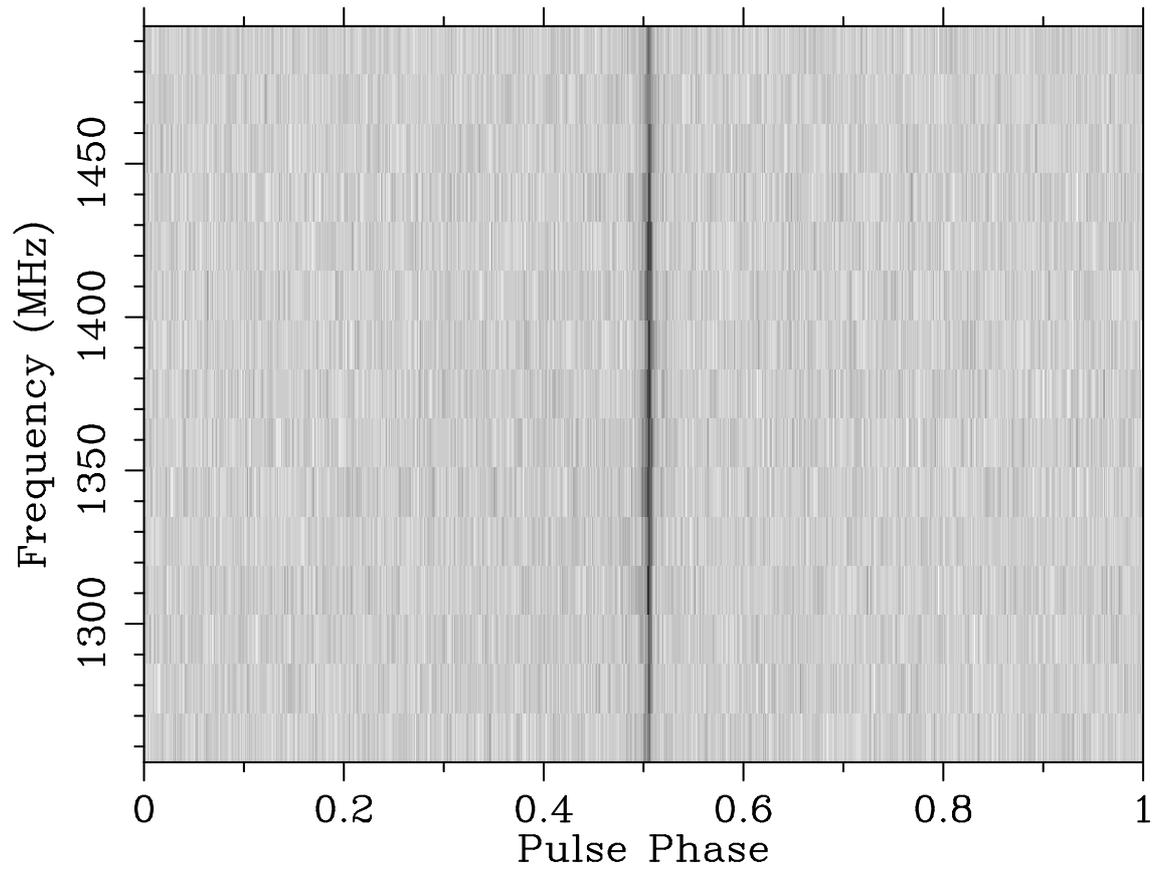


Figure 3.4: Greyscale image of a single-pulse from PSR J0437–4715 as a function of pulse phase and frequency taken with the APSR instrument. We stress that (a) the single-pulse persists across the whole available band, and (b) that each single-pulse is very different from the average profile. See Fig. 5.4 for a comparison of a single-pulse with the average profile.

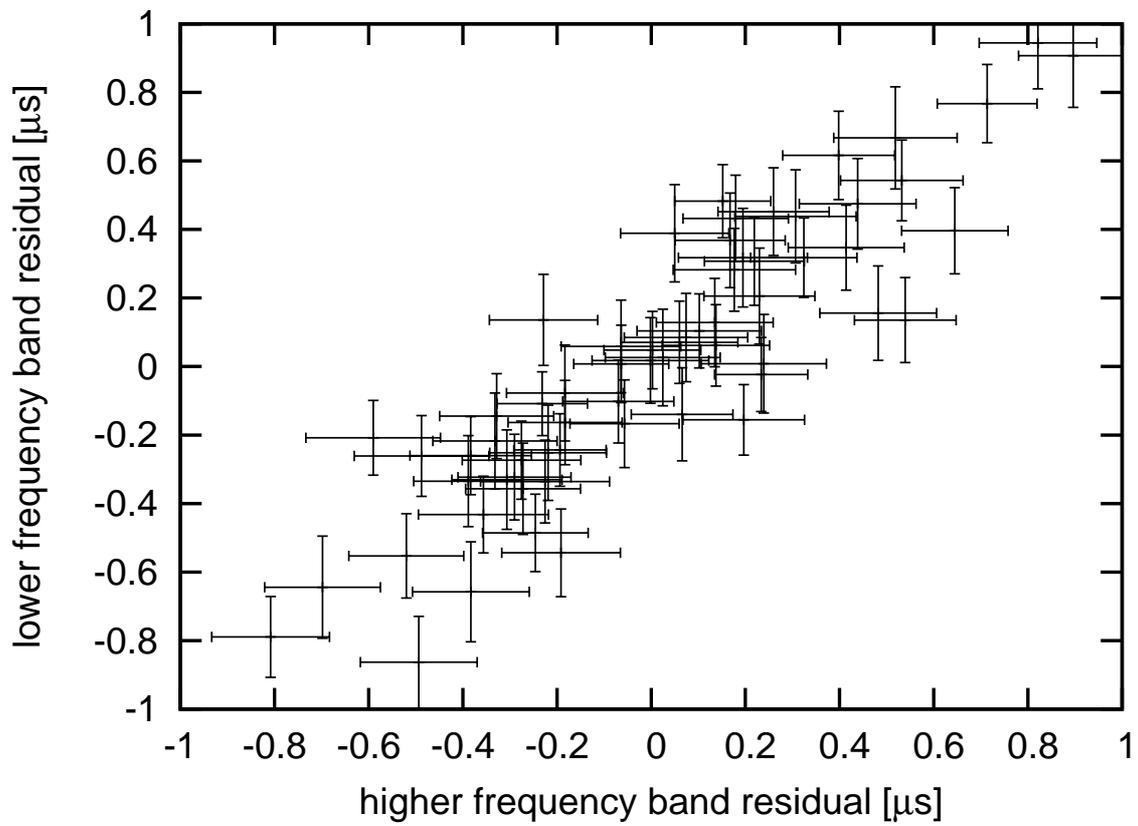


Figure 3.5: Timing residuals derived from two independent sub-bands every 60 seconds plotted against each other to demonstrate their remarkable degree of correlation.

The heteroscedastic and both temporally and spectrally correlated properties of SWIMS can lead to significant random error⁴ in arrival time estimates derived from sources with amplified flux densities comparable to the system equivalent flux density. The following sections report on an investigation of one possible method of correcting these random errors.

3.4 Method

As explained in the previous section, failure to account for the statistical characteristics of SWIMS leads to measurement error and underestimation of arrival time uncertainty. In this section, we explore the use of principal component analysis (PCA) to correct the random error through a series of simplified simulations. These simulations do not model the large number of impulsive intensity fluctuations; rather, the simulations demonstrate that the PCA model corrects only the arrival time random error due to profile shape variations and that no other sources of phase noise are incorrectly mitigated. An overview of PCA is given in Hyvarinen et al. (2001). We extend the analysis introduced by Demorest (2007) and then present a number of illustrative simulations that demonstrate the validity of our method and its implementation. A very similar methodology has been under development by Cordes and his collaborators since the 1990s (private communication, Cordes, 1993).

The PCA method provides a rigorous and unbiased statistical method for analysing temporally correlated variations in total intensity. For each one-minute observation of PSR J0437–4715 $\sim 15 \times 10^6$ samples are integrated in each of the N_{bin} pulse phase bins; therefore, by the central limit theorem, the fluctuations in total intensity are well described by a multivariate normal distribution. If the distribution of these fluctuations was strongly non-normal, better performance might be achieved by a similar method using independent component analysis (Hyvarinen et al., 2001). We note that the number of pulses integrated in each minute is approximately an order of magnitude larger than the number of pulses considered in previous studies of profile stability (Helfand et al., 1975; Rathnasree & Rankin, 1995).

Assume that N such observations have been made of a given pulsar. We describe the profile for the i 'th observation as a column vector⁵ \vec{p}^i . The j 'th element, p_j^i , is the

⁴SWIMS causes a systematic error in the estimate of ToA, however, this bias will be different for every ToA as it will have a different SWIMS realisation, resulting in random errors in all ToAs

⁵Our notation is defined as follows: All matrices are denoted by a font with a hat (e.g. \hat{M}). All vectors are denoted by a font with an arrow (e.g. \vec{v}). An element of a matrix is denoted by M_{ij} . An i 'th column of a matrix is denoted by \hat{M}_i . An i 'th element of a vector is denoted by v_i . If there are multiple vectors of a given type, we denote the i 'th vector by superscripting, e.g. \vec{v}^i . The indices i and j always are in range $[1, N]$ and $[1, N_{\text{bin}}]$, respectively.

amplitude of the j 'th bin in the i 'th profile. The covariance matrix is typically computed after subtracting the mean of all observations from each observation. Here, we assume that the template, \vec{s} , is a good estimate of the mean profile and, before subtraction, each observation is first adjusted to match the template using the best-fit phase shift, scale and offset as derived from the template-matching procedure used for pulsar timing (see section 3.3). We then form the covariance matrix of the dataset by computing the outer product of template matched profiles:

$$\hat{C} = \frac{\sum_{i=1}^N w_i (\vec{p}^i - \vec{s}) (\vec{p}^i - \vec{s})^T}{\sum_{i=1}^N w_i}, \quad (3.1)$$

where w_i is the S/N of the i 'th profile, the T superscript denotes transposition and the prime superscript signifies that the profiles have been matched to the template. The resulting covariance matrix, \hat{C} , is a symmetric matrix with the number of rows and columns equal to the number of bins, N_{bin} , in each profile. We note that at least N_{bin} observations are necessary for \hat{C} to have full rank. Furthermore, the data set should be large enough so that all potential modes of profile variation are represented.

Template matching before subtracting the standard profile removes three degrees of freedom from the shape fluctuations that are intrinsic to the pulsar signal. For example, to first order, the best-fit phase shift removes all variations that correlate with the derivative of the standard profile with respect to pulse phase. Removing variations with a certain profile shape is equivalent to projecting the N_{bin} -dimensional vector space of the total intensity fluctuations onto the $N_{\text{bin}}-1$ -dimensional subspace that is orthogonal to the axis defined by that profile shape. A significant amount of the fluctuation information may be lost by this projection. However, if the best-fit phase shift were not first removed, any actual phase shifts would be misinterpreted as shape variations; therefore, this dimension must be excluded from the analysis. Similarly, the best-fit scale and offset remove variations in pulsar flux and system temperature, respectively; these fluctuations are not the focus of this chapter. In practice, all the data are fit for phase shift, flux scale and baseline offset and these three dimensions are always projected out of the available vector space in which the pulse profiles are described. The eigenvectors corresponding to these three dimensions all have the same eigenvalue (zero) and hence together form an eigenspace; we refer to the template matching eigenspace as the fit-space throughout the remainder of the chapter.

To characterise the remaining fluctuations of the total intensity, we solve the eigenproblem of the covariance matrix. The eigenvectors \vec{e}^j define the principal axes in the N_{bin} -dimensional vector space of profile shape variations along which the intensity fluctuations are correlated as a function of pulse phase. Sorting the eigenvectors in order of decreasing eigenvalue, λ , allows us to determine the most significant variations. The variance corresponding to each eigenvector is equal to the corresponding eigenvalue.

The eigenvectors form an orthonormal basis onto which each residual difference profile can be projected. The projection coefficient, α_{ij} , of the i 'th difference profile onto j 'th eigenvector is

$$\alpha_{ij} = (\vec{p}^i - \vec{s})^T \vec{e}^j. \quad (3.2)$$

These coefficients can be thought of as the residual of the i 'th pulse profile in the basis spanned by the eigenvectors \vec{e} and are often referred to as the principal components.

After the subspace projection that removes the phase shift, scale, and offset dimensions, the remaining projection coefficients for each residual profile are uncorrelated; i.e.,

$$\sum_{i=1}^N \alpha_{ij} \alpha_{ik} \propto \delta_{jk}, \quad (3.3)$$

where δ_{jk} is the Kronecker delta. However, these coefficients may possibly be correlated with unobservable variations in the three d.o.f. that have been removed. These correlations are exploited by the technique developed in section 3.4.1, where we introduce a new method for using these projection coefficients to correct the timing residuals for the random error introduced by pulse shape variability. Simulations to confirm our algorithm are presented in section 3.4.2.

3.4.1 Correcting the timing residuals: multiple regression

Demorest (2007) measured the correlation between the first projection coefficient (corresponding to the largest variance in the data) and the arrival time residuals. This was subsequently used to detect corrupted data and he has shown that it could be used to remove their deleterious effect on the timing residuals. However, his method only used the information stored in the first projection coefficient. Here we apply a multi-variate statistics method of multiple regression to simultaneously remove the effects of multiple varying components.

To predict the random error in ToA estimate, I_i , caused by SWIMS we assume that there is a linear function relating this random error to the projection coefficients. We use

the observed timing residuals, R_i , to determine the best-fit parameters for this function.

We wish to predict I_i using the linear predictor

$$I_i = a + \vec{A}^T \hat{\alpha}_i \quad (3.4)$$

where a and \vec{A} are the regression coefficients of the linear predictor. These are determined from the observed residuals by minimising the mean squared error between the predicted and observed residuals: $\sum_i (R_i - I_i)^2$. The analytic solution is given by (Johnson, 1998):

$$\vec{A} = \hat{D}^{-1} \gamma, \quad (3.5)$$

and

$$a = \nu + \vec{A}^T \mu, \quad (3.6)$$

where

$$\hat{D} = \frac{\sum_{i=1}^N (\hat{\alpha}_i - \langle \hat{\alpha} \rangle) (\hat{\alpha}_i - \langle \hat{\alpha} \rangle)^T}{N}, \quad (3.7)$$

is the covariance matrix of the projection coefficients, γ is the vector of covariances between the residuals and the projection coefficients

$$\gamma_j = \frac{1}{N} \left((\hat{\alpha}^T)_j - \langle (\hat{\alpha}^T)_j \rangle \right) (\vec{R} - \nu), \quad (3.8)$$

and ν is a vector with each element equal to the mean of the observed residuals. The elements μ_i are the mean values of $\hat{\alpha}_i$ and $\langle . \rangle$ denotes average of vectors. The values of \vec{A} and a allow us to predict the random error in ToA estimation induced by pulse shape variations. This predictor has minimum mean square error and maximum correlation with the R_i . Subtracting the predicted error from the estimated arrival times has the potential to reduce the post-fit arrival time residual rms.

The expected improvement in timing residuals can be calculated from the projection coefficients and the observed residuals as

$$\frac{\sigma'}{\sigma} = \sqrt{1 - \rho^2}, \quad (3.9)$$

where

$$\rho = \sqrt{\frac{\gamma^T \hat{D}^{-1} \gamma}{\sigma^2}}. \quad (3.10)$$

Here σ' is the rms timing residual for ToAs with the SWIMS-induced random error removed

and ρ is called the population multiple correlation coefficient.

It is necessary to restrict the number of eigenvectors to model only pulse variability. Many approaches have been put forward in the literature for determining the number of significant principal axes (Johnson, 1998). We introduce a new parameter, ξ_j , which is the Pearson’s product moment correlation coefficient between the timing residuals, R , and the projection coefficients onto the j ’th eigenvector, that is the j ’th row of $\hat{\alpha}$. The standard deviation of non-significant ξ values is determined. Significant values are identified using a Tukey’s bi-weighting scheme starting with an initial guess of the standard deviation obtained from the median absolute deviation. The resulting standard deviation of the ξ values is a robust and resistant estimator. For more details see Andrews (1972) and Hoaglin et al. (1983). Starting from the last correlation coefficient we search for three consecutive ξ values that are more than three times this measured standard deviation. The number of the eigenvectors used in all subsequent processing is equal to the index of last of these three values, when counting from one. For cases in which fewer than three ξ values are significant, we compare the results of using only one eigenvector with those of using the first five.

An implementation of this method is publicly available as a part of the PSRCHIVE suite. The relevant application is called “psrpca” and it requires the GNU Scientific Library⁶ to work.

3.4.2 Simulations

To confirm that our method correctly detects pulse shape variations that can be used to correct random error in ToA estimation, we carry out three simulations of data with noise in the third regime (i.e., SWIMS). In every simulation, each observed profile is replaced by a copy of the template profile plus white noise and additional varying components. The amount of white noise is set such that in many random realisations of the simulated observation the mean S/N of the simulated profile would match that of the given observation. Although there can be several sub-pulses per pulse period, we illustrate the technique using a simple model in which only a single pulse is added per minute of observation. Note that this procedure does not affect the observing parameters (such as frequency, observation time etc.) which are held fixed at the values in the actual observations. The resulting simulated observations are cross-correlated with the template to form ToAs (and hence residuals) in exactly the same manner as the actual observations.

We initially tested the trivial cases of simulated data with white radiometer noise only

⁶<http://www.gnu.org/software/gsl/>

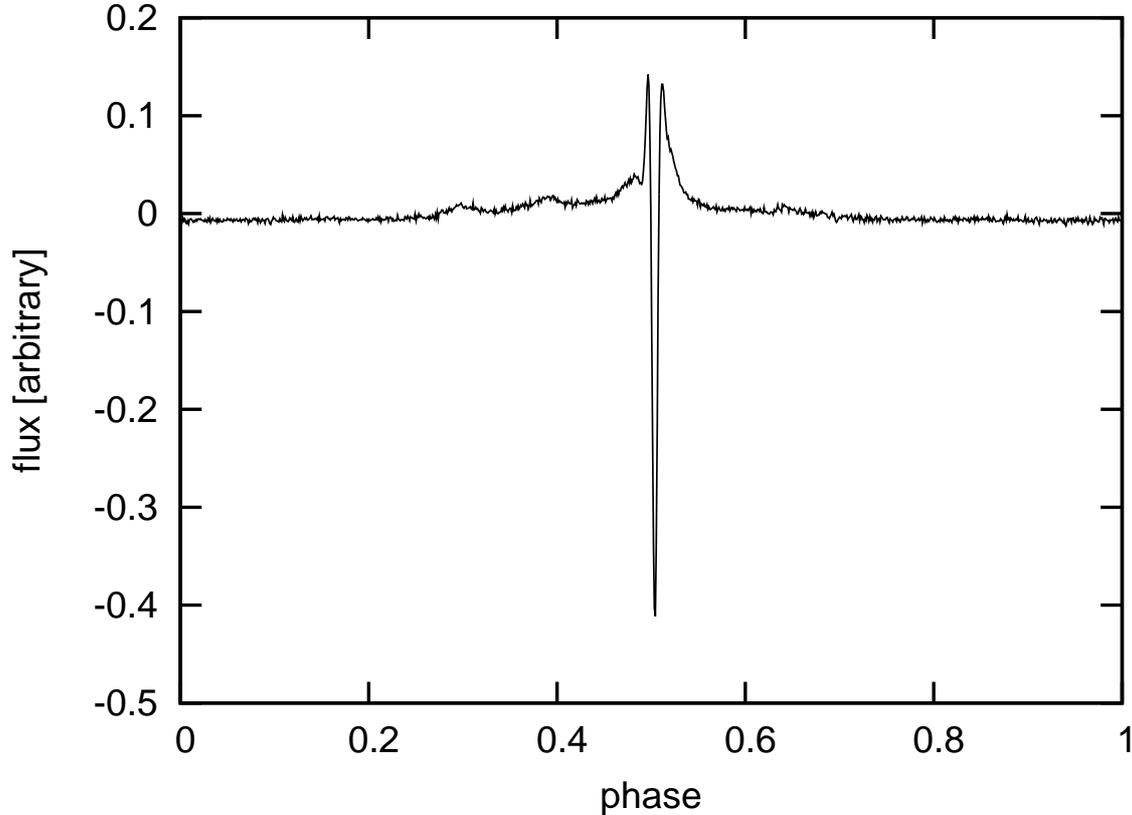


Figure 3.6: First eigenvector for simulation with white noise, arbitrary shifts and a von Mises component on top of the template profile. Notice that this eigenvector does not look purely like von Mises distribution as it has to be orthogonal to the fit-space.

with or without arbitrary phase shifts applied to the data. As expected, no significant eigenvectors were detected in either case. Attempting to correct the residuals regardless of that yields no significant improvement in the rms timing residual. These two cases demonstrate that our method will not artificially decrease the arrival time residual arising from white noise or arbitrary phase shifts. The latter could arise from any phase shift such as that due to a gravitational wave and it is important that such signal is not removed by the PCA.

Simulation 1: Single, fixed component

Pulsar emission is often modelled as consisting of multiple Gaussian components (e.g. Kramer et al., 1999). Many pulsars show mode changing where one or more components are active for only a finite amount of time (Wang et al., 2007). To verify that the random error introduced by a single “mode changing” component can be detected and corrected, we now include an extra component in the profile that varies in amplitude. This compo-

ment is created from a von Mises distribution (which is a periodic analog to a Gaussian distribution) defined by:

$$v\left(\frac{j}{N_{\text{bin}}}\right) = A_v \cdot e^{\kappa_v(\cos((j/N_{\text{bin}} - \mu_v)) \cdot 2\pi) - 1}, \quad (3.11)$$

where j is the phase bin index for which the value of the component is calculated, A_v is the amplitude of von Mises component, κ_v is the concentration parameter, and the component is centered on the phase μ_v . The amplitude is drawn from a normal distribution with zero mean and an rms of 3% of the peak template flux. The component added in this simulation is centred at pulse phase $\mu_v = 0.504$ and κ_v is set to 0.1131. The resulting weighted rms timing residual was 255 ns and $\chi^2/\text{d.o.f.} = 14.4$, while without the additional component the same realisation of white noise leads to an rms of 59 ns and a $\chi^2/\text{d.o.f.}$ of 1.0. We emphasise that this increased rms residual and high $\chi^2/\text{d.o.f.}$ is due solely to the pulse shape variations and therefore should be detected and corrected using our method.

We obtain a significant first eigenvalue. The automatic determination of useful eigenvectors fails, as there is only one significant eigenvector, associated with the introduced profile variation. Fig. 3.6 shows the first eigenvector, which is different from the introduced component as it is projected onto the $N_{\text{bin}} - 3$ dimensional space, in which the three dimensions corresponding to the fit-space are removed. Although these degrees of freedom have been removed, the residual shape variation is still highly correlated with the ToA residual, as discussed above. Correcting the residuals using the one significant eigenvector reduces the rms residual to 67 ns and the $\chi^2/\text{d.o.f.}$ to 1.0. We note that this improvement agrees very well with prediction from equation 3.9. We calculate reference residuals by simulating and timing data without any additional components but with the same realisation of white noise as in the simulation with the varying component. As expected the corrected residuals are highly correlated with the reference residuals (correlation coefficient of 0.93). This indicates that we have recovered most of the signal in the original residuals and have removed the effect of the pulse variability.

We note that after removing the random error in ToA residuals the weighted rms timing residual is almost the same as in the case of reference residuals; however, the $\chi^2/\text{d.o.f.}$ can be smaller than that of the reference residuals because the ToA uncertainties have increased. The total intensity fluctuations decrease the cross-correlation between the observation and the template, thereby increasing the estimated uncertainty (see equation A10 from Taylor, 1992).

In order to check the effect of using a different number of eigenvectors to correct the

Table 3.1: Parameters of the multi-component simulation.

	1	2	3	4	5	6
μ_v	0.504	0.496	0.512	0.524	0.509	0.520
A_v	0.74%	0.36%	0.86%	1.4%	0.7%	0.4%
κ_v	0.452	0.164	0.098	0.050	0.577	0.104

timing residuals, we re-analysed the data using five eigenvectors. In this case the rms residual and $\chi^2/\text{d.o.f.}$ were 64 ns and 0.92 respectively. The corrected residuals are still pleasingly highly correlated with the reference residuals. We therefore conclude that our result is not highly sensitive to the number of eigenvectors used, but care needs to be taken when choosing that number.

In addition, we also tested an extended version of this simulation, where arbitrary phase shifts were included to simulate, e.g., the effect of gravitational waves. As expected, only the random error in the timing residuals due to the introduced profile variation is removed and the arbitrary phase shifts remain unaffected. Again, the residuals are highly correlated with reference residuals where the same realisations of white noise and arbitrary shifts were introduced.

Simulation 2: Multiple, fixed components

The previous simulation dealt with only a single varying component while many pulsars can emit radiation simultaneously from several components. Even if only one of the multiple components was present in each rotation of the pulsar, several components would be present after integration over multiple pulse periods. To demonstrate again that PCA and multiple regression do not remove phase shifts that are not caused by pulse shape variations, we now introduce six von Mises functions whose amplitudes are allowed to vary independently. The parameters of these components are presented in Table 3.1. We also apply arbitrary phase shifts after adding the varying components. This leads to a weighted rms residual of 372 ns and $\chi^2/\text{d.o.f.}$ of 32.1. Note that the reference residuals obtained from a simulation with the same realisation of white noise and arbitrary phase shifts, but no additional components, have an rms of 206 ns and $\chi^2/\text{d.o.f.}$ of 12.2 due to the arbitrary phase shifts; i.e., we do not expect the rms to be of the same order as the ToA measurement error after random error removal.

Our method gives corrected residuals with weighted rms residual of 216 ns ($\chi^2/\text{d.o.f.} = 10.8$) and they are highly correlated with the reference residuals. With multiple components the random error is not completely removed because more than one projection coefficient correlates with the arrival time residual and these projection coefficients may

not be statistically independent of each other. Nevertheless, significant improvement in the rms residual and in the $\chi^2/\text{d.o.f.}$ is apparent and much random error is removed. The additional shifts, corresponding to unmodelled timing noise processes are still unaffected; i.e., the post-correction residuals are highly correlated (correlation coefficient of 0.87) with the reference residuals.

Simulation 3: Single, random component

We now consider a possibly more realistic case of the pulsar emission being erratic and distributed in phase over the whole region in which the average profile is visible; this case corresponds to the stochastic nature (Cordes & Downs, 1985; Cordes, 1993) of modulated pulses. In this simulation we allow a single von Mises component per simulated profile to be centred anywhere in within the central peak (central phase uniformly distributed between 0.479 and 0.518). The concentration parameter of the component is uniformly distributed between 0.055 and 1.386 $1/\text{bin}^2$. The amplitude of this component is normally distributed with an rms equal to 2.3% of the intensity of the template at the centre phase of the component. This value is chosen in order that the resulting rms timing residual is 380 ns and $\chi^2/\text{d.o.f.}$ value of 35.1, both very close to the observed values for PSR J0437–4715.

Multiple significant eigenvectors are detected, with the number varying between different realisations from 10 to 50. Correcting the residuals reduces the rms to 266 ns and $\chi^2/\text{d.o.f.}$ to 17.2. We note that in this simulation no arbitrary shifts were included and therefore we conclude that the PCA method has failed to completely remove the random error in ToAs (and hence in residuals) induced by this kind of erratic shape variation. As explained in section 3.4 some fraction of the fluctuation power is lost during the vector subspace projection effected by removing the best-fit shift, scale, and offset. As before, for data integrated over multiple pulse periods, more than one varying component per pulse profile would be present and this could make it even more difficult to remove the random error in ToAs.

3.5 Results

Applying our method to the observed dataset of 1145 profiles leads to:

- the detection of significant pulse shape variations with at least ten significant eigenvectors,
- a reduction in rms timing residual from 372 ns to 294 ns and a reduction in $\chi^2/\text{d.o.f.}$ from 33.8 to 21.1.

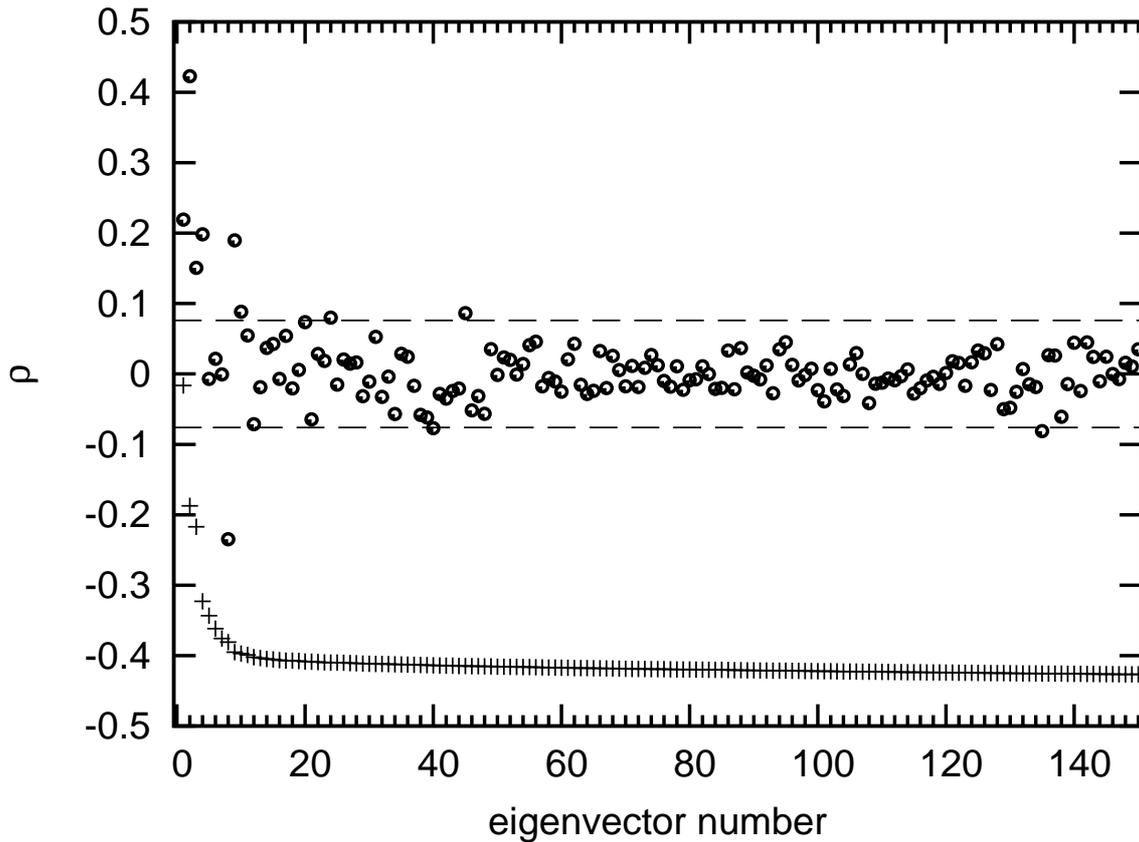


Figure 3.7: Distribution of correlation coefficients between the residuals and projection onto eigenvectors for the actual observations, shown by open circles. Several correlation coefficients are more than 3 standard deviations (denoted by the dashed lines) above the background noise level. Only the first 150 correlation coefficients are shown, the rest of the correlation coefficients looking very similar to the last few shown. The crosses are showing the corresponding eigenvalues normalised by their sum, multiplied by a factor of 10 and offset by -0.45 for clarity of the figure.

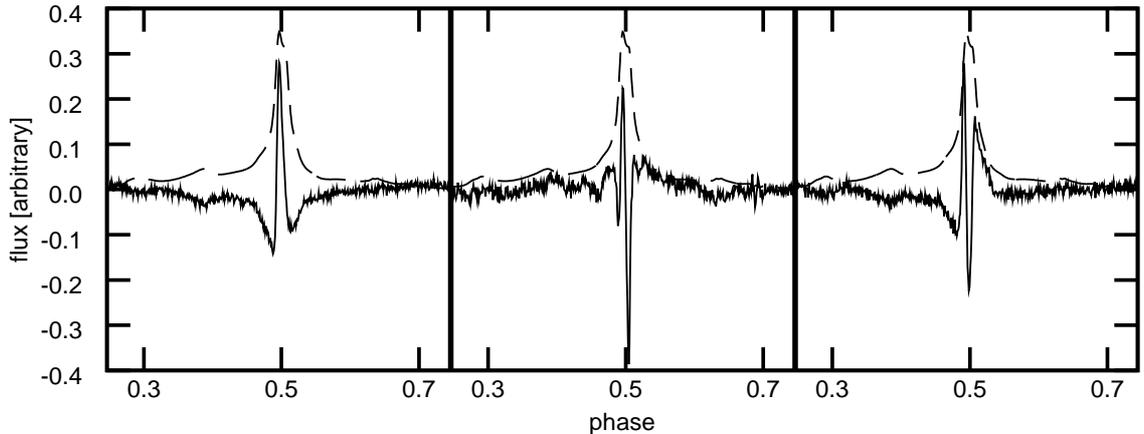


Figure 3.8: First three eigenvectors for the observed data. We show only the central part of the profile for clarity. The vertical thick lines separate the eigenvectors. The dashed line represents the scaled template added for reference.

The first 150 values of ξ are plotted in Fig. 3.7 which indicates that around 10 eigenvectors significantly correlate with the observed residuals while the rest are consistent with being white noise. The choice of exact number of eigenvectors would be difficult to make without the rigorous criteria described before. In Fig. 3.8, we present the three most significant eigenvectors overlaid on the pulsar template profile. The detection of significant eigenvalue-eigenvector pairs is a direct consequence of temporally correlated fluctuations in total intensity; i.e., significant shape variations are detected. As discussed in sections 3.1, 3.2, and 3.3, it is likely that such variations originate from the pulse-to-pulse variations of the pulsar emission. The most significant variation occurs at phase 0.497, in the main peak of the pulse profile. The majority of other significant eigenvectors peak around the main and second highest peak in the mean profile. One exception is the 10th eigenvector that peaks at phase 0.385, that is in the local peak on the left hand side of the main peak.

Using the 10 most significant eigenvectors to correct the random error in arrival time estimates, the rms timing residual was reduced by 20% and $\chi^2/\text{d.o.f.}$ was reduced by 36%. Only 3% of the variance in the timing residuals can be attributed to the most significant eigenvector; therefore, multiple regression is required to provide the best estimate of SWIMS-induced random error. The third simulation has demonstrated that pulse profile variability can introduce random error in ToAs that cannot be removed completely using our method and the timing residuals after random error correction are still affected by these errors. Therefore, even though only 20% of the rms timing residual is corrected, there need not be another explanation for the remaining timing noise. For completeness, other effects that might contribute to the rms of timing residuals are discussed in the next

section.

To investigate if the intensity fluctuations arise from a stationary stochastic process, at least in the wide sense, we used five hours of observations of PSR J0437–4715 during July 2009 obtained using the same observing system as described in section 3.3. These additional data were processed to form pulse profiles, arrival times and timing residuals in exactly the same way as our main dataset. We used the eigenvectors and regression coefficients obtained above to correct these July 2009 timing residuals. We again achieved a 22% decrease in the rms timing residual from 380 ns to 296 ns and the $\chi^2/\text{d.o.f.}$ was reduced by 36% from 39.1 to 24.9. The fractional improvement is similar to before, implying that the covariance matrix and hence profile variability is stationary in time and that after the regression coefficients and eigenvectors have been determined for one dataset they can subsequently be applied to other observations of the same pulsar obtained with the same instrumentation and observing parameters.

3.6 Discussion

We first discuss the method that we have introduced to search for and correct pulse shape variability. Second, we discuss the astrophysical implications of broad-band profile shape variations intrinsic to the pulsar.

3.6.1 Discussion of the method

It is clear from the simulations that our method successfully detects significant pulse shape variations and partially corrects the random error induced in timing residuals due to such variations in many cases. The method has certain limitations. First, in order for the covariance matrix, \hat{C} , to have full rank, the method requires at least N_{bin} observations. With fewer observations than phase bins, it may be necessary to average adjacent phase bins, or, if prior knowledge suggests that the pulse variability occurs in only a restricted region then only these bins could be included in the analysis. If full phase resolution is required then the covariance matrix can be determined, but not all of the eigenvectors can be calculated. Alternatively, following Demorest (2007), the PCA method can be developed in the frequency domain using only the significant harmonics thus reducing the number of required observations.

A large number of observations is desirable for two reasons. Firstly, dividing a fixed observing time into smaller intervals provides a greater number of estimates of the pulse profile, thereby increasing the S/N of the covariance matrix estimate. Secondly, even

for $N > N_{\text{bin}}$ our method is limited by the SEFD present in all observations. Such noise will reduce the precision with which the eigenvectors may be determined, reducing their ability to fully describe the shape variations. For instance, in the timing residuals of simulations with only white radiometer noise added, even though no significant eigenvalues were measured, the rms timing residual is reduced by a negligible amount when using just a few eigenvectors. When all of the eigenvectors are included in the random error removal, the effects of white noise are artificially reduced. This occurs because the white radiometer noise affects the regression coefficients a and \vec{A} in equation 3.4. Since the eigenvectors are measured using the data that are being “corrected”, the noise in the eigenvectors correlates with the noise in the data. This is similar to the “self-standardizing” effect that arises when the mean of a set of observed profiles is used as the template to derive arrival times from the same data, as described in Appendix A of Hotan et al. (2005). Applying the eigenvectors to a completely independent data set leads to no significant change to the timing residuals as the noise in the data no longer correlates with noise in the eigenvectors. The degree of correlation between individual residual profiles and the white-noise eigenvectors may also be reduced by increasing the number of observations from which the covariance matrix \hat{C} is estimated. If the eigenvectors are obtained from the data to be corrected, then it is essential to apply the rigorous criteria described in section 3.4.1 to choose only the significant eigenvectors.

Many pulsar observations are affected by radio frequency interference (RFI) and the measured eigenvectors are extremely sensitive to the presence of RFI in the data. As pointed out by Demorest (2007), PCA is also a sensitive and robust method for detecting RFI and other types of data corruption and distortions. It is therefore essential that the data used in forming the eigenvectors are unaffected by RFI. In our case nearly 32% of the pulse profiles had to be rejected. RFI might also be mitigated through the use of a robust estimator of the covariance matrix, such as the minimum covariance determinant estimator (Hubert & Debruyne, 2010).

In contrast to more traditional applications of the PCA method, our method relies on first aligning each pulse profile to the template using the best-fit phase shift, scale, and baseline offset. As a result of this fit, the last three eigenvalues are several orders of magnitude smaller (i.e., close to zero) and the corresponding eigenvectors spanning the fit-space are highly correlated with the template profile, its phase derivative, and the baseline offset or their linear combinations. In other words, the profiles are originally described in an N_{bin} dimensional vector space and the fit projects the data onto $N_{\text{bin}} - 3$ dimensional vector space in which the fit-space has been removed. This removes three degrees of freedom from

the remaining eigenvectors and limits the efficacy of the correction scheme presented in this chapter because the fit-space component of the intrinsic shape fluctuations has been removed. Consequently, the random error due to any intrinsic shape fluctuations that correlate with these eigenvectors cannot be fully corrected. This is especially important in the case of profile variations correlated with the template derivative as such variations will introduce most random error in the ToAs. Only the total intensity fluctuations that are orthogonal to the fit-space contribute to the predictor computed in equation 3.4. The degree to which such fluctuations are correctable depends on how strongly they correlate with these three eigenvectors and the degree of correlation between the remaining projection coefficients and the arrival time residual.

We note that our work has implications for any relevant template matching algorithms. Such algorithms normally assume that the errors in the measurements of intensity are homoscedastic and uncorrelated and are based on the least squares minimisation. Our detection of profile variations implies that the errors are, in fact, heteroscedastic and correlated. A template matching algorithm based on generalised least squares, i.e., one that could include the covariance matrix, which carries the information about the correlation and heteroscedasticity of the noise, into the template matching algorithms will yield more realistic estimates of ToA uncertainty. It may also remove the necessity of correcting the residuals by the means described in this chapter because any random errors caused by the intensity fluctuations might be removed at the time of ToA determination. This is similar to the Cholesky decomposition which can remove bias in estimation of various parameters, such as parallax, when estimating from residuals with red noise present (Coles et al., 2011). Examples of template matching algorithms are the standard ToA derivation algorithm presented by Taylor (1992) and the matrix template matching algorithm that allows all Stokes parameters to be used in the ToA estimation (van Straten, 2006). An unbiased generalisation of template matching would be a logical extension of this work.

3.6.2 Application to PSR J0437–4715

Using one-minute integrations of PSR J0437–4715, we have detected shape fluctuations that we attribute to the stochastic single-pulse structure of the pulsar emission. However, for many timing applications, the exact cause of the pulse shape variations is irrelevant. This technique can be used to correct random errors in ToA estimates and improve sensitivity to any phenomena that do not induce shape variations. For instance, in order to place a limit on the existence of a gravitational wave background (e.g. Jenet et al., 2006) some methods compare the amount of power in the timing residuals with the power pre-

dicted to be induced by gravitational waves. Such waves will not affect the pulse profile shape. Therefore, reducing the rms timing residuals by accounting for pulse shape variations allows an improved limit on the existence of a gravitational wave signature to be obtained. We note that our correction method does not remove any of the signal induced by gravitational waves or any other phase shift of the pulsar profile, as verified by the simulations.

The long term timing of PSR J0437–4715 shows significant low frequency structure present in the residuals (e.g. Verbiest et al., 2008). Such red timing noise is a common feature in the timing residuals. It is important to determine the best predictor for correcting the residuals based on a short data span, as it will be less affected by any non-Gaussian noise. In the presence of non-Gaussian components, the best predictor will be affected as the correlation coefficients ξ between the residuals and projection coefficients will be biased toward zero. Presence of Gaussian noise (or any other symmetrically distributed noise) will increase the uncertainty in the estimate of ξ values but will not bias them.

The timing residuals for our Parkes observations are only partially corrected by our new method. We have demonstrated that partial correction does not necessarily imply the existence of other sources of scatter in ToA residuals that do not affect the pulse shape. At the same time it is not possible to completely exclude the existence of other effects such as hardware or software errors that are affecting the timing residuals at a lower level. Other possible effects that can increase the rms timing residual are described in detail by Cordes & Shannon (2010). Such issues are beyond the scope of this chapter, which concentrates solely on the pulse shape variability. We note that if any other non-white process affecting the residuals could be corrected, the residuals induced by profile shape variations could be removed more completely as the estimates of ξ are affected by the presence of phase noise.

We stress that the intrinsic shape variations lead to the heteroscedastic and correlated component of the SWIMS⁷. The uncorrelated component, originating from the standard radiometer noise in the weak source limit and the self-noise, is described by the diagonal of the covariance matrix \hat{C} while the temporally correlated part is described by the off diagonal elements. However, whatever phenomena contribute to the correlated component will naturally also contribute to the diagonal of the covariance matrix. As the self-noise and single-pulse modulation are all measured and described simultaneously by the covariance matrix, it is natural to consider them all as one phenomena, which we called SWIMS throughout this chapter. Two details of covariance matrix we would like to stress are

⁷We postpone the calculation of the expected amount of ToA scatter introduced by SWIMS (based on the measured covariance matrix) to following publications.

that: a) it does not contain any information about the spectral correlation of the noise and b) the measured covariance matrix also contains a contribution from the SEFD which can be subtracted if necessary. Extrinsic sources can also introduce a heteroscedastic and correlated component of noise, such as lightning strikes, other types of RFI or some instrumental effects, such as non-linearity of the backends. The effects of impulsive interference should be present across the whole pulse profile as they occur randomly in pulse phase. After careful removal of RFI, the measured eigenvectors are consistent with white noise outside the mean profile (see section 3.5).

As shown in Fig. 3.6, interpretation of the eigenvectors derived from the covariance matrix is complicated by the fact that the fit-space has been projected out of the N_{bin} -dimensional space of the shape variations. To investigate the structure of the shape variations prior to this projection, one can make the assumption that the pulsar timing model accurately predicts pulse phase (at least over the time-scale of the observations) and compute the covariance matrix of observed profile residuals after template matching by varying only the scale and offset (i.e., no phase shift). In this case, the covariance matrix contains the cyclic, phase-resolved auto-correlation (ACF) of the intensity fluctuations. The mean ACF computed by summing elements along the diagonals of this matrix is plotted in Fig. 3.9. The characteristic width of ACF, as determined by fitting a Laplace function, is equal to $67 \mu\text{s}$, which is consistent with the average width of microstructure events reported by Jenet et al. (1998). For comparison, the mean ACF formed from the covariance matrix used for random error removal (i.e., best-fit phase shift removed) is also plotted with a dashed line. A large fraction of the fluctuation power is removed by fitting for the phase shift during template matching. The phase shift fit removes variations that correlate with the template derivative, and the auto-correlation of the template derivative reaches its first minimum (below zero) at roughly the width of the pulse profile (around $140 \mu\text{s}$). This corresponds with the first dip seen in the ACF formed from data in which no phase shift has been removed; no dip is present in the ACF formed after fitting for phase shift. The measured auto-correlation function may be affected by other non-intrinsic effects, especially those related to propagation through ISM (Smits et al., 2003). We argue below that the ISM is not an important factor in considerations of PSR J0437–4715. The ACF shows a periodic ripple (most apparent at large phase lags), which is believed to be an instrumental artefact; it has a period of roughly 18 bins or $100 \mu\text{s}$. Its origin is currently unknown. Preliminary analysis of data from another instrument (Caltech Parkes Swinburne Recorder 2, CPSR2; Bailes, 2003; Hotan, 2006) confirms that this effect is intrinsic to PDFB3 as the ACF calculated from CPSR2 has no periodic ripple present.

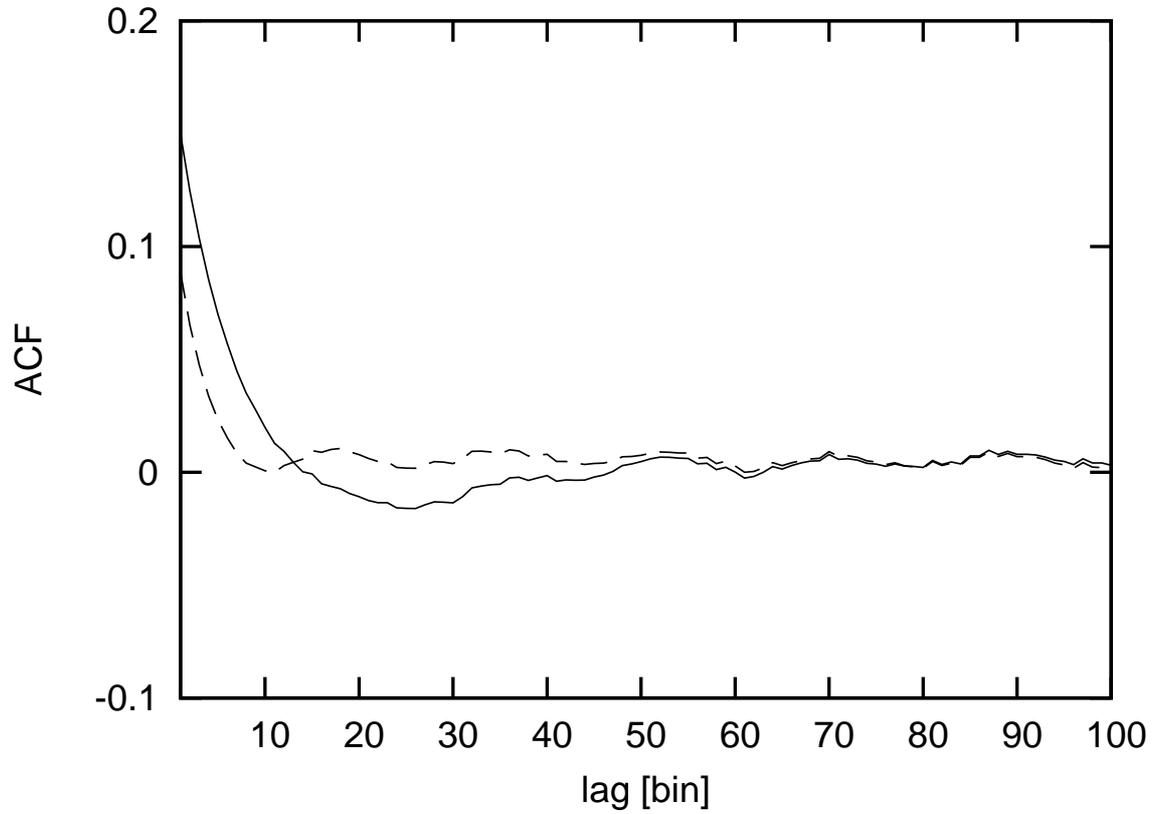


Figure 3.9: The ACF calculated with the assumption of the timing model accurately predicting the pulse phase over the course of observations (solid line) and without this assumption, i.e., when performing a full template matching (dashed line). The width of the ACF is determined by the characteristic width of single pulses (Rickett, 1975). The SEFD contributes an unresolved spike at zero lag, which is scaled to unit height for the solid line and the dashed line uses the same scaling factor.

The pulse profile variations are still detected thus confirming their origin as intrinsic to the pulsar.

The interstellar medium can also introduce shape variations, such as broadening of the pulse profile, which increases quadratically with DM and decreases quartically with frequency. Given the very low $DM = 2.64 \text{ cm}^{-3}\text{pc}$ and the observing wavelength $\sim 20 \text{ cm}$ used in this work, the expected variations in the pulse width for PSR J0437–4715 due to broadening are of the order of 1 ns (Bhat et al., 2004; Gwinn et al., 2006). Interstellar scintillation combined with frequency dependence of the pulse profile can also lead to fluctuations of ToAs. As shown in Fig. 3.5 the residuals are highly correlated when estimating the ToAs from two separate frequency channels. This high correlation implies that diffractive narrow-band interstellar scintillation is not responsible for the additional scatter. As we observe variations on a time-scale of minutes it is unlikely that broad-

band refractive scintillation is a contributing factor to the observed pulse shape changes as the time-scale for such variations is of the order of 1000 s (Gwinn et al., 2006). Based on the high degree of correlation between arrival times measured in separate bands (Fig. 3.5) and the observed broadband nature of single pulses (Fig. 3.4), we conclude that the intensity fluctuations are correlated over wide bandwidths. Consequently, increasing the bandwidth does not increase the timing precision. Only active mitigation or longer integrations can reduce the timing rms if the fluctuations of total intensity are indeed the main cause of the arrival time variations that greatly exceed the predicted uncertainty. Simultaneous multifrequency observations of PSR J0437–4715 would help to determine if the shape variations are persistent over very wide bands. Some observations to date have shown that the giant pulses from Crab extend over GHz bandwidths (Sallmen et al., 1999; Hankins et al., 2003; Hankins & Eilek, 2007). Another group has shown that the single pulses in PSR B0329+54 persist across 1.3 GHz (Karastergiou et al., 2001) but this may not be true for all single-pulse structure in the general population of pulsars.

It is also worth considering the impact of polarization variations on arrival time estimates. Emission from PSR J0437–4715 is highly polarised and a sudden change in the position angle of the linear polarisation occurs near the peak in the total intensity profile (as noted by Navarro et al., 1997). This implies that poor polarisation calibration will lead to significant profile changes, as quantified by van Straten (2006). For a single dish that is not equatorially mounted, these variations occur on time-scale of hours. Pulse profile shape changes detected by Vivekanand et al. (1998) were argued to be caused by calibration errors (Sandhu et al., 1997). Even though we detect fluctuations of the total intensity pulse profile on much shorter time-scales, we investigated if this was the case in our data as well. We studied the random error in timing residuals induced by the measured pulse shape variation as a function of parallactic angle. We found that there was no dependence between these two quantities. We also applied the PCA to uncalibrated data and a correlation between the induced random error I and the parallactic angle was readily apparent.

We conclude that the shape variations are more likely to originate at the pulsar rather than in the observing hardware or from interference. The polarisation calibration has been performed sufficiently to alleviate at least minute-time-scale fluctuations and does not introduce detectable pulse shape variations. The interstellar medium is also unlikely to cause such variations. Even if the detected variability is not intrinsic, the presented methodology remains valid. The intrinsic variation is expected from the stochastic single-pulse structure and will be detectable if the pulsar is bright enough.

Since the detected variations are likely to be intrinsic to the pulsar, a question arises whether the profile variations in PSR J0437–4715 are related only to SWIMS or if they are due to mode changing. We searched for clustering in the space spanned by the projection coefficients onto the ten significant eigenvectors by applying a friends-of-friends algorithm known from n-body simulations to identify dark matter haloes (Davis et al., 1985). We did not find any evidence of clustering in this space and hence conclude that the pulse profile variations are not due to mode changing.

We would like to stress the importance of our work for the next generation telescopes, which are likely to provide more sensitivity than currently available. With its huge collecting area of 1 km^2 , the Square Kilometre Array (SKA) is expected to revolutionise pulsar astronomy. One of the Key Science Projects of the SKA requires pulsar observations with the highest possible timing precision (Kramer et al., 2004). It is assumed that the SKA will observe of the order of 100 MSPs with an rms timing precision better than 100 ns. With the SKA’s phenomenal sensitivity, the S/N of a pulse profile should be >1000 on a time-scale of only minutes for many pulsars (compared with many hours with the Parkes telescope). The short observing times required to achieve such high S/N ratios would allow the SKA to observe multiple pulsars in a short time. However, the increased sensitivity of next-generation telescopes will also increase the relative importance of SWIMS as the radiometer noise is decreased. If the intrinsic pulse shape variations that we have detected for PSR J0437–4715 are typical of many MSPs at the observing frequency being used, they will induce a floor on timing precision that can be ameliorated only with longer integration times and active mitigation using methods such as the one presented in this chapter. Cordes & Downs (1985) demonstrated that, for the majority of their sample of 24 pulsars, timing precision is likely to be limited by phase jitter. No MSPs were included in this sample but later Cordes et al. (1990) detected SWIMS in a millisecond pulsar PSR B1937+214. Shannon & Cordes (2010) later argue that the timing noise of MSPs is similar to that of classical pulsars, only much smaller. In MSPs the phase jitter is also limited by averaging of a larger number of pulse periods in an observation of a given length as compared to an observation of a classical pulsar with the same duration of observation. Although SWIMS has not yet been detected in most millisecond pulsars, it is likely to be revealed with better instrumentation, more sensitive telescopes or longer data spans. Since the detected pulse shape variations are likely to be very broadband, increasing bandwidth will not reduce the random error introduced by SWIMS. This must be considered when predicting the potential science of current and future pulsar timing array projects and the observing time and strategy necessary to achieve the stated goals. For example, with an

SKA-like telescope, if many pulsars in a timing campaign are limited by SWIMS, then it is better to observe multiple pulsars simultaneously with fraction of the array for longer time rather than using full sensitivity to observe pulsars one by one for a short time. Some proposed astrophysical experiments demand extremely accurate ToAs over short intervals of the order of minutes, such as when a pulsar passes behind black hole in a close binary. SWIMS will make such experiments difficult or impossible.

The relative importance of the correlated component of SWIMS is also expected to vary between pulsars as it depends on intrinsic single-pulse emission properties and the shape of the average pulse profile. The fractional improvement in rms timing residual is expected to vary from case to case and it can be hoped that for pulsars with simpler and/or narrower profiles, variability in the single-pulse structure will be less severe (in agreement with equation 5 from Cordes & Shannon (2010)). As demonstrated by our first simulation, in simple cases our method works very well and can completely remove the random error in ToAs for some pulsars. Our method can be used to identify pulse profile modes which can lead to improved timing.

We note that, for current telescopes, equation 1.10 is a good approximation for the vast majority of MSPs, which have time averaged mean flux densities an order of magnitude smaller than PSR J0437–4715. For example, in the PPTA sample, fluxes vary between 1.3 and 13.8 mJy (see Table 2 of Yan et al., 2011) with a median value of 2.4 mJy, compared to the mean value of 150 mJy for PSR J0437–4715. Consequently, the profile variability arising from SWIMS has been neglected to date. We note that future telescopes like the SKA are likely to perform timing array experiments at higher frequencies to avoid some of the problems caused by the interstellar medium. Whether the self-noise will be a crucial limitation to precision timing for the majority of pulsars at all observing frequencies remains to be seen.

3.7 Conclusion

We have developed an extended principal component analysis method that is applicable to searching for pulse shape variations in pulse profiles. Applying this method to PSR J0437–4715 shows the presence of pulse profile variability that is likely to be intrinsic to the pulsar. The statistics of this variability are consistent over many months. The detection of significant intensity fluctuations implies that self-noise may be a limiting factor for timing precision of PSR J0437–4715 for current generation of telescopes. Future technological developments including construction of larger antennae and increased instrumental bandwidth will not improve timing precision as the single-pulse structure is a

source-intrinsic broadband phenomenon. However, the effects of SWIMS can be partially corrected by the method presented in this chapter and the proposed generalised template matching.

4

Improved methods of dealing with SWIMS

This chapter contains further developments in dealing with the stochastic wideband impulse modulated self-noise. The first part of this chapter presents an improved method of removing the random error in the times of arrival, based upon and extending the methodology presented in the previous chapter.

At the highest levels of pulsar timing precision achieved to date, experiments are limited by noise intrinsic to the pulsar when its peak flux approaches the system equivalent flux density. This stochastic wideband impulse modulated self-noise limits pulsar timing precision by introducing additional scatter in the measured times of arrival. We discuss an improved methodology of removing this random error in the measured times of arrival for pulsars by including information about polarised radiation. Observations over a one-week session of PSR J0437–4715 are used to demonstrate a nearly 40 per cent improvement in the root mean square timing residual with this extended analysis. In this way, based on the observations over a 64 MHz bandwidth centred at 1341 MHz with integrations over 16.78 s we achieve a 476 ns rms timing residual. The applicability of this methodology to achieve extrapolated rms timing residual of 30 ns with one hour integrations is prevented by the impact of Earth’s ionosphere. This improvement provides an opportunity to increase the sensitivity in various pulsar timing experiments, for example pulsar timing arrays that pursue the detection of the stochastic background of gravitational waves. The fractional improvement is highly dependent on the properties of the pulse profile and the stochastic wideband impulse modulated self-noise of the pulsar in question.

4.1 Introduction

Radio pulsars are very stable rotators whose radio emission is received on Earth as series of radio pulses. The measurement of the average pulse train time of arrival (ToA) is

referred to as pulsar timing and leads to the majority of derived pulsar astrophysics. At the heart of this methodology lies the comparison of the estimated ToA to a model of the pulsar's properties, including astrometric, spin-down, and binary (both Keplerian and relativistic) parameters. MSPs, with their short periods, narrow features in their phase-resolved light curves (pulse profiles) and very low spin frequency derivatives, are exceptionally stable (Matsakis et al., 1997) and provide the best opportunity for high precision timing experiments. A number of authors have achieved timing precision of a few hundred nanoseconds over time-scales of five to ten years (e.g., Verbiest et al., 2009, and references therein). With such high precision, a number of experiments are possible, such as testing theories of relativistic gravity (e.g. Hulse & Taylor, 1975; Weisberg & Taylor, 2005; Kramer et al., 2006b); detection of planets orbiting a pulsar (Wolszczan & Frail, 1992; Bailes et al., 2011) and the most precise published mass measurements of planets in the Solar System (Champion et al., 2010); determining properties of the interstellar medium (ISM, e.g. You et al., 2007; Demorest, 2011); constraining the equation of state of dense matter (e.g. Demorest et al., 2010); or detecting irregularities in terrestrial time-scales and developing a pulsar based one (Petit & Tavella, 1996; Rodin, 2008; Hobbs et al., 2012).

One particular application of the pulsar timing methodology is the detection of gravitational waves, i.e., perturbations of the space-time metric predicted by general relativity. These waves carry energy and so their emission can be inferred by observing the decrease of the semi-major axis of an emitting binary system of the order of 1 cm a day, as detected with pulsar timing of binaries (e.g., Taylor & Weisberg, 1982; Stairs et al., 1998; Kramer et al., 2006b; Jacoby et al., 2006; Bhat et al., 2008; Freire et al., 2012). A method of direct detection of gravitational waves relies on their impact on the relative distance between Earth and the pulsar with a quadrupolar spatial signature as the gravitational wave passes through space-time (e.g., Jenet et al., 2006).

A number of collaborations around the world are observing ~ 30 MSPs in the pursuit of detecting gravitational waves with so called pulsar timing array experiments. These experiments, based on early ideas of Sazhin (1978) and Foster & Backer (1990), exploit the correlation between the timing residuals for pairs of pulsars in different directions (Hellings & Downs, 1983) in the sky to measure the spatial signature of the gravitational wave background and thus distinguish it from the timing noise, interstellar medium (ISM) effects, clock or Solar System ephemeris errors, which have different multipole signatures or are uncorrelated between different objects. We note that recently van Straten (2013) pointed out that instrumental distortions due to imperfect polarimetric calibration can

corrupt the quadrupolar signature of the gravitational waves.

A number of authors have demonstrated that in order to detect gravitational waves, a high precision of pulsar timing observations needs to be achieved. For example, Jenet et al. (2006) showed that 20 pulsars need to be timed at a precision of 100 to 500 nanoseconds over a period of 10 years. While the desired precision has been attained for a small number of objects, the timing precision of the MSPs is worse than the theoretically predicted values for objects expected to be the best timers. A prime example of this is the closest and brightest MSP, PSR J0437–4715 (Johnston et al., 1993). Despite numerous attempts, the timing precision of this pulsar has been significantly worse than expected from the pulse profile and signal-to-noise ratio. Sandhu et al. (1997); van Straten et al. (2001); Verbiest et al. (2008) all achieved progressively better timing precision with the latest results being 199 ns rms timing residual over 10 years with integration lengths between 10 and 120 minutes, still a factor of 4 or 5 worse than theoretical predictions. An important factor in the improvement of the timing precision was related to improved methods of dealing with polarimetric calibration (Britton, 2000; van Straten, 2004, 2006).

The long standing problem of PSR J0437–4715 under-performing in terms of timing precision has been understood by Osłowski et al. (2011b, chapter 3 in this thesis) which demonstrated that the theoretical predictions neglect the contribution of the stochastic wideband impulse modulated self-noise (SWIMS), also referred to as pulse (or phase) jitter (e.g., Cordes & Shannon, 2010; Liu et al., 2012). For pulsars whose peak flux approaches the system equivalent flux density this contribution cannot be neglected and it randomly biases the estimates of the times of arrival (ToA) of the pulse train. Chapter 3 also presented a method based on decomposition of the residual profiles onto eigenvectors of the noise covariance matrix to help remove the effects of SWIMS in post processing of timing residuals. We applied this methodology to PSR J0437–4715 and achieved a 20 per cent reduction in the rms timing residual. This work is concerned with improvements to this methodology, allowing even better recovery of unbiased ToAs by analysing the polarised flux of the pulsar.

This chapter is structured as follows: section 4.2 reviews the concept of SWIMS and its impact on attainable timing precision. Section 4.3 describes the data used and details of processing techniques applied. In section 4.4 we present an extension to the methodology allowing for removal of the random error introduced by SWIMS. The results of applying this method to archival observations of PSR J0437–4715 are presented in section 4.5. We discuss these results in detail in section 4.6 before drawing our conclusions in section 4.7.

The work presented in this chapter is a logical extension of the work presented in

chapter 3 and it should be consulted in order to fully understand the current work. In the following section we briefly summarise the key concepts of chapter 3 for the reader's convenience.

4.2 Pulsar timing and the stochastic wideband impulse modulated self-noise

The pulsar timing methodology is based on cross-correlation of the observed pulse profile with a template, typically done in the Fourier domain (Taylor, 1992). The derived ToAs and the pulsar model are used to form timing residuals as implemented by the TEMPO2 software package (Hobbs et al., 2006). As outlined in section 4.1, in the case of PSR J0437–4715 the rms timing residual is much larger than what is theoretically expected based on the considerations of the pulse profile effective width and the signal-to-noise ratio (S/N) of the data. In chapter 3 we provided evidence that this is due to the presence of SWIMS, i.e., the pulsar itself becomes a source of significant noise that is normally not accounted for in the noise balance and predictions of timing precision. SWIMS will affect not only the total intensity measurement but also the polarised flux. Fig. 6 of McKinnon & Stinebring (2000) demonstrates presence of SWIMS in Stokes U parameter for PSR B2020+28.

The amplitude of SWIMS is phase dependent and thus the variance of noise is also phase dependent, that is the noise is heteroscedastic or non-stationary in pulse phase. In addition, due to the emission properties of the pulsar, SWIMS exhibits temporal and spectral correlations. Both types of correlation have adverse consequences. The temporal correlation implies that the noise in the pulse profile is not only heteroscedastic but also not independent between different phase bins. While this correlation can amplify the random error introduced by SWIMS into the ToA estimation, it allows for removal of this random error in post-processing. The spectral correlations neutralise the benefit of increasing the instrumentation bandwidth by inducing correlations the ToAs estimated from separate parts of the observing radio band. Using wideband receivers is still beneficial for other reasons as it allows to address many issues related to the interstellar medium such as dispersion measure variations or scattering (e.g., You et al., 2007).

The temporal correlations and heteroscedasticity of SWIMS can be fully characterised by the covariance matrix of the pulse profiles after subtracting the template profile that was used during the ToA estimation. This allows us to study the noise properties at any integration time longer than the pulse period as the covariance matrix scales inversely

with the integration length¹ and the relative contribution of SWIMS to the covariance matrix is integration length independent.

For a more detailed discussion of SWIMS we refer the reader to section 3.2.

4.3 Observations and data processing

The dual-polarisation data were recorded at the Parkes 64 m radio telescope using the central beam of the 20 cm multibeam receiver (Staveley-Smith et al., 1996) and the second generation of Caltech Swinburne Parkes baseband recorder (Bailes, 2003; Hotan, 2006, CPSR2). A total of ~ 36.5 hours were recorded between the 19th and 27th of July 2003 using two bit digitisation with a bandwidth of 64 MHz centred at 1341 MHz and split over 128 channels, each 0.5 MHz wide, resolved into 1024 phase bins, and integrated over 16.78 seconds. CPSR2 adjusts the sampling thresholds twice every minute to maintain linearity of the digitiser response. The incident radiation was dedispersed coherently to preserve the narrow features of the pulse profile. The data were processed as in van Straten (2013) and we refer the reader to this work for details. The Stokes parameters in the recorded data follow the definitions from van Straten et al. (2010). Here, we present a brief summary of the applied processing.

As explained in chapter 3 and section 4.6.1 below, it is beneficial to have a large number of sub-integrations available for the analysis. Therefore, we demonstrate our methodology using a different dataset than in chapter 3. We describe here an archival dataset ample enough for our purposes. This dataset can be obtained from the CSIRO's Data Access Portal² (Hobbs et al., 2011b).

We generated a high S/N, calibrated full Stokes template profile using an independent dataset, as follows: 50 hours of observations of PSR J0437–4715 were calibrated using the measurement equation modelling technique (van Straten, 2004). This technique makes use of observations of a pulsed noise diode coupled to the receptors that injects a polarised reference signal into the feed horn. Observations of the radio source 3C 218 assumed to have a constant flux and negligible circular polarisation, were used to constrain the boosts along the Stokes V axis. Then, the five best quality of the 8 hour observations were integrated together to form an initial pulse profile template. The remaining 45 hours were matched to it using the matrix template matching technique (van Straten, 2006). A final pulse profile template was created from the best 41.7 hours of data processed in the same

¹This statement assumes that we are already looking at integration times longer than correlation time of any possible correlations between different pulses.

²<http://datanet.csiro.au/dap/>

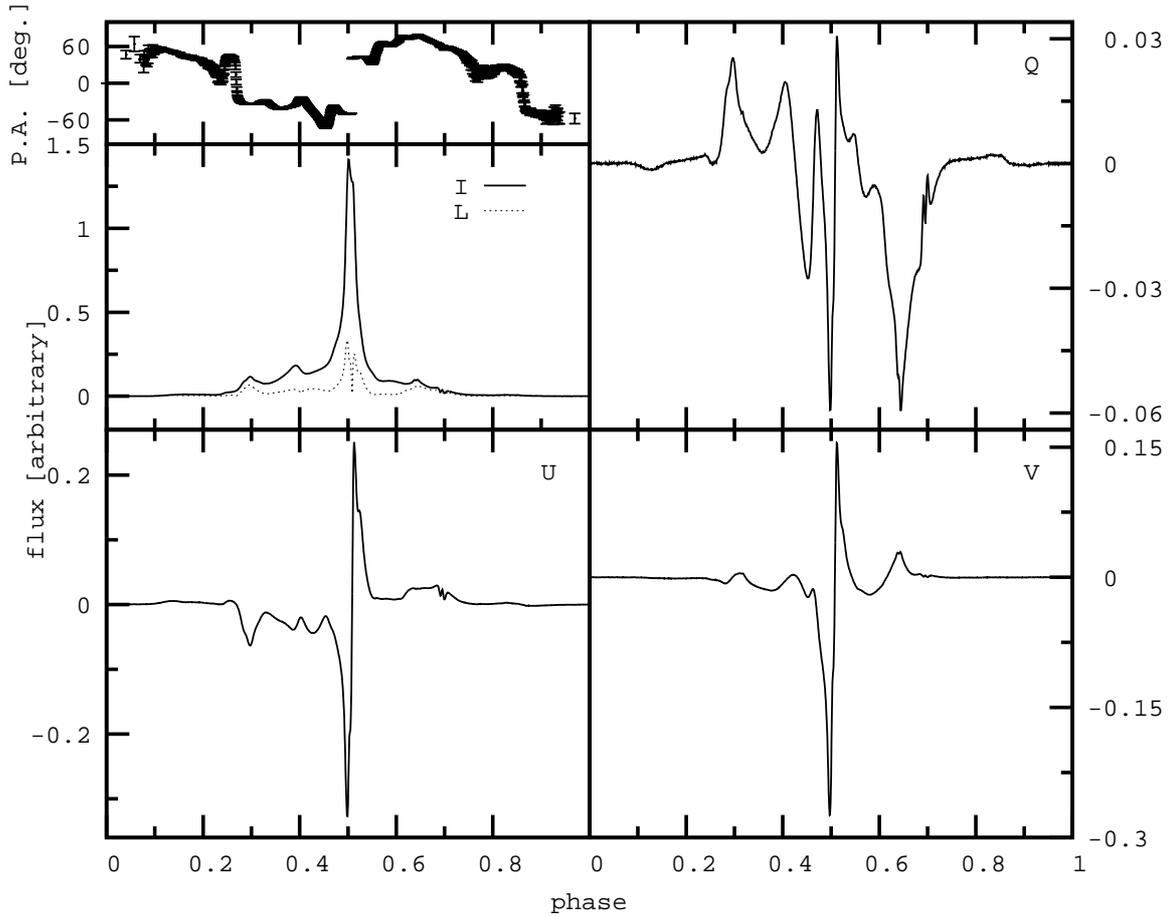


Figure 4.1: A high S/N template of PSR J0437–4715 in the 20cm observing band formed from 41.7 hours of data. The top left short panel shows the position angle swing and is situated directly above a panel showing the total intensity pulse profile in solid line and the linear polarisation in a dotted line. The top right panel shows the Stokes Q parameter of the template, while the bottom left and right panels correspond to Stokes U and V, respectively. The panels on the left use the ordinate axes on the left edge while the right hand side panels use the far right ordinate axes. All Stokes parameters are on the same flux scale.

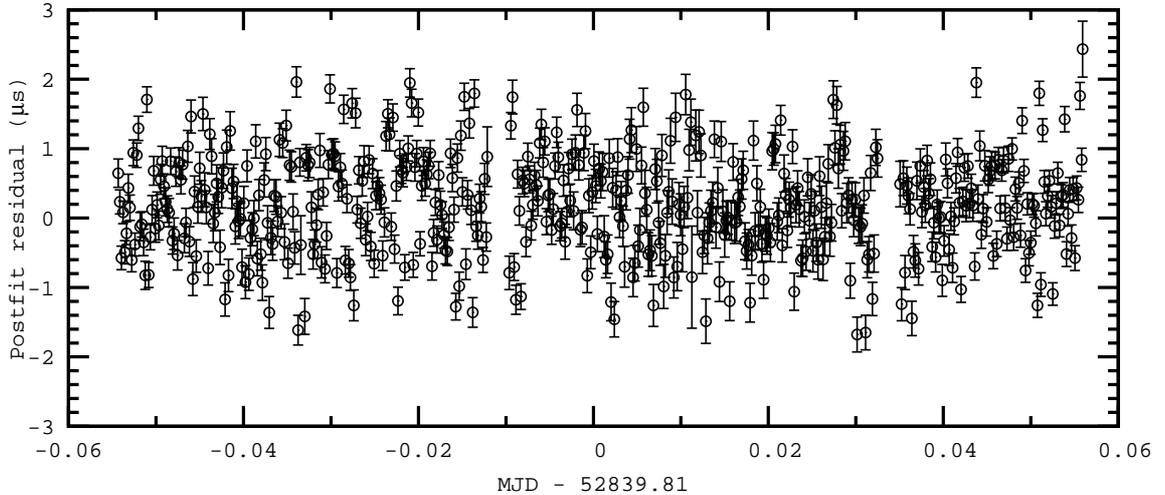


Figure 4.2: Timing residuals estimated for 16.78 s sub-integrations of data taken during one of the days timed against the standard from Fig. 4.1. The mean ToA estimation error in the whole dataset is 270 ns, whereas the weighted rms of the residual ToAs is 774 ns. The fit has $\chi^2/\text{d.o.f.}$ of 11.8. For clarity, we have plotted the residuals as a function of MJD.

manner and integrated hierarchically³ to minimise quantisation errors into a profile with S/N equal to 18.6×10^3 . The full polarisation information was retained and is shown in Fig. 4.1. Throughout this publication we use the same template in all contexts.

The generated template is used for polarimetric calibration of the primary dataset spanning a week of observations. The calibration was performed with the measurement equation template matching technique (van Straten, 2013). This method yields superior results by using observations of PSR J0437–4715 at different epochs as a polarised reference source. It combines the measurement equation modelling with the matrix template matching to break certain degeneracies encountered during the polarimetric calibration.

In both datasets used for creating the template profile and estimating the covariance matrix, 15 per cent of the frequency channels on each side of the band were rejected to avoid quantisation noise and aliasing⁴. CPSR2 is a 2-bit recorder which therefore introduces significant distortions into the measured pulse profile (Jenet & Anderson, 1998). A 2-bit correction scheme has been applied and the data were searched for cases of imperfect correction. We achieve this by measuring the $\chi^2/\text{d.o.f.}$ of the fit of the pulse profile to a high signal to noise ratio (S/N) template, where d.o.f. refers to the number of degrees

³That is, we first form integrations of a small number of sub-integrations at a time for the whole dataset and once all the longer integrations are formed we repeat this process with them and continue doing so until all the data are added together.

⁴In this context by aliasing we mean reflection of the harmonics beyond the Nyquist frequency corresponding to the phase resolution of the observations back into the pulse profile.

of freedom of the fit. The profiles with worst and uncorrectable 2-bit distortions have distinctly higher values of χ^2 which allows for their easy identification and rejection. All of the processing steps were performed using the PSRCHIVE⁵ data processing and analysis suite (Hotan et al., 2004b; van Straten et al., 2012).

The timing residuals for a subset of the 7845 profiles based on the 16.78 sub-integrations used in the analysis are presented in Fig. 4.2 as a function of MJD. The ToAs were estimated using the Fourier domain with Markov chain Monte Carlo algorithm in PSRCHIVE. The weighted rms timing residual is 774 ns and the $\chi^2/\text{d.o.f.}$ equals 11.8. The unweighted rms timing residual is slightly higher and equals 806 ns. The mean ToA estimation error is 270 ns (calculated as a square root of the unweighted mean of the squared ToA uncertainties), and the mean S/N is 302. Based on the findings of chapter 3, we attribute the higher-than-expected rms timing residual to SWIMS and the induced ToA estimation random error and $\chi^2/\text{d.o.f.}$ on the invalid assumptions used in the ToA uncertainty estimation. The rms timing residual scales as the square root of integration time for pulsars without temporal correlations between single pulses and without long-term timing noise, or when studied over short periods of time. Based on this extrapolation, which we have shown to be true for PSR J0437–4715 (see Fig. 3.3) the data would yield an rms timing residual of 53 ns when derived from 60 minute integrations.

4.4 Method

The method described here is a direct extension of what was presented in section 3.4 and continues the ideas from Demorest (2007). To summarise, this method consisted of fitting for a scale, phase shift and offset between the observed total-intensity pulse profiles and a template profile and calculating the covariance matrix of the residual profiles. Next, the eigenvectors of the covariance matrix are calculated and used as a new basis in which the residual profiles are presented. This allows a natural identification of the observations with most SWIMS by analysing the profiles in this new basis, as the highest variance in the dataset is associated with the eigenvectors corresponding to the highest eigenvalue. As a final step, a correlation between the new presentation of the residual profiles and timing residuals is found using multiple regression and enables a post-processing correction of the random error in ToA.

The method presented in 3.4 yielded a 20 per cent reduction of the rms timing residual by exploiting this correlation. However, it is unable to correct the random error introduced by SWIMS correlated with any of the three degrees of freedom that are removed from

⁵<http://psrchive.sourceforge.net/>

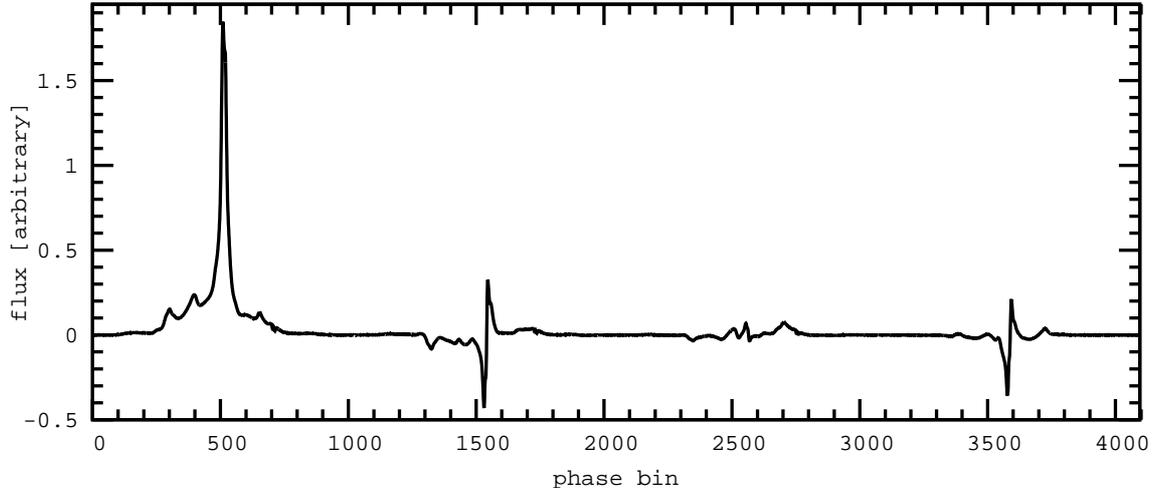


Figure 4.3: An example of a profile used as input for the methodology presented in this work. The profile is formed from all four Stokes parameters spread over $4N_{\text{bin}}$ phase bins instead of four separate profiles with N_{bin} phase bins each. We used the same profile as in Fig. 4.1. The input profile is resolved into 1024 pulse phase bins and after combing all the Stokes parameters into one profile we obtain a pulse profile with 4096 phase bins.

the data during the template matching fit. Unfortunately, any components of SWIMS correlated with one of these degrees of freedom, namely the template time derivative, also have a high impact on the random error of the ToA. In this work, we extended the previously published method by incorporating the polarised flux of the pulsar. As we demonstrate, this allows to recover some of the signal lost in the previous attempt at the mitigation of additional ToA scatter introduced by SWIMS.

The simplest way to implement this extended methodology is by considering each pulse profile resolved into N_{bin} phase bins with all four Stokes parameters as a $4N_{\text{bin}}$ profile with only one polarisation. The additional, artificially introduced, phase bins are populated with the polarised flux as measured in the Stokes parameters. That is, the first N_{bin} phase bins represent the information contained in the total intensity measurement; the bins from $N_{\text{bin}} + 1$ to $2N_{\text{bin}}$ represent the Stokes Q measurements; $2N_{\text{bin}} + 1$ to $3N_{\text{bin}}$ represent the Stokes U flux; and the last N_{bin} phase bins are populated with the circular polarisation, i.e., the Stokes V parameter. An example of such a profile is shown in Fig. 4.3, where the mean off-pulse values were chosen to equal zero for each Stokes parameter independently.

Before forming the $4N_{\text{bin}}$ profiles, we performed template fitting, which is necessary to estimate the profile covariance matrix \hat{C} . As we considered all four Stokes parameters, we have six free parameters of the fit: a scale factor to account for flux variation of the pulsar

(intrinsic, related to the effects of interstellar medium, or instrumental effects); a phase shift between the pulse profile and the template; and four offsets between the pulse profile and the template for each of the Stokes parameters. The phase shift and scale factor are determined based on the total intensity but applied to all four Stokes parameters. The baseline offset is calculated and applied separately for each of the Stokes parameters. The fitting procedure results in six out of $4N_{\text{bin}}$ eigenvectors sharing the eigenvalue equal zero and form what we refer to as the fit-space.

After the template fitting we form the $4N_{\text{bin}}$ profiles and directly follow the method presented in chapter 3 in equations 3.1 through to 3.9. When exploiting the additional information available in all the Stokes parameters, the sizes of some of the mathematical quantities involved are different than previously, namely: the covariance matrices \hat{C} and \hat{D} are both $4N_{\text{bin}}$ by $4N_{\text{bin}}$ square matrices; there are $4N_{\text{bin}}$ eigenvalues associated with as many eigenvectors, each of length $4N_{\text{bin}}$; the matrix of projections of residual profiles onto eigenvectors has now $4N_{\text{bin}}$ columns; the vector of covariances between the residuals and the projection coefficients γ has $4N_{\text{bin}}$ elements; and there are $4N_{\text{bin}}$ values of ξ , the Pearson's product moment correlation coefficients between the timing residuals and the projection coefficients onto one of the eigenvectors. After the analysis is finished the $4N_{\text{bin}}$ profiles and eigenvectors are converted back to N_{bin} profiles with four Stokes parameters.

Similar to the methodology described in chapter 3, we have performed a number of simulations to test if the extended methodology does not remove arbitrary phase shifts not related to pulse profile deviation from the template profile. We refer the reader to our previous work for a description of how such simulations are realised. Here we report that no removal of arbitrary shift occurs and thus the method is valid and useful for astrophysical experiments.

An implementation of this method is publicly available as a part of the PSRCHIVE suite as an upgraded version of the previously available application “psrpca”. It requires the GNU Scientific Library⁶ to work and can be significantly sped up by using the CULA library⁷ (Humphrey et al., 2010) if an NVIDIA[®] CUDA[™] enabled Graphics Processing Unit is available.

4.5 Results

Applying our method to the observed dataset of 7845 profiles leads to the following:

⁶<http://www.gnu.org/software/gsl/>

⁷<http://www.culatools.com/>

- the detection of significant pulse shape variations with at least 21 significant eigenvectors,
- a reduction in rms timing residual from 776 ns to 473 ns, a 39 per cent reduction; and a reduction in $\chi^2/\text{d.o.f.}$ from 11.8 to 4.4.

The $\chi^2/\text{d.o.f.}$ remains large as the errors of ToA residuals are estimated with the PSRCHIVE software suite under the assumption of SWIMS not being present in the data. Other factors may be contributing to the post-correction rms timing residual, see section 3.6 for a discussion of possibilities.

The first three, most significant eigenvectors are plotted in in Fig. 4.4, 4.5, and 4.6. Each panel represents a different Stokes parameter, from total intensity on top, via Q and U to V at the bottom. We note that the Stokes Q and U parameters in the first eigenvector are highly correlated with the Stokes U and Q, respectively, of the template profile shown in Fig. 4.1. The Stokes U parameter in the first eigenvector has an additional significant peak near the leading edge of the total intensity profile as compared to the Stokes U of the template profile. The majority of the detected pulse profile variance is located in the central parts of the profile.

We plot the first 150 values of ξ in Fig. 4.7. In chapter 3, we chose the number of significant eigenvectors by calculating a resistant and robust estimate of the standard deviation of ξ and searched for three subsequent values of ξ above three times the measured standard deviation. Applying the same criteria for automatic selection of the number of significant eigenvector results in 301 significant eigenvectors. This value prompted a closer investigation of the correlation coefficients ξ and revealed that the automatic selection process does not work very well in this case. With a larger number of eigenvectors available, more stringent criteria are necessary or the number needs to be chosen by investigating the ξ values manually. In this case, we arrive at a reasonable number of significant eigenvectors by increasing the required number of consecutive points above a chosen threshold to six.

4.6 Discussion

We first discuss the details of the extended method and stress the differences between this new method and the one presented in chapter 3. We then discuss the impact of the extended method on the timing precision of PSR J0437–4715 and on the prospects of precision timing.

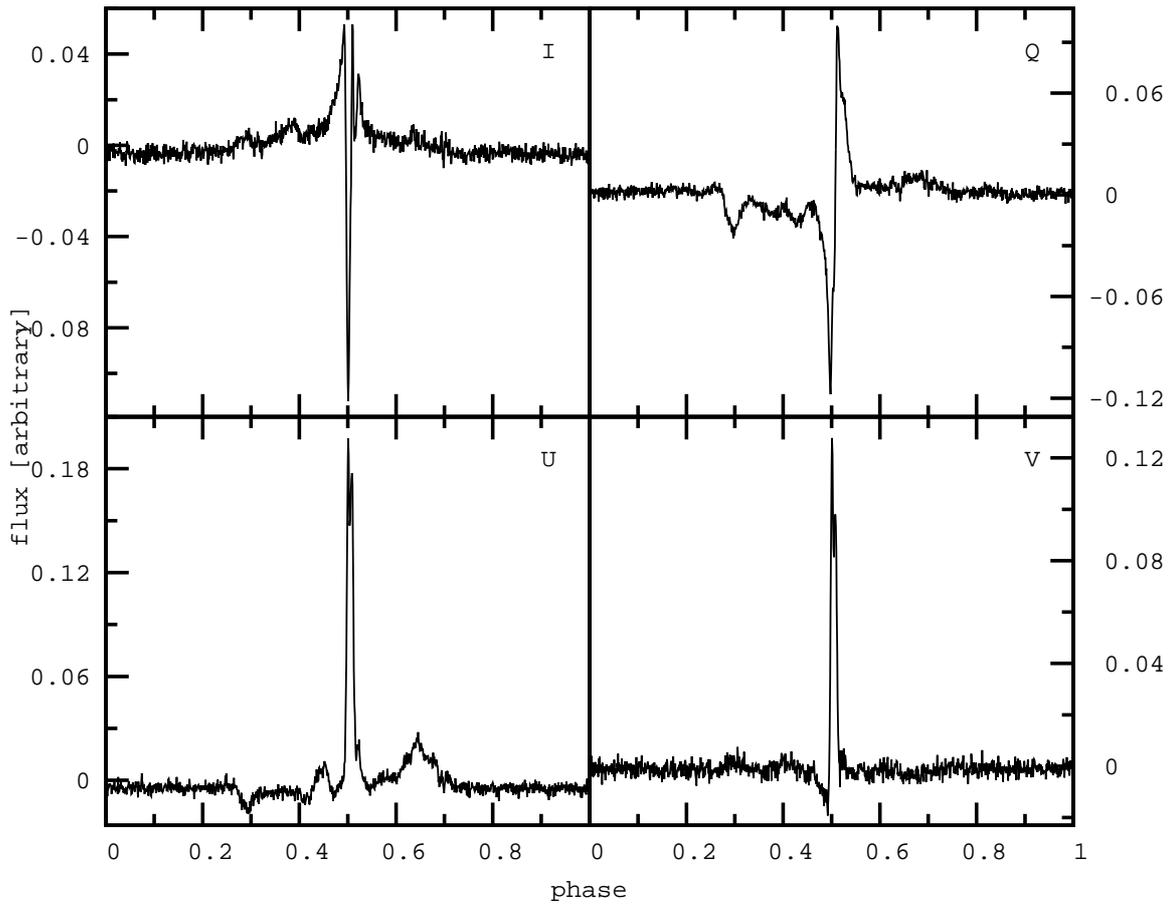


Figure 4.4: First, most significant eigenvector, presented in the same way as the template profile in Fig. 4.1, except we do not show the polarisation angle for the eigenvector. Note that the flux in total intensity is smaller than flux in the polarised signal.

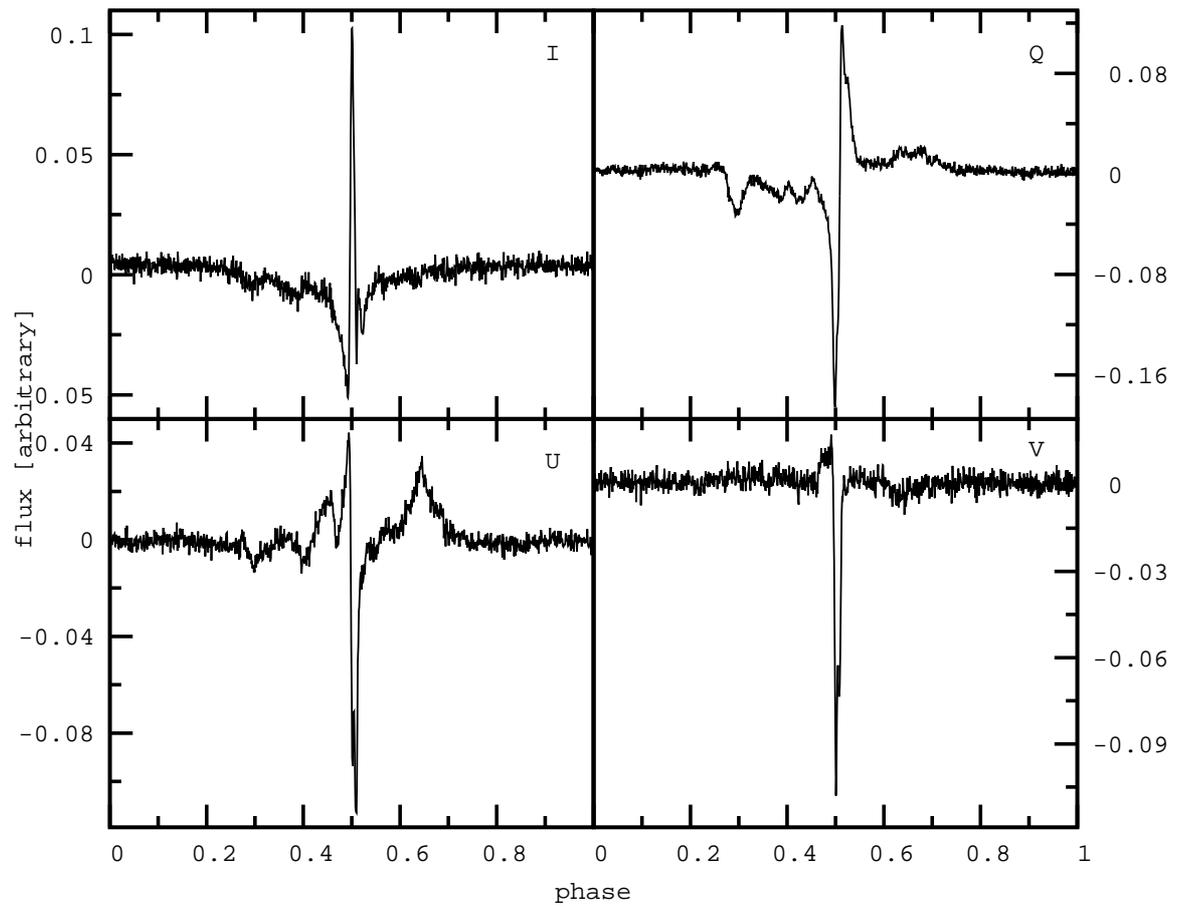


Figure 4.5: The second eigenvector.

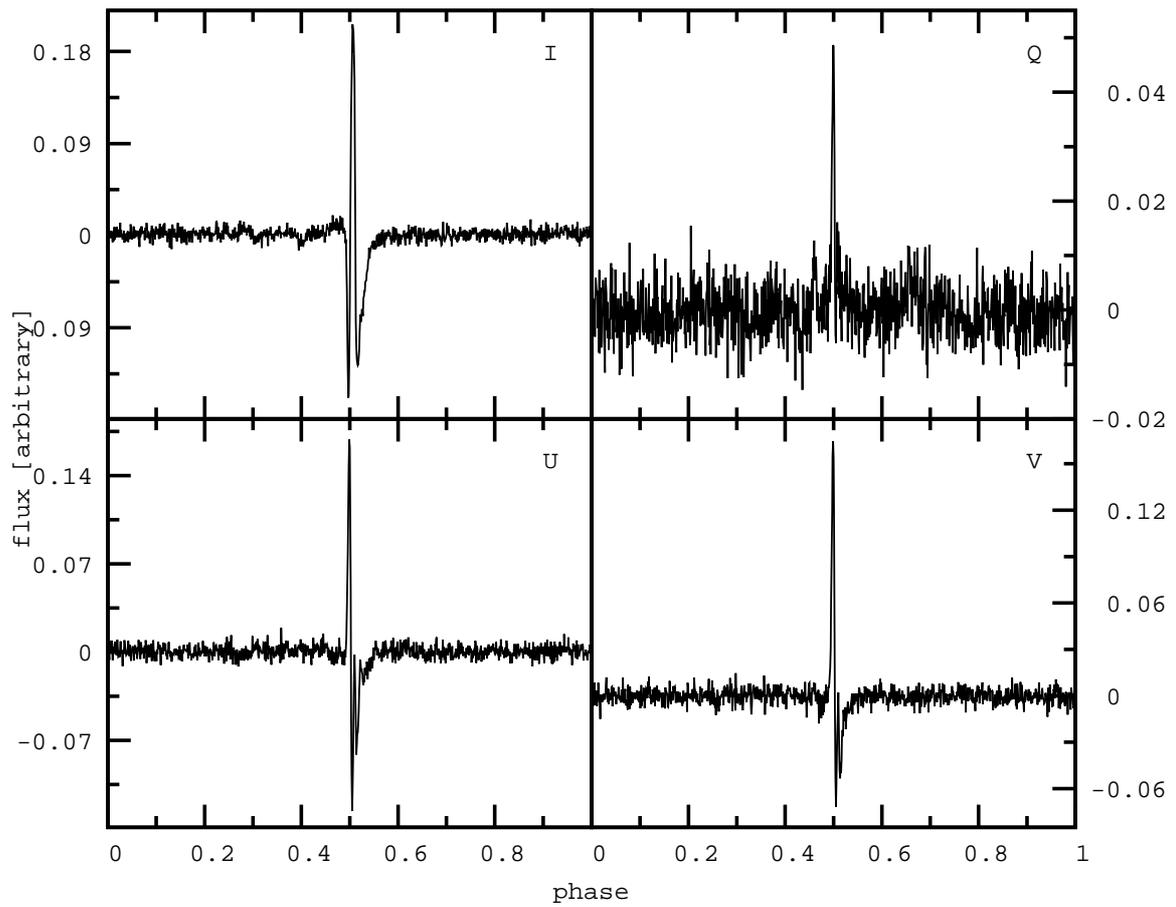


Figure 4.6: The third eigenvector.

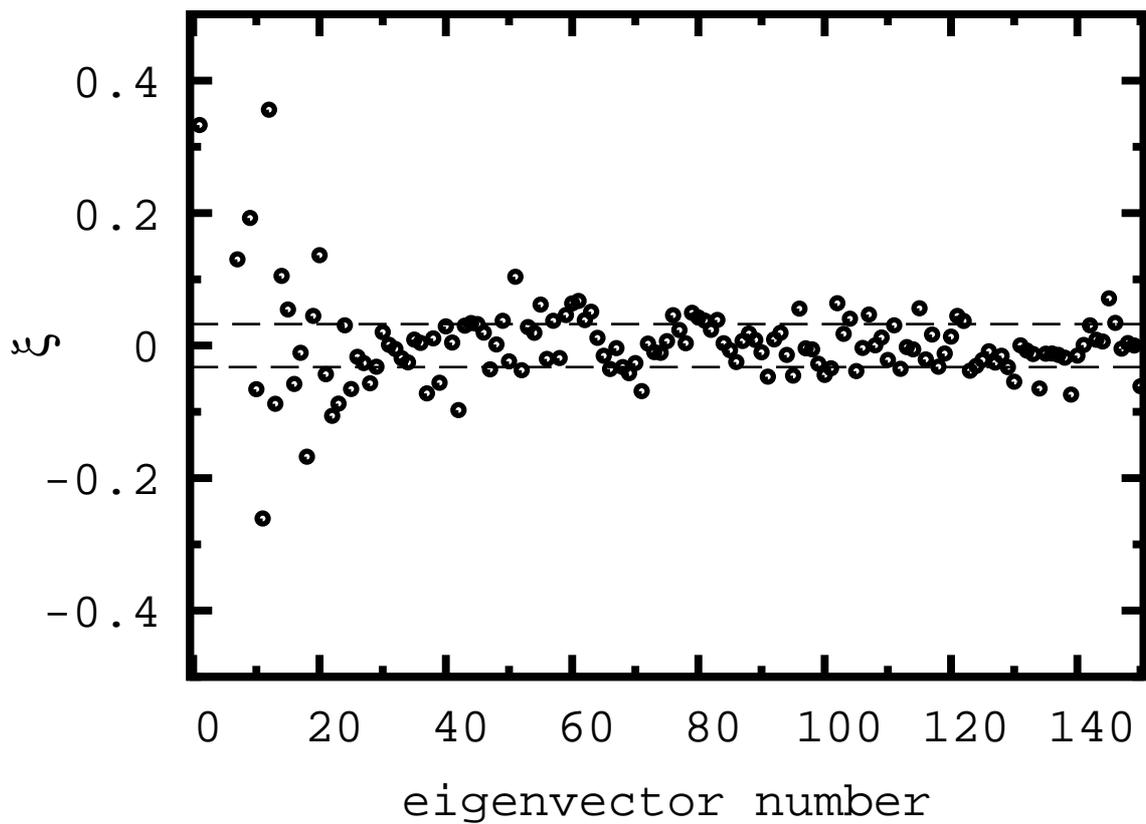


Figure 4.7: Distribution of correlation coefficients between the residuals and projection on to eigenvectors for the actual observations, shown by open circles. The threshold level for choosing the number of eigenvectors of three standard deviations of ξ is marked with dashed lines. The correlation coefficients are highly scattered, thus requiring more stringent criteria for selection of the significant eigenvectors.

4.6.1 Discussion of the method

The previously published method used pulse profiles resolved into N_{bin} phase bins but only exploited one Stokes parameter. In the extension presented in this publication we employ all four Stokes parameters with the same phase resolution as before. This implies that in order for the covariance matrix to be fully determined we now need a minimum of $4N_{\text{bin}}$ observations. Our conclusions from chapter 3 still hold and we briefly reiterate them here: a) when the number of available observations is lower than $4N_{\text{bin}}$, the number of eigenvectors with non-null eigenvalues will be limited to number of observations and care needs to be taken when choosing the number of significant eigenvectors; b) the number of required observations can be reduced by either gating the pulse profile or by reducing phase resolution; c) reducing the phase resolution of the pulse profile may wash away the correlation of the noise between the phase bins which will render the presented methodology less efficient; d) in addition, reducing the phase resolution can lead to aliasing of the observations as, by definition, we are dealing with high S/N observations; e) high number of observations is highly desirable in order to increase the S/N of the covariance matrix estimate and to ensure complete measurement of the SWIMS representations in the data; f) given a constant total integration time of all available observations, a larger number of separate sub-integrations will yield a higher S/N of the covariance matrix. This follows from variance of sample variance being inversely proportional to the number of samples (Kenney & Keeping, 1964). By considering the case of fixed total integration time and related change of S/N when changing the number of samples, remembering that S/N is related to the ratio of sample variance to the square root of variance of sample variance, and assuming we are considering timescales beyond any correlations between pulses, we conclude that S/N of covariance matrix is inversely proportional to the square root of number of sub-integrations.

Similar to the chapter 3, the eigenvectors are sorted by the corresponding eigenvalues in the decreasing order, that is, most variance in the data is associated with the first eigenvector and the eigenvectors corresponding to the fit-space are the last ones as all the variance in the data in these dimensions is removed during the template fitting procedure. Despite this monotonicity of the eigenvalues, the regression coefficients used in the ToA estimate correction scheme are not monotonic as the highest variance in profile does not directly translate into the highest random biasing of the ToA estimate. For this reason, applying multiple regression is crucial for a successful correction scheme and projections onto more than one eigenvector need to be used.

We note that our method calculates the eigenvectors of the covariance matrix assuming

Euclidean geometry, i.e., the measured eigenvectors are orthogonal only in a classical sense. The Stokes parameters span the Minkowski space and to allow any physical interpretations of the measured eigenvectors, their orthogonality must be ensured in the Minkowski space instead (see, e.g., Renardy, 1996). Four dimensional Euclidean orthogonal transformations may not preserve the time-like interval of the Stokes parameters. A prime example of this effect is visible in Fig. 4.4, on which the polarised flux of the first eigenvector exceeds its flux in total intensity. This corresponds to the variance of polarised flux being larger than that of the total intensity.

As was the case in the method based on only total intensity, the obtained eigenvectors are affected by the template matching process, this time with six degrees of freedom. This means that if one was to perform a simulation with a known pulse profile distortion inserted into data, the measured eigenvector would not correspond directly to the inserted noise (see Fig. 3.6 and the relevant text).

The method presented here is superior to the technique presented in chapter 3 as the latter was limited primarily by any variance in the profile correlated with the time derivative of the profile. In chapter 3, the template matching procedure ensured that no variance is left in the dimensions corresponding to the template profile, its time derivative and the baseline offset. Variations that correlate with the time derivative of the template profile can easily introduce very significant random error into the ToA estimate without adding much power to the profile, as can be demonstrated by simulations. By including the information about the statistics of polarised noise, as presented in this work, some part of the information previously completely removed during the template fitting is now retained in the Stokes Q, U and V parameters.

The covariance matrix constructed in the way described in section 4.4 is related to the covariance matrix of the Stokes parameters at a given phase bin as described by van Straten (2009) in equation 28. In van Straten (2010) the author generalised this matrix to include cross covariances between different polarised states, such as the Stokes parameters observed in different phase bins. Contrary to both these publications, the covariance matrix we use is not based on any assumptions about the distribution of the Stokes parameters, but merely on the conjecture that the template represents the mean Stokes parameters accurately. We note that the covariance matrix \hat{C} used here contains more information than either of the two above formulations as it contains information both about covariances between the Stokes parameters at any phase bins as well as between different phase bins.

4.6.2 Application to PSR J0437–4715

For comparison, the method presented in chapter 3 yields timing residual with an rms of 642 ns, a 17 per cent reduction; and $\chi^2/\text{d.o.f.}$ of 9.23. The improvement based on this method is somewhat worse than for the dataset used in chapter 3, most likely due to the lower S/N of the observations used in this publication. By exploiting the information available in the polarisation of the incident radiation we are able to improve the results of the random error removal from ToA estimates by more than a factor of two compared with the previously published approach. We note that despite the significant improvement in the attained rms timing residual, the results may vary greatly for different pulsars depending on their pulse profile and the nature of SWIMS.

As we are employing the information contained in the polarised incident radiation, it is useful to compare the rms timing residual obtained using the methodology presented in this chapter with other methods that use not only the total intensity but the other three Stokes parameters as well. Two such methods are the matrix template matching (van Straten, 2006) and timing of the invariant interval pulse profile (Britton, 2000). The first allows us to achieve higher timing precision by exploiting the fact that the polarised signal contains more power in the higher harmonics of the Fourier transform of the pulse profile and can compensate for some of the systematic errors introduced during the calibration. The latter minimises the errors arising from polarimetry related errors by calculating an invariant interval and is for example typically used for the precision timing of PSR J0437–4715 within the Parkes Pulsar Timing Array project (Manchester et al., 2013).

The ToAs derived using these two methods have an rms value of 588 and 771 nanoseconds and $\chi^2/\text{d.o.f.}$ of 4.9 and 5.2, respectively. The latter $\chi^2/\text{d.o.f.}$ value is significantly better than the corresponding value for ToAs derived from total intensity as calculation of the invariant interval significantly reduces the S/N of the data for highly polarised pulsars, such as PSR J0437–4715, thus making the random error introduced by SWIMS less obvious by increasing the uncertainties in the estimated ToAs. Furthermore, systematic errors can be present in the invariant interval based on observations with CPSR2 due to the automatic gain control used by this baseband system.

It is possible to correlate the ToA residuals derived from matrix template matching with the residual profiles expressed in the basis spanned by the eigenvectors to obtain an rms timing residual of 449 ns and $\chi^2/\text{d.o.f.}$ of 2.9. The rms timing residual is similar to the value presented in section 4.5 and 23.5 per cent lower than the input value. The $\chi^2/\text{d.o.f.}$ value is smaller than presented in section 4.5 as the ToA uncertainty estimate based on matrix template matching is more realistic because it is the formal least squares

uncertainty, which takes into account some of the covariances between the fit parameters.

The method presented here is very sensitive to polarimetric calibration of the data and any imperfections in calibration will be detectable by presence of characteristic signals in the eigenvectors. Any effect that affects the polarisation of the pulsar should be detectable using this method. The fact that the first eigenvector for Stokes Q and U are correlated with the Stokes U and Q parameters in the template, respectively, can be explained by invoking the Earth’s ionosphere. Based on the International Reference Ionosphere⁸ (IRI) and International Geomagnetic Reference Field⁹ (IGRF) models, we derived a strong correlation with correlation coefficient of 0.83 between the absolute position angle derived in the calibration solutions and the position angle change induced by the ionosphere, as shown in Fig. 4.8. Note that the ionosphere has been quite stable during the week of observations, however the calibration solutions were derived at different times of the day, corresponding for different lines of sight through the atmosphere and thus different angles between the line of sight and the Earth’s magnetic field. The co-incidence of a spike in Stokes I, U and V and the mixing of Stokes Q and U is most likely a purely geometrical effect. It does not imply that these two types of variation are correlated as we project the observations onto a subspace of the N_{bin} dimensional space with the fit-space removed.

For pulsar timing experiments, it is often desirable to derive ToAs based on hour long observations. However, the temporal correlation of the mixing of Stokes Q and U implies that the covariance matrix may no longer scale as the inverse of integration time. This in turn renders the predictor derived from one time-scale not applicable to other time-scales. In principle, without red noise present and with ionosphere related effects removed, we should achieve rms timing residual of 30 ns with one hour integrations with the dataset presented. Note that this is comparable to what modern observing systems with four times larger bandwidth achieve (Manchester et al., 2013), stressing the spectral correlation of SWIMS. To render the predictors useful at all scales, calibration solutions need to be derived at higher time resolution, or ultimately the state of ionosphere can be monitored in near real-time using the dual-frequency Global Positioning System satellites (see e.g., Garner et al., 2008). Both these options should allow to remove the mixing of Stokes Q and U and enable higher precision experiments.

As noted in chapter 3 and Liu et al. (2011), PSR J0437–4715 is about 100 times brighter than other known MSPs. Assuming its properties are not very unusual compared to other pulsars used for timing experiments, we can thus expect to encounter similar problems when timing other MSPs with the next generation of radio telescopes such as

⁸<http://iri.gsfc.nasa.gov>

⁹<http://www.ngdc.noaa.gov/IAGA/vmod/igrf.html>

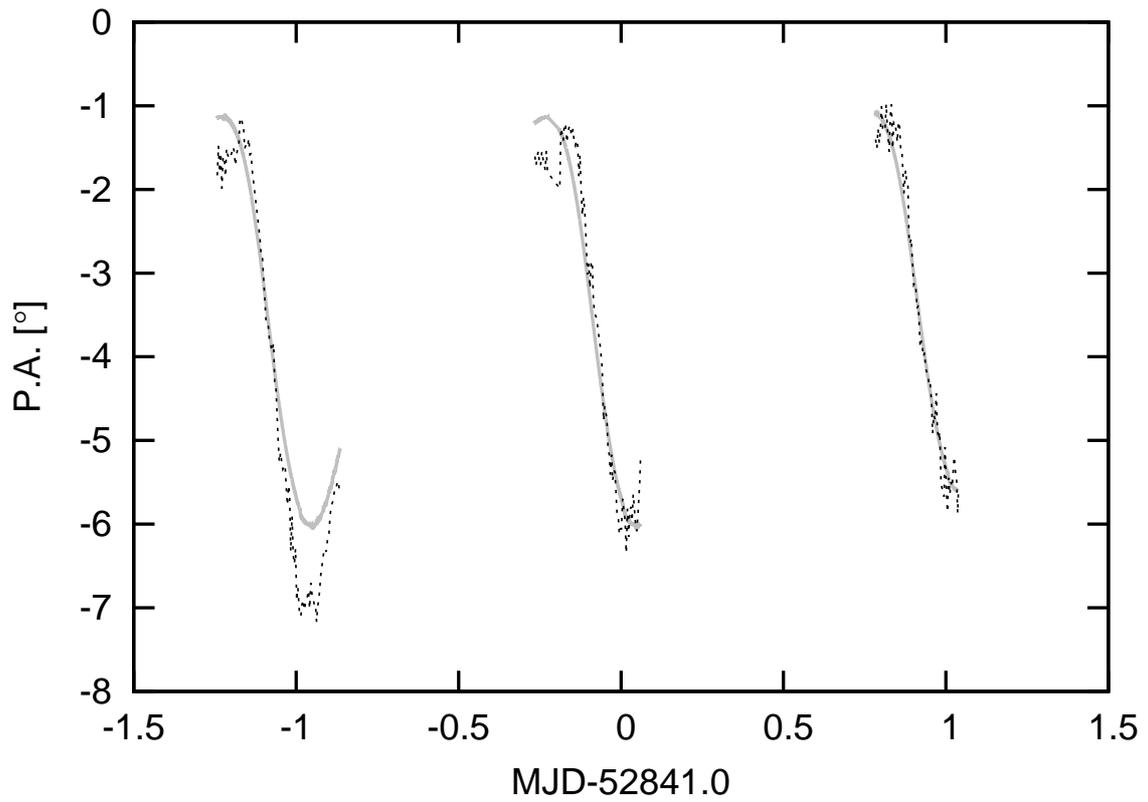


Figure 4.8: The position angle over three chosen days of observations. Grey solid line shows the combined predictions of IRI and IGRF while black dashed line corresponds to the measurements of the absolute position angle. We note that while the overall agreement is very good, at some epochs there are significant differences between the models and data.

the Square Kilometre Array or Five hundred metre Aperture Spherical Telescope. Based on the expected collecting areas and system temperatures of these instruments (Nan, 2006; Schilizzi et al., 2010), one can predict the timing precision of future timing experiments. Cordes et al. (2005), Cordes & Shannon (2010) and Liu et al. (2011) concluded that the timing precision of MSPs will be greatly improved with these instruments; however, it will likely be limited by the impact of SWIMS on the ToA estimation.

4.7 Conclusion

A solution to the long-standing problem of failing to achieve the expected timing precision for PSR J0437–4715 was presented in chapter 3. In the same chapter, a method for improving the precision of the timing of this pulsar was presented. Here, we address one limitation of the previous method and subsequently attain a more than 40 per cent reduction of the rms timing residual, more than a factor of two better than the method published in chapter 3. The reduced rms timing residual equals 476 ns, and corresponds to about 30 ns in one hour integration. If such a precision was achieved for a number of pulsars, the detection of gravitational waves would be much easier than with current precision attained. This corresponds to reduction of observing time necessary to achieve a given precision by almost a factor of three. Thus this method is likely to be very important for many pulsars when observed with the next generation of radio telescopes. The rms timing residual achieved with the method presented here is lower than can be attained with any other available method, including matrix template matching and timing of invariant interval pulse profiles. Currently, the applicability of this method, and thus the timing precision of PSR J0437–4715, is limited by ignoring the Earth’s ionosphere in the data analysis, thus providing a natural next step in the pursuit of ultimate timing precision and remains to be addressed in future work.

5

PSR J0437–4715 timing and polarimetry - a single-pulse perspective

Single pulses contain a wealth of information about a pulsar; including internal emission properties and insight into timing properties. It was shown recently that single-pulse emission is responsible for limiting the timing stability of the brightest of MSPs. In this chapter, we report on an analysis of single pulses from PSR J0437–4715 and present many of their properties including polarimetry. We present the S/N distribution, timing and polarimetry of average profiles integrated from single pulses with chosen S/N cut-offs, modulation properties of the emission, phase-resolved statistics of S/N , and two dimensional spherical histograms of polarisation vector orientation. These indicate the presence of orthogonally polarised modes (OPMs). Combined with dependence of polarisation fraction on the S/N and polarimetry of the brightest pulses, the existence of OPMs constrains pulsar emission mechanisms and plasma physics in the magnetosphere.

5.1 Introduction

So far this thesis has studied PSR J0437–4715 by analysing the statistics of the noise in the pulse profile and its relation to the estimated times of arrivals (ToAs) and their formal uncertainty. This chapter focuses on another way of studying pulsars, i.e., by analysing the single pulses emitted by the pulsar. While the average pulse profile is generally stable over time (Helfand et al., 1975), every single pulse is different (Drake & Craft, 1968b). Their amplitudes can be modulated temporally with or without a drift in phase (Cole, 1970; Edwards & Stappers, 2003). The emission can exhibit two orthogonally polarised modes (Manchester et al., 1975; Backer et al., 1976; Cordes et al., 1978; Backer & Rankin, 1980; Stinebring et al., 1984). Some pulsars give off giant pulses (Staelin & Reifenstein,

1968; Comella et al., 1969) and micropulses (Johnston et al., 2001), and some have high nulling fractions, that is fraction of time when the emission of the pulsar is not detectable (Backer, 1970b). The latter effect may be connected with the magnetospheric state and the spin-down rate (Kramer et al., 2006a). MSPs are typically less luminous and hence fainter than more slowly rotating pulsars. For this reason, single-pulse studies of MSPs are much harder. Such studies can shed light on the elusive problem of the pulsar emission mechanism, continuity of properties of MSPs and slower pulsars, plasma physics, or even affect the attainable timing precision.

Empirical and theoretical considerations of single-pulse properties and studies of slow pulsars' single-pulse emission has been undertaken in the past (e.g., most recently, Karastergiou et al., 2001, 2002, 2003a; Kramer et al., 2003; Karastergiou et al., 2003b; Melrose & Luo, 2004; Melrose et al., 2006; Bhat et al., 2007; Weltevrede et al., 2012). The sample of MSP single-pulse studies is more limited but still numerous (Cognard et al., 1996; Jenet et al., 2001a; Edwards & Stappers, 2003; Jenet & Gil, 2004; Knight et al., 2006; Burke-Spolaor et al., 2012). PSR J0437–4715, discovered in the 70-cm Parkes Survey (Johnston et al., 1993), is the brightest MSP known, making it a perfect candidate for single-pulse studies and a number of such studies has been undertaken previously (Johnston et al., 1993; Ables et al., 1997; Jenet et al., 1998; Vivekanand et al., 1998; Vivekanand, 2000).

In this chapter, we focus on the relation between the single-pulse properties and the timing of PSR J0437–4715 as well as on the polarimetric measurement properties of the single pulses. The structure is as follows: in section 5.2 we describe our observations including hardware used and details of data processing; section 5.3 focuses on an analysis of the total intensity properties of PSR J0437–4715; the polarimetric properties are presented in section 5.4; and we summarise our findings in section 5.5.

5.2 Observations and data processing

Observations of PSR J0437–4715 were recorded in three approximately hour-long sessions (two in November 2011 and one in March 2012) using the Parkes 64-m radio telescope and the central beam of the 20-cm multibeam receiver (Staveley-Smith et al., 1996). We used the CASPER Parkes Swinburne Recorder (CASPSR), a GPU-based digital signal processing system, capable of real-time phase-coherent dispersion removal in a 400 MHz wide band, and a baseband recorder. The latter mode was used to obtain the 8-bit real-sampled voltages from a 300 MHz wide band¹ centred at 1382 MHz. In total 177.3 minutes of data were recorded, or about 1.8 million pulsar rotations. A 512 channel filterbank

¹The remaining 100 MHz are filtered out before digitisation due to strong radio frequency interference

was created offline with the phase-coherent dispersion removal during its formation using DSPSR (van Straten & Bailes, 2011). The inter channel dispersion delays were removed as well before recording the data on a Redundant Array of Independent Disks (RAID array). All the subsequent processing was performed using the PSRCHIVE software suite (Hotan et al., 2004b; van Straten et al., 2012). The data were converted to PSRFITS format (Hotan et al., 2004b) and checked for incorrect header parameters.

Narrow-band radio frequency interference (RFI) was removed by applying a median filter by comparing the total flux in each channel with that of its 49 neighbouring channels. We attempted several methods of post-folding removal of impulsive RFI but all of them remove some of the brightest pulses, rendering automatic removal impossible. Jenet et al. (1998) demonstrated that single pulses of PSR J0437–4715 are much narrower than the average pulse profile, thus allowing to pre-select potential RFI by searching for wide pulses and visually inspecting all the pulses broader than 54 phase bins, the expected maximum width of a sub-pulse² (Jenet et al., 1998). Also, the top 1% of pulses sorted by signal-to-noise ratio (S/N) were visually scrutinised as the data with unexpectedly high S/N is likely to be RFI. This inspection found that for two of our three observing sessions only 0.01% of data is affected by strong impulsive RFI and thus we will ignore its presence throughout the analysis, unless otherwise noted. One hour of data recorded in November 2011 is corrupted by the presence of quasi-periodic impulsive RFI and is completely discarded from further analysis.

In order to perform polarimetric calibration using measurement equation template matching (METM) we apply the following procedure: An eight-hour long observing track of PSR J0437–4715 observed within a month of the observing session in November 2011 was calibrated using the measurement equation modelling technique (van Straten, 2004). A high S/N calibrated template was obtained and presented in Fig. 5.1. For each hour session of our main dataset, five minute integrations were formed by integrating single pulses in a hierarchical manner to avoid quantisation noise. These five minutes integrations were used to derive properties of the receiver at the time of observation. This information was used to calibrate the single pulses. See van Straten (2013) for details of the outlined procedure. Observations of the Hydra A radio galaxy, which is assumed to have a constant flux of 43.1 Jy at 1400 MHz and a spectral index of -0.91 (Baars et al., 1977), were used to calibrate the flux density scale.

²By a sub-pulse we refer to a single emission event within a single pulse, that is a single pulse can consist of multiple sub-pulses

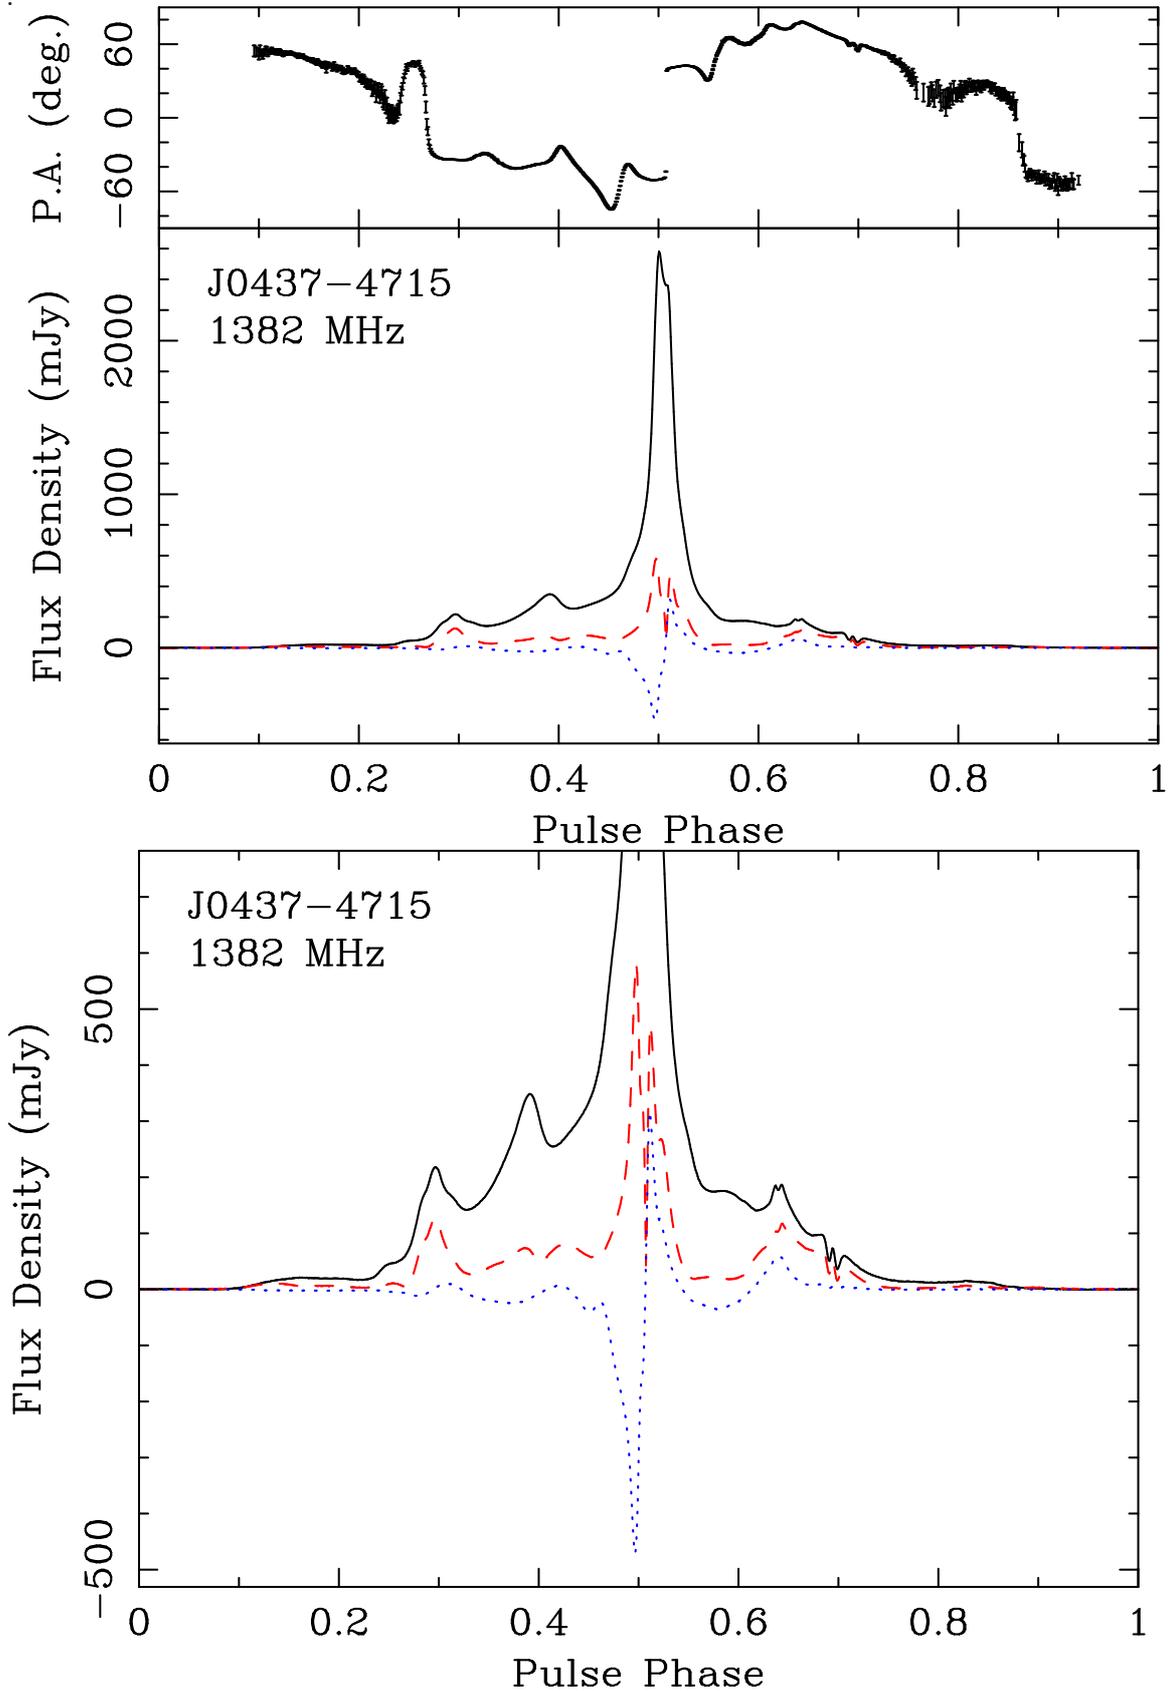


Figure 5.1: The high S/N ($\sim 20,000$) template for PSR J0437-4715 created from 8.5 hours of observations is plotted in the top panel. The solid black line represents total intensity, dashed red line corresponds to linear polarisation while the dotted blue line represents the circular polarisation. Bottom panel shows the same profile but zoomed in to facilitate examination of the details of the polarisation curves in the low flux regions.

5.3 Total intensity properties

We begin the data analysis with a focus on the total intensity of single pulses. We present the distribution of the single-pulse intensity, present dependence of the average pulse profile on the S/N threshold for data inclusion, discuss the impact of the pulse energy distribution on timing properties of PSR J0437–4715, analyse the implications of phase resolved flux distributions on plasma physics in the emitting region, and discuss the quasi-periodic stationary and drifting intensity modulation.

5.3.1 Distribution of instantaneous signal-to-noise ratio

The flux distribution of PSR J0437–4715 single pulses has been presented previously by Jenet et al. (1998). Here we present the distribution of instantaneous signal-to-noise ratio, as measured according to two separate definitions. The first one is:

$$S/N = \frac{\sum_{i=1}^{N_{\text{on}}} (F_i - B)}{\sqrt{N_{\text{on}} \sigma_{\text{off}}}}, \quad (5.1)$$

where F_i is the pulse flux in the i -th on-pulse bin; B is the mean off-pulse flux; N_{on} is the number of on-pulse phase bins; and σ_{off} is the off-pulse root mean square (rms) flux. The signal-to-noise ratio defined in this way is independent of the number of the phase bins used to resolve the pulse period. We note that the S/N of a narrow pulse with high flux can be the same as a broad pulse with smaller flux, making the definition of S/N sensitive to RFI which is often broader than the single pulses. Therefore, we checked our results using an alternative definition:

$$S/N_{\text{peak}} = \frac{F_{\text{max}} - B}{\sigma_{\text{off}}}, \quad (5.2)$$

where F_{max} is the maximum flux; S/N_{peak} is a phase resolution-dependent quantity but it is more useful for finding high flux spikes in the data. We note that both of these definitions are based on the off-pulse noise estimates and as such, do not take SWIMS into account in the definition of noise as measured by σ_{off} . We use 1024 phase bins throughout this chapter, unless otherwise noted.

The pulsar signal travels through the interstellar medium on the way to Earth. This has several consequences, such as dispersion and interstellar scintillation, the latter of which causes the apparent flux to vary significantly. By chance, the apparent flux was more stable during observations presented in this work and variations were within a factor of two. To facilitate comparison of single pulses with similar intrinsic S/N we use the

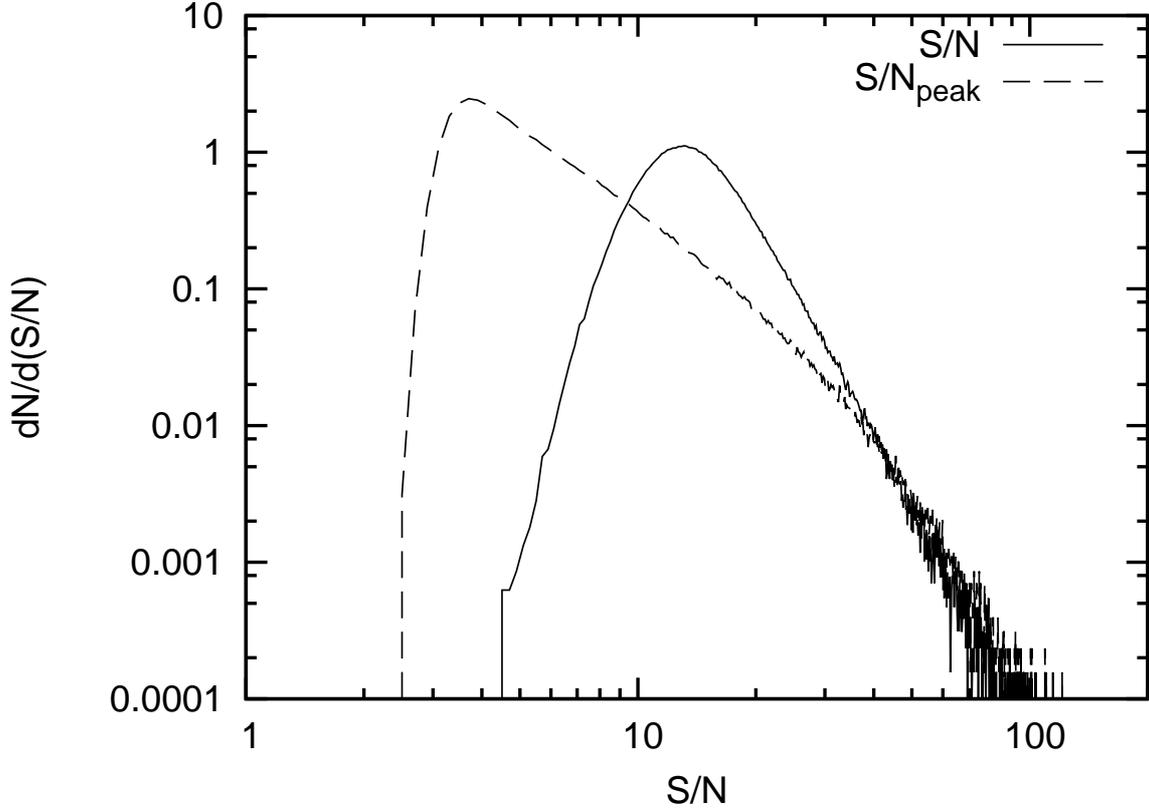


Figure 5.2: Observed instantaneous signal to noise ratio. The solid line is the distribution of S/N while the dashed line represents the distribution of S/N_{peak} . Overlap of the two distributions at high S/N range suggests that the brightest pulses are also most narrow. The distribution of normalised S/N values are very similar but shifted to lower values of S/N .

normalised quantity $S/N_n = \frac{S/N}{\langle S/N \rangle}$, where $\langle S/N \rangle$ is the mean S/N in a minute centred on the given single pulse.

The peak of S/N_{peak} distribution, as seen in Fig. 5.2, corresponds to a random peak in the noise as we have 1024 phase bins and expect a 3 standard deviations peak in each observation. The instantaneous median unnormalised S/N is 10. This immediately implies that SWIMS will be significant in observations of this pulsar (Osłowski et al., 2011b). Due to heteroscedasticity and temporal correlation of SWIMS, the ToA estimates of average pulse profiles of PSR J0437–4715 are randomly biased (Cordes & Downs, 1985; Cordes, 1993; D’Alessandro et al., 1993; Cordes et al., 2005; Cordes & Shannon, 2010; Osłowski et al., 2011b; Liu et al., 2012) increasing the root means square (rms) timing residuals (i.e. the difference between the modelled ToAs and the observations) by a factor of four. Chapters 3 and 4 have presented a statistical methods of correcting this random error by studying the noise in the pulse profiles. In the following two subsections we present an

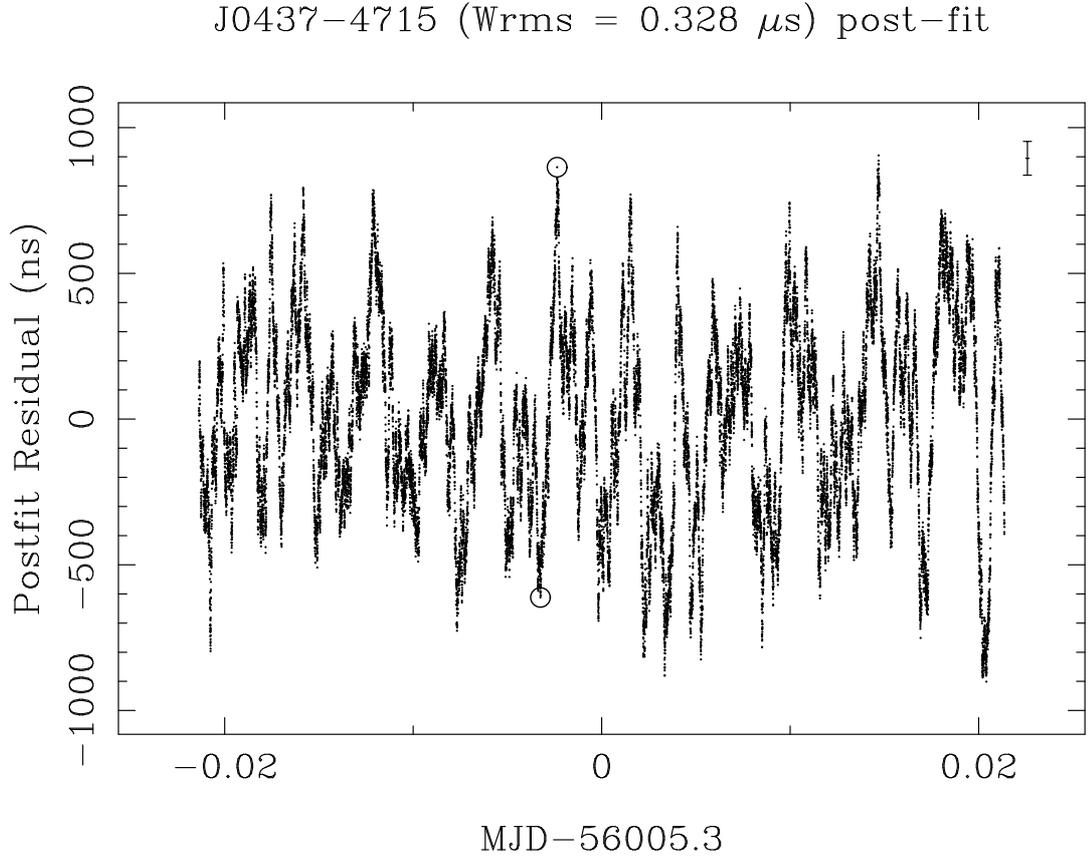


Figure 5.3: Running ToA estimate for one minute integrations. For practical plotting reasons, only every 25th point is plotted and no measurement errors are included. The two points marked with circles are taken 76.4 s apart and the difference in residual ToA is 1500 ns while the mean ToA estimation error is 58 ns. To facilitate comparison, we added a point with an error bar of this size in the top right corner of the figure. The apparent correlation is due to lack of independence between the data points.

investigation of a different approach to detecting the impact of SWIMS.

5.3.2 Impact of a single pulse on the time of arrival estimate

When the instantaneous signal to noise ratio (S/N) is larger than one, the effect of a single pulse on the ToA estimated from an integration of some length should be directly measurable. In order to measure this effect we implement the following procedure: we integrated one minute of data and formed an average pulse profile; this profile was timed against the template profile presented in Fig. 5.1; we then removed the first single pulse from this minute of data and added the next single pulse that follows the original minute of

integration; finally we go back to the first step and average the profile. We have analysed in this way all of the single pulses in one hour of data, thus forming a “running ToA” presented in Fig. 5.3. A typical ToA error estimate equals 58 ns. Note that we present only every 25th resulting ToA for practical reasons. The data points in this plot are not independent as two neighbouring ToAs are derived from data with 99.76% of overlap, producing the visible correlation in the figure.

We highlighted two points in this plot around the -0.0025 abscissa value. These two ToAs are separated by only 76.4 s and yet the difference between their estimates is 1500 ns, nearly 26 times the typical ToA error estimate. We can use these values to directly measure the effect of SWIMS on the ToA estimate. Every single pulse between the two highlighted ToAs has affected the estimated time of arrival by 0.11 ns on average. While this number seems negligible, it implies that if an observer decided to postpone both the start and finish of a 1-minute observation by one second, the estimated ToA would change by 19.1 ns. With next-generation radio telescopes, such as FAST or the Square Kilometre Array (SKA) (Schilizzi et al., 2010) an observer might be satisfied with the achieved S/N ratio after a one-minute integration; however the effect of SWIMS may dominate the ToA estimate.

We performed a set of simulations, similar to those described in section 3.4.2 and implemented with the “psrover” application available as part of PSRCHIVE software suite. In these simulations we tested the impact of a single pulse with varying width and flux on the ToA estimate. They show that a single pulse can affect the ToA even if it has only a fraction of the flux of the average profile, especially if it is a wide pulse. In case of PSR J0437–4715 the single pulses are quite narrow but can have a very high flux, as demonstrated in Fig. 5.4. For the sake of visibility of the average pulse profile, the template’s flux is artificially increased by a factor of 45. The amount of energy in a single sub-pulse can be larger than in the average pulse profile and is concentrated over a much narrower phase range. The average pulse profile will be measurably different depending on the inclusion of this sub-pulse in the average, impacting the template matching algorithm, used to derive the ToAs.

5.3.3 Average pulse profile dependence on sub-pulse S/N

We first present the dependence of the pulse profiles averaged over 1-minute on the thresholds of S/N_n applied to the data used for creation of these profiles. Extending the ideas of

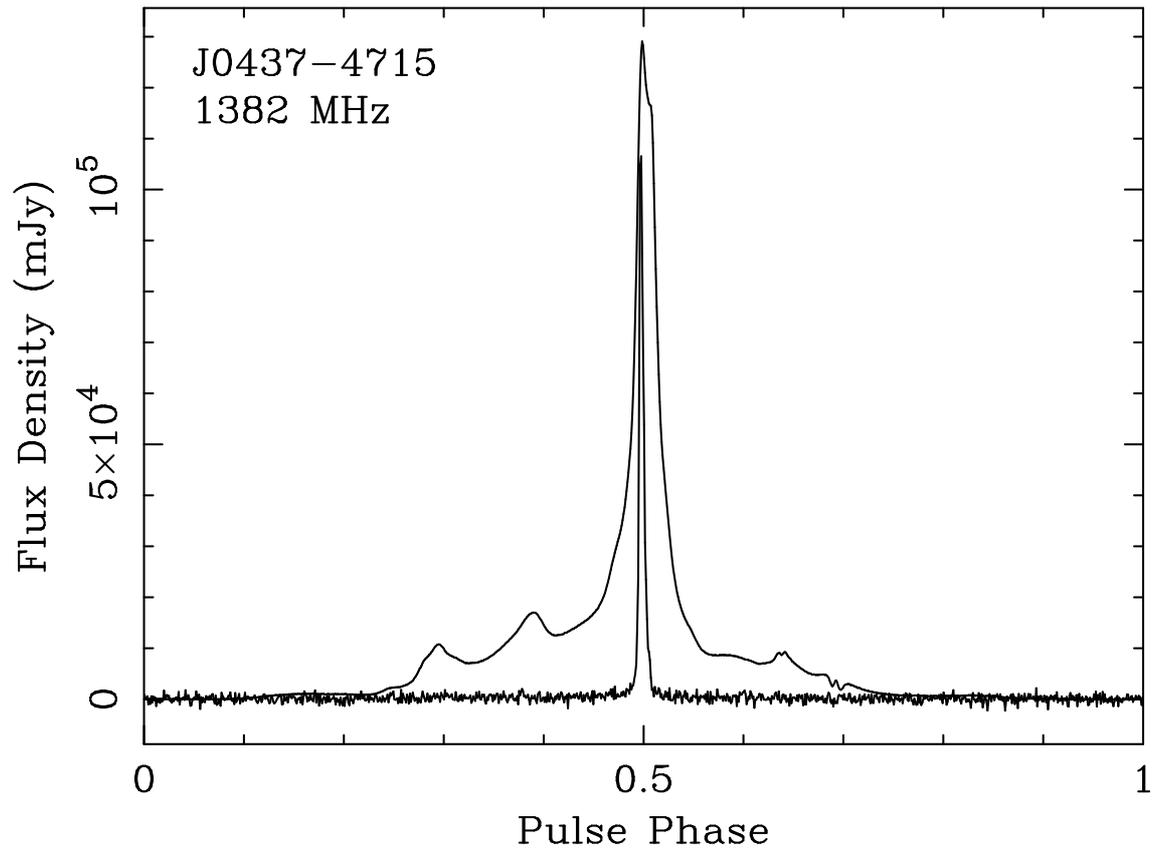


Figure 5.4: Comparison of a single pulse with an average pulse profile. The flux scale matches the single pulse and the flux of the average pulse profile has been multiplied by a factor of 45 for clarity. Total flux contained in this sub-pulse equals 636.8 Jy compared to the total flux in the average pulse profile of the observing session this pulse comes from of 155.8 Jy.

Krishnamohan & Downs (1983)³, we divide our data into following S/N_n ranges: [0.0; 0.5); [0.5, 0.6); [0.6, 0.7); [0.7, 0.8); [0.8, 0.9); [0.9, 1.0); [1.0, 1.1); [1.1, 1.2); [1.2, 1.3); [1.3, 1.4); [1.4, 1.5); [1.5, 2.0); [2.0; 4.0); [4.0 : 8.0); [8.0, 16.0); and > 16.0 . The average pulse profiles in a subset of these S/N_n ranges are plotted in Fig. 5.5. The averaged profiles were scaled to have the same peak intensity to facilitate comparison. The brightest pulses tend to arrive near the peak of the template profile intensity which is consistent with this region of the profile exhibiting the highest modulation. The peak of the average pulse profiles falls in an earlier pulse phase for the profiles constructed from brighter sub-pulses. Shannon & Cordes (2012) have seen similar behaviour for PSR J1713+0747, which exhibits a correlation between the S/N of the sub-pulses and their time of arrival. In the average pulse profile formed from the weakest pulses an additional component on the leading edge becomes visible. This component may correspond to the thickening of the average pulse profile seen on the leading edge of the template profile, see Fig. 5.1, phases around 0.47.

We note that all the profiles, even those constructed from the brightest pulses, exhibit the same structure of “wings” as the average profile. This effect persists in all the S/N ranges. Visual inspection of a number of the brightest sub-pulses, conducted during the search for RFI described in section 5.2, reveals that not a single one of them has any visible emission in the wings phase region, and vast majority of the sub-pulses are detected only in 20 phase bins near phase 0.5. The presence of wings in the average pulse profile of these brightest pulses implies the presence of low level emission, which is only revealed by averaging away the radiometer noise. This suggests that the pulsar magnetosphere is local, i.e., the presence of a bright sub-pulse does not affect other regions of magnetosphere, which continue to emit without being affected by any physical process which generated the bright emission.

5.3.4 Timing properties as a function of sub-pulse S/N

Because SWIMS begins to impact the ToA when the instantaneous S/N exceeds unity, it may be possible to limit the amount of such noise in the data by rejecting the brightest sub-pulses as identified by their normalised S/N . While such a procedure will reduce the S/N of the average pulse profiles, it may be beneficial as the noise will be closer to having equal variance in each pulse phase bin and not being temporally correlated, that is, presence of SWIMS will be limited. We note that this will not remove any of the spectral correlations in the noise. In Table 5.1, we present the rms timing residuals

³The authors referred to the process of selective integration as gating which nowadays is often used to describe selecting a phase range of the pulse profile.

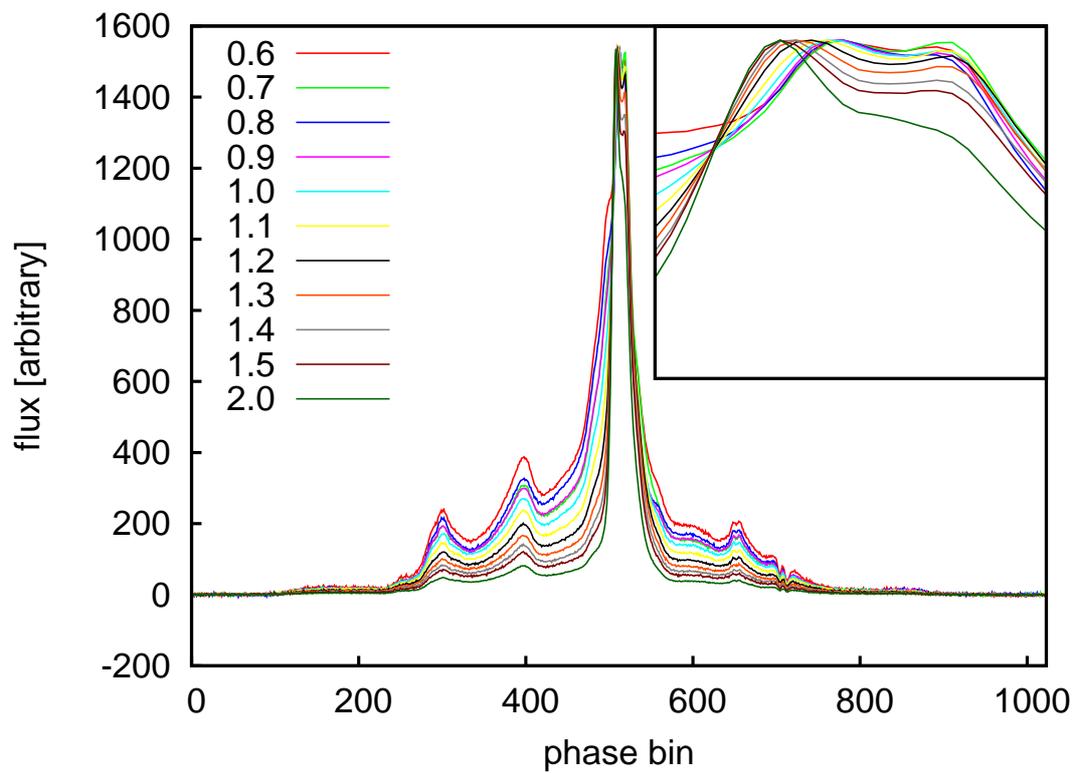


Figure 5.5: Average pulse profile integrated using only data within given instantaneous normalised S/N ranges. The profiles were normalised to have the same peak flux. The numbers in the label denote to the top end of the normalised S/N range. The lower limit of any range is the upper limit of the previous range, or 0.0 in the case of range denoted as 0.6. The inset in top left corner zooms in on the peak region.

Table 5.1: Rms timing residual for data formed by excluding sub-pulses in given S/N ranges. The rms timing residual is given in nanoseconds.

Included range	rms timing residual	$\chi^2/\text{d.o.f.}$	Fraction of pulses
$0.0 < S/N < 0.5$	4998	22.3	0.0054
$0.0 < S/N < 1.0$	449	14.3	0.6011
$0.0 < S/N < 2.0$	331	31.1	0.9476
$0.0 < S/N < 4.0$	314	34.7	0.9673
$0.0 < S/N < 8.0$	312	34.5	0.9679
$S/N > 0.0$	338	39	1.000
$S/N > 0.5$	316	34.4	0.9625
$S/N > 1.0$	340	44.9	0.3669
$S/N > 2.0$	931	181.3	0.0204
$S/N > 4.0$	3790	157	0.0007

for data formed by including sub-pulses with normalised S/N in the following ranges: $0 < S/N_n < 0.5$; $0 < S/N_n < 1.0$; $0 < S/N_n < 2.0$; $0 < S/N_n < 4.0$; $0 < S/N_n < 8.0$ and the complementing ranges: $[0.0; 0.5)$; $[0.0; 1.0)$; $[0.0; 2.0)$; $[0.0; 4.0)$; $[0.0; 8.0)$ and the complementing ranges: > 0.5 ; > 1.0 ; > 2.0 ; > 4.0 ; > 8.0 ; In this case we have taken the identification of RFI into account and used only the data that was not flagged as affected by RFI. The table also contains the fraction of pulses which have been classified as belonging in the respective range.

The rms timing residual is marginally reduced for the data formed by rejecting the sub-pulses of S/N_n greater than two, four or eight. The $\chi^2/\text{d.o.f.}$ is lower in these cases as well as for any other case when upper limit of S/N_n is imposed. As the high instantaneous S/N pulses are removed from data, the relative contribution of SWIMS is reduced along with a reduction of the average S/N . This causes the ToA measurement uncertainty estimates to be both larger and more realistic, thus reducing the $\chi^2/\text{d.o.f.}$. On the opposite end of the spectrum, the brighter the included sub-pulses are, the ToA estimates become dominated by SWIMS and the timing is strongly affected, visible in both the rms timing residual and the $\chi^2/\text{d.o.f.}$ of the fit. While it is not possible to significantly improve the rms timing residual by selectively rejecting sub-pulses in this dataset, it may be possible for other pulsars with a different distribution of sub-pulse fluxes.

5.3.5 Observed electric field intensities and plasma physics

The pulsar emission mechanism remains elusive despite many years of research in this area. Many possibilities have been presented. Following Cairns et al. (2003b) we list the proposed linear processes including, but not limited to linear acceleration and maser cur-

vature emission (Luo & Melrose, 1995; Melrose, 1996, 2003), relativistic plasma emission (Melrose, 1996; Asseo, 1996), and a streaming instability into an escaping mode (Gedalin et al., 2002). Non-linear processes have been proposed by several groups. The possibilities include direct conversion of plasma turbulence into electromagnetic emission (Weatherall, 1997, 1998) soliton collapse (Asseo & Porzio, 2006, and references therein), and antenna mechanism (Pottelette et al., 1992; Cairns & Robinson, 2000). All these models have trouble explaining all of the rich pulsar emission phenomenology, including time variability and microstructure. It is therefore plausible that either none of these mechanisms provide an accurate description of pulsar magnetosphere or that a combination of these mechanisms is occurring simultaneously.

As pointed out in a series of papers by Cairns and collaborators (Cairns et al., 2001, 2003a,b, 2004), field statistics provide insight into plasma behaviour in the pulsar's magnetosphere. In these papers, two competing theories of interactions between plasma waves, driving particles and the background plasma are summarised. In brief, self-organised criticality (SOC Bak et al., 1988; Bak, 1996) describe systems interacting self-consistently without any preferred distance- or time-scales and predicts power law distributions of intensities with indices typically in the range 0.5 to 2.0. Stochastic growth theory (SGT) (Robinson, 1992, 1995; Robinson et al., 1993; Cairns & Menietti, 2001) in turn describes self-consistently interacting systems where the interactions take place in an independent homogeneous medium and introduce distance- and time-scales. This theory predicts a log-normal distribution of electric field.

The field statistics have been studied so far for the Vela pulsar in the aforementioned series of papers by Cairns and collaborators, for PSR J1644–4559 and PSR J0953+0755 by Cairns et al. (2004) and Smirnova (2006) demonstrated that the intensity is distributed log-normally at certain phases. The search for giant pulses, as seen in the Crab pulsar (Staelin & Reifenstein, 1968), in four MSPs was conducted by Knight et al. (2006). The authors found that PSR J0218+4232 emits giant pulses with power law distribution of energies. Burke-Spolaor et al. (2012) studied the field statistics for a sample of 315 pulsars, including two MSPs PSR J1439–5501 and PSR J1744–1134 but these distributions were not phase-resolved. The MSPs were found to have non-Gaussian and non-log-normal distribution of energies. Another related study of MSPs was done previously by Cognard et al. (1996) who studied the giant pulses from PSR J1939+2134. They found that the amplitudes of these pulses follow power law distributions that agree for the pulses from the main pulse and inter-pulse.

While power law distributions of electric field have been interpreted in the past as

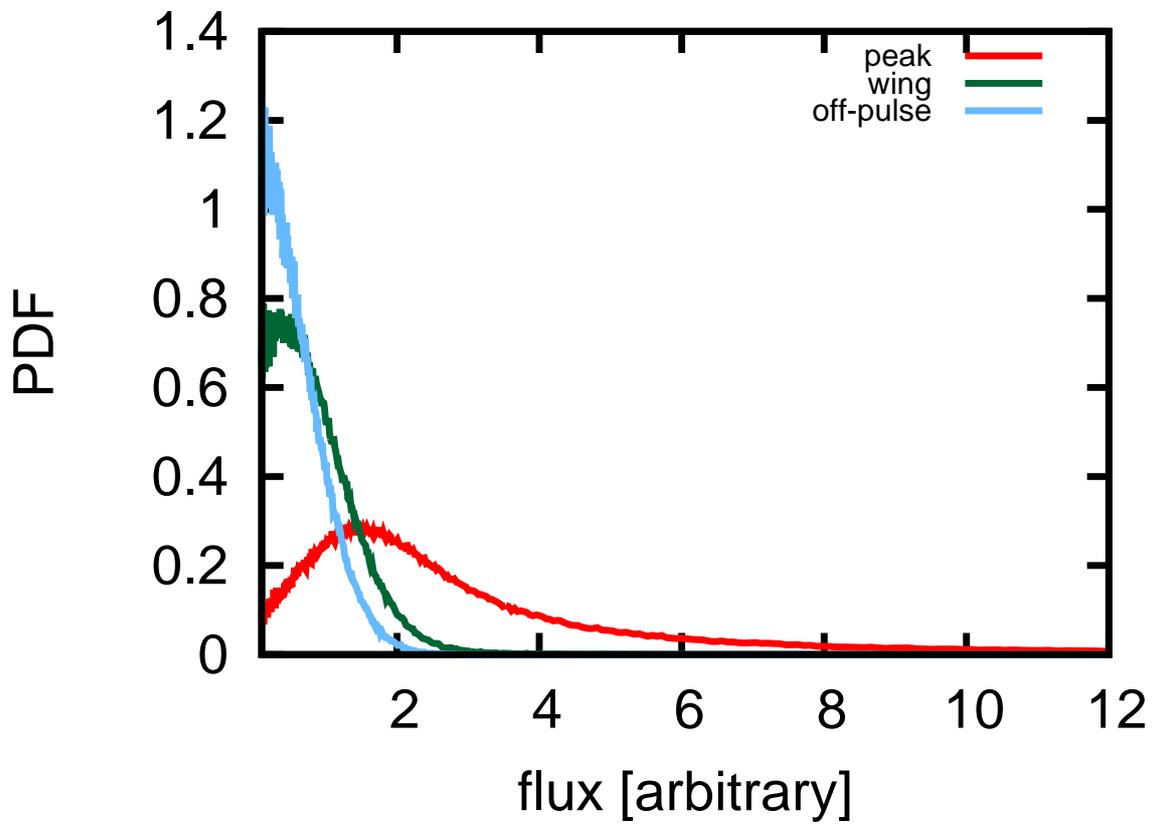


Figure 5.6: Phase-resolved intensity probability density function.

evidence of non-linear processes in the pulsar's magnetosphere, the observed indices did not agree with theoretical predictions. Cairns et al. (2003a) proposed an explanation for this discrepancy. They argue that power law distributions can be mimicked by two plasma wave populations that follow either log-normal statistics, as predicted by SGT, or normal statistics. Cairns et al. (2002) present a convolution technique that also applies to electromagnetic fields originating from different source regions. This explains the unexpected indices of power law distributions without a necessity for non-linear processes to be involved.

To calculate the distribution of the electric field's magnitude, we used the formula:

$$P(\log E) = 2 \log 10 E^2 P(I) , \quad (5.3)$$

where I is the intensity and E is the wave electric field. We note that in practice, to avoid rejecting very low flux samples, arbitrary offset is added to the observed flux. This corresponds to recovering the average flux of the pulsar and causes a shift and scaling of the abscissa in the plot of field statistics. See e.g., Cairns et al. (2003a) for details.

The measured phase resolved flux distributions are shown in Fig. 5.6 for a few chosen phase ranges. The blue curve shows the noise distribution in the off-pulse region where we expect only radiometer noise; the green curve shows the distribution in the low flux region around phase 0.4; and the red curve corresponds to the peak of the average pulse profile. Both the off-pulse and low flux regions are well fit by a Gaussian distribution; however, the wing region exhibits a small excess of pulses with high energies. It is not possible to obtain an acceptable fit of the flux distribution near the peak intensity and a log-normal fit is necessary, consistent with SGT predictions. An investigation of both the low and high flux tails of the distribution reveals that a better fit may be in fact obtained by a convolution of two populations, as suggested by Cairns et al. (2003a). This behaviour, already demonstrated for a number of pulsars, provides further support for the hypothesis that plasma in the pulsar magnetosphere is indeed well modelled by SGT.

5.3.6 Intensity modulation

Many pulsars exhibit periodic intensity modulations. Sometimes the modulation function is stationary while in other cases it drifts with respect to pulse phase. The drifting phenomenon was first discovered by Drake & Craft (1968a) by visual inspection of sequences of sub-pulses. We have searched our data for evidence of such phenomena by calculating the longitude resolved fluctuation spectra (Backer, 1970a,c, 1973; Backer et al., 1975) and

two dimensional fluctuation spectra (Edwards & Stappers, 2002). An implementation of both of these spectrum calculations is available as the “drifting_sub-pulses” application, developed by the candidate, within the PSRCHIVE software suite.

We found no evidence of either stationary or drifting quasi-periodic intensity modulation. This places constraints on some models of PSR J0437–4715. Gil & Krawczyk (1997) predict that PSR J0437–4715 should manifest stationary modulation of intensity. The lack of sub-pulse drift is somewhat surprising as Weltevrede et al. (2006) found that it’s quite a common behaviour of pulsars. They point out that 68 of their luminosity biased sample of 187 pulsars exhibit drifting sub-pulses, with older pulsars more likely to do so. After considerations of their selection effects they conclude that 55 or more per cent pulsars are drifters and another 5 per cent exhibit stationary modulation. In one of the models, lack of drifting in PSR J0437–4715 implies that this pulsar spins so that the relative orientation of rotation spin and polar-cap magnetic flux density is positive as it is the opposite direction of spin provides physical basis for phenomena such as nulling, drifting, and mode switching (Jones, 2012).

We note that the search for intensity modulation in PSR J0437–4715 performed using data from the ATNF Parkes Swinburne Recorder (APSR van Straten & Bailes, 2011), a frequency-division multiplexing baseband recorder and a coherent dedispersion system produced a different result. Both the longitude resolved fluctuation spectrum and the two dimensional fluctuation spectrum showed a clear signature of intensity modulation. Fig. 5.7 presents a high resolution two dimensional fluctuation spectrum. The power spectrum is colour coded using the cubehelix colour palette (Green, 2011) in which the brightness perceived by human eye changes linearly on the scale on the right hand side. The detected intensity modulation occurs across the full phase range, including the off-pulse region. This implies either the presence of errors in the software used for calculation of the fluctuation spectra, artefacts in data introduced by the down-conversion system or the pulsar-specific backends used for this observation, or intensity-modulated low level emission of the pulsar are present across the whole phase range.

Based on numerous tests of the code and the fact that data obtained with CASPSR does not exhibit this modulation, we conclude that this must be an instrumental effect. The baseband data from APSR and CASPSR presented in this chapter are processed with the same software. However, APSR receives its data stream not directly from the down-conversion system, but it uses Pulsar Digital Filterbank (PDFB3) which employs a two-stage analysis/synthesis polyphase filterbank to perform the frequency-division multiplexing. The artefacts must be therefore either introduced by PDFB3 or during the

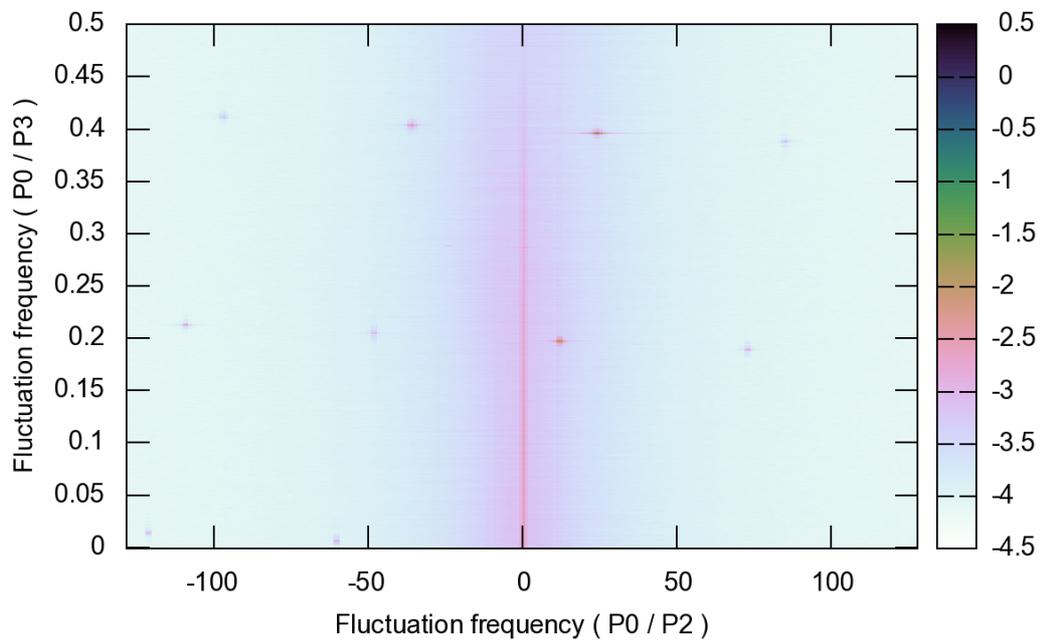


Figure 5.7: Two dimensional fluctuation spectrum obtained based on data recorded with APSR. The pulsar period here is denoted by P_0 , P_2 is the horizontal drift band separation, and P_3 is the vertical drift band separation, see Fig. 1 in Weltevrede et al. (2006) for details of these definitions. The brighter band around the fluctuation frequency equal to zero is an artefact introduced by the average profile. Other coloured areas in the plot are most likely of an instrumental origin, see text.

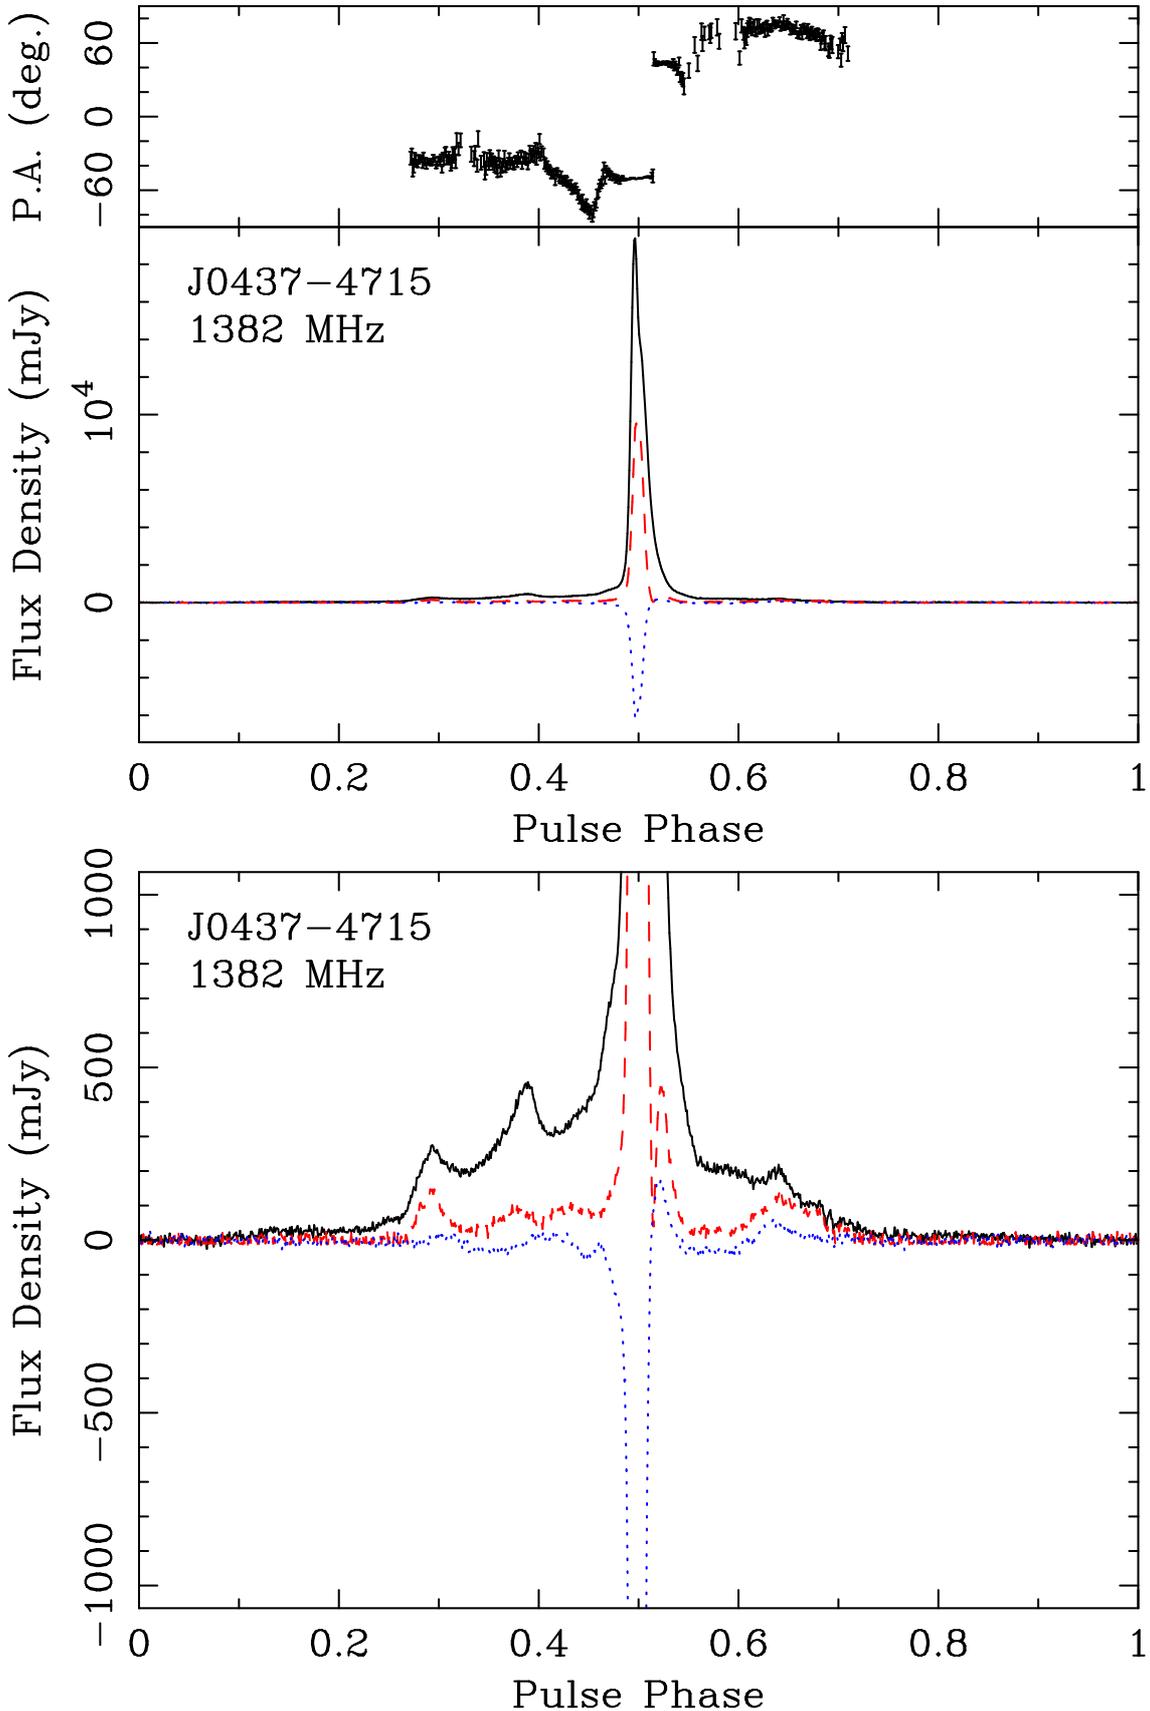


Figure 5.8: The average pulse profile of the brightest 1% of sub-pulses. The solid black, dashed red, and dotted blue lines correspond to Stokes I, linear polarisation, and Stokes V, respectively. Bottom panel shows the same pulse profile but zoomed in to demonstrate the details of polarisation in the “wings” of the profile.

synthesis stage. It is unclear which of these two is true, however we have already demonstrated that PDFB3 introduces some artefacts in the data (see Fig. 3.9) which occur on the time-scales corresponding to the low frequency signal in Fig. 5.7.

5.4 Polarimetry

We now move on to discuss some properties of the polarised radiation observed. We begin by discussing the dependence of polarisation on the signal-to-noise ratio. Afterwards we describe the detection of orthogonally polarised modes.

5.4.1 Polarisation dependence on S/N

As in section 5.3.3, we now investigate the polarisation properties of the sub-pulses in chosen S/N_n ranges. As previously noted, when average profiles are formed by selectively integrating sub-pulses, the structure of the “wings” of the profile remains fixed with only the ratio of the peak flux to flux in the wings changing. An example of the pulse profile constructed from the 1% brightest pulses as defined by their S/N_{peak} is shown in Fig. 5.8. We note that the polarisation intensity in all the Stokes parameters of the profile in Fig. 5.1 are consistent within noise with those of the template in Fig. 5.1. However, the brightest pulses in the centre of the averaged pulse profile have a different polarisation profile and degree of polarisation as readily visible in the bottom panel of Fig. 5.8. The low level emission appears independent of the emission in the central part as its properties are consistent with those of the template profile. The PA curve of the brightest pulses also follows the relevant curve of the template profile.

The change in the polarisation degree is also visible when we investigate it in the data formed from selective integration in the same S/N_n regimes as considered previously. Fig. 5.9 shows how the degree of linear and circular polarisation changes across the integration ranges within 40 phase bins from phase 0.459 to 0.498, i.e. where the biggest difference in polarised pulsed profiles in Fig. 5.8 is visible. Note that the degree of linear polarisation is calculated as $\sum_i \sqrt{Q_i^2 + U_i^2}/I_i$ and that of circular polarisation as $\sum |V_i|/I_i$, where summation is over phase bins in the aforementioned phase range and no bias removal was performed. Bias is likely to be responsible for increasing the degree of polarisation when measured in the lowest S/N_n regimes. The steady rise of polarised flux in the central part of the pulse profile with increasing S/N_n threshold is apparent.

Mitra et al. (2009) analysed the polarisation of single pulses from 10 classical pulsars. They found that all the objects in their sample emit bright pulses that have a high degree

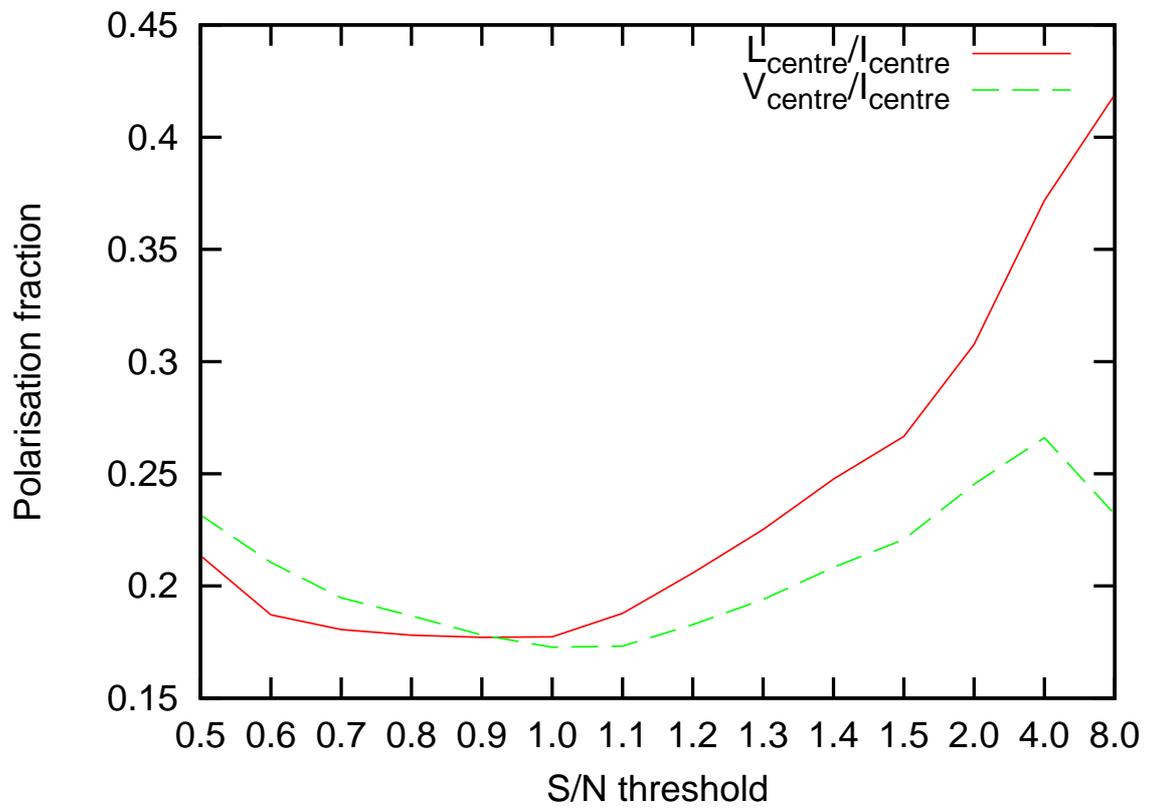


Figure 5.9: Polarisation fraction as a function of S/N_n . The red solid and green dashed lines correspond to the degree of linear and circular polarisation, respectively.

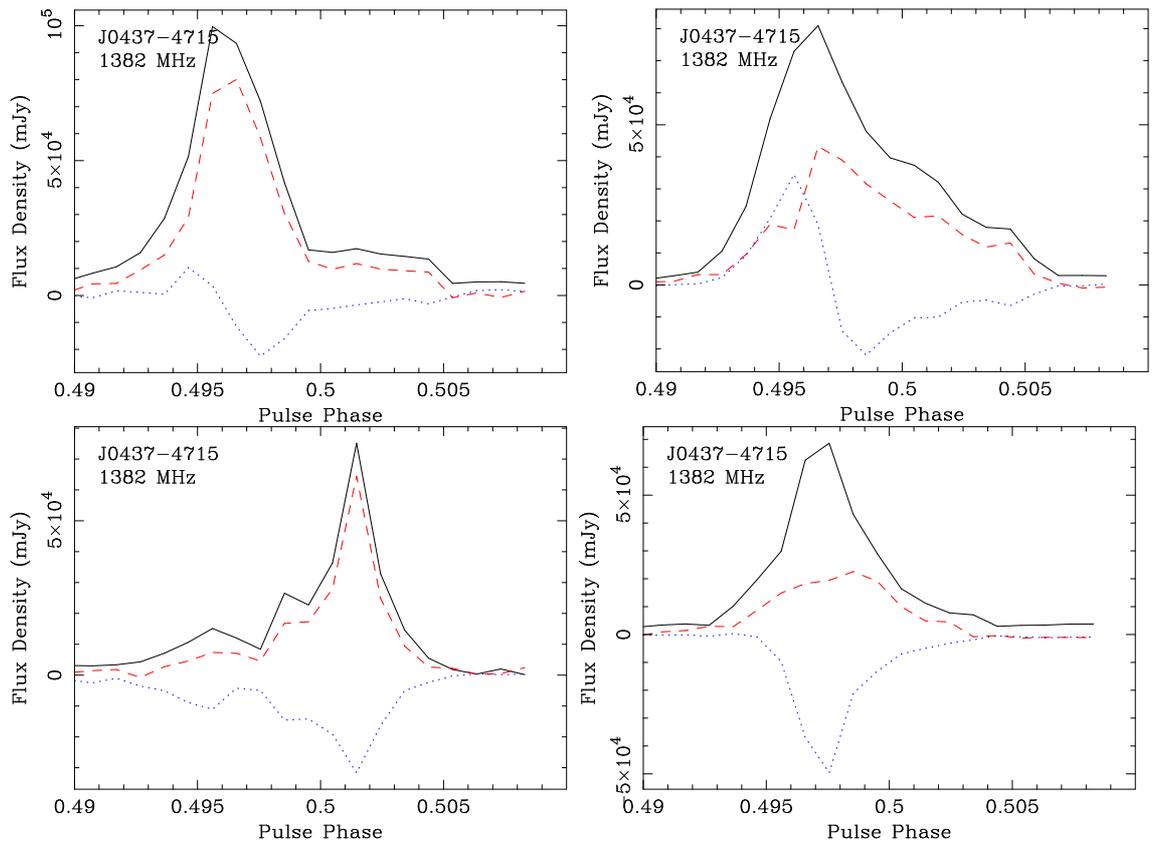


Figure 5.10: Examples of bright sub-pulses. Note that about half of the 1 per cent brightest sub-pulses before the phase 0.5 exhibit the swing of circular polarisation while none of those after this reference phase do not.

of linear polarisation and a change of the sense of circular polarisation across the sub-pulse. These features originate from the coherent curvature radiation from a single charged particle travelling at relativistic speed in the presence of magnetic field (Michel, 1987), producing an escaping extraordinary plasma wave (Arons & Barnard, 1986). Mitra et al. (2009) argue that a soliton is a good candidate for such a particle. A sample of bright pulses, as selected from the top 1% of brightest pulses according to their S/N_{peak} , is shown in Fig. 5.10. The two top pulses exhibit the same characteristics that Mitra et al. (2009) discussed. The example in the bottom left panel shows a pulse that has a high degree of linear polarisation but the Stokes V does not change the sense across the pulse while bottom right panel is an example of a sub-pulse with a lower degree of polarisation. The data indicate that at least some of PSR J0437–4715 emission can originate from coherent curvature radiation of solitons. We note that only the pulses occurring before the phase 0.5 exhibit a swing of circular polarisation.

5.4.2 Phase resolved polarisation histograms

We analysed the polarimetric properties of PSR J0437–4715 by studying phase resolved histograms of Stokes parameters on a Poincaré sphere. We implemented this using the Hierarchical Equal Area isoLatitude Pixelation of a sphere (HEALPix, Górski et al., 2005) with the resolution parameter set to three, corresponding to 768 pixels on the sphere. The generation of Poincaré sphere histograms was performed with the “psrpol” application written by the doctoral candidate within the PSRCHIVE framework. This application reads in every sub-pulse and measures the Stokes vector at every rotation phase. It computes the appropriate bin on the sphere weighing the contribution to that pixel by the polarised flux at the given spin phase. We present a few chosen Poincaré spheres in Fig. 5.11, 5.12, 5.13, and 5.14⁴ with colour coding showing the value of the spherical histogram in given bin. This representation is similar to that of Edwards & Stappers (2004) but we used the Mollweide projection of the sphere onto a two dimensional image. We note that the mean of the polarisation vector varies significantly across the pulse phase, well beyond the variance of the polarisation at any phase.

Orthogonally polarised modes

The most notable feature of the polarisation patterns is the presence of orthogonally polarised modes (OPM Stinebring et al., 1984) near the peak flux of the total intensity;

⁴An animation showing the evolution of the Poincaré sphere with the rotation phase can be obtained from the author.

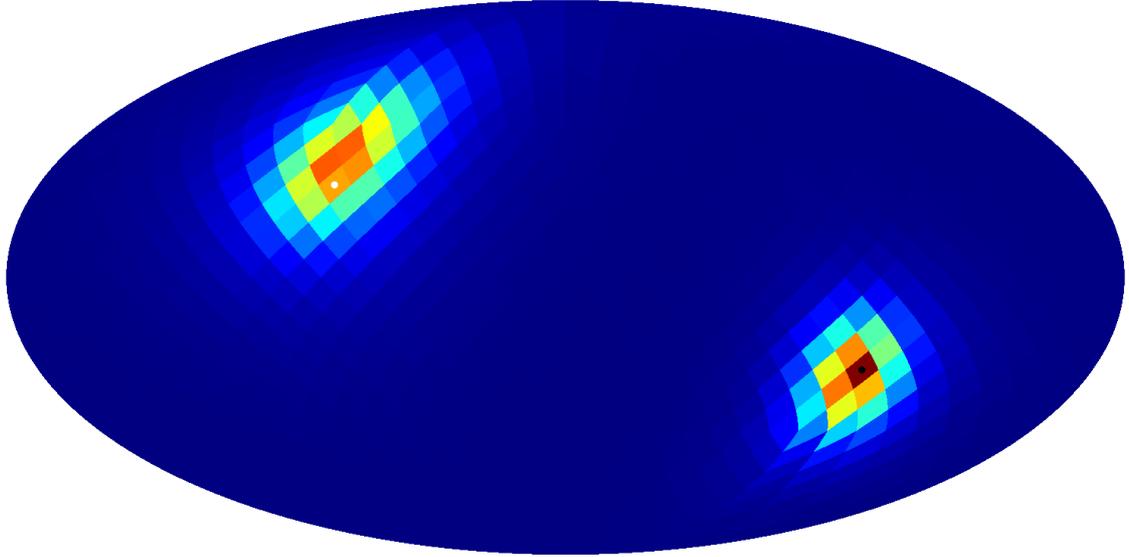


Figure 5.11: Mollweide projection of the Poincaré sphere at phase 0.512. Orthogonally polarised modes are clearly visible. The poles and equator correspond to purely circular polarisation and pure linear polarisation, respectively. The black and white dots correspond to the dominant polarised mode and its orthogonal counterpart.

to our knowledge this is the first direct detection of OPM in an MSP. We do not see annuli⁵ as seen by Edwards & Stappers (2004) in the object they studied. The modes are consistent with not deviating from orthogonality as can be seen with the visual guide of black and white dots in Fig. 5.11, with the former centred in the peak of one of the modes and latter on the opposite side of the sphere. Detection of the modes near the peak intensity is expected, as this is the region with highest modulation, as shown in chapters 3 and 4. The association of OPMs with highly modulated emission has been first noted by McKinnon (2004) and explained statistically by van Straten (2009). Furthermore, Dyks et al. (2007) and Dyks et al. (2010) argue existence of orthogonally polarised modes in PSR J0437–4715 and that the trailing side of the pulse profile must be dominated by the extraordinary polarisation mode as only this can explain the creation of notches in, i.e., the “w” shaped feature of the average pulse profile visible around the phase 0.7.

5.5 Conclusions

We have presented a study of PSR J0437–4715’s sub-pulse properties. We attempted to reduce the post-fit arrival time residual by minimising SWIMS via rejection of the sub-pulses with the highest flux while taking scintillation into account. While we have not

⁵We note that our low resolution on the Poincaré sphere, dictated by S/N, may be the cause of this.

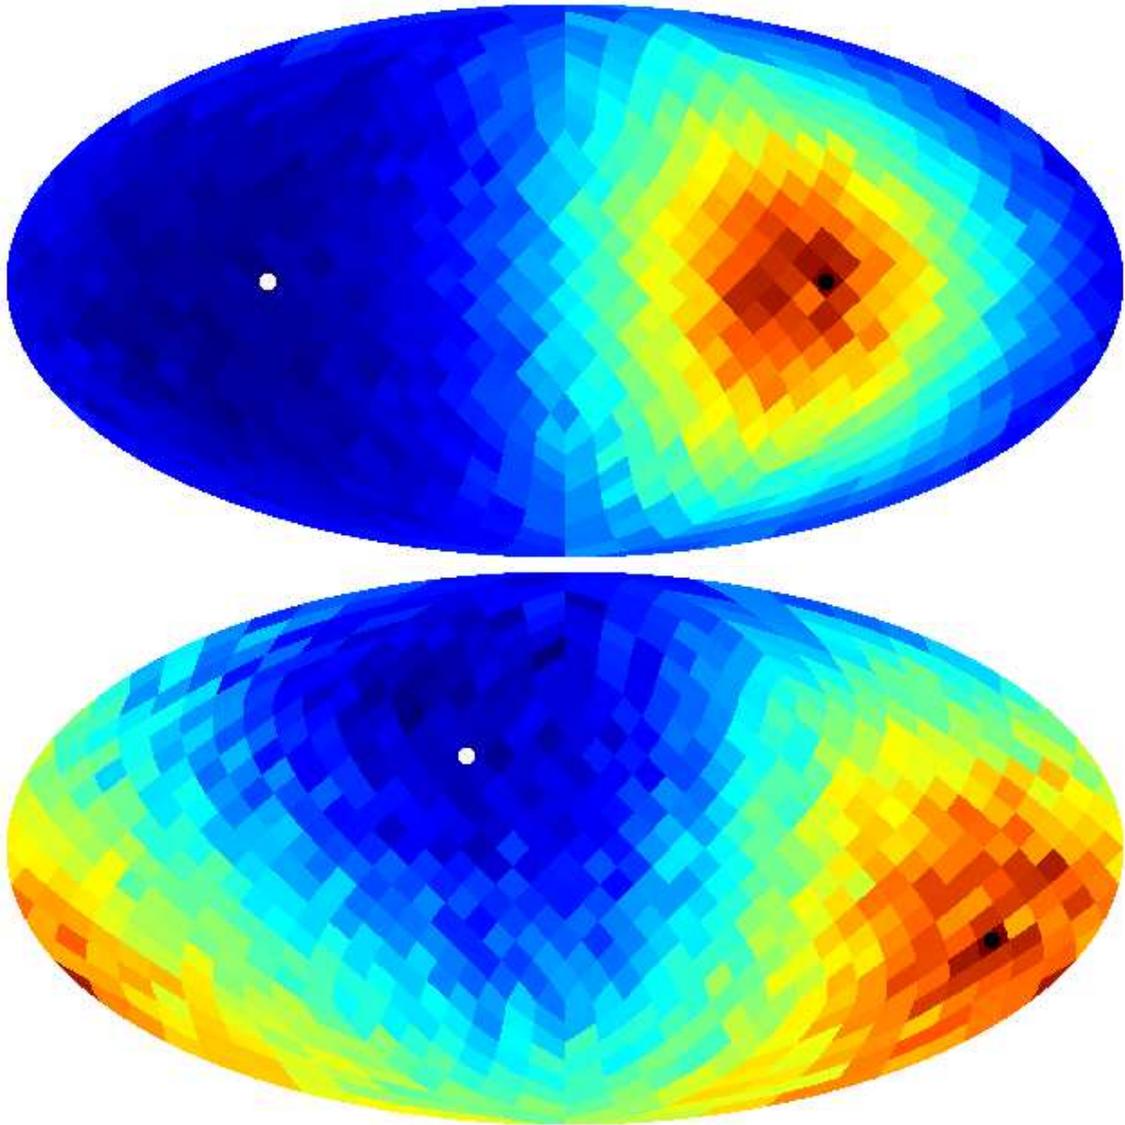


Figure 5.12: Mollweide projection of the Poincaré sphere near the phases 0.293 (top) and 0.449 (bottom). The black and white dots correspond to the dominant and the secondary polarised mode and its orthogonal counterpart.

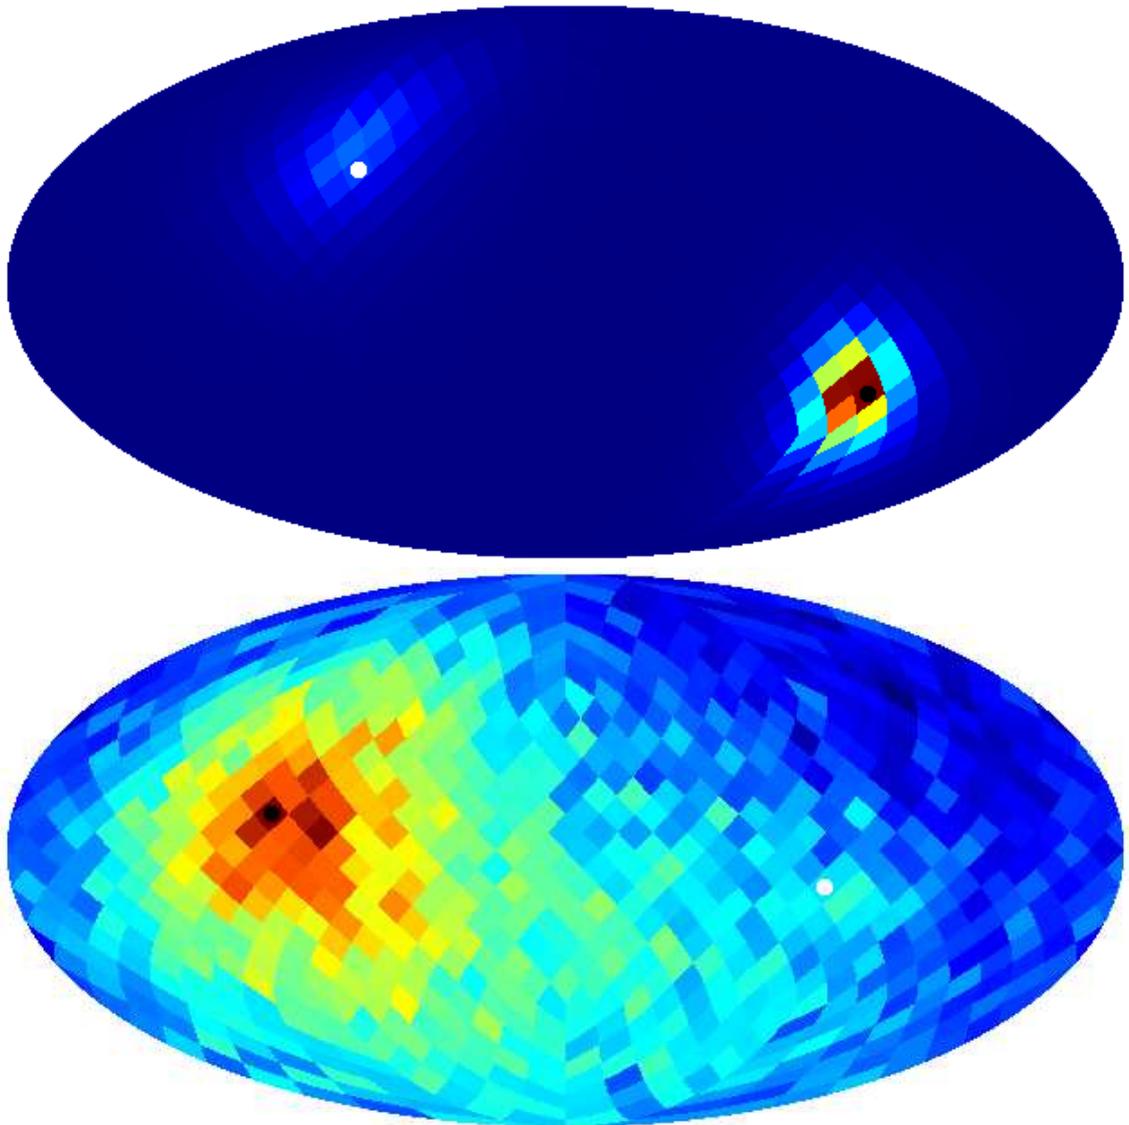


Figure 5.13: Mollweide projection of the Poincaré sphere near the phases 0.504 (top) and 0.547 (bottom). The polarisation vector visible at phase 0.547 may have a bimodal orientation distribution within a narrow angle on the sphere, however the data S/N is too low to draw decisive conclusions. The black and white dots correspond to the dominant and the secondary polarised mode and its orthogonal counterpart.

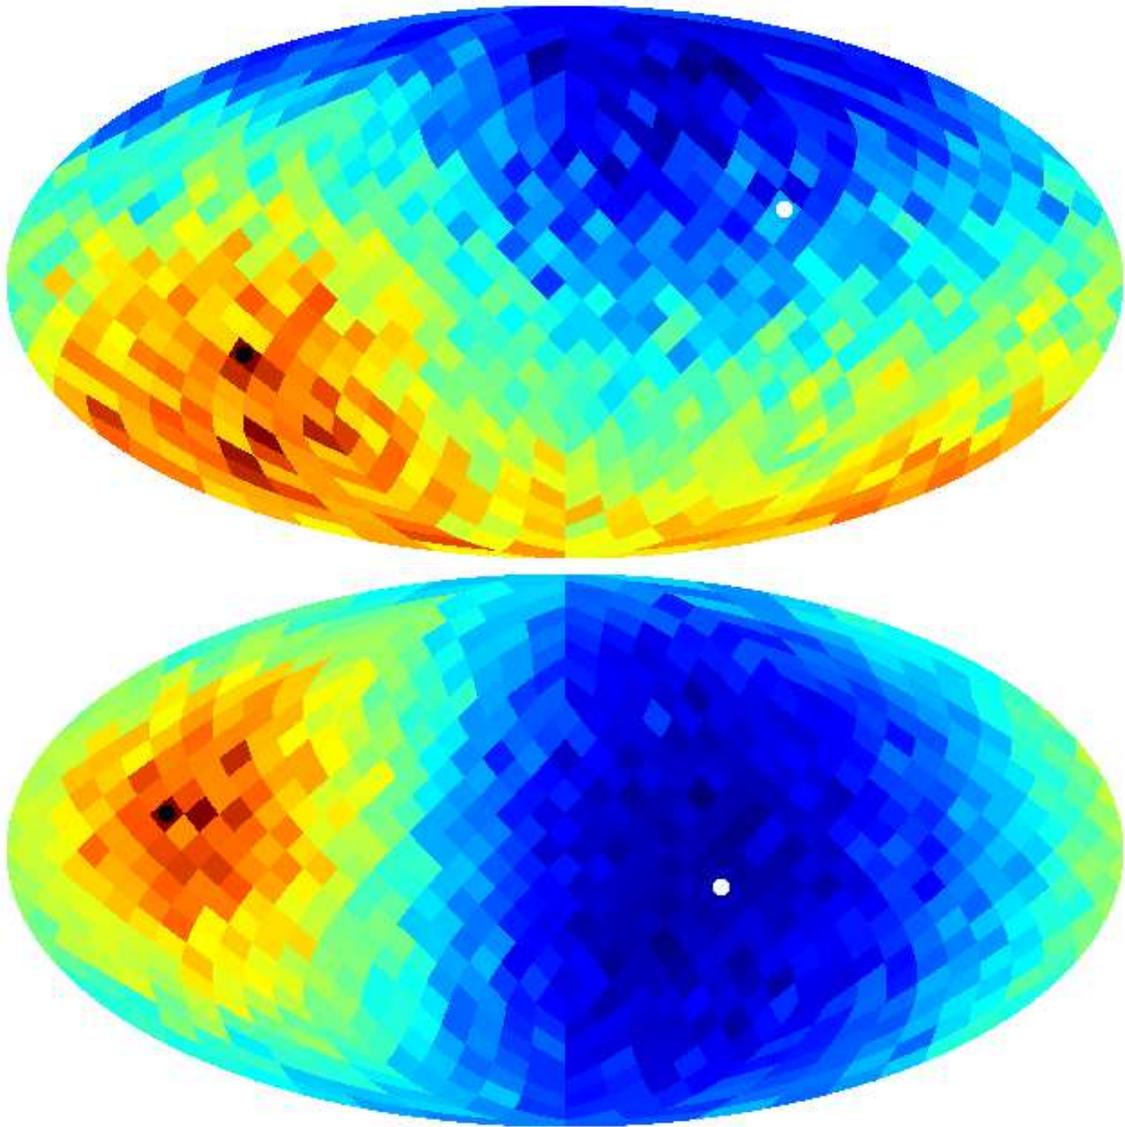


Figure 5.14: Mollweide projection of the Poincaré sphere near the phases 0.574 (top) and 0.684 (bottom). In the top panel, the polarisation vector is strongly circularly polarised and occupies a very large region near the southern pole of the Poincaré sphere. The black and white dots correspond to the dominant and the secondary polarised mode and its orthogonal counterpart.

achieved a significant improvement in the timing precision, we demonstrated that the rms timing residual and the statistical goodness of timing model fit can be improved by rejecting the brightest sub-pulses and demonstrated a correlation between the average ToA and the S/N of the sub-pulses. Where possible, we have drawn conclusions about the physics of PSR J0437–4715. Most notably, we were able to comment on the alignment of the pulsar’s angular momentum vector with respect to its magnetic flux vector; coherent curvature emission as a plausible mechanism generating the brightest sub-pulses; and direct detection of the orthogonally polarised modes based on the phase-resolved histograms of single-pulse polarisation.

6

Conclusions

This thesis is concerned with the highest precision pulsar timing of the brightest MSP known, PSR J0437–4715. By pushing the limits of precision timing and addressing the issues that arise when studying this particular pulsar, it lays the ground work for future generations of telescopes which are likely to encounter similar issues for all the other pulsars whose peak flux will be comparable to the system equivalent flux densities of next generation of instruments. Currently, PSR J0437–4715 is one of the few windows that allow us to see into the future of pulsar timing. This thesis therefore contributes indirectly to all the areas of science outlined briefly in chapters 1 and 2. We begin this final chapter (section 6.1) with a concise summary of the state-of-the-art precision timing techniques before this thesis commenced. The contributions of this thesis are described in section 6.2 along with their impact on modern pulsar astrophysics. Future directions of work stemming from this thesis are then discussed in section 6.3.

6.1 Pulsar timing before this thesis

Here we present the best precision timing techniques commonly in use before this thesis commenced. As was described in more detail in section 1.2.4, pulsar timing methodology requires deriving a physical model of the pulsar, including its spin, astrometric and binary parameters. Such models are dependent upon models of the Solar System ephemeris, Earth’s rotation, the state of the intervening ISM and many other factors. These models are initially constructed from, and subsequently compared to, the observed mean ToAs of the radio pulse train. Taylor (1992) described the most commonly used ToA estimation technique, which makes use of the total intensity pulse profile and determines the phase shift in the Fourier domain. This technique is part of the PSRCHIVE software suite.

Foster et al. (1991) developed the basis for describing pulse profiles as composed of multiple Gaussian components to enable modelling of pulse profile evolution with frequency. Later, Kramer et al. (1994) and Kramer (1994) extended and formalised this methodology to constrain the emission geometry of pulsars. After discovering apparent pulse profile variations in observations of PSR J1022+1001, a modification of this technique was used to allow variations of the components at one radio frequency (Kramer et al., 1999; Stairs et al., 2000a)

It has been argued that the pulse profile variations observed in PSR J1022+1001 are likely to be of polarimetric origin (Hotan et al., 2004a; van Straten, 2013). Earlier, Sandhu et al. (1997) and Britton (2000) realised that lack or incorrect polarimetric calibration can introduce systematic error in ToA estimation. The latter author suggested use of the invariant interval to mitigate this problem. This solution is not free of its own problems as the invariant interval pulse profiles suffer from lower S/N, going to zero for fully polarised sources; are biased by the variance of the noise in the data; and their intensities are not normally distributed. To avoid these problems, polarimetric calibration techniques have been continually improved over the last decade (van Straten, 2004, 2006, 2013). The second paper in this series also provided a template matching technique that exploits all four Stokes parameters to achieve higher precision and accuracy of arrival time estimates, as suggested by Kramer (1996) who noted that the polarised pulse profiles typically have sharper features than the total intensity.

Other efforts before the commencement of this thesis related to ToA derivation include the work by Hotan et al. (2005). These authors improved the traditional time domain algorithm by using a Gaussian interpolation function, which is necessary for sub-phase bin time resolution, instead of interpolation with a three point parabola. We note that an optimal interpolation function is a sinc function and thus the Fourier domain implementation is superior to time-domain approach with the aforementioned interpolation functions. Demorest (2007) addressed the question of deriving the optimal template profile and suggested using a noise-free template formed from data with a cut-off in the number of Fourier harmonics, similarly to the matrix template matching derived by van Straten (2006). Other widely used techniques include forming the template from a separate dataset, as used throughout this thesis, or forming an analytic template (Lommen, 2001).

After arrival time estimation using any of the above methods, timing residuals were calculated. If the goodness of fit of data to the timing model, as measured by $\chi^2/\text{d.o.f.}$, exceeded unity, sometimes the ToA estimation uncertainties were artificially multiplied.

This was done in one of two ways: the ToA uncertainties were all multiplied by the same factor, so called EFAC (error factor); or by adding noise in quadrature (EQUAD) to improve the goodness of fit (see e.g., Verbiest et al., 2009). The additional uncertainty was often attributed to instrumental effects, using templates derived from observations made with a different backend or other unknown effects.

6.2 Summary of the work accomplished in this thesis

As the precision attained in the pulsar timing experiments increased with time, it became clear that further progress was being inhibited by some new unknown effects, whose contributions were harder to previously detect and could have been ignored. Although it was initially predicted that the Parkes Pulsar Timing Array project would detect the stochastic gravitational wave background as early as 2010, only upper limits have been provided (most recently by Jenet et al., 2006) and no detection has been made (Yardley et al., 2011), along with many other scientific results (see Manchester et al., 2013, for a recent summary).

It became clear that in order to achieve the required sensitivity (Jenet et al., 2005), more effort needed to be placed on resolving the remaining issues related to pulsar timing. In particular, a natural improvement was to stop introducing arbitrary additional uncertainties but instead try to understand the origin of the discrepancies. PSR J0437–4715 provided a natural candidate for studies of what affects the attained timing precision and why it is worse than expected. We note that simultaneously other issues affecting the long term pulsar timing precision are being addressed by other members of the Parkes Pulsar Timing Array (e.g., Keith et al., 2012, addressed the variations of dispersion measure and their effect on timing precision).

Several causes of worse than expected timing precision are plausible, including, but not limited to, systematic instrumental effects, impact of the interstellar medium, data processing software deficiencies, and processes intrinsic to the pulsar. This thesis has been primarily concerned with the latter possibility. In chapters 3 and 4 we demonstrated that a large fraction of the additional scatter in the timing residuals on short time-scales can be attributed to stochastic wideband impulse modulated self-noise. In the same chapters, we provided two new methodologies that allow us to reduce the rms timing residual by 20 and nearly 40%, respectively. We demonstrated that SWIMS is not yet an important factor for majority of the pulsar population, as the contribution of the pulsar flux to the noise balance can be safely neglected for most pulsars. As the next generation of telescopes becomes available, the conceptual understanding of the encountered difficulties as well

as the devised methodology will be increasingly important. The importance of stochastic wideband impulse modulated self-noise (SWIMS) for future projects has also been realised by other groups (Cordes & Downs, 1985; Cordes, 1993; D’Alessandro et al., 1993; Cordes et al., 2005; Cordes & Shannon, 2010; Liu et al., 2011, 2012; Shannon & Cordes, 2012), however the work presented in this thesis is the first to not only clearly demonstrate the presence and wideband nature of SWIMS in millisecond pulsar observations but also to develop techniques of mitigating its deleterious effects by extending the ideas of Cordes (1993) and Demorest (2007). These techniques were applied to achieve a 40 per cent reduction of rms timing residual.

Even with the improved methods of dealing with SWIMS presented in chapter 4, the goodness of fit of the timing residuals is still not satisfactory. In chapters 3 and 4 we argue that this is, to some extent, due to the template matching algorithms making false assumptions about homoscedastic and uncorrelated nature of the noise in pulse profiles. As outlined in section 3.6.1, it is possible to derive a generalised template matching algorithm which will enable realistic estimates of the ToA uncertainties by taking the heteroscedasticity and correlations of noise into account. At the moment we cannot exclude the presence of smaller magnitude effects that are also contributing to the error budget. For example, it is likely that a number of instrumental effects are still present in the data, see Fig. 3.9 and 5.7, and e.g., apparent over-polarisation of the Pulsar Digital Filterbank generations 3 and 4 (not discussed in this thesis).

Finally, while pursuing the goal of highest precision timing, we analysed a plethora of single-pulse data in chapter 5. Pulsars are unusual objects in astrophysics in the sense that we are able to obtain over-abundant data and many seemingly disconnected facts about these stars. Single-pulse data are the best example of this, where the variety of observed phenomena is overwhelming, as is the sheer volume of the data. We attempted to make as many connections to various models present in the pulsar literature as possible. The highlights of this chapter were: selective rejection of bright sub-pulses allows us to achieve as good a timing precision as all the data, as expected from SWIMS properties. In one-minute sub-integrations, a single pulse can bias the ToA estimate by 0.11 ns on average. Brighter pulses tend to arrive earlier than their dimmer counterparts. We found support for the stochastic growth theory as the correct description of plasma behaviour in the magnetosphere over certain regions of pulse longitude. Lack of stationary and modulated intensity modulation refuted some models of PSR J0437–4715’s emission and constrains the relative orientation of its spin and magnetic flux. A first direct detection of orthogonally polarised modes (OPMs) in an MSP was presented. Brighter pulses were

found to be more polarised, with half of them showing characteristic features of coherent curvature radiation from solitons.

6.3 Future directions

As it is often the case, many of the problems solved in this thesis provoke further questions. While we managed to achieve a 40% reduction in the rms timing residual in chapter 4, we were unable to apply the same correction scheme to data gathered with different integration times. Furthermore, we studied only short time-scales, but many of the MSPs show long term timing noise that limits their usefulness for pulsar timing arrays and other timing experiments.

Fig. 4.8 demonstrates that there is a need to take the ionosphere into account, when performing the highest precision polarimetric calibration. Once this is included, application of the correction scheme based on one time-scale should be applicable to separate datasets with similar or different integration times. After removing the short term variations a better measurement of the low frequency timing noise can be obtained and addressed in subsequent work. An interesting question is whether the covariance matrix measured in equation 3.1 is constant for a given pulsar or whether the emission statistics change in time. In chapter 3 we demonstrated its stability over a few months, however it may become unstable on longer time-scales, thus contributing to the low frequency timing noise.

Recently, Shannon & Cordes (2012) showed that PSR J1713+0747's timing is also limited by SWIMS, however they were unable to remove the effects of the additional noise from the ToA estimates. While more investigation is necessary, it is possible that this is because of the template profile of this pulsar. Namely, any pulse profile variability that is correlated with the template derivative is not accessible by the method attempted by the authors, which is fundamentally similar to the technique described in chapter 3. We can possibly circumvent this problem by extending their analysis and studying the polarised pulse profiles, as in chapter 4.

Before analysing long term timing noise, it would be worthwhile to finalise the generalised template matching algorithm. The realistic ToA uncertainties will help answer the question if more processes are affecting the short term timing and correctly estimate the magnitude of long term variations.

More work can be undertaken with the baseband data gathered for the purposes of chapter 5. An example of a potential study is analysing the 4-th order moments of the electric field at nanosecond timescales (e.g., Hankins, 1972; Cordes & Hankins, 1973; van

Straten, 2009), a capability of the DSPSR suite currently under development. Many interesting polarisation properties have been revealed by analysing the single pulses and further study is warranted. An interesting possibility would be a study of the polarisation at much higher time resolutions which could give further insight into the pulsar magnetosphere.

Finally, it would be folly to take a solitary pulsar (PSR J0437–4715) and expect all pulsars to precisely mimic its emission properties and precision timing performance. As larger telescopes come online, the techniques developed in this thesis for PSR J0437–4715 should be applied to define the ultimate precision each pulsar can attain. In turn, this will enable realistic predictions about the science that can be achieved with pulsar timing array experiments and the potential for direct detection of gravitational waves.

Bibliography

- Abbott B. & the LIGO collaboration *ApJ* 659:918–930 April 2007.
- Ables J. G. & Manchester R. N. *A&A* 50:177–184 July 1976.
- Ables J. G. McConnell D. Deshpande A. A. & Vivekanand M. *ApJ* 475:L33 January 1997.
- Abramovici A. Althouse W. E. Drever R. W. P. Gursel Y. Kawamura S. Raab F. J. Shoemaker D. Sievers L. Spero R. E. & Thorne K. S. *Science* 256:325–333 April 1992.
- Adams W. S. *PASP* 27:236 December 1915.
- Adams W. S. *Proceedings of the National Academy of Science* 11:382–387 July 1925.
- Aharonian F. A. Hofmann W. Konopelko A. K. & Völk H. J. *Astroparticle Physics* 6: 343–368 March 1997.
- Allen B. & Romano J. D. *Phys. Rev. D* 59(10):102001 May 1999.
- Alvarez L. W. *Physical Review* 52:134–135 July 1937.
- Andrews D. *Robust Estimates of Location: Survey and Advances*. Princeton University Press Princeton 1972. ISBN 0691081131.
- Ardavan H. Ardavan A. Singleton J. & Perez M. R. *MNRAS* 388:873–883 August 2008.
- Arnison G. et al. (UA1 collaboration) *Physics Letters B* 122:103–116 February 1983a.
- Arnison G. et al. (UA1 collaboration) *Physics Letters B* 126:398–410 July 1983b.
- Arons J. & Barnard J. J. *ApJ* 302:120–137 March 1986.
- Arzoumanian Z. *Radio Observations of Binary Pulsars: Clues to Binary Evolution and Tests of General Relativity*. PhD thesis Princeton University 1995.
- Asseo E. *Plasma Instabilities : Sources for Coherent Radio Emission*. In Johnston S. Walker M. A. & Bailes M. editors *IAU Colloq. 160: Pulsars: Problems and Progress* volume 105 of *Astronomical Society of the Pacific Conference Series* page 147 1996.
- Asseo E. & Porzio A. *MNRAS* 369:1469–1490 July 2006.
- Baade W. & Zwicky F. *Proceedings of the National Academy of Science* 20:254–259 May 1934a.

- Baade W. & Zwicky F. *Proceedings of the National Academy of Science* 20:259–263 May 1934b.
- Baade W. & Zwicky F. *Phys. Rev.* 45(2):138 Jan 1934c.
- Baars J. W. M. & Swenson G. *Physics Today* 61(7):070000 2008.
- Baars J. W. M. Genzel R. Pauliny-Toth I. I. K. & Witzel A. *A&A* 61:99–106 October 1977.
- Backer D. C. *Nature* 227:692–695 August 1970a.
- Backer D. C. *Nature* 228:42–43 October 1970b.
- Backer D. C. *Nature* 228:1297–1298 December 1970c.
- Backer D. C. *ApJ* 182:245–276 May 1973.
- Backer D. C. & Rankin J. M. *ApJS* 42:143–173 January 1980.
- Backer D. C. & Sallmen S. T. *AJ* 114:1539 October 1997.
- Backer D. C. Rankin J. M. & Campbell D. B. *ApJ* 197:481–487 April 1975.
- Backer D. C. Rankin J. M. & Campbell D. B. *Nature* 263:202–207 September 1976.
- Backer D. C. Kulkarni S. R. Heiles C. Davis M. M. & Goss W. M. *Nature* 300:615–618 December 1982.
- Bagnaia P. et al. (UA2 collaboration) *Physics Letters B* 129:130–140 September 1983.
- Bailes M. *Precision Timing at the Parkes 64-m Radio Telescope*. In M. Bailes, D. J. Nice, & S. E. Thorsett editor *Radio Pulsars* volume 302 of *Astronomical Society of the Pacific Conference Series* page 57 2003.
- Bailes M. *The art of precision pulsar timing*. In S. A. Klioner, P. K. Seidelmann, & M. H. Soffel editor *IAU Symposium* volume 261 of *IAU Symposium* pages 212–217 January 2010. doi: 10.1017/S1743921309990421.
- Bailes M. Bates S. D. Bhalerao V. Bhat N. D. R. Burgay M. Burke-Spolaor S. D’Amico N. Johnston S. Keith M. J. Kramer M. Kulkarni S. R. Levin L. Lyne A. G. Milia S. Possenti A. Spitler L. Stappers B. & van Straten W. *Science* 333:1717– September 2011.
- Bak P. *How nature works : the science of self-organized criticality*. 1996.

- Bak P. Tang C. & Wiesenfeld K. *Phys. Rev. A* 38:364–374 July 1988.
- Banner M. et al. (UA2 collaboration) *Physics Letters B* 122:476–485 March 1983.
- Barker B. M. & O’Connell R. F. *Phys. Rev. D* 12:329–335 July 1975.
- Bartel N. Morris D. Sieber W. & Hankins T. H. *ApJ* 258:776–789 July 1982.
- Bartel N. Chandler J. F. Ratner M. I. Shapiro I. L. Pan R. & Cappallo R. J. *AJ* 112:1690
October 1996.
- Baym G. Pethick C. & Sutherland P. *ApJ* 170:299–+ December 1971.
- Bekenstein J. D. *Phys. Rev. D* 12:3077–3085 November 1975.
- Belczynski K. Kalogera V. Rasio F. A. Taam R. E. Zezas A. Bulik T. Maccarone T. J. &
Ivanova N. *ApJS* 174:223–260 January 2008.
- Beringer J. on behalf of Particle Data Group. *The Review of Particle Physics* volume 86.
August 2012.
- Bertotti B. Iess L. & Tortora P. *Nature* 425:374–376 September 2003.
- Beskin V. S. *Soviet Physics Uspekhi* 42:1071–1074 November 1999.
- Bessarab A. V. Martynenko S. P. Prudkoi N. A. Soldatov A. V. & Terekhin V. A. *Radiation
Physics and Chemistry* 75:825–831 August 2006.
- Bhat N. D. R. Cordes J. M. Camilo F. Nice D. J. & Lorimer D. R. *ApJ* 605:759–783 April
2004.
- Bhat N. D. R. Gupta Y. Kramer M. Karastergiou A. Lyne A. G. & Johnston S. *A&A* 462:
257–268 January 2007.
- Bhat N. D. R. Bailes M. & Verbiest J. P. W. *Phys. Rev. D* 77(12):124017 June 2008.
- Bhattacharya D. *Detection of Radio Emission from Pulsars: a Pulsar Observation Primer*.
In Bucccheri R. van Paradijs J. & Alpar A. editors *NATO ASIC Proc. 515: The Many
Faces of Neutron Stars*. page 103 1998.
- Blandford R. & Teukolsky S. A. *ApJ* 205:580–591 April 1976.
- Boesgaard A. M. & Steigman G. *ARA&A* 23:319–378 1985.

- Born M. & Wolf E. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. 1999.
- Boynton P. E. Groth III E. J. Partridge R. B. & Wilkinson D. T. *ApJ* 157:L197+ September 1969.
- Boynton P. E. Groth E. J. Hutchinson D. P. Nanos Jr. G. P. Partridge R. B. & Wilkinson D. T. *ApJ* 175:217+ July 1972.
- Bracewell R. N. *The Fourier transform and its applications*. 2000.
- Bradaschia C. Calloni E. Cobal M. Del Fabbro R. di Virgilio A. Giazotto A. Holloway L. E. Kautzky H. Michelozzi B. Montelatici V. Passuello D. & Velloso W. *VIRGO: a ground based interferometric antenna for gravitational wave detection above 10 Hz*. In *Problems of Fundamental Modern Physics II* pages 341–356 1991.
- Brans C. & Dicke R. H. *Physical Review* 124:925–935 November 1961.
- Britton M. C. *ApJ* 532:1240–1244 April 2000.
- Burgay M. D’Amico N. Possenti A. Manchester R. N. Lyne A. G. Joshi B. C. McLaughlin M. A. Kramer M. Sarkissian J. M. Camilo F. Kalogera V. Kim C. & Lorimer D. R. *Nature* 426:531–533 December 2003.
- Burke-Spolaor S. Bailes M. Johnston S. Bates S. D. Bhat N. D. R. Burgay M. D’Amico N. Jameson A. Keith M. J. Kramer M. Levin L. Milia S. Possenti A. Stappers B. & van Straten W. *MNRAS* 416:2465–2476 October 2011.
- Burke-Spolaor S. Johnston S. Bailes M. Bates S. D. Bhat N. D. R. Burgay M. Champion D. J. D’Amico N. Keith M. J. Kramer M. Levin L. Milia S. Possenti A. Stappers B. & van Straten W. *MNRAS* 423:1351–1367 June 2012.
- Cairns I. H. *ApJ* 610:948–955 August 2004.
- Cairns I. H. & Menietti J. D. *J. Geophys. Res.* 106:29515–29530 December 2001.
- Cairns I. H. & Robinson P. A. *Roles Played by Electrostatic Waves in Producing Radio Emissions*. In Stone R. G. Weiler K. W. Goldstein M. L. & Bougeret J.-L. editors *Radio Astronomy at Long Wavelengths* page 27 2000.
- Cairns I. H. Johnston S. & Das P. *ApJ* 563:L65–L68 December 2001.
- Cairns I. H. Robinson P. A. & Das P. *Phys. Rev. E* 66(6):066614 December 2002.

- Cairns I. H. Das P. Robinson P. A. & Johnston S. *MNRAS* 343:523–532 August 2003a.
- Cairns I. H. Johnston S. & Das P. *MNRAS* 343:512–522 August 2003b.
- Cairns I. H. Johnston S. & Das P. *MNRAS* 353:270–286 September 2004.
- Castillo G. *Differentiable Manifolds: A Theoretical Physics Approach*. Birkhäuser Boston 2011. ISBN 9780817682705.
- Champion D. J. Hobbs G. B. Manchester R. N. & +20 other authors *ApJ* 720:L201–L205 September 2010.
- Chandrasekhar S. *ApJ* 74:81 July 1931.
- Clemence G. M. *AJ* 50:126–127 July 1943.
- Cognard I. Shrauner J. A. Taylor J. H. & Thorsett S. E. *ApJ* 457:L81 February 1996.
- Cole T. W. *Nature* 227:788–791 August 1970.
- Coles W. Hobbs G. Champion D. J. Manchester R. N. & Verbiest J. P. W. *MNRAS*, *accepted* August 2011.
- Collins J. C. & Perry M. J. *Physical Review Letters* 34:1353–1356 May 1975.
- Comella J. M. Craft H. D. Lovelace R. V. E. & Sutton J. M. *Nature* 221:453–454 February 1969.
- Cordes J. M. *ApJ* 237:216–226 April 1980.
- Cordes J. M. *The detectability of planetary companions to radio pulsars*. In J. A. Phillips, S. E. Thorsett, & S. R. Kulkarni editor *Planets Around Pulsars* volume 36 of *Astronomical Society of the Pacific Conference Series* pages 43–60 January 1993.
- Cordes J. M. & Chernoff D. F. *ApJ* 482:971 June 1997.
- Cordes J. M. & Downs G. S. *ApJS* 59:343–382 November 1985.
- Cordes J. M. & Hankins T. H. *Subpulse Time Scales in Several Pulsars*. In *Bulletin of the American Astronomical Society* volume 5 of *Bulletin of the American Astronomical Society* page 18 January 1973.
- Cordes J. M. & Helfand D. J. *ApJ* 239:640–650 July 1980.
- Cordes J. M. & Lazio T. J. W. *ApJ* 475:557 February 1997.

- Cordes J. M. & Shannon R. M. *ArXiv e-prints* October 2010.
- Cordes J. M. Rankin J. & Backer D. C. *ApJ* 223:961–972 August 1978.
- Cordes J. M. Wolszczan A. Dewey R. J. Blaskiewicz M. & Stinebring D. R. *ApJ* 349: 245–261 January 1990.
- Cordes J. M. Kramer M. Backer D. C. Lazio T. J. W. & Science Working Group for the Square Kilometer Array Team. *Key Science with the Square Kilometer Array: Strong-field Tests of Gravity using Pulsars and Black Holes*. In *American Astronomical Society Meeting Abstracts* volume 37 of *Bulletin of the American Astronomical Society* page 137.04 December 2005.
- Cui W. *ApJ* 482:L163+ June 1997.
- Cumming A. Zweibel E. & Bildsten L. *ApJ* 557:958–966 August 2001.
- D’Alessandro F. McCulloch P. M. King E. A. Hamilton P. A. & McConnell D. *MNRAS* 261:883 April 1993.
- Damour T. & Deruelle N. *Ann. Inst. Henri Poincaré Phys. Théor., Vol. 43, No. 1, p. 107 - 132* 43:107–132 1985.
- Damour T. & Deruelle N. *Ann. Inst. Henri Poincaré Phys. Théor., Vol. 44, No. 3, p. 263 - 292* 44:263–292 1986.
- Damour T. & Ruffini R. *Academie des Sciences Paris Comptes Rendus Serie Sciences Mathematiques* 279:971–973 December 1974.
- Damour T. & Taylor J. H. *Phys. Rev. D* 45:1840–1868 March 1992.
- D’Angelo C. R. & Spruit H. C. *MNRAS* 406:1208–1219 August 2010.
- Davis M. Efstathiou G. Frenk C. S. & White S. D. M. *ApJ* 292:371–394 May 1985.
- de Broglie L. *Recherches sur la théorie des quanta*. PhD thesis Paris University 1924.
- Demorest P. B. *Measuring the gravitational wave background using precision pulsar timing*. PhD thesis University of California, Berkeley August 2007.
- Demorest P. B. *MNRAS* 416:2821–2826 October 2011.
- Demorest P. B. Pennucci T. Ransom S. M. Roberts M. S. E. & Hessels J. W. T. *Nature* 467:1081–1083 October 2010.

- Demorest P. B. Ferdman R. D. Gonzalez M. E. Nice D. Ransom S. Stairs I. H. Arzoumanian Z. Brazier A. Burke-Spolaor S. Chamberlin S. J. Cordes J. M. Ellis J. Finn L. S. Freire P. Giampanis S. Jenet F. Kaspi V. M. Lazio J. Lommen A. N. McLaughlin M. Palliyaguru N. Perrodin D. Shannon R. M. Siemens X. Stinebring D. Swiggum J. & Zhu W. W. *ArXiv e-prints* January 2012.
- Detweiler S. *ApJ* 234:1100–1104 December 1979.
- Dewey R. J. Taylor J. H. Weisberg J. M. & Stokes G. H. *ApJ* 294:L25–L29 July 1985.
- Dirac P. A. M. *Royal Society of London Proceedings Series A* 114:243–265 March 1927.
- Dirac P. A. M. *Proc. Roy. Soc.* 113:432 1926.
- Dowd A. Sisk W. & Hagen J. *WAPP — Wideband Arecibo Pulsar Processor*. In Kramer M. Wex N. & Wielebinski R. editors *IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond* volume 202 of *Astronomical Society of the Pacific Conference Series* page 275 2000.
- Downs G. S. & Reichley P. E. *ApJS* 53:169–240 September 1983.
- Drake F. D. & Craft H. D. *Nature* 220:231–235 October 1968a.
- Drake F. D. & Craft Jr. H. D. *Science* 160:758–760 May 1968b.
- Duncan R. C. & Thompson C. *ApJ* 392:L9–L13 June 1992.
- Durbin J. M. Large M. I. Little A. G. Manchester R. N. Lyne A. G. & Taylor J. H. *MNRAS* 186:39P–41P March 1979.
- Dyks J. Rudak B. & Rankin J. M. *A&A* 465:981–991 April 2007.
- Dyks J. Rudak B. & Demorest P. *MNRAS* 401:1781–1795 January 2010.
- Dyson F. W. Eddington A. S. & Davidson C. *Royal Society of London Philosophical Transactions Series A* 220:291–333 1920.
- Eddington A. S. *The Internal Constitution of the Stars*. 1926.
- Edwards R. T. & Stappers B. W. *A&A* 393:733–748 October 2002.
- Edwards R. T. & Stappers B. W. *A&A* 407:273–287 August 2003.
- Edwards R. T. & Stappers B. W. *A&A* 421:681–691 July 2004.

- Edwards R. T. Hobbs G. B. & Manchester R. N. *MNRAS* 372:1549–1574 November 2006.
- Einstein A. *Annalen der Physik* 322:132–148 1905.
- Einstein A. *Annalen der Physik* 354:769–822 1916.
- Einstein A. & Fokker A. D. *Annalen der Physik* 349:321–328 1914.
- Englert F. & Brout R. *Physical Review Letters* 13:321–323 August 1964.
- Erdmann K. & Wildon M. *Introduction to Lie Algebras*. Springer Undergraduate Mathematics Series. Springer 2007. ISBN 9781846280405.
- Everett J. E. & Weisberg J. M. *ApJ* 553:341–357 May 2001.
- Everitt C. W. F. Debra D. B. Parkinson B. W. Turneure J. P. Conklin J. W. Heifetz M. I. Keiser G. M. Silbergleit A. S. Holmes T. Kolodziejczak J. Al-Meshari M. Mester J. C. Muhlfelder B. Solomonik V. G. Stahl K. Worden Jr. P. W. Bencze W. Buchman S. Clarke B. Al-Jadaan A. Al-Jibreen H. Li J. Lipa J. A. Lockhart J. M. Al-Suwaidan B. Taber M. & Wang S. *Physical Review Letters* 106(22):221101 June 2011.
- Faraday M. *Faraday's diary* volume 4. September 1845.
- Faucher-Giguère C.-A. & Kaspi V. M. *ApJ* 643:332–355 May 2006.
- Feliachi R. . PhD thesis University of Orléans, France 2010.
- Ferdman R. D. van Haasteren R. Bassa C. G. Burgay M. Cognard I. Corongiu A. D'Amico N. Desvignes G. Hessels J. W. T. Janssen G. H. Jessner A. Jordan C. Karuppusamy R. Keane E. F. Kramer M. Lazaridis K. Levin Y. Lyne A. G. Pilia M. Possenti A. Purver M. Stappers B. Sanidas S. Smits R. & Theureau G. *Classical and Quantum Gravity* 27 (8):084014 April 2010.
- Fermi E. *Reviews of Modern Physics* 4:87–132 January 1932.
- Fermi E. *Rend. Acc. Lincei* 3:145–149 1926.
- Foster R. S. & Backer D. C. *ApJ* 361:300–308 September 1990.
- Foster R. S. Fairhead L. & Backer D. C. *ApJ* 378:687–695 September 1991.
- Fourier M. *Théorie analytique de la chaleur*. Chez Firmin Didot, Père et fils 1822.
- Fowler R. H. *MNRAS* 87:114–122 December 1926.

- Freire P. C. Kramer M. Lyne A. G. Camilo F. Manchester R. N. & D'Amico N. *ApJ* 557: L105–L108 August 2001.
- Freire P. C. C. Wex N. Esposito-Farèse G. Verbiest J. P. W. Bailes M. Jacoby B. A. Kramer M. Stairs I. H. Antoniadis J. & Janssen G. H. *MNRAS* 423:3328–3343 July 2012.
- Gabor D. *J. Inst. Electr. Engrs.* 93:429–457 1946.
- Gaensler B. M. & Slane P. O. *ARA&A* 44:17–47 September 2006.
- Gardner W. & Cadzow J. A. *Appl. Opt.* 29:1399 April 1990.
- Garner T. W. Gaussiran II T. L. Tolman B. W. Harris R. B. Calfas R. S. & Gallagher H. *Advances in Space Research* 42:720–726 August 2008.
- Gedalin M. Gruman E. & Melrose D. B. *MNRAS* 337:422–430 December 2002.
- Georgi H. & Glashow S. L. *Physical Review Letters* 32:438–441 February 1974.
- Geppert U. & Urpin V. *MNRAS* 271:490–+ November 1994.
- Ghosh P. & Lamb F. K. *ApJ* 232:259–276 August 1979a.
- Ghosh P. & Lamb F. K. *ApJ* 234:296–316 November 1979b.
- Ghosh P. Pethick C. J. & Lamb F. K. *ApJ* 217:578–596 October 1977.
- Gil J. & Krawczyk A. *MNRAS* 285:561–566 March 1997.
- Glashow S. *Nucl. Phys. A* 22:579–588 February 1961.
- Glendenning N. K. *ApJ* 293:470–493 June 1985.
- Gold T. *Nature* 218:731–732 May 1968.
- Goldreich P. & Julian W. H. *ApJ* 157:869–+ August 1969.
- Gonthier P. L. Ouellette M. S. Berrier J. O'Brien S. & Harding A. K. *ApJ* 565:482–499 January 2002.
- Gonthier P. L. Van Guilder R. & Harding A. K. *ApJ* 604:775–790 April 2004.
- Gonthier P. L. Nagelkirk E. Stam M. & Harding A. K. *ArXiv e-prints* December 2009.

- Gonthier P. L. Billman C. Harding A. K. Grenier I. A. & Pierbattista M. *Population Synthesis of Radio and Gamma-ray Pulsars in the Fermi Era*. In Burgay M. D'Amico N. Esposito P. Pellizzoni A. & Possenti A. editors *American Institute of Physics Conference Series* volume 1357 of *American Institute of Physics Conference Series* pages 245–248 August 2011. doi: 10.1063/1.3615126.
- Górski K. M. Hivon E. Banday A. J. Wandelt B. D. Hansen F. K. Reinecke M. & Bartelmann M. *ApJ* 622:759–771 April 2005.
- Granet C. Zhang H. Z. Forsyth A. R. Graves G. R. Doherty P. Greene K. J. James G. L. Sykes P. Bird T. S. Sinclair M. W. Moorey G. & Manchester R. N. *IEEE Antennas Propagation Magazine* 47:13–19 June 2005.
- Green D. A. *Bulletin of the Astronomical Society of India* 39:289–295 June 2011.
- Greenstein J. L. Oke J. B. & Shipman H. L. *ApJ* 169:563 November 1971.
- Greiner W. Schramm S. Stein E. & Bromley D. *Quantum Chromodynamics*. Springer 2007. ISBN 9783540485346.
- Groth E. J. *ApJS* 29:431–442 November 1975a.
- Groth E. J. *ApJS* 29:443–451 November 1975b.
- Groth E. J. *ApJS* 29:453–465 November 1975c.
- Guillemot L. Freire P. C. C. Cognard I. Johnson T. J. Takahashi Y. Kataoka J. Desvignes G. Camilo F. Ferrara E. C. Harding A. K. Janssen G. H. Keith M. Kerr M. Kramer M. Parent D. Ransom S. M. Ray P. S. Saz Parkinson P. M. Smith D. A. Stappers B. W. & Theureau G. *MNRAS* 422:1294–1305 May 2012.
- Guralnik G. S. Hagen C. R. & Kibble T. W. *Physical Review Letters* 13:585–587 November 1964.
- Gwinn C. *PASP* 116:84–97 January 2004.
- Gwinn C. R. *ApJ* 554:1197–1209 June 2001.
- Gwinn C. R. *PASP* 118:461–477 March 2006.
- Gwinn C. R. & Johnson M. D. *ApJ* 733:51–+ May 2011.
- Gwinn C. R. Hirano C. & Boldyrev S. *A&A* 453:595–599 July 2006.

- Gwinn C. R. Johnson M. D. Smirnova T. V. & Stinebring D. R. *ApJ* 733:52–+ May 2011.
- Halzen F. & Klein S. R. *Review of Scientific Instruments* 81(8):081101 August 2010.
- Hamaker J. P. *A&AS* 143:515–534 May 2000.
- Han J. *Progress in Astronomy* 19:201–204 June 2001.
- Hankins T. H. *ApJ* 169:487 November 1971.
- Hankins T. H. *ApJ* 177:L11 October 1972.
- Hankins T. H. & Eilek J. A. *ApJ* 670:693–701 November 2007.
- Hankins T. H. & Rickett B. J. *Methods in Computational Physics* 14:55–129 1975.
- Hankins T. H. Moffett D. A. Novikov A. & Popov M. *ApJ* 417:735 November 1993.
- Hankins T. H. Kern J. S. Weatherall J. C. & Eilek J. A. *Nature* 422:141–143 March 2003.
- Hasert F. J. Kabe S. Krenz W. von Krogh J. Lanske D. Morfin J. Schultze K. Weerts H. Bertrand-Coremans G. H. Sacton J. van Doninck W. Vilain P. Camerini U. Cundy D. C. Baldi R. Danilchenko I. Fry W. F. Haidt D. Natali S. Musset P. Osculati B. Palmer R. Pattison J. B. M. Perkins D. H. Pullia A. Rousset A. Venus W. Wachsmuth H. Brisson V. Degrange B. Haguenaue M. Kluberg L. Nguyen-Khac U. Petiau P. Belotti E. Bonetti S. Cavalli D. Conta C. Fiorini E. Rollier M. Aubert B. Blum D. Chounet L. M. Heusse P. Lagarrigue A. Lutz A. M. Orkin-Lecourtois A. Vialle J. P. Bullock F. W. Esten M. J. Jones T. W. McKenzie J. Michette A. G. Myatt G. & Scott W. G. *Physics Letters B* 46:138–140 September 1973.
- Hawking S. W. *Particle creation by black holes*. In Isham C. J. Penrose R. & Sciama D. W. editors *Quantum Gravity* pages 219–267 1975.
- Hawking S. W. & Penrose R. *Royal Society of London Proceedings Series A* 314:529–548 January 1970.
- Hecht E. *Optics 4th edition*. 2001.
- Helfand D. J. Manchester R. N. & Taylor J. H. *ApJ* 198:661–670 June 1975.
- Helfand D. J. Taylor J. H. Backus P. R. & Cordes J. M. *ApJ* 237:206–215 April 1980.
- Hellings R. W. & Downs G. S. *ApJ* 265:L39–L42 February 1983.

- Hertz H. *Annalen der Physik* 272:769–783 1889.
- Hessels J. W. T. Ransom S. M. Stairs I. H. Freire P. C. C. Kaspi V. M. & Camilo F. *Science* 311:1901–1904 March 2006.
- Hetherington N. S. *QJRAS* 21:246–252 September 1980.
- Hewish A. Bell S. J. Pilkington J. D. H. Scott P. F. & Collins R. A. *Nature* 217:709–713 February 1968.
- Higgs P. W. *Physical Review Letters* 13:508–509 October 1964.
- Hoaglin D. C. Mosteller F. & Tukey J. W. editors. *Understanding robust and exploratory data analysis*. Wiley series in probability and mathematical statistics. Wiley-Interscience 1983.
- Hobbs G. Hollow R. Champion D. Khoo J. Yardley D. Carr M. Keith M. Jenet F. Amy S. Burgay M. Burke-Spolaor S. Chapman J. Danaia L. Homewood B. Kovacevic A. Mao M. McKinnon D. Mulcahy M. Osłowski S. & van Straten W. *Publ. Astron. Soc. Aust.* 26:468–475 November 2009.
- Hobbs G. Archibald A. Arzoumanian Z. Backer D. Bailes M. Bhat N. D. R. Burgay M. Burke-Spolaor S. Champion D. Cognard I. Coles W. Cordes J. Demorest P. Desvignes G. Ferdman R. D. Finn L. Freire P. Gonzalez M. Hessels J. Hotan A. Janssen G. Jenet F. Jessner A. Jordan C. Kaspi V. Kramer M. Kondratiev V. Lazio J. Lazaridis K. Lee K. J. Levin Y. Lommen A. Lorimer D. Lynch R. Lyne A. Manchester R. McLaughlin M. Nice D. Osłowski S. Pilia M. Possenti A. Purver M. Ransom S. Reynolds J. Sanidas S. Sarkissian J. Sesana A. Shannon R. Siemens X. Stairs I. Stappers B. Stinebring D. Theureau G. van Haasteren R. van Straten W. Verbiest J. P. W. Yardley D. R. B. & You X. P. *Classical and Quantum Gravity* 27(8):084013 April 2010.
- Hobbs G. Miller D. Manchester R. N. Dempsey J. Chapman J. M. Khoo J. Applegate J. Bailes M. Bhat N. D. R. Bridle R. Borg A. Brown A. Burnett C. Camilo F. Cattalini C. Chaudhary A. Chen R. D’Amico N. Kedziora-Chudczer L. Cornwell T. George R. Hampson G. Hepburn M. Jameson A. Keith M. Kelly T. Kosmynin A. Lenc E. Lorimer D. Love C. Lyne A. McIntyre V. Morrissey J. Pienaar M. Reynolds J. Ryder G. Sarkissian J. Stevenson A. Treloar A. van Straten W. Whiting M. & Wilson G. *ArXiv e-prints* May 2011a.

- Hobbs G. Miller D. Manchester R. N. Dempsey J. Chapman J. M. Khoo J. Applegate J. Bailes M. Bhat N. D. R. Bridle R. Borg A. Brown A. Burnett C. Camilo F. Cattalini C. Chaudhary A. Chen R. D'Amico N. Kedziora-Chudczer L. Cornwell T. George R. Hampson G. Hepburn M. Jameson A. Keith M. Kelly T. Kosmynin A. Lenc E. Lorimer D. Love C. Lyne A. McIntyre V. Morrissey J. Pienaar M. Reynolds J. Ryder G. Sarkissian J. Stevenson A. Treloar A. van Straten W. Whiting M. & Wilson G. *Publ. Astron. Soc. Aust.* 28:202–214 August 2011b.
- Hobbs G. Coles W. Manchester R. N. Keith M. J. Shannon R. M. Chen D. Bailes M. Bhat N. D. R. Burke-Spolaor S. Champion D. Chaudhary A. Hotan A. Khoo J. Kocz J. Levin Y. Osłowski S. Preisig B. Ravi V. Reynolds J. E. Sarkissian J. van Straten W. Verbiest J. P. W. Yardley D. & You X. P. *MNRAS* 427:2780–2787 December 2012.
- Hobbs G. B. Edwards R. T. & Manchester R. N. *MNRAS* 369:655–672 June 2006.
- Hotan A. W. *High-Precision Observations of Relativistic Binary and Millisecond Pulsars*. PhD thesis Centre for Astrophysics and Supercomputation, Swinburne University of Technology 2006.
- Hotan A. W. Bailes M. & Ord S. M. *MNRAS* 355:941–949 December 2004a.
- Hotan A. W. van Straten W. & Manchester R. N. *Publ. Astron. Soc. Aust.* 21:302–309 2004b.
- Hotan A. W. Bailes M. & Ord S. M. *MNRAS* 362:1267–1272 October 2005.
- Hotan A. W. Bailes M. & Ord S. M. *MNRAS* 369:1502–1520 July 2006.
- Hubert M. & Debruyne M. *Wiley Interdisciplinary Reviews: Computational Statistics* 2 (1):36–43 2010. ISSN 1939-0068.
- Hulse R. A. & Taylor J. H. *ApJ* 195:L51–L53 January 1975.
- Humphrey J. R. Price D. K. Spagnoli K. E. Paolini A. L. & Kelmelis E. J. *SPIE DEfense and Security Symposium* pages 770502–770502–7 2010.
- Hyvarinen A. Karhunen J. & Oja E. *Independent Component Analysis*. John Wiley and Sons, New York 2001.
- Illarionov A. F. & Sunyaev R. A. *Astron. Astrophys.* 39:185–196 1975.
- Jackson J. D. *Classical electrodynamics*. 1975.

- Jacoby B. A. Cameron P. B. Jenet F. A. Anderson S. B. Murty R. N. & Kulkarni S. R. *ApJ* 644:L113–L116 June 2006.
- Jansky K. G. *Nature* 132:66 July 1933.
- Jaschek C. & Gómez A. E. *PASP* 82:809 August 1970.
- J Jeans J. H. *Royal Society of London Philosophical Transactions Series A* 199:1–53 1902.
- Jenet F. Finn L. S. Lazio J. Lommen A. McLaughlin M. Stairs I. Stinebring D. Verbiest J. Archibald A. Arzoumanian Z. Backer D. Cordes J. Demorest P. Ferdman R. Freire P. Gonzalez M. Kaspi V. Kondratiev V. Lorimer D. Lynch R. Nice D. Ransom S. Shannon R. & Siemens X. *ArXiv e-prints* September 2009.
- Jenet F. A. & Anderson S. B. *PASP* 110:1467–1478 December 1998.
- Jenet F. A. & Gil J. *ApJ* 602:L89–L92 February 2004.
- Jenet F. A. Anderson S. B. Kaspi V. M. Prince T. A. & Unwin S. C. *ApJ* 498:365 May 1998.
- Jenet F. A. Anderson S. B. & Prince T. A. *ApJ* 546:394–400 January 2001a.
- Jenet F. A. Anderson S. B. & Prince T. A. *ApJ* 558:302–308 September 2001b.
- Jenet F. A. Hobbs G. B. Lee K. J. & Manchester R. N. *ApJ* 625:L123–L126 June 2005.
- Jenet F. A. Hobbs G. B. van Straten W. Manchester R. N. Bailes M. Verbiest J. P. W. Edwards R. T. Hotan A. W. Sarkissian J. M. & Ord S. M. *ApJ* 653:1571–1576 December 2006.
- Johnson J. D. *Classical Electrodynamics*. Wiley 3. ed edition August 1998. ISBN 047130932X.
- Johnson R. *Applied multivariate statistical analysis*. Prentice Hall Upper Saddle River, NJ [u.a.] 4. ed edition 1998. ISBN 013834194X.
- Johnston S. *Publ. Astron. Soc. Aust.* 19:277–281 2002.
- Johnston S. Lorimer D. R. Harrison P. A. Bailes M. Lyne A. G. Bell J. F. Kaspi V. M. Manchester R. N. D’Amico N. & Nicastro L. *Nature* 361:613–615 February 1993.
- Johnston S. van Straten W. Kramer M. & Bailes M. *ApJ* 549:L101–L104 March 2001.

- Jones P. B. *MNRAS* 423:3502–3511 July 2012.
- Jordan P. *Zeitschrift fur Physik* 157:112–121 February 1959.
- Ju L. Blair D. G. & Zhao C. *Reports on Progress in Physics* 63:1317–1427 September 2000.
- Kaluza T. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite p. 966-972* pages 966–972 1921.
- Kaplan D. B. & Nelson A. E. *Physics Letters B* 175:57–63 July 1986.
- Karastergiou A. von Hoensbroech A. Kramer M. Lorimer D. R. Lyne A. G. Doroshenko O. Jessner A. Jordan C. & Wielebinski R. *A&A* 379:270–278 November 2001.
- Karastergiou A. Kramer M. Johnston S. Lyne A. G. Bhat N. D. R. & Gupta Y. *A&A* 391: 247–251 August 2002.
- Karastergiou A. Johnston S. & Kramer M. *A&A* 404:325–332 June 2003a.
- Karastergiou A. Johnston S. Mitra D. van Leeuwen A. G. J. & Edwards R. T. *MNRAS* 344:L69–L73 October 2003b.
- Karastergiou A. Roberts S. J. Johnston S. Lee H. Weltevrede P. & Kramer M. *MNRAS* pages 590–+ April 2011.
- Kaspi V. M. Taylor J. H. & Ryba M. F. *ApJ* 428:713–728 June 1994.
- Keith M. J. Jameson A. van Straten W. Bailes M. Johnston S. Kramer M. Possenti A. Bates S. D. Bhat N. D. R. Burgay M. Burke-Spolaor S. D’Amico N. Levin L. McMahon P. L. Milia S. & Stappers B. W. *MNRAS* 409:619–627 December 2010.
- Keith M. J. Coles W. Shannon R. M. Hobbs G. Manchester R. N. and Bailes M. Bhat N. D. R. Burke-Spolaor S. G. Champion D. J. Chaudhary A. Hotan A. W. Khoo J. Kocz J. Osłowski S. Ravi V. Reynolds J. R. Sarkissian J. M. van Straten W. & Yardley D. *MNRAS*, *submitted* 2012.
- Kennedy J. F. & Keeping E. S. *Mathematics of statistics*. Van Nostrand 3. ed edition 1964.
- Kerr R. P. & Schild A. *A new class of vacuum solutions of the Einstein field equations*. In *IV Centenario Della Nascita di Galileo Galilei, 1564-1964* page 222 1965.
- Kiel P. D. Hurley J. R. Bailes M. & Murray J. R. *MNRAS* 388:393–415 July 2008.

- Kinkhabwala A. & Thorsett S. E. *ApJ* 535:365–372 May 2000.
- Knight H. S. Bailes M. Manchester R. N. Ord S. M. & Jacoby B. A. *ApJ* 640:941–949 April 2006.
- Kopeikin S. M. *ApJ* 439:L5–L8 January 1995.
- Kopeikin S. M. *ApJ* 467:L93 August 1996.
- Kramer M. *A&AS* 107:527–539 November 1994.
- Kramer M. *ApJ* 509:856–860 December 1998.
- Kramer M. 1996. Poster presented at Pulsar timing, General Relativity and the Internal Structure of Neutron Stars, 24 - 27 September, Amsterdam.
- Kramer M. Wielebinski R. Jessner A. Gil J. A. & Seiradakis J. H. *A&AS* 107:515–526 November 1994.
- Kramer M. Xilouris K. M. Camilo F. Nice D. J. Backer D. C. Lange C. Lorimer D. R. Doroshenko O. & Sallmen S. *ApJ* 520:324–334 July 1999.
- Kramer M. Karastergiou A. Gupta Y. Johnston S. Bhat N. D. R. & Lyne A. G. *A&A* 407:655–668 August 2003.
- Kramer M. Backer D. C. Cordes J. M. Lazio T. J. W. Stappers B. W. & Johnston S. *New Astronomy Reviews* 48:993–1002 December 2004.
- Kramer M. Lyne A. G. O’Brien J. T. Jordan C. A. & Lorimer D. R. *Science* 312:549–551 April 2006a.
- Kramer M. Stairs I. H. Manchester R. N. McLaughlin M. A. Lyne A. G. Ferdman R. D. Burgay M. Lorimer D. R. Possenti A. D’Amico N. Sarkissian J. M. Hobbs G. B. Reynolds J. E. Freire P. C. C. & Camilo F. *Science* 314:97–102 October 2006b.
- Krauss J. D. *Radio Astronomy*. McGraw-Hill, New York 1966.
- Krishnamohan S. & Downs G. S. *ApJ* 265:372–388 February 1983.
- Kuiper G. P. *PASP* 47:15–+ February 1935.
- Kuiper G. P. *ApJ* 93:133 January 1941.
- Kulkarni S. R. *AJ* 98:1112–1130 September 1989.

- Kulkarni S. R. & Heiles C. *AJ* 85:1413–1420 October 1980.
- Kurkela A. Romatschke P. & Vuorinen A. *Phys. Rev. D* 81(10):105021 May 2010.
- Lada C. J. *ApJ* 640:L63–L66 March 2006.
- Landau L. D. *Phys. Z. Sowjetunion* 1:285 1932.
- Lange C. Camilo F. Wex N. Kramer M. Backer D. C. Lyne A. G. & Doroshenko O. *MNRAS* 326:274–282 September 2001.
- Lattimer J. M. & Prakash M. *Science* 304:536–542 April 2004.
- Le Verrier U. *Comptes rendus hebdomadaires des sances de l’Acadmie des sciences* 49: 379–383 1859.
- Lee K. J. Guillemot L. Yue Y. L. Kramer M. & Champion D. J. *ArXiv e-prints* May 2012.
- Liu K. Verbiest J. P. W. Kramer M. Stappers B. W. van Straten W. & Cordes J. M. *MNRAS* 417:2916–2926 November 2011.
- Liu K. Keane E. F. Lee K. J. Kramer M. Cordes J. M. & Purver M. B. *MNRAS* 420: 361–368 February 2012.
- Lommen A. N. *Precision multi-telescope timing of millisecond pulsars: New limits on the gravitational wave background and other results from the pulsar timing array*. PhD thesis University of California, Berkeley 2001.
- Lommen A. N. & Backer D. C. *ApJ* 562:297–302 November 2001.
- Lorimer D. R. & Kramer M. *Handbook of Pulsar Astronomy*. November 2004.
- Luo Q. & Melrose D. B. *MNRAS* 276:372–382 September 1995.
- Lyne A. Hobbs G. Kramer M. Stairs I. & Stappers B. *Science* 329:408– July 2010.
- Lyne A. G. Pritchard R. S. & Graham-Smith F. *MNRAS* 321:67–70 February 2001.
- Lyne A. G. Burgay M. Kramer M. Possenti A. Manchester R. N. Camilo F. McLaughlin M. A. Lorimer D. R. D’Amico N. Joshi B. C. Reynolds J. & Freire P. C. C. *Science* 303: 1153–1157 February 2004.
- Manchester R. N. & Taylor J. H. *ApJ* 191:L63+ July 1974.
- Manchester R. N. Taylor J. H. & Huguenin G. R. *ApJ* 196:83–102 February 1975.

- Manchester R. N. Lyne A. G. Camilo F. Bell J. F. Kaspi V. M. D'Amico N. McKay N. P. F. Crawford F. Stairs I. H. Possenti A. Kramer M. & Sheppard D. C. *MNRAS* 328:17–35 November 2001.
- Manchester R. N. Hobbs G. B. Teoh A. & Hobbs M. *VizieR Online Data Catalog* 7245: 0–+ August 2005.
- Manchester R. N. Hobbs G. Bailes M. Coles W. A. van Straten W. Keith M. J. Shannon R. M. Bhat N. D. R. Brown A. Burke-Spolaor S. G. Champion D. J. Chaudhary A. Edwards R. T. Hampson G. Hotan A. W. Jameson A. Jenet F. A. Kesteven M. J. Khoo J. Kocz J. Maciesiak K. Osłowski S. Ravi V. Reynolds J. R. Sarkissian J. M. Verbiest J. P. W. Wen Z. L. Wilson W. E. Yardley D. Yan W. M. & You X. P. *Publ. Astron. Soc. Aust.* 30:e017 June 2013.
- Matsakis D. N. Taylor J. H. & Eubanks T. M. *A&A* 326:924–928 October 1997.
- Maxwell J. C. *Phil. Man. and Journal of Science* 21:161–175 March 1861.
- Maxwell J. C. *Phil. Trans. Roy. Soc.* 155:459–512 1865.
- McKinnon M. M. *ApJ* 606:1154–1166 May 2004.
- McKinnon M. M. & Stinebring D. R. *ApJ* 529:435–446 January 2000.
- McMaster W. H. *Reviews of Modern Physics* 33:8–27 January 1961.
- Melrose D. Miller A. Karastergiou A. & Luo Q. *MNRAS* 365:638–646 January 2006.
- Melrose D. B. *Royal Society of London Philosophical Transactions Series A* 341:105–115 October 1992.
- Melrose D. B. *Whence the pulses?* In Johnston S. Walker M. A. & Bailes M. editors *IAU Colloq. 160: Pulsars: Problems and Progress* volume 105 of *Astronomical Society of the Pacific Conference Series* page 139 1996.
- Melrose D. B. *Plasma Physics and Controlled Fusion* 45:523–534 May 2003.
- Melrose D. B. & Luo Q. *MNRAS* 352:915–923 August 2004.
- Michel F. C. *ApJ* 322:822–830 November 1987.
- Mitra D. Wielebinski R. Kramer M. & Jessner A. *A&A* 398:993–1005 February 2003.
- Mitra D. Gil J. & Melikidze G. I. *ApJ* 696:L141–L145 May 2009.

- Moore G. E. *Electronics* 38(8) April 1965.
- Moore J. H. *PASP* 40:229 August 1928.
- Naghizadeh-Khouei J. & Clarke D. *A&A* 274:968 July 1993.
- Nan R. *Science in China G: Physics and Astronomy* 49:129–148 March 2006.
- Navarro J. Kulkarni S. R. Thorsett S. E. & Phillips J. A. *Pulsar observations with the Caltech Fast Pulsar Timing Machine*. In *American Astronomical Society Meeting Abstracts* volume 24 of *Bulletin of the American Astronomical Society* page 1276 December 1992.
- Navarro J. Manchester R. N. Sandhu J. S. Kulkarni S. R. & Bailes M. *ApJ* 486:1019 September 1997.
- Newman E. T. & Janis A. I. *Journal of Mathematical Physics* 6:915–917 June 1965.
- Noether E. *Transport Theory and Statistical Physics* 1:186–207 1971.
- Nordström G. *Annalen der Physik* 347:533–554 1913.
- Nordström G. *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences* 20:1238–1245 1918.
- Noutsos A. *Space Sci. Rev.* 166:307–324 May 2012.
- Oppenheimer J. R. & Volkoff G. M. *Physical Review* 55:374–381 February 1939.
- Ord S. M. Johnston S. & Sarkissian J. *Sol. Phys.* 245:109–120 September 2007.
- Ośłowski S. Bulik T. Gondek-Rosińska D. & Belczyński K. *MNRAS* 413:461–479 May 2011a.
- Ośłowski S. van Straten W. Hobbs G. B. Bailes M. & Demorest P. *MNRAS* 418:1258–1271 December 2011b.
- Ostriker J. P. & Gunn J. E. *ApJ* 157:1395–+ September 1969.
- Pacini F. *Nature* 219:145–146 July 1968.
- Paczynski B. *ARA&A* 9:183 1971.
- Pauli W. *Zeitschrift für Physik* 31:765–783 February 1925.
- Penrose R. *Nuovo Cimento Rivista Serie* 1:252 1969.

- Peters P. C. *Physical Review* 136:1224–1232 November 1964.
- Petit G. & Tavella P. *A&A* 308:290–298 April 1996.
- Phinney E. S. *Royal Society of London Philosophical Transactions Series A* 341:39–75 October 1992.
- Planck M. *Verhandlungen der Deutschen phys. Gessellschaft* 2(13):202–204 October 1900a.
- Planck M. *Verhandlungen der Deutschen phys. Gessellschaft* 2(17):237–245 1900b.
- Pottelette R. Treumann R. A. & Dubouloz N. *J. Geophys. Res.* 97:12029 August 1992.
- Pound R. V. & Rebka G. A. *Phys. Rev. Lett.* 4(7):337–341 Apr 1960.
- Pound R. V. & Snider J. L. *Phys. Rev. Lett.* 13:539–540 Nov 1964.
- Rajagopal M. & Romani R. W. *ApJ* 446:543 June 1995.
- Rasio F. A. *ApJ* 427:L107–L110 June 1994.
- Rathnasree N. & Rankin J. M. *ApJ* 452:814–+ October 1995.
- Reber G. *Proc. IRE* 28:68–70 1940.
- Rees M. J. & Gunn J. E. *MNRAS* 167:1–12 April 1974.
- Renardy M. *Linear Alebra and its Applications* 236:53–58 1996.
- Rickett B. J. *ApJ* 197:185–191 April 1975.
- Rickett B. J. *ARA&A* 28:561–605 1990.
- Robinson P. A. *Sol. Phys.* 139:147–163 May 1992.
- Robinson P. A. *Physics of Plasmas* 2:1466–1479 May 1995.
- Robinson P. A. Cairns I. H. & Gurnett D. A. *ApJ* 407:790–800 April 1993.
- Rodin A. E. *MNRAS* 387:1583–1588 July 2008.
- Rohlf K. & Wilson T. L. *Tools of radio astronomy.* 2004.
- Romani R. W. *Nature* 347:741–743 October 1990.
- Romanova M. M. Ustyugova G. V. Koldoba A. V. & Lovelace R. V. E. *ApJ* 616:L151–L154 December 2004.

- Roy A. E. *Orbital motion*. 1988.
- Rudak B. & Ritter H. *MNRAS* 267:513–+ April 1994.
- Ruze J. *Suppl. al Nuovo Cimento* 9:364–380 1952.
- Salam A. & Strathdee J. *Nuovo Cimento A Serie* 11:397–435 September 1972.
- Sallmen S. Backer D. C. Hankins T. H. Moffett D. & Lundgren S. *ApJ* 517:460–471 May 1999.
- Sandhu J. S. Bailes M. Manchester R. N. Navarro J. Kulkarni S. R. & Anderson S. B. *ApJ* 478:L95 April 1997.
- Sathyaprakash B. S. & Schutz B. F. *Living Reviews in Relativity* 12:2 March 2009.
- Sazhin M. V. *Soviet Ast.* 22:36–38 February 1978.
- Scalo J. M. *Fundamentals of Cosmic Physics* 11:1–278 May 1986.
- Scheck L. Plewa T. Janka H.-T. Kifonidis K. & Müller E. *Physical Review Letters* 92(1): 011103 January 2004.
- Schilizzi R. T. Dewdney P. E. F. & Lazio T. J. W. *The square kilometre array*. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* volume 7733 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* July 2010. doi: 10.1117/12.856344.
- Schwarzschild K. *Abh. Konigl. Preuss. Akad. Wissenschaften Jahre 1906,92, Berlin,1907* pages 189–196 1916.
- Sesana A. *ApJ* 719:851–864 August 2010.
- Shannon R. M. & Cordes J. M. *ApJ* 725:1607–1619 December 2010.
- Shannon R. M. & Cordes J. M. *ArXiv e-prints* October 2012.
- Shapiro I. I. *Physical Review Letters* 13:789–791 December 1964.
- Shapiro I. I. Pettengill G. H. Ash M. E. Stone M. L. Smith W. B. Ingalls R. P. & Brockelman R. A. *Physical Review Letters* 20:1265–1269 May 1968.
- Shapiro S. L. & Teukolsky S. A. *Black holes, white dwarfs, and neutron stars: The physics of compact objects*. 1983.

- Shklovskii I. S. *Soviet Ast.* 13:562 February 1970.
- Smarr L. L. & Blandford R. *ApJ* 207:574–588 July 1976.
- Smirnova T. V. *Astronomy Reports* 50:915–925 November 2006.
- Smits J. M. Stappers B. W. Macquart J. Ramachandran R. & Kuijpers J. *A&A* 405:795–801 July 2003.
- Staelin D. H. & Reifenstein III E. C. *Science* 162:1481–1483 December 1968.
- Stairs I. H. *Living Reviews in Relativity* 6:5 September 2003.
- Stairs I. H. Arzoumanian Z. Camilo F. Lyne A. G. Nice D. J. Taylor J. H. Thorsett S. E. & Wolszczan A. *ApJ* 505:352–357 September 1998.
- Stairs I. H. Lyne A. G. & Shemar S. L. *Nature* 406:484–486 August 2000a.
- Stairs I. H. Thorsett S. E. Taylor J. H. & Arzoumanian Z. *Geodetic Precession in PSR B1534+12*. In M. Kramer, N. Wex, & R. Wielebinski editor *IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond* volume 202 of *Astronomical Society of the Pacific Conference Series* page 121 2000b.
- Stairs I. H. Faulkner A. J. Lyne A. G. Kramer M. Lorimer D. R. McLaughlin M. A. Manchester R. N. Hobbs G. B. Camilo F. Possenti A. Burgay M. D’Amico N. Freire P. C. & Gregory P. C. *ApJ* 632:1060–1068 October 2005.
- Stanley M. *Isis. Journal of the History of Science Society* 94:57–89 March 2003.
- Stappers B. W. Kramer M. Lyne A. G. D’Amico N. & Jessner A. *Chinese Journal of Astronomy and Astrophysics Supplement* 6(2):020000–303 December 2006.
- Staveley-Smith L. Wilson W. E. Bird T. S. Disney M. J. Ekers R. D. Freeman K. C. Haynes R. F. Sinclair M. W. Vaile R. A. Webster R. L. & Wright A. E. *Publ. Astron. Soc. Aust.* 13:243–248 November 1996.
- Stinebring D. R. Cordes J. M. Rankin J. M. Weisberg J. M. & Boriakoff V. *ApJS* 55:247–277 June 1984.
- Stoner E. C. *Phil. Mag.* 48:719 1924.
- Tartaglia A. Ruggiero M. L. & Capolongo E. *ArXiv e-prints* September 2011.

- Taylor J. H. *Royal Society of London Philosophical Transactions Series A* 341:117–134 October 1992.
- Taylor J. H. & Weisberg J. M. *ApJ* 253:908–920 February 1982.
- Taylor J. H. Manchester R. N. & Huguenin G. R. *ApJ* 195:513–528 January 1975.
- Thorne K. S. *Reviews of Modern Physics* 52:299–340 April 1980.
- Tout C. A. *Contemporary Physics* 47:145–155 June 2006.
- Turlo Z. Forkert T. Sieber W. & Wilson W. *A&A* 142:181–188 January 1985.
- van den Heuvel E. P. J. *A&A* 291:L39–L42 November 1994.
- van Haasteren R. Levin Y. Janssen G. H. Lazaridis K. Kramer M. Stappers B. W. Desvignes G. Purver M. B. Lyne A. G. Ferdman R. D. Jessner A. Cognard I. Theureau G. D’Amico N. Possenti A. Burgay M. Corongiu A. Hessels J. W. T. Smits R. & Verbiest J. P. W. *MNRAS* 414:3117–3128 July 2011.
- van Straten W. *ApJ* 568:436–442 March 2002.
- van Straten W. *ApJS* 152:129–135 May 2004.
- van Straten W. *ApJ* 642:1004–1011 May 2006.
- van Straten W. *ApJ* 694:1413–1422 April 2009.
- van Straten W. *ApJ* 719:985 August 2010.
- van Straten W. *ApJS* 204:13 January 2013.
- van Straten W. & Bailes M. *Publ. Astron. Soc. Aust.* 28:1–14 January 2011.
- van Straten W. Bailes M. Britton M. Kulkarni S. R. Anderson S. B. Manchester R. N. & Sarkissian J. *Nature* 412:158–160 July 2001.
- van Straten W. Manchester R. N. Johnston S. & Reynolds J. E. *Publ. Astron. Soc. Aust.* 27:104–119 March 2010.
- van Straten W. Demorest P. & Osłowski S. *Astronomical Research and Technology* 9(3): 237–256 May 2012.
- van Vleck J. H. & Middleton D. *Proc. I.E.E.E.* 54:2 1966.

- Verbiest J. P. W. Bailes M. van Straten W. Hobbs G. B. Edwards R. T. Manchester R. N. Bhat N. D. R. Sarkissian J. M. Jacoby B. A. & Kulkarni S. R. *ApJ* 679:675–680 May 2008.
- Verbiest J. P. W. Bailes M. Coles W. A. Hobbs G. B. van Straten W. Champion D. J. Jenet F. A. Manchester R. N. Bhat N. D. R. Sarkissian J. M. Yardley D. Burke-Spolaor S. Hotan A. W. & You X. P. *MNRAS* 400:951–968 December 2009.
- Vivekanand M. *ApJ* 543:979–986 November 2000.
- Vivekanand M. *MNRAS* 326:L33–L36 September 2001.
- Vivekanand M. Narayan R. & Radhakrishnan V. *Journal of Astrophysics and Astronomy* 3:237–247 September 1982.
- Vivekanand M. Ables J. G. & McConnell D. *ApJ* 501:823 July 1998.
- Wang N. Manchester R. N. & Johnston S. *MNRAS* 377:1383–1392 May 2007.
- Weatherall J. C. *ApJ* 483:402 July 1997.
- Weatherall J. C. *ApJ* 506:341–346 October 1998.
- Webbink R. F. *ApJ* 277:355–360 February 1984.
- Weinberg S. *Physical Review Letters* 19:1264–1266 November 1967.
- Weinberger P. *Philosophical Magazine Letters* 88:897–907 December 2008.
- Weinreb S. *A digital spectral analysis technique and its application to radio astronomy*. MIT Research Laboratory of Electronics 1963.
- Weisberg J. M. *Kinematic distances for southern pulsars: GUM and Carina*. In B. K. S. J. & J. W. editors *Pulsars: Problems and Progress* volume 105 of *IAU Colloquium 160* page 479 1996.
- Weisberg J. M. & Taylor J. H. *The Relativistic Binary Pulsar B1913+16: Thirty Years of Observations and Analysis*. In Rasio F. A. & Stairs I. H. editors *Binary Radio Pulsars* volume 328 of *Astronomical Society of the Pacific Conference Series* page 25 July 2005.
- Weisberg J. M. Romani R. W. & Taylor J. H. *ApJ* 347:1030–1033 December 1989.
- Weltevrede P. Edwards R. T. & Stappers B. W. *A&A* 445:243–272 January 2006.

- Weltevrede P. Wright G. & Johnston S. *MNRAS* 424:843–854 August 2012.
- Wesemael F. *QJRAS* 26:273–278 September 1985.
- Wex N. *Classical and Quantum Gravity* 12:983–1005 April 1995.
- Wex N. Johnston S. Manchester R. N. Lyne A. G. Stappers B. W. & Bailes M. *MNRAS* 298:997–1004 August 1998.
- Wick G. C. *Zeitschrift fur Physik* 84:799–800 November 1933.
- Will C. *Living Reviews in Relativity* 4:4 May 2001.
- Wolf E. *Nuovo Cimento* 13:1165–1181 1959.
- Wolszczan A. & Frail D. A. *Nature* 355:145–147 January 1992.
- Yan W. M. Manchester R. N. van Straten W. Reynolds J. E. Hobbs G. Wang N. Bailes M. Bhat N. D. R. Burke-Spolaor S. Champion D. J. Coles W. A. Hotan A. W. Khoo J. Osłowski S. Sarkissian J. M. Verbiest J. P. W. & Yardley D. R. B. *MNRAS* 414:2087–2100 July 2011.
- Yardley D. R. B. Coles W. A. Hobbs G. B. Verbiest J. P. W. Manchester R. N. van Straten W. Jenet F. A. Bailes M. Bhat N. D. R. Burke-Spolaor S. Champion D. J. Hotan A. W. Osłowski S. Reynolds J. E. & Sarkissian J. M. *MNRAS* 414:1777–1787 June 2011.
- You X. P. Hobbs G. Coles W. A. Manchester R. N. Edwards R. Bailes M. Sarkissian J. Verbiest J. P. W. van Straten W. Hotan A. Ord S. Jenet F. Bhat N. D. R. & Teoh A. *MNRAS* 378:493–506 June 2007.
- Young M. D. Manchester R. N. & Johnston S. *Nature* 400:848–849 August 1999.
- Zee A. *Physical Review Letters* 42:417–421 February 1979.
- Zhang C. M. & Kojima Y. *MNRAS* 366:137–143 February 2006.

A

Some details on general relativity

A.1 Derivation of the Einstein's field equation

Since the curvature of space-time is a central premise of general relativity, and it can be fully described by measuring the distances along paths, we begin by introducing the metric tensor. This tensor defines the infinitesimal interval (equivalent of the distance in Euclidean geometry) squared:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \quad (\text{A.1})$$

where Einstein's summation convention is used. All the indices go from 0 to 3 for Greek letters and from 1 to 3 for Latin letters. While the interval is invariant, the metric depends on the choice of coordinates.

Following nomenclature is used for the interval defined in equation A.1. If $ds^2 < 0$ then we call that interval time-like; for $ds^2 > 0$ it is called space-like; while in the border case we refer to it as light-like. In standard Euclidean geometry the spatial part of the metric would simply have all elements equal to unity on the diagonal with all the others equal zero in Cartesian coordinates. In special relativity the metric included time coordinate and, in Cartesian coordinates, it is a diagonal matrix with all elements equal to ± 1 with the three spatial coordinates all having the same sign and the time-like component equal ∓ 1 . The choice of signs depends on the chosen convention.

According to the principle of equivalence, for a particle moving freely and affected only by gravitational forces, there is a coordinate system in which the equation of motion for this particle is:

$$\frac{d^2 \xi^\alpha}{d\tau^2} = 0, \quad (\text{A.2})$$

where $d\tau$ is the proper, or coordinate, time, defined in the same manner as in special relativity: $d\tau^2 = -g_{\alpha\beta} d\xi^\alpha d\xi^\beta$, assuming the same sign convention as above. Note that

here we have effectively presented gravitation as a curvature affecting motion of a test particle. If we convert this equation to another coordinate system x^μ we arrive at:

$$\frac{d^2 x^\kappa}{d\tau^2} + \Gamma_{\mu\nu}^\kappa \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = 0, \quad (\text{A.3})$$

where Γ are Christoffel symbols, i.e. the affine connection:

$$\Gamma_{\beta\gamma}^\alpha = \frac{\partial dx^\alpha}{\partial \xi^\kappa} \frac{\partial^2 \xi^\kappa}{\partial x^\beta \partial x^\gamma}. \quad (\text{A.4})$$

It turns out the Christoffel symbols are closely related to the curvature of the space-time. The Riemann curvature tensor is constructed from Christoffel symbols and measures the change of a vector under parallel transport around an infinitesimal closed paths. This tensor is uniquely defined for any metric and linear in the second derivatives. This tensor is given by:

$$R_{\sigma\mu\nu}^\chi = \frac{\partial}{\partial \mu} \Gamma_{\nu\sigma}^\chi - \frac{\partial}{\partial \nu} \Gamma_{\mu\sigma}^\chi + \Gamma_{\mu\lambda}^\chi \Gamma_{\nu\sigma}^\lambda - \Gamma_{\nu\lambda}^\chi \Gamma_{\mu\sigma}^\lambda, \quad (\text{A.5})$$

Of the 256 elements of the tensor $R_{\sigma\mu\nu}^\chi$ only 20 are independent from symmetry and commutation constraints. We can form other useful quantities by using metric to contract the Riemann curvature tensor. The two most useful of these linear combinations of the $R_{\sigma\mu\nu}^\chi$ are the Ricci tensor and the curvature scalar defined respectively by:

$$R_{\sigma\nu} = R_{\sigma\chi\nu}^\chi \quad (\text{A.6})$$

and

$$R = g^{\sigma\nu} R_{\sigma\nu}. \quad (\text{A.7})$$

We have now introduced all the quantities useful to describe the curvature of space-time and its effects on motion of physical systems within the frame of this thesis. To fully present GR we will now describe how to derive how these physical systems generate the space-time curvature. We do this by introducing the energy-momentum tensor $T^{\mu\nu}$ which maps an infinitesimal volume in space-time to the energy and momentum in this volume. This tensor is already present in the theory of special relativity. For completeness, we present its most important properties here.

In order to preserve the Newtonian mechanics in the weak static field regime, the first element of the energy-momentum tensor is $T^{00} = \rho$. The T^{0i} elements represent the momentum density (or energy flux density) while T^{ik} represent the momentum flux densities and stresses. To ensure that energy and momentum are conserved, the covariant

divergence of the energy-momentum tensor must vanish everywhere: $T_{;\nu}^{\mu\nu} = 0$ where the index $;\nu$ denotes the covariant derivative¹. The energy-momentum tensor is symmetric.

By considering these properties of the energy momentum tensor, Einstein found a tensor that could be equated to it by assuming the sought tensor will only contain terms linear in first the second derivatives or quadratic in the first derivatives of metric. This assumption originates from the desire to maintain uniformity of gravitational field with respect to scale. It turns out there is a unique class of tensors with these properties, and it is now known as the Einstein's tensor $G^{\mu\nu}$:

$$G^{\mu\nu} = R_{\mu\nu} - 0.5g_{\mu\nu}, \quad (\text{A.8})$$

If the Poisson's equation (defining the relation between the Newtonian potential and the matter density distribution) is to remain valid in the weak field approximation, the field equations must be:

$$G_{\mu\nu} = 8\pi T_{\mu\nu}. \quad (\text{A.9})$$

Due to symmetry, only 10 of these equations are independent. In this view gravity is no longer a force but the curvature of space-time caused by the energy-momentum distribution.

A.2 Solutions to the Einstein's field equations

After introducing his field equations, Einstein could not find a solution that would correspond to a static Universe, which he favoured for philosophical reasons. He therefore relaxed the assumptions about the properties of the sought tensor on the left hand side of equation A.9 and introduced the cosmological constant Λ . This constant enables a stationary solution to these field equations by replacing the Einstein's tensor on the left hand side by $G^{\mu\nu} + \Lambda g^{\mu\nu}$. This modification preserves the symmetry and vanishing of covariant divergences of the left hand side but does not yield Poisson's equation in the weak static field limit. This naturally allows constraints on Λ by measuring the deviations from Newton's dynamics in this limit. After Hubble's discovery of cosmological redshift of other galaxies, Einstein dropped the idea of the cosmological constant. In a twist of fate, after the discovery of the Universe's expansion acceleration, recently awarded with Nobel prize, the cosmological constant has gained on popularity as one of the possible mechanisms to explain this acceleration, i.e. a possible form of the "dark energy".

¹Covariant derivative of a tensor $A^{\alpha\beta}$ is defined as: $\frac{\partial A^{\alpha\beta}}{\partial x^\mu} + \Gamma_{\gamma\mu}^\alpha A^{\gamma\beta} + \Gamma_{\gamma\mu}^\beta A^{\alpha\gamma}$.

The first exact solution to the Einstein's field equations was obtained by a German physicist Karl Schwarzschild soon after the publication of these equations. He found a solution for the metric around an uncharged point mass (Schwarzschild, 1916). Very soon this solution was extended to allow the point mass to be electrically charged (Nordström, 1918). This type of metric is not very interesting from astrophysical point of view as astronomical bodies are generally not electrically charged. The more interesting case of a rotating spherical mass remained elusive for many more years. Finally in 1965 Schwarzschild's solution was generalised to allow rotation of the central object. This unique and complete solution was a contribution from Roy Kerr (Kerr & Schild, 1965). In the same year a fully generalised solution which allowed the rotating mass to have an electrical charge was derived by Ezra Newman (Newman & Janis, 1965). We present here the most astrophysically relevant Kerr metric in the spherical coordinate system;

$$\begin{aligned}
 ds^2 = & \left(1 - \frac{r_s r}{r^2 + a^2 \cos^2 \theta}\right) dt^2 - \frac{r^2 + a^2 \cos^2 \theta}{r^2 - r_s r + a^2} dr^2 - (r^2 + a^2 \cos^2 \theta) d\theta^2 + \\
 & - \left(r^2 + a^2 + \frac{r_s r a^2}{r^2 + a^2 \cos^2 \theta} \sin^2 \theta\right) \sin^2 \theta d\phi^2 + \frac{2r_s r a \sin^2 \theta}{r^2 + a^2 \cos^2 \theta} dt d\phi, \tag{A.10}
 \end{aligned}$$

where r_s is the Schwarzschild radius equal $2M$, where M is the mass of the central object. The parameter a relates to the rotation of the central mass and is defined as $a = \frac{J}{M}$, where J is the angular momentum of the central mass.

This solution can be used to describe the space-time around a planet, a star or a black hole. Because this solution, as expected, yields results indistinguishable from Newton's predictions at large distance from the star, it is most interesting to study it for objects where the behaviour of a test particle outside the central object is different from the classical expectations. The ratio M/R , where R is the radius of the object in question, is a good indicator of how relativistic an object is. For a black hole this ratio is 0.5 while for a main sequence star like our Sun it is of the order of 2×10^{-6} . We note that a neutron star is a highly relativistic object with this ratio equal ~ 0.2 , with the exact value depending on the equation of state of dense matter (Shapiro & Teukolsky, 1983). In the most extreme case, some effects unique to relativistic dynamics are strongly pronounced in the vicinity of a Kerr black-hole. There are three distinct regions of space around such an object where the test particle will deviate strongly from classical behaviour. Far away from the object, the space-time appears flat but when we approach the curvature becomes more significant. At radii smaller than $0.5 \left(r_s + \sqrt{r_s^2 - 4a^2 \cos^2 \theta}\right)$, the rotation of the black hole is causing the Lense-Thirring effect (also known as the frame dragging). This means that any object that enters this region of space, known as ergosphere, is forced to co-rotate

with the black hole. The name originates from the fact that reference frames connected with such an object are the only inertial frames possible within the ergosphere. The dragging is caused by the reversal of sign of the “temporal” component of the metric at the ergosphere radius. Any non-co-rotating geodesics within this region become time-like (i.e. correspond to trajectories of tachyons, or faster-than-light particles) and are not physically allowed. In general this effect allows particles to gain energy at the expense of the black hole and escape from the ergosphere (this is known as the Penrose mechanism Penrose, 1969). Gravity B probe explored the Lense-Thirring effect. The satellite contains and monitors four cryogenically cooled gyroscopes on a low Earth orbit. While the systematic errors have been larger than expected, the mission successfully confirmed both the Lense-Thirring effect and the Earth-Moon system drift in the orbit plane caused by the curvature of space-time (Everitt et al., 2011).

When a test particle comes closer to the black holes, it encounters the event horizon at $r = 0.5 \left(r_s + \sqrt{r_s^2 - 4a^2} \right)$. This is the surface at which the radial component of the metric becomes infinite. At this surface, the escape velocity exceeds that of the speed of light and hence the name black hole. A phenomenon related to the Penrose process is the Hawking radiation (Hawking, 1975; Bekenstein, 1975). Due to quantum fluctuations, a pair of particles can be created in the ergosphere in a configuration in which one of the particles will escape the black hole to the flat space region while the other one will enter the event horizon. The black holes are emitting photons and particles in this way with a black body spectrum with intensity inversely proportional to the black hole’s mass.

A.3 Classical tests of general relativity

As the predictions of the general relativity were quite different than those of Newtonian gravity, it is necessary to test them experimentally. Are accelerating reference frames really equivalent? Einstein himself suggested three tests that could confirm or refute the GR. Firstly, the orbits of planets around stars are no longer conic sections, in particular there are no closed elliptical orbits like there were in Newtonian gravity. This enabled a test of GR using the planets in the Solar System. Secondly, as predicted by one of the Einstein’s famous thought experiments, massive objects affect the trajectory of light and Einstein predicted that the apparent positions of stars will change by 1’45” when the line of sight towards them will graze the Sun. Thirdly, gravitational redshift predicted by Einstein’s thought experiment could be tested both astrophysically and physically.

The first test was passed thanks to Einstein. In 1859 French mathematician Urbain Le Verrier, also known for the discovery of Neptune, pointed out that the precession of the

Mercury's orbit does not match the predictions of the celestial mechanics based on Newtonian gravity. Classical dynamics predicted additional precession due to perturbation from other planets in the solar system to be $531'$ per century whereas Le Verrier's analysis of 150 years of data showed the rate of $569'$ per century (Le Verrier, 1859). While his measurement was off by $5'$ per century, it pointed out a problem with the Newtonian gravity. Einstein calculated that the general relativity will cause orbit of Mercury to precess at a rate increased by $43'$ per century, in perfect agreement with observations (Einstein, 1916; Clemence, 1943). This was the first success of the new theory of gravitation.

The second test was performed soon after the first one by sir Arthur Eddington, after whom the famous luminosity limit for astrophysical sources is named. In 1919 he and his collaborator Astronomer Royal Frank Dyson organised expeditions to various places in western Africa to watch the stars near the Sun during the total eclipse on May 29th. Not only the bending of light matched Einstein's predictions but the expedition and Dyson's support gained Eddington an exemption for conscription. Before the expedition he had refused to fight in the World War I and was frowned upon by the British Government and risked being arrested (Stanley, 2003). The scientific results of the experiment were published soon after the expedition (Dyson et al., 1920). Similar tests have been since repeated a few times and have given consistent results.

The last "classical" test, exploiting the gravitational redshift, proved most difficult. It was first performed by Walter Adams, an American astronomer, director of the Mount Wilson Observatory. In 1915 he measured the spectrum of Sirius B thus discovering the second known white dwarf (Adams, 1915). 10 years later he used this white dwarf to measure the gravitational redshift and obtained results consistent with Einstein's prediction (Adams, 1925), also confirmed by an independent work (Moore, 1928). This result has been contested several times, both by pointing out that this measurements were contaminated by the Sirius A star (Greenstein et al., 1971) and by arguing that his result was an effect of confirmation bias (Hetherington, 1980). Adam's scientific integrity was defended by Wesemael (1985) who supported the thesis of data contamination by Sirius A star. In the meantime, Robert Pound and Glen Rebka based at Harvard University managed to confirm gravitational redshift to within 10% by implementing Einstein's thought experiment using the Mössbauer effect (Pound & Rebka, 1960). Further improvements to the experimental setup allowed the precision to be improved to 1% as early as 1964 (Pound & Snider, 1964). Global Positioning System relies heavily on the predictions of GR - if the gravitational time dilation was not accounted for, the reported locations would be off by kilometers within one day.

A.4 Gravitational waves

For a small perturbation of the metric $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, where $\eta_{\mu\nu}$ is the unperturbed metric far away from any mass and $h_{\mu\nu} \ll 1$ is the small perturbation, the Ricci tensor can be presented as:

$$R_{\mu\nu} = \frac{\partial}{\partial x^\nu} \Gamma_{\sigma\mu}^\sigma - \frac{\partial}{\partial x^\sigma} \Gamma_{\mu\nu}^\sigma + o(h^2), \quad (\text{A.11})$$

where $o(h^2)$ represents all the elements of the order of h^2 and higher. We can also simplify the affine connection with this metric:

$$\Gamma_{\mu\nu}^\sigma = 0.5\eta^{\sigma\lambda} \left[\frac{\partial}{\partial x^\mu} h_{\lambda\nu} + \frac{\partial}{\partial x^\nu} h_{\lambda\mu} - \frac{\partial}{\partial x^\lambda} h_{\mu\nu} \right] + o(h^2). \quad (\text{A.12})$$

Note that since we are interested only in terms linear with respect to the metric perturbation, we use the flat metric η for all operations of lowering and raising the indices. With this assumption and from equations A.11 and A.12, we obtain the fully linearised Ricci tensor:

$$R_{\mu\nu} = 0.5 \left(\square^2 h_{\mu\nu} - \frac{\partial}{\partial x^\sigma \partial x^\mu} h_\mu^\sigma - \frac{\partial}{\partial x^\sigma \partial x^\nu} h_\nu^\sigma + \frac{\partial}{\partial x^\mu \partial x^\nu} h_\sigma^\sigma \right), \quad (\text{A.13})$$

where \square^2 is the d'Alembert operator for flat space, i.e., divergence of the gradient in Minkowski space. We can now subtract the part of the field equations A.9 related to the flat metric η and obtain field equations:

$$\square^2 h_{\mu\nu} - \frac{\partial}{\partial x^\sigma \partial x^\mu} h_\mu^\sigma - \frac{\partial}{\partial x^\sigma \partial x^\nu} h_\nu^\sigma + \frac{\partial}{\partial x^\mu \partial x^\nu} h_\sigma^\sigma = -16\pi (T_{\mu\nu} - 0.5\eta_{\mu\nu} T_\sigma^\sigma). \quad (\text{A.14})$$

Here the momentum-energy tensor is also linearised and so is independent of $h_{\mu\nu}$. This simplifies the conservations laws to only standard derivatives: $\frac{\partial}{\partial x^\mu} T_\nu^\mu = 0$. By exploiting these conservation laws and choosing a convenient gauge (i.e. the harmonic coordinate system in which $g^{\mu\nu} \Gamma_{\mu\nu}^\sigma = 0$ holds everywhere) we arrive at the final linearised form of the field equations:

$$\square^2 h_{\mu\nu} = -16\pi (T_{\mu\nu} - 0.5\eta_{\mu\nu} T_\sigma^\sigma). \quad (\text{A.15})$$