Golam Shaifullah

Timing & Properties of Recycled Pulsars

A dissertation submitted
to the Faculty of Physics,
Bielefeld Universität.

All of time

would not have sufficed.

yet now I know

the heart of darkness

trembling

from memory of light

Flickering like fireflies

on a summer's night

or glimmering through eyes

that know the colour of nine.

In loving memory of Saheda Begum, my mother and Ishita Maity, my partner.

Thesis advisor
JProf. Dr. Joris PW Verbiest
Faculty of Physics
Bielefeld Universität

Examination board
Prof. Dr. Thomas Huser
Faculty of Physics
Bielefeld Universität

Prof. Dr. Dominik Schwarz
Faculty of Physics
Bielefeld Universität

Prof. Dr. Thomas Tauris
Argelander Institut für Astronomie
Universität Bonn

This document was typeset using LaTeX, with a modified version of the tufte-latex package. The bibliography was processed by Bibtex.

This thesis uses the Vollkorn type-face by Friedrich Althausen for the basic text and display. Additionally, some mathematical symbols use the Asana Math type-face by Apostolos Syropoulos. Chinese text is typeset using the WenQuanYi Micro Hei type-face, Japanese text uses the GenEi Antique type-face and Bengali text is typeset using the Hind Siliguri type-face.

Unless otherwise mentioned, all figures in this thesis are original work by the author. Non-original figures are credited to the original creators and used subject to their terms and licenses.

Cover art is original work by the author.

This research was supported by the Max-Planck-Institute für Radioastronomie, Bonn.

First printing, August 2017

Printed on non-ageing paper ISO 9706

Contents

List of Figures

List of Plates

List of Tables

This thesis has been prepared using the International System of Units (SI) of units. However, the following International Astronomical Union (IAU) approved non-SI units are also used:

Quantity	Unit: Name	Symbol	Value
time	minute	min or ”	60 s
time	hour	h	3600 s =60 min
time	day	d	86 400 s =24 h
angle	second of arc	$''$	$\pi/648\,000$ rad
angle	minute of arc	$'$	$\pi/10\,800$ rad
angle	degree	$^\circ$	$\pi/180$ rad
length	astronomical unit	au	$0.149\,598 \times 10^{12}$ m
length	parsec	pc	30.857×10^{16} m
length	solar radius	$R_\odot$	6.957×10^{8} m
mass	solar mass	$M_\odot$	1.9891×10^{30} kg
flux density	Jansky	Jy	10^{-26} W m^{-1} Hz^{-2}

See the IAU recommendations page for detailed explanations.

In addition, the following centimetre–gram–second (CGS) units are also used:

Quantity	Unit: Name	Symbol	Value
magnetic field strength	gauss	G	10^{-4} T

Finally the notation of $\mathrm{i} = \sqrt{-1}$ is used to avoid confusion with running indices.

Abstract

Recycled pulsars are old pulsars which have been spun up to very high rotational frequencies through the transfer of angular momentum by stellar material accreted from a companion via Roche-lobe overflow. These pulsars consist of matter at extreme densities, subject to some of the strongest electromagnetic fields and while the study of these objects has remained a challenging task, recycled pulsars are also extremely sensitive probes for studying fundamental physics. In this thesis I present results from three investigations related to the observation of recycled pulsars. I present an overview of the most significant artifacts that affect digital data recording systems for pulsar observations. I show that for the special case of a coherent dedispersion pulsar backend, the action of the dedispersion filter is non-linear in phase and therefore the resultant signal cannot be reconstructed perfectly by an analysis-synthesis filterbank which is built using simple digital filters, if the channel bandwidths are very high. I review a least-mean-squares based filter optimisation algorithm with the aim of addressing this issue. I then present an updated pulsar timing solution for PSR J2051–0827, which was the second black-widow pulsar to be discovered. For the first time for this system, we are able to detect a decrease in the dispersion measure of $\sim 2.5 \times 10^{-3}\,\mathrm{cm^{-3}\,pc}$. The extended timing also results in the most precise measurements to date of the mean proper motion ($6.1(1)\,\mathrm{mas\,yr^{-1}}$) and the 2-D transverse velocity ($30(9)\,\mathrm{km\,s^{-1}}$) of this system. Secular variations in the orbital period are recovered for more than one complete cycle and we detect previously unknown short-term variations, on timescales of $\sim$150 days. The 21-year dataset results in a weighted timing residual of of $\sim 5\,\mu\mathrm{s}$, which is comparable to that of sources already in PTAs suggesting some black-widow pulsars *may* indeed be useful PTA sources. Finally, I measure the spectral indices of 12 recycled pulsars for 11 of which flux densities at less than two frequency bands were known, using multi-epoch flux density measurements at three frequency bands carried out with the Arecibo Radio Observatory. We add rederived spectral indices for 43 recycled pulsars which had pre-existing spectral index measurements by adding flux density measurements at other frequency bands and 19 recycled pulsars for which only flux densities were available in literature to increase the sample size to 74, which is almost two-thirds of the known Galactic population of 195. The measured spectral indices suggest that while no obvious difference exists between isolated recycled pulsars and those in binaries, redback systems appear to have steeper spectral indices. Recycled pulsars which are also visible in the γ-ray regime appear to have a steeper spectral index than those which are visible only radio frequencies. This may be the reason why targeted searches for such sources at high radio frequencies have been less successful compared to concurrent searches at lower frequencies. The two-tailed Kolmogorov-Smirnov test shows that the spectral index distribution of recycled pulsars is similar to that of classical pulsars, as well.

I

Introduction

立ち別れ
いなばの山の
峰に生ふる
まつとしきかば
今かへりこむ

And should I then presume?
And how should I begin?

— **T. S. Eliot;** *The Love Song of J. Alfred Prufrock*

Pulsars or pulsating stars are exotic stars which are created in the aftermath of violent explosions called supernovae. The most characteristic property of pulsars is their tightly beamed radiation, often detectable only in the radio frequency regime. This beam is typically the only detectable component of the pulsar and appears and disappears at the rate of rotation of the pulsar, leading to the misleadingly named 'pulsations', which are merely the result of the beam crossing the detector. This chapter presents a brief summary of pulsars and their properties, defines the various terms used in the following chapters and introduces the specific class of pulsars which form the basis of the studies carried out as part of this thesis.

1.1 What are pulsars?

Stars do not die. Or rather, stars being giant collections of primarily atomic gas ignited into radiation of energy released via the fusion of elemental Hydrogen into Helium under the influence of the immense pressures at their cores due to gravity, do not experience anything akin to human lives. They do, however, proceed through well defined stages of evolution almost entirely driven by the contest between gravity and outward pressure resulting from the arrangement (and state) of stellar matter.

A star like our Sun or even up to 8 times as massive as it, will eventually exhaust all of the Hydrogen at its core and gravity will overcome the radiation pressure from the release of photons due to nuclear fusion to start squeezing the stellar material at the core into a denser state. This will lead first to a state of expansion driven by burning of Hydrogen in a shell around the core, while Helium is fused at the core to form Carbon. After the Helium at the core is also depleted, the inner core collapses while two shells of Helium and Hydrogen begin to burn. This is followed by a stage of an ejection of the outer layers of stellar material to form a planetary nebula. The inner layers will continue to contract until the matter is packed so tight that the electrons of the elemental matter begin to repel each other following the rules of quantum mechanics. The star is now called a white dwarf (WD). If the star at the start of its collapse had a mass of between ~8 to 20 times the mass of the Sun, even the electron degeneracy pressure cannot sustain the material and the collapse continuous on till the remaining matter is squeezed into almost purely nucleonic states (i.e., protons and neutrons) and the star is now called a neutron star (NS). If the NS is magnetised it emits large amounts of radiation. A rotating NS of this kind was first detected in the radio frequency regime by Jocelyn Bell, Anthony Hewish and their collaborators (Hewish et al., 1968). These objects were first hypothesised to be the end result of supernovae (SNe) by Baade and Zwicky (1934a), although the theoretical work of Chandrasekhar (1931a) and Landau (1932) had predicted the existence of NS. That a NS could be the powerhouse which drove the brilliant Crab Supernova was a hypothesis first tendered by Pacini (1967) and the theory of a compact stellar object with a dipolar magnetic field co-rotating with the plasma near its surface was first presented by Gold (1968), while Goldreich and Julian (1969) built on the work of Ostriker and Gunn (1969) to provide a competing theory for the source of pulsar emission. Ruderman and Sutherland (1975) introduced superfluidity as the possible origin of the intense magnetic fields and Cheng et al. (1976) and subsequent work by the same authors improved upon the Goldreich-Julian model by introducing the concepts of polar cap current flow, electron-positron production, and the effect of the magnetoactive plasma around the rotating pulsar. Alternatively, Michel (1973b) attempts to solve the *pulsar problem* by studying the magnetic flux of a freely rotating charged, magnetosphere. However, the theory of the pulsar phenomenon remains an actively researched field and a firm conclusion on the theory has yet to be made.

Pulsar astronomy, however, has continued to grow in leaps and bounds through the discovery of ever more exotic and unique systems. In 1975, Hulse and Taylor discovered a pulsar in a binary system with a NS (Hulse and Taylor, 1975). By measuring the change in the orbital period of the binary, Taylor and Weisberg (1982) provided the first evidence for Gravitational Wave emission. Less than fifteen years after Bell's original discovery, Backer et al. (1982) discovered the first millisecond pulsar (MSP), an isolated object whose existence was difficult to explain with known evolution scenarios, leading to the *recycling* scenario being proposed by Bhattacharya and van den Heuvel (1991), even as alternative theories were put forward by Henrichs and van den Heuvel (1983); Ruderman and Shaham (1983). Although the discovery of a pulsar which is in the process of ablating its companion (Fruchter et al., 1988) was initially thought to support the idea that recycling would ultimately lead to the production of an isolated MSP, it is now thought that the efficiency of the ablation process is typically too low in the systems we observe to lead to this. The study of recycled pulsars (RPs) and the effects of artefacts in pulsar receivers used to observe them form the subject of this thesis.

However, even more exotic pulsars continue to be discovered, such as the first MSP in a hierarchical triple system (Ransom et al., 2014), the double pulsar system PSR J0737+3039 (Burgay et al., 2003) which now provides up to five tests of General Relativity (Kramer et al., 2006), the *magnetars* (Duncan and Thompson, 1992), pulsars which have surface magnetic field strengths far greater than the typical radio bright pulsars, the *intermittent* pulsars or rapidly rotating radio transients (RRATs; see e.g., McLaughlin et al., 2006) as well as the *transition* binary PSR J1023+0038 (Stappers et al., 2014) which appears to switch between a radio bright pulsar state and an accretion powered X-ray bright state.

1.1.1 *Before pulsars*

NS are the collapsed inner-most layers of a massive star after it has gone through a SN stage; a colossal explosion that marks the end of nuclear fusion. The first recorded SN was probably SN 185, from 185 AD when Chinese astronomers documented a mysterious "guest star" which appeared in the night sky and remained visible for about 8 months.

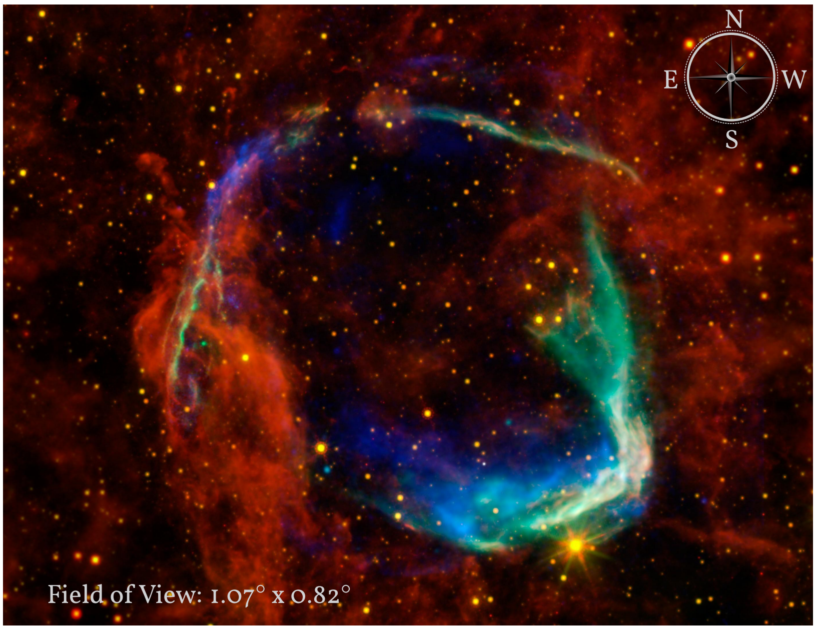

Figure 1.1: An image of the oldest documented example of a SN, RCW86, created using X-ray images from NASA's *Chandra* X-ray Observatory (coloured in blue) and the European Space Agency's *XMM* -Newton Observatory (coloured in green) of interstellar gas that has been heated to $\sim 10^6$ K by passing shock waves from the SN. Also shown are infrared data from NASA's *Spitzer* Space Telescope (yellow) and *WISE*, the Wide-Field Infrared Survey Explorer (red), showing dust radiating at a temperature of 100 K. RCW86 is approximately ~2.5 kpc away. At about ~25 pc in diameter, it occupies a region of the sky in the southern constellation of Circinus that is slightly larger than the full moon. This image was compiled in October 2011.
Image Credit: X-ray: NASA/CXC/SAO and ESA; Infared: NASA/JPL-Caltech/B. Williams (NCSU). Obtained from the image gallery at NASA.gov.
The original image has been modified and custom annotations added.

Figure 1.1 shows a multiband image of RCW86, the supernova remnant (SNR) associated with SN 185.X-ray images from NASA's Chandra X-ray Observatory and the European Space Agency's XMM-Newton Observatory were combined to form the blue and green colours in the image. The X-rays show the interstellar gas that has been heated to millions of degrees by the passage of the shock wave from the SN.

One of the most well-studied SNRs in recent times however, is probably that of SN 1054, also known as the Crab Nebula. Records of the appearance of a star bright enough to be visible in the daytime sky can be found in the Chinese[1] and Japanese[2] astronomical texts dating this event to the first quarter of 1054 AD. The nebula of relatively cool gas and dust formed from the outer layers of the SN is even visible with reasonably well-constructed amateur optical telescopes and is located about 2.0(5) kpc (Kaplan et al., 2008) away in the constellation Taurus.

Figure 1.2 shows an image from the Hubble Space Telescope, created from three separate, high resolution images taken over 30 years, which reveals the central core of this fascinating object. At its very heart lies a star between ten and twelve kilometres in diameter, with a mass slightly

[1] The first mention is found in *Xù zīxùn tōng jiàn chángbiān* (續資治通鑑長編), literally; 'Extended Continuation to The Comprehensive Mirror in Aid of Governance' a historical record of the Northern Song dynasty from ~976 to ~1126 by Li Tao (李燾) (1115–1184)

[2] *Meigetsuki* (明月記), literally; 'The Record of the Clear Moon'; a personal diary from 1180 to ~1241 maintained by Fujiwara Sadaie (藤原定家), better-known as Fujiwara no Teika

Figure 1.2: Peering deep into the core of the Crab Nebula, this close-up image reveals the beating heart of one of the most historic and intensively studied remnants of a SN. This image was created from three separate, high resolution images taken with the Hubble Space Telescope over 30 years, which reveals the central core of this fascinating object. At its very heart lies a star between ten and twelve kilometres in radius, with a mass about one-and-a-half times that of the sun and rotating roughly 30 times every second.
Image Credit: Original image by ESA/Hubble. Obtained from spacetelescope.org.
The original image has been modified and custom annotations added.

less than one-and-a-half times that of the sun and rotating roughly 30 times every second. This star emits relatively little light in the optical regime and yet is one of the brightest point sources in the radio-frequency regime, emitting colossal numbers of low-energy photons directly in the radio[3] and perhaps gamma ray frequencies, while secondary processes in the nebula like re-emission by gas heated by direct emission and collisions of the surrounding gas with highly energetic streams of particles, etc., lead to emission across multiple frequency domains from the infra-red radiation (IR) to γ rays.

[3] which is still a small fraction of the total energy emitted.

Gas, coloured in red, that was ejected at the time of the SN surrounds the NS at the centre and is shaped into an intricate web of filaments and cavities by a stream of particles blowing outward from the NS. The hazy blue glow is due to radiation given off by electrons trapped in the magnetic field of the NS as they spiral around the magnetic field lines, called synchrotron radiation. A zoomed out image would show the blue 'jets' from the centre being propelled in opposite directions out to a few parsecs.

The dimensions and mass of the central object imply densities even higher than the average value in the nucleus of an atom and hence, it is believed that the star must be entirely composed of either neutrons or neutrons and other fundamental particles like quarks. The star is therefore called a NS. Due to conservation of angular momentum, this NS is born as a rapidly spinning object. Having inherited magnetic fields from the pre-NS star, this star also emits an extremely tight beam of coherent radiation. This radiation, which is predominantly *detected* in the radio-frequency regime, sweeps over the earth once with every rotation of the NS, much like a lighthouse beam seen from a ship at sea.

SNe like that associated with the Crab Nebula have continued to have been recorded through history, although those bright enough to be visible to the naked eye are typically quite few. Tycho Brahe presented an immensely detailed investigation of a SN when in 1572 he documented the appearance of a new star in the Cassiopeia constellation. Figure 1.3 shows a page from Brahe's notebooks, marking the position of SN 1572. Kepler, who built on Brahe's painstaking records of the motion of the Solar System planets to provide the first laws of planetary motion, also recorded a similar event in 1604.

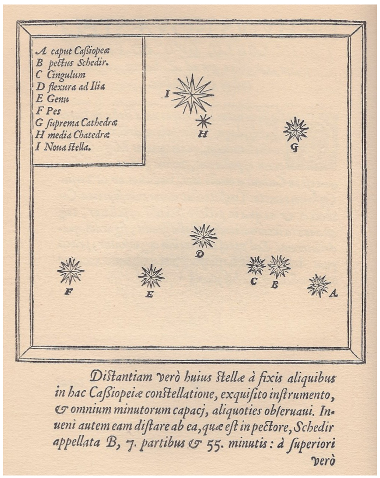

Figure 1.3: Tycho Brahe's drawing showing the position of SN 1572 in his notes published under the title *Tychonis Brahe Dani, Epistolarum Astronomicarum Libri*. **Image Credit:** SLUB Dresden, under (CC-BY-SA 4.0). Obtained from the archives of SLUB Dresden.

It was not until the late 20th century that the first SN was actually observed in the process of tearing itself apart, when astronomers, first at the Las Campanas Observatory in Chile (Kunkel et al., 1987) and then across the world were able to observe SN 987A, an event in the Large Magellanic Cloud (LMC) which had probably exploded 168,000 years earlier[4]

[4] That is to say the light had taken that amount of time to reach the Earth.

While SNe themselves are fascinating astronomical events, we press on to that which they leave behind; a NS like in the case of the Crab SN.

1.1.2 *A deeper understanding of the stars*

Even with the detailed observations of Brahe, Kepler and many other astronomers, it took almost 250 years for us to understand the real nature of these cosmic explosions. In particular, after the birth of Quantum Theory and the General Theory of Relativity fuelled a renaissance in physics and astronomy, astronomers and physicists quickly realised that many of the predictions of these theories did rather well in explaining more than a few of the strange astronomical phenomena they had encountered.

Even though the foundations of both these theories were already well laid out within the first decade of the 20th century, the 1930s marked a special period in the growth of our understanding of SNe and their consequences. In 1931, Chandrasekhar presented a theory for the evolution of self-gravitating systems (Chandrasekhar, 1931a,b) leading to the formation of WD stars once thermal energy was no longer sufficient to sustain them against gravitational collapse. In 1932, the neutron was discovered by Chadwick (1932). However, there was little evidence yet to suggest the existence of a NS and in fact, astronomers were not fully convinced of the need to organise novae into different classes (Shapley and Curtis, 1921). Although Chandrasekhar's theory indicated that the collapsed core of these SNe would have densities greater than that of either the electron or proton, little was known about the behaviour of matter at such high densities. In spite of these gaps in the prevalent theories, in a pair of remarkably prescient papers (Baade and Zwicky, 1934a,b), Walter Baade and Fritz Zwicky coined the term neutron-star and posited that these may be the result of SNe like SN 1572. To quote their original words; "*With all reserve we advance the view that a super-nova* *represents the transition of an ordinary star into a neutron star, consisting mainly of neutrons. Such a star may possess a very small radius and an extremely high density. As neutrons can be packed much more closely than ordinary nuclei and electrons, the "gravitational packing" energy in a cold neutron star may become very large, and, under certain circumstances, may far exceed the ordinary nuclear packing fractions. A neutron star would therefore represent the most stable configuration of matter as such."*(Baade and Zwicky, 1934a)

It is worth reflecting that the theory of Lev Landau (Landau, 1932) discussing the behaviour of stellar cores denser than WDs was not quite as well publicised, although his original article had appeared in 1932. It was only in 1938 (Landau, 1938) that Landau presented his result showing that a stellar core composed of only neutrons was possible. In 1939, Oppenheimer and Volkoff independently derived the minimum stellar mass required for the collapsed core of SNe to form a NS (Oppenheimer and Volkoff, 1939) based on the analytical results by Tolman (Tolman, 1934). However, both Landau, as well as Oppenheim and Volkoff were able to show for the first time that this mass was $\sim$1.5 $M_{\odot}$.

At about the same time as Chandrasekhar, Chadwick and Landau were presenting results and discoveries that would alter our ideas of stellar behaviour, Karl Jansky, a radio engineer with Bell Labs, was trying to characterise the various sources of noise plaguing the Trans-Atlantic radio communication systems. In doing so he 'serendipitously' discovered radio-frequency signals which originated from the Milky Way Galaxy itself (Jansky, 1933). This was perhaps one of the most significant events of this decade and would lead to remarkable changes in our perception of the Universe.

1.1.3 *The discovery of pulsars*

Jansky's discovery caused considerable excitement before the Second World War interrupted this period of rapid scientific growth. However, the heavy use of radio communications and radar technology during the war led to a large surplus of radio technology and researchers eager to apply their techniques to understanding the Universe. One of the subsequent, surprising discoveries was that the space between the planets was filled with a stream of particles, which appeared to be originate from and whose column-densities appeared to be linked to the activity of the Sun itself. Several groups started experiments to probe this inter-planetary medium and its interaction with the Solar System planets.

In early 1965, Anthony Hewish and his research group started building a large, low-frequency array at the Mullard Radio Astronomy Observatory of Cambridge University. The goal of this interferometric array operating at a centre frequency of 81.5 MHz was to study the properties of the diffuse medium between the Solar System planets using distant quasars as a reference source. This technique relies on the fluctuations of the radio flux of the quasars as the density variations of the inter-planetary medium pass in the foreground of the distant, stationary quasars and is called inter-planetary scintillation (IPS). In the same year, Jocelyn Bell, a fresh graduate student at Cambridge, joined Hewish's IPS team and along with a few other graduate students and interns, helped to construct the ~5 acre large interferometric array.

By July 1967, the telescope had started regular operations even though it would take several more months to complete its construction. Using pen-chart recorders and paper spools, the group collected several months[5] of data, among which Bell noticed the appearance of some rather unusual 'scruff' (Bell Burnell, 1977) consisting of several regular ticks with a period of about 1.337 s, each tick being about 0.3 s wide.

[5] Although Hewish's 1968 article in Nature used only ~3 hours of data from August, 1967.

These appeared to originate from an unknown source at a right ascension (R.A.) of $19^{h}19^{m}38(3)^{s}$ and a declination (DEC) of $22°0(3)'$ (Hewish et al., 1968) with a reference epoch of B1950 (Newcomb, 1895). Bell, who had already been trained by Hewish and others in identifying radio frequency interference (RFI) from terrestrial sources and had a deep and practical knowledge[6] of the behaviour of the array and its receiving setup was reasonably convinced this was of extra-terrestrial origin. Hewish, on hearing Bell's first report, was somewhat sceptical and believed these were some RFI which they had failed to identify. However, over the fall of 1967, Bell detected another source with a similar regular period of 1.2 s at an R.A. of $11^{h}33^{m}$. Finally in January 1968 Bell detected another couple of sources at R.A.s $8^{h}34^{m}$ and $9^{h}50^{m}$ which showed extremely regular 'pulsations'.

[6] Since the entire telescope was constructed by hand by Hewish's group, including Bell.

While this left little room for doubt that these signals were of astronomical origin, a much stronger confirmation was provided by looking at these sources with an independent telescope by Scott and Collins. Pilkington looked at the swept nature of the signals and calculated the 'dispersion'[7] of the source and placed it outside the Solar System but within the Galaxy.

[7] Dispersion is defined as the frequency-dependent quadratic delay introduced due to the fact that the photons are not travelling through free space but rather, through a rarified medium and therefore travel with a group velocity less than their free space velocity.

In February 1968, Hewish, with Bell and the others as co-authors presented their results in the now famous article in Nature (Hewish et al., 1968). The article made the correct inference that NSs and not WDs were the likely sources. However, Hewish et al. (1968) assumed these pulsations were linked to radial pulsations (Meltzer and Thorne, 1966) of the entire star leading to 'stellar-flares' occurring over the entire star, once per oscillation. As it turned out, this was not the correct interpretation. However, using this idea the term 'pulsars' was coined by the science correspondent of The Daily Telegraph [8]. Hewish was subsequently awarded the Nobel prize in Physics 1974 for the discovery of pulsars[9].

[8] "Pulsating Star Traced", by Dr. Anthony Michaelis, The Daily Telegraph, 5th March, 1968

[9] The prize was actually shared between Hewish and Sir Martin Ryle, who received it for his contributions to radio astronomy, specifically radio interferomtery.

1.2 Radio emission from pulsars and Pulsar Magnetospheres

While Deutsch in 1955 was the first to put forward a theory of the behaviour of strong magnetic fields associated with normal stars, two weeks before the remarkable discovery by Bell, Franco Pacini (Pacini, 1967) hypothesised that a magnetised NS rotating about a misaligned axis at the heart of the Crab nebula could output enough energy to explain the luminosity of the Crab nebula. He also suggested that the majority of this emission could be in the radio frequency regime.

Shortly after the announcement of the discovery of pulsars, using the fact that pulsar rotation periods are only observed to lengthen with time ($\dot{P} \equiv dP/dt < 0$, although rare increases in the spin-period do occur and are called glitches (see e.g. Alpar et al., 1981; Espinoza et al., 2011a; Lyne, 1999).) and rarely the other way round, Ostriker and Gunn (1969) presented a simple model for the spin-down process, modelling the NS as a magnetised body of moment of inertia I, rotating in vacuum with angular velocity Ω. It loses rotational energy due to the time-variation of its magnetic dipole moment vector μ, inclined with respect to the rotation axis by an angle α:

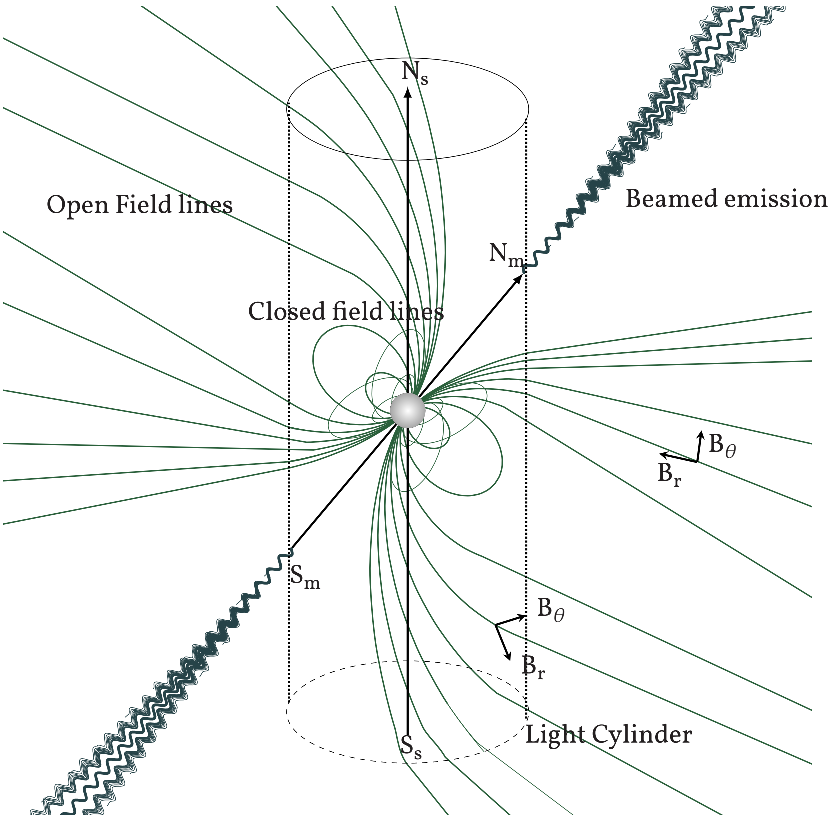

Figure 1.4: Sketch (not to scale) showing the primary features of the lighthouse model of the pulsar. The figure shows a pulsar indicated by the black sphere at the centre, which is inclined at 40.1 degrees with respect to the axis of rotation. The green lines show the magnetic field lines and dark blue lines show a possible scenario of beamed emission. Close to the surface of the pulsar, the magnetic fields are so intense that it 'rips' particles off of the surface and forces them to co-rotate with the pulsar magnetosphere, locked into the magnetic lines of force. At a distance $\vec{\mathbf{r}}_c$ where the trapped particles would need to travel faster than the speed of light to co-rotate with the magnetosphere we define the 'light cylinder' where co-rotation breaks down and particles are now free to escape at a relativistic velocity. While there is significant lack of agreement as to where the actual emission region resides, most modern models agree that along with beamed radiation, a stream of particles is continuously escaping outwards from the pulsar. This is known as the pulsar wind.

$$
\begin{aligned}
-\frac{d}{dt}\left(\frac{1}{2}I\Omega^2\right) &= \frac{2}{3c^3}|\ddot{\vec{\mu}}|^2 \\
&= \frac{1}{6c^3}B^2R^6\Omega^4 sin^2\alpha
\end{aligned}
\tag{1.1}
$$

which provides an estimate of the magnetic field (in units of Gauss) $B[\mathrm{G}] \approx 3.2\times 10^{19}\sqrt{P[s]\dot{P}}$ for a radius R = 10 km, I = 1×10^{45} g cm^2 and an orthogonal rotator, $\alpha = 90°$.

To overcome the shortcomings of this "dipole in vacuum" model, Goldreich and Julian (1969), presented a slightly more realistic model which attempted to explain the origin of the pulsar emission due to sparking across a massive potential that builds up due to charges being developed on the particles that are trapped in and rotate with the rotation of the magnetic field. While there are several criticisms of the Goldreich-Julian model and particularly the stability of the plasma structures, it is still useful as a starting point for understanding pulsar emission. A parallel development to the Goldreich-Julian model was that of Thomas Gold, who in 1968 (Gold, 1968) had argued that due to the strong magnetic fields and high rotational rates of NSs, any plasma in the surrounding magnetosphere will be forced to behave relativistically and lead to radiation in the pattern of a rotating beacon. The Gold (1968) model associates the radiation with neutral surfaces and invokes circulating bunches of particles as the origin of the emission.

Ruderman and Sutherland (1975) improved upon the work of Sturrock (1971) to present a model of the pulsar with a superfluid core, which is responsible for the generation of intense magnetic fields at the surface of the NS. This rotating magnetic field sets up an electric field and subsequently a potential drop that can be of the order of 10^{12} V. A polar magnetospheric gap is formed that spans the open field lines from the stellar surface up to an altitude of about 10 km. The scalar product of the electric and magnetic fields is non-zero in this gap, although it vanishes essentially everywhere else in the near magnetosphere. Electrons that lie in this region, shaded in Figure 1.5, are accelerated almost instantaneously. These electrons stay bound to the magnetic lines of force and emit curvature radiation as they move along it. When the electron has an energy of $\sim 10^{12}$ eV, the emitted photons have energies of several times $m_e c^2$ and can travel across the field lines. In doing so, they decay into electron-positron pairs almost immediately. As these newly generated particles lose energy once again through curvature radiation, they give off more photons. The positrons move out along the open field lines and electrons flow to the stellar surface to close the pulsar's homopolar generator circuit, as sketched in Figure 1.6.

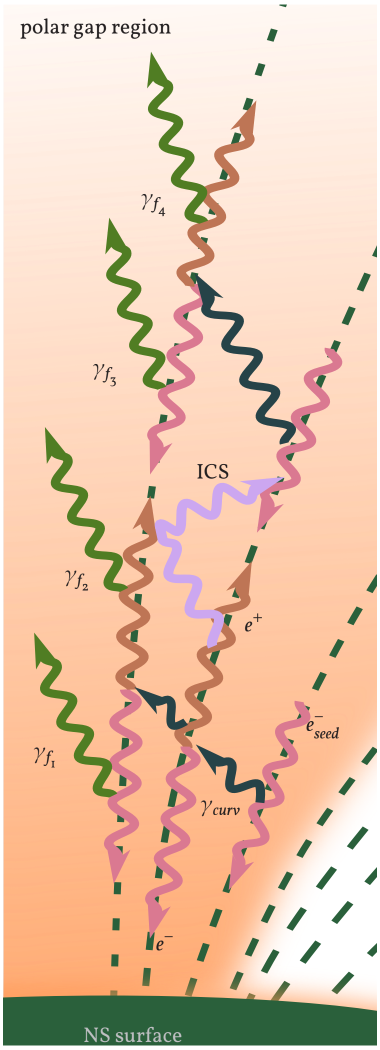

Figure 1.5: Schematic of pair cascades in the strongly magnetised field of an NS, showing the seed electron e^-_{seed} and other electrons and positrons moving along the magnetic lines of force, while photons generated from curvature emission create new pairs when they cross other lines of force. Also shown are inverse Compton scattering (ICS) processes which also add to the pair cascade process leading to pulsar emission. A final detail is the emission of different frequencies at different heights as per the radius-to-frequency mapped emission that is expected to occur near the polar gap of the NS.

The end result of this is a "pair cascade", in which a single seed electron generates a large numbers of charges. The flowing of all these charges shorts out the electric field that initiated the whole process, after which they travel relativistically out into space. Their high density makes maser activity favourable, so intense radio waves are generated. The cycle resumes after the charges have moved away.

In a series of papers, Cheng, Ruderman and Sutherland (Cheng et al., 1976; Cheng and Ruderman, 1977b,a; Cheng and Ruderman, 1979, 1980) introduced improvements for polar cap current flow, electron-positron production, and subsequent growth of a two-stream bunching instability. They also analysed the effect of the magnetoactive plasma on the subpulse polarization patterns presuming that the radiation originates from highly relativistic plasma streaming out along open field lines.

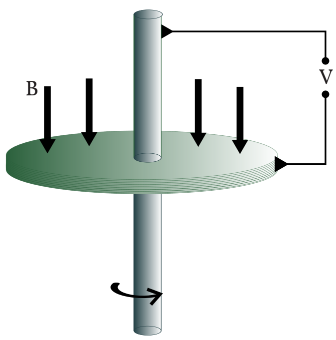

Figure 1.6: Schematic of a homopolar generator.

While the above model is perhaps most commonly discussed, alternative models like Michel (1973b) also exist which build on Gold (1968). Michel (1973b) constructs the so-called 'pulsar equation'; a function of the magnetic flux Ψ where the poloidal field is $B_p = \nabla\Psi \times \phi / r_c$ and the co-rotating electric field is defined as $E \equiv \frac{r_c \Omega}{c} B_p \times \hat{\phi}$ such that the force-free constraint becomes

$$(1 - x^2)\left[\frac{\partial^2 \Psi}{\partial x^2} + \frac{\partial^2 \Psi}{\partial z^2}\right] - \frac{1 + x^2}{x}\frac{\partial \Psi}{\partial x} = -I(\Psi)\, I'(\Psi) \qquad (1.2)$$

where $x \equiv r_c/R_L$ and $z \equiv z_c/R_L$ are scaled cylindrical coordinates, r_c is the radius of the light cylinder (see Figure 1.4) and $I(\Psi)$ is an unknown function proportional to the poloidal current enclosed by the flux surface Ψ. This is a nonlinear second-order elliptic equation with a regular singularity at the light cylinder $x = 1$. Analytical solutions for the monopolar field (Michel, 1973b) and a current-less, corotating dipole magnetosphere (Michel, 1973a) were established early on and the model was improved further by the work of Mestel et al. (1979); Mestel and Wang (1979) who introduced small gaps between ions and electrons along the surface at which the electric field falls to zero and Holloway and Pryce (1981) who introduced finite temperature limits on the vacuum gaps while Okamoto (1975) computed the magnetic field arrangement in the case of a non-corotating plasma. The problem of inertial particles was investigated by Scharlemann and Wagoner (1973) and Schmalz et al. (1979, 1980) while Beskin et al. (1983) obtained solutions for an arbitrarily inclined rotator. Numerical solutions for dipolar fields with currents were obtained for the condition where the last open field line corresponds to $\Psi_{open} = 1.36\Psi_{pc}$, where $\Psi_{pc} \equiv \mu/R_L$ is the flux through the polar cap for the unperturbed dipole field by Contopoulos et al. in 1999(Contopoulos et al., 1999). Independently, Ogura and Kojima (2003) and work by Spitkovsky and collaborators (See e.g., Spitkovsky and Arons, 2002; Spitkovsky, 2004) show numerical solutions exist for $\Psi_{open} = 1.66\Psi_{pc}$.

Apart from the first-principles approaches, there exists a large body of empirical work, (see e.g., Rankin, 2015, and references therein) which try to address the pulsar emission problem by investigating the properties of the emission beam. Such empirical theory currently is able to model the emission as the result of rotating dipolar regions embedded in the emission region, popularly known as the 'carousel' model. Radius-to-frequency mapping (RFM, see e.g., Cordes, 1978; Phillips, 1992, and others), where the pulsar emission originates at a specific height above the NS surface determined directly by the emitted frequency, has recently been shown to agree quite well with the model of Dyks and Rudak (2015) which is an improved version of the streaming model proposed by Michel (1987), although this solution does not explicitly rule out the carousel model or the patchy-cone model of Karastergiou and Johnston (2007) either.

While a single or unified solution to the pulsar problem is yet to be generally accepted, the field of pulsar magnetospheric research is evolving rapidly and has grown to include extremely diverse approaches. A more thorough review than is possible here should include approaches inspired by studies of plasma physics, as well as those involving true magnetohydrodynamic (MHD) solutions. The reviews by Spitkovsky (2008) and Pétri (2016) offer excellent summaries of pulsar magnetosphere research and the interested reader is directed to those works.

1.3 *Fundamental properties of pulsars*

The birth masses of pulsars can range from 1.28 $M_\odot$–~1.7 $M_\odot$ (see Tauris et al., 2012; Timmes et al., 1996; Özel et al., 2012) starting from progenitor stars of masses between 8 $M_\odot$–20 $M_\odot$. These are the suggested birth masses of pulsars that have been produced by the method described earlier in the text. Pulsars which are in binaries with other stars do not follow the same channel of formation and can have greater masses while there are indications that at least some pulsars may even be born as more massive objects. If the binary companion is another NS, the expected birth mass is 1.33(5) $M_\odot$ and for pulsars with lighter companions this can be higher, at 1.48(20) $M_\odot$ (Özel et al., 2012).

The gravitational collapse leading to the birth of the NS forces them to be born with high rotation rates and the majority of NSs are believed to be ejected from their site of formation as a reaction to any asymmetry in the SN. This is also called the 'natal kick' effect, which is believed to be responsible for the high velocities associated with the majority of known young pulsars.

NSs are not perfect, solid spheres and therefore, neither are pulsars. The average mass density of the core of the NS is about (Lorimer and Kramer, 2005):

$$\langle \rho \rangle = 6.7 \times 10^{14}\ \mathrm{g\,cm^{-3}} \tag{1.3}$$

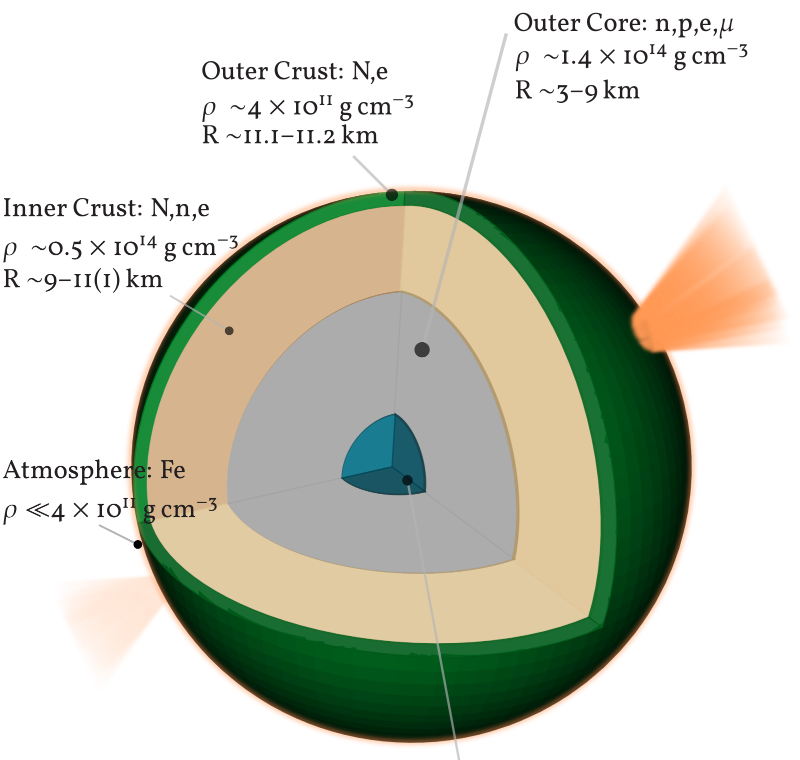

Figure 1.7: Schematic of the structure of a NS following Chamel and Haensel (2008). Fe denotes elemental iron, while N stands for nuclei while n,p and e represent neutrons, protons and electrons, respectively.

This turns out to be higher than even the density of nuclear matter, 2.7×10^{14} g cm^{-3}. Given our lack of experimental knowledge of such extreme states of matter the values presented in the following text should be treated as representative only, since they are dependent on models of the structure of NSs and there are significant differences among competing models (see e.g. Chamel et al., 2015; Becker, 2009, and references therein).

Near the surface, a thin atmosphere surrounds the NS. The density of this atmosphere is about 10^6 g cm^{-3} and it is composed mainly of elemental Helium and Hydrogen or iron depending on the model being used. The outermost region is a thin crust composed mainly of iron nuclei and a sea of degenerate electrons, which carries about 1.4 percent of the total inertia and extends for ~100 to 300 m. The next layer, the inner crust, has a density of $\sim 0.5 \times 10^{14}$ g cm^{-3} and spans about 0.7 to 3 km in thickness. A few hundred meters inside the inner crust, the density rises to above 4×10^{11} g cm^{-3} to 8×10^{11} g cm^{-3}, at the point known as the neutron drip point. After this the outer core dissolves fully at 2×10^{14} g cm^{-3} into a neutron superfluid and a small percentage (~5 percent) of superconducting electrons and protons as well as a small fraction of muons. The extremely high density forces the material to occupy a number of specific geometries, where the neutron and proton superfluids are forced into non-spherical localisations resembling rods, tubes, bubbles and sheets which are often colloquially referred to as the *nuclear pasta* phase (see e.g., Chamel and Haensel, 2008). Beyond this layer the density increases even further until the only states of matter that can possibly exist are exotic forms of matter, which may consist of a pion or kaon condensate or even quark matter at the inner core of the NS.

1.3.1 *Spin periods and spin-down*

Pulsars are known to spin with frequencies as high as 716.35 Hz (Hessels et al., 2006) to about as slow as 0.0848 Hz(Dib and Kaspi, 2014) and represent some of the most precisely measured astronomical quantities, with examples of attosecond level precision (See, e.g., Verbiest et al., 2008, for PSR J0437−4715)[10]

Just as incredible as the precise measurement of spin frequencies is perhaps the spread of the rate at which they decay. Figure 1.8 shows a plot of the spin period plotted as a function of the time rate of decay of the spin period, commonly called the P–$\dot{P}$ diagram. This figure encodes a wealth of information connected to the possible evolution of pulsars. Several kinds of clustering can be identified, with slow period pulsars [11] forming a large group towards the middle and the right of the plot and a smaller group of rapidly rotating pulsars forming an island of MSPs in the bottom left.

[10] This high precision is a property that makes pulsars some of the most accurate 'celestial clocks' and lets us measure incredibly small variations in the propagation path and is hoped will lead to independent detection and measurement of gravitational waves from some of the most massive black holes believed to exist, the fundamental technique for which is described in section 2.5.1.

[11] which I refer to as 'classical' pulsars throughout

1.3.2 *Spin-down luminosity*

As the pulsar rotates, it dissipates energy at a rate

$$\dot{E} = -\frac{dE_{rot}}{dt}. \tag{1.4}$$

by converting rotational energy into radiation. This is called the spin down luminosity of the pulsar. If the radiation is assumed to be generated due to the rotation of a purely dipolar magnetic field, we can quantify this in terms of the spin period and the time derivative of the

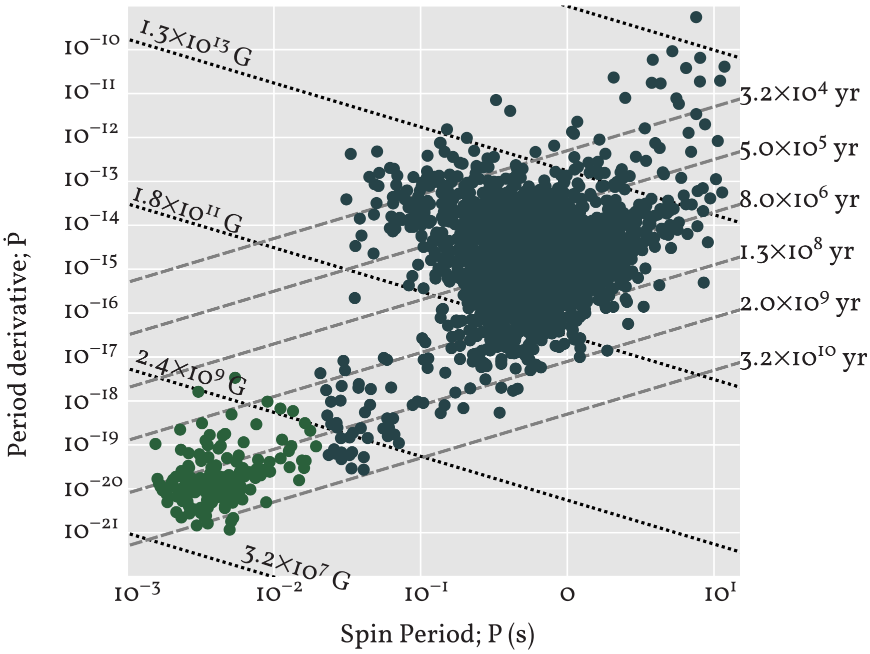

Figure 1.8: P − Ṗ diagram for all currently known pulsars. Dark blue points mark the classical pulsars and green, the MSPs. The gray dashed lines are lines of constant age, calculated from eqn. (1.12) while the black dotted lines mark constant surface magnetic field strengths estimated from eqn. (1.16).

spin period as (Lorimer and Kramer, 2005)

$$\dot{E} = -\frac{d(I\Omega^2/2)}{dt} = -I\Omega\dot{\Omega} = 4\pi^2 I\dot{P}P^{-3}. \tag{1.5}$$

where $\Omega = 2\pi/P$ is the angular frequency. Assuming the moment of inertia is 10^{45} g cm^{-3}, the spin-down luminosity is about (Lorimer and Kramer, 2005)

$$\dot{E} = 3.95 \times 10^{24}\,\mathrm{J\,s^{-1}} \left(\frac{\dot{P}}{10^{-15}}\right)\left(\frac{P}{s}\right)^{-3}. \tag{1.6}$$

1.3.3 *Braking index*

By restricting the pulsar to have a purely dipolar magnetic field, we can write, for a magnet of dipole moment $|\mu|$ (Lorimer and Kramer, 2005)

$$\dot{E}_{dipole} = \frac{2}{3c^3}|\mu|^2\Omega^4 sin^2\alpha, \tag{1.7}$$

$$\dot{\Omega} = \frac{2}{3Ic^3}|\mu|^2\Omega^3 sin^2\alpha. \tag{1.8}$$

If we were to express this in terms of a rotation frequency, $\nu = 1/P$ and write it as a power law, we would get

$$\dot{\nu} = -k\nu^n \tag{1.9}$$

where n is the *braking index* and $k = \frac{2}{3Ic^3}|\mu|^2 sin^2\alpha$.

If we were to take a second derivative of ν we could write the expression for the braking index as

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2} \tag{1.10}$$

If the moment of inertia is constant, then for purely dipolar radiation the canonical estimate of n is 3. However, for the few studies where significant measurements of n have been possible, these values do not converge to 3 (see table 1.1).

Source	n	Reference
J0534+2200 (Crab)	2.51(1)	Lyne et al. (1993)
J0540-6919	2.140(9)	Livingstone et al. (2007)
J0835-4510 (Vela)	1.4(2)	Lyne et al. (1996)
J1119-6127	2.684(2)	Weltevrede et al. (2011)
J1513-5908	2.839(1)	Livingstone et al. (2007)
J1734-3333	0.9(2)	Espinoza et al. (2011b)
J1833-1034	1.8569(10)	Roy et al. (2012)
J1846-0258	2.65(10)	Livingstone et al. (2007)
	2.16(13)	Livingstone et al. (2011)

Table 1.1: Braking indices of 8 pulsars.

1.3.4 *Characteristic age and spin period at birth*

FROM eqn. (1.9) and eqn. (1.10), we can write the time derivative of the period of rotation as

$$\dot{P} = kP^{2-n} \tag{1.11}$$

This gives the age of the pulsar, assuming a *spin period at birth* P_0 and that the spin-down is described by a simple continuous function

$$T = \frac{P}{(n+1)\dot{P}}\left[1 - \left(\frac{P_0}{P}\right)^{n-1}\right]. \tag{1.12}$$

From this we can now define a *characteristic age* for the pulsar as

$$\tau_c \equiv \frac{P}{2\dot{P}} \hat{=} 15.8 Myr \left(\frac{P}{s}\right)\left(\frac{\dot{P}}{10^{-15}}\right) \tag{1.13}$$

where we have assumed that $P_0 << P$ and n=3 for a spin down due to magnetic dipole radiation only.

As a consequence of the conservation of angular momentum, the gravitational collapse of the core of the SN progenitor forces the resulting NS to rotate very rapidly. This birth period is estimated to be much higher than the observed spin period of the pulsar. This is can be estimated if the true age of the pulsar is known (e.g., from SN association) and a braking index has been measured for the pulsar. Substituting eqn. (1.13) in eqn. (1.12), we obtain

$$P_0 = P\left[1 - \frac{(n-1)}{2}\frac{T}{\tau_c}\right]^{\left(\frac{1}{n-1}\right)} \tag{1.14}$$

This implies that the spin period at the birth of the Crab pulsar is ~19 ms (Lyne et al., 1993). Theoretical considerations suggest that for any given pulsar this must lie between ~ 11 ms to 150 ms (see e.g., Faucher-Giguère and Kaspi, 2006).

1.3.5 *Surface magnetic field strength*

While it is not possible to obtain direct measurements of the surface magnetic field strength of a pulsar, under the approximation of a purely dipolar magnetic field, we can rewrite eqn. (1.8) to obtain an expression for the surface magnetic field strength

$$B_s \equiv B_{(r=R)} = \sqrt{\frac{3c^3}{8\pi^2}\frac{I}{R^6 sin^2\alpha}P\dot{P}}, \tag{1.15}$$

For an assumed I=10^{45} g cm^2, a radius of 10 km and assuming that α = 90°, the surface magnetic field becomes

$$B_s = 3.2 \times 10^{19} G\sqrt{P\dot{P}} \simeq 10^{12} G\left(\frac{\dot{P}}{10^{-15}}\right)^{1/2}\left(\frac{P}{s}\right)^{1/2}. \tag{1.16}$$

given the spin period P and spin down $\dot{P}$ of the pulsar.

1.3.6 True radius of a pulsar

Due to the strong gravitational field of the NS, the observed radius of a pulsar is larger than the intrinsic radius and therefore, the surface temperature appears to be lower than the intrinsic value. The observed radius, R_{obs} and the intrinsic radius, R, are related by

$$R_{obs} = \frac{R}{\sqrt{1 - GM/Rc^2}} = \frac{R}{\sqrt{1 - R_s/R}} \qquad (1.17)$$

where R_{obs} is the observed radius and R_s is the Schwarzchild radius given by

$$R_s = \frac{GM}{c^2} \simeq 4.2 \frac{M}{1.4\,M_\odot} \qquad (1.18)$$

Assuming that the NS has a mass of $1.4\,M_\odot$ (see e.g., Lattimer and Prakash, 2004) we can limit the minimum radius to

$$R_{min} = 1.5R_s = 3\frac{GM}{c^2} = 6.2 \times \frac{M}{1.4\,M_\odot}\,\mathrm{km} \qquad (1.19)$$

While the simplified discussion presented above touches upon some of the fundamental properties of pulsars, detailed summaries can be found in Lorimer and Kramer (2005) or Lyne and Graham-Smith (2012) and other specific reviews mentioned above.

1.4 Recycled pulsars

Much of the earlier discussion has focused on the more commonly found class of pulsars, the classical pulsars. These are usually objects which do not have gravitationally bound companions and typically have rotation periods $\sim 10^{-1}$ s to 10 s. If the pulsar is gravitationally bound to a companion then it is said to be in a binary system. If the companion is another massive star, then it must also follow a similar evolutionary route as the pulsar's progenitor. Evidently, this evolution is affected by the properties of both the stars and the end results might also differ, depending on the manner in which both the stars evolve and interact. In most cases the end result of the evolution of the binary results in at least one rapidly rotating pulsar, along with a lighter companion. In some special systems, the companion might also evolve into a pulsar, e.g., PSR J0737+3039 or a radio-quiet NS. These double neutron star (DNS) systems typically possess spin periods of ~30 ms–100 ms. If the spin period of the pulsar is less than ~30 ms it is called an MSP, although this limit is not well-defined.

The Australia Telescope National Facility (ATNF) pulsar catalogue[12] (Manchester et al., 2005) shows the current number of known MSPs to be ~321 of which 195 are located within the Galaxy and 126 are found in globular clusters. These are plotted in figure 1.9, where green dots show the projected positions of MSPs and the red crosses mark the positions of classical pulsars, estimated from the Taylor and Cordes (1993) model of free electron distribution in the Galaxy. Of the 321 MSPs, about 66 appear to lack a companion.

[12] www.atnf.csiro.au/research/pulsar/psrcat, ver. 1.54

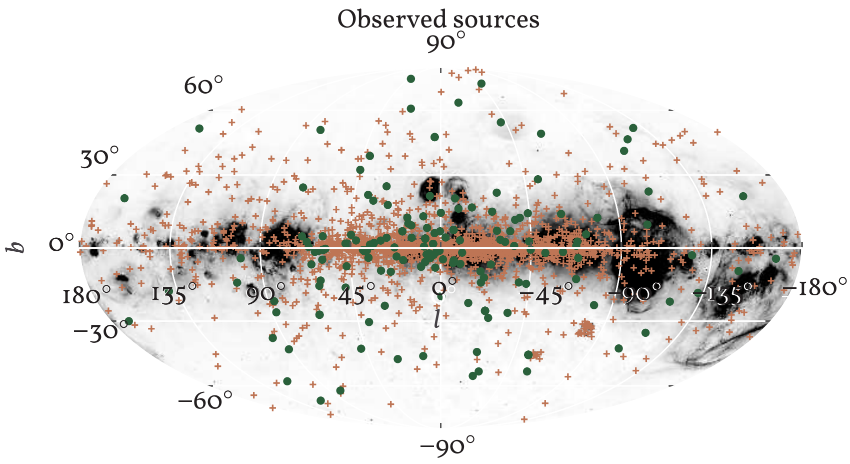

Figure 1.9: An all-sky plot of all the pulsars currently known, from the pulsar catalogue (Manchester et al., 2005). Brown '+' symbols mark the locations of young pulsars, in galactic coordinates and filled green circles mark the positions of the known MSPs. The positions are overlaid on a Mollweide projection of three merged H-α surveys. The H-α data (Finkbeiner, 2003) were obtained from the NASA/Goddard Legacy Archive for Microwave Background Data (LAMBDA) data products page.

1.4.1 *The process of recycling pulsars*

Although the first MSP to be discovered, PSR J1939+2138 (Backer et al., 1982), was an isolated system it is believed that all MSPs originate in binary or tertiary systems, where the pulsar accretes matter from the (inner) companion. Infalling matter from the companion transfers angular momentum to the pulsar once the surface magnetic field of the rotating NS falls low enough to allow efficient accretion, initially prevented due to either the magnetodipole radiation pressure or propeller effects (Illarionov and Sunyaev, 1975). The accretion torque acting on the spinning NS is due to a dominant material term, a magnetic term and a viscous stress term (Ghosh and Lamb, 1992; Shapiro and Teukolsky, 1983, etc.).

The exchange of angular momentum at the magnetospheric boundary eventually increases the spin angular momentum of the NS, which depends on the acting torque as

$$N \approx \sqrt{GMr_A}\dot{M}\upsilon \tag{1.20}$$

where $\upsilon \simeq 1$ is a numerical factor which depends on the flow pattern (Ghosh and Lamb, 1978, 1992) and

$$\begin{aligned} r_A &\simeq \left(\frac{B^2R^6}{\dot{M}\sqrt{2GM}}\right)^{2/7} \\ &\simeq 22\,\mathrm{km}\cdot B_8^{4/7}\left(\frac{\dot{M}}{0.1\dot{M}_{\rm Edd}}\right)^{-2/7}\left(\frac{M}{1.4\,\mathrm{M}_\odot}\right)^{-5/7}, \end{aligned} \tag{1.21}$$

is the Alfvén radius[13](see e.g., Pringle and Rees, 1972), a typical value for which could be ~ 40 km assuming a surface magnetic field of 1×10^8 G and a mass-loss rate of $0.01\dot{M}_{\rm Edd}$[14]. This process is known as recycling (Bhattacharya and van den Heuvel, 1991; Tauris and van den Heuvel, 2006).

[13] Defined as the location at which the infalling material couples with the magnetic field lines emanating from the NS magnetosphere and co-rotates with it.

[14] $M_{\rm Edd}$ is called the Eddington mass, the theoretical maximum mass a star or accretion disk can have before its luminosity begins to blow away the outer layers.

1.4.2 *The companions of recycled pulsars*

While there are many classifications in use for MSPs, the most common bases are the mass and the nature of the companion, which can include degenerate stars like Helium white dwarfs (He-WDs) and Carbon-Oxygen white dwarfs (CO-WDs), semi-degenerate brown dwarfs or non-degenerate low-mass dwarf stars or gas-giant planets. Of these the last two can often experience severe mass loss and ablation by the pulsar wind, leading to the formation of the so-called black-widow pulsar (BWP) and red-back pulsar (RBP) systems (Roberts, 2013; Chen et al., 2013).

X-ray binaries are believed to be the precursors of all MSPs with low mass X-ray binaries (LMXBs) forming the bulk of the progenitors and low mass X-ray binaries (IMXBs) leading to Carbon-Oxygen/Oxgen-Neon-Magnesium white dwarf (CO/ONeMg-WD) companions. Low mass X-ray binaries (HMXBs) lead to the formation of slower DNS systems. Figure 1.10 shows a schematic of the many possible paths by which an MSP binary may be produced.

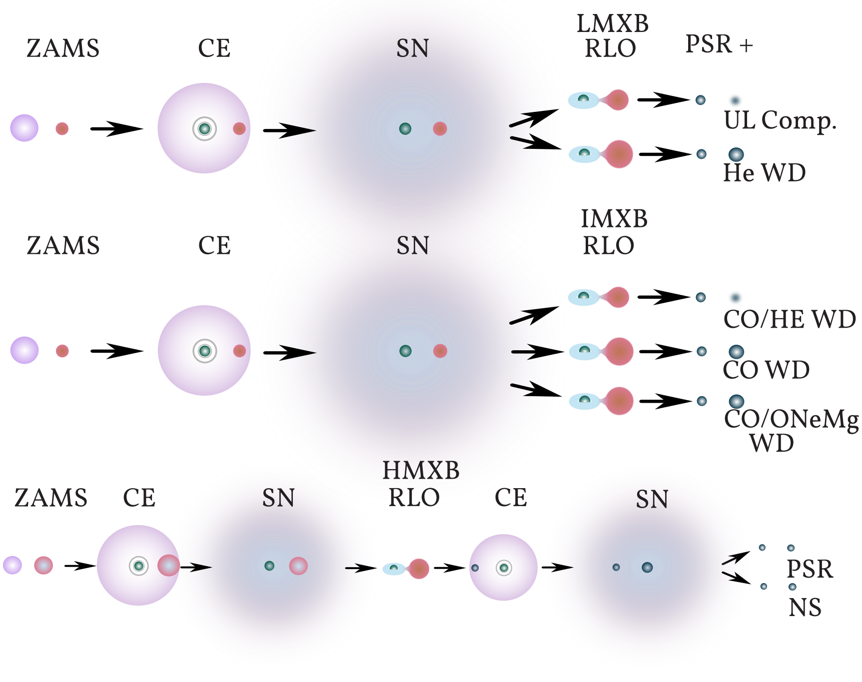

Figure 1.10: Schematic showing the formation of an MSP binary starting from a binary of zero-age main-sequences (ZAMSs), one of which is enters the red-giant phase first and expands till its outer layers encompass both the stars in what is known as the common envelope (CE) phase, after which the bloated star explodes in a SN, leaving behind a pulsar (PSR) and a companion orbiting each other around a common centre of mass. The pair now form an X-ray binary which is classified as a low-, intermediate- or high-mass X-ray binary (L/I/HMXB) as the second star bloats up and starts to undergo Roche-lobe overflow (RLO). This leads to an accretion disk around the pulsar, which forces it to rotate faster via transfer of angular momentum via accretion onto the pulsar. Depending on the progenitor mass of the companion leaves behind an evolved companion which may be ultra-light (UL) or a Helium, Carbon-Oxygen or Oxygen-Neon-Magnesium WD (He/CO/ONeMg WD. In special cases, the evolution can also produce a double-pulsar binary like PSR J0737+3039 or a radio quiet NS-pulsar binary like PSR J1915+1606 (B1913+16). Figure following Tauris (2011). Sizes and distances are not to scale.

Of these three 'channels' of MSP binary evolution, the first leads to several exotic systems. The progenitors in this case are typically two ZAMS stars, one of which has a mass $8\,M_{\odot}$–$20\,M_{\odot}$ while the second is typically much lighter, $M \leq 8\,M_{\odot}$. The more massive companion exhausts nuclear fuel first and bloats to the CE state where the stellar material of the bloated star envelops the companion. It subsequently undergoes a SN and produces a classical pulsar or NS. The less massive star, after it has ended its nuclear burning phase, bloats up and undergoes RLO which causes matter to be accreted onto the NS and forces it to rotate faster by transferring angular momentum. During this phase the binary emits copious amounts of X-rays and γ-rays and forms an LMXB. Finally, the stellar material in the companion is significantly depleted so that it can no longer overflow and the companion contracts to produce an UL companion. If the remaining mass is sufficient, the resulting star then forms an ultracool He-WD. In the special cases where the companion has too little mass to produce a He-WD, either due to excessive mass loss via accretion onto the NS or due to the outer layers of the companion being stripped away by the pulsar wind, the resulting star is a semi-degenerate object, possibly a 'brown' dwarf star. If these binaries have very short orbital periods and the mass of the companion is $\leq 0.05\,M_\odot$ then the system is called a BWP system while a more massive companion, between $0.1\,M_\odot$–$0.5\,M_\odot$, leads to the system being labelled a RBP system. Although the recycling process is believed to be the most likely process for producing MSPs, it has been argued that there may be cases where pulsars are born as MSPs (see e.g., Tauris and Takens, 1998; Miller and Hamilton, 2001; Freire et al., 2008). While conventional formation scenarios do not favour the formation of such objects, the mounting observational evidence against the production of isolated MSPs like PSR J1939+2134 via the companion evaporation route provides some support to this proposed route.

1.4.3 *Too little, too varied?*

RPs differ from their classical counterparts in a number of ways, apart from their intrinsically higher rotation rates, which range from 1.4 ms–185 ms (Hessels et al., 2006; Swiggum et al., 2015, respectively). Perhaps the most significant difference lies in the masses of the pulsars in binary systems. Unlike the limited spread expected for isolated pulsars, MSP masses appear to be distributed over a fairly wide range, from $1.17(1)\,M_\odot$ to $2.01(4)\,M_\odot$ (Martinez et al., 2015; Antoniadis et al., 2013, respectively) as shown by observations. Empirical evidence also suggests that the surface magnetic field strengths of MSPs are much lower than for young pulsars, although the exact process by which the pulsar loses its magnetic field is not understood very well (Bhattacharya, 2002). At the same time, MSPs are also more stable rotators as compared to their classical counterparts, although recent evidence (Cognard and Backer, 2004; McKee et al., 2016) shows more of the MSPs may experience glitches than previously thought.

While a large number of MSPs have been discovered in the last two decades, the expected number in the Galaxy is highly uncertain and successive surveys have met with less success than originally predicted. The discovery of MSPs has been traditionally burdened by two major factors. The first relates to the high computing cost which is the result of the lengthy Fast Fourier Transforms (FFT, Cooley and Tukey, 1965) that must be constructed to search for objects with such short spin periods. The second major factor has been the typically small-number statistics that were available for estimation of the properties of MSPs. The first hurdle has been significantly addressed with the advent of high speed computing infrastructure being available for significantly lower costs. The second hurdle can only be addressed through greater studies of MSPs themselves, part of which would involve studying the properties and behaviour of these systems, as carried out in chapter 4 while another part would involve studying the entire population (and sub-populations) of MSPs to estimate their emission characteristics (i.e., their spectra), as carried out in chapter 5 to allow future surveys to predict discovery numbers more accurately.

1.5 *Structure and organisation of this thesis*

Having reviewed the properties and formation of classical pulsars and their older counter-parts, MSPs, we proceed in the next chapter to introduce the three aspects of pulsar astronomy that form the body of this thesis. Specifically, in Chapter 2 I introduce some fundamental concepts related to pulsar astronomy. I discuss the techniques and tools used to record pulsar observations, the properties of the data and their analysis. I introduce the concepts of flux and phase calibration for pulsar data and discuss the technique of generating pulse times of arrival from pulsar observation. Finally, I review the technique of pulsar timing.

In Chapter 3 I present an overview of the most important artifacts that can be found in software based data recording systems. I review the theory of polyphase filterbanks and apply a least mean square optimisation scheme to the special case of a coherent dedispersion pulsar backend. I show that the action of the dedispersion filter is non-linear in phase and therefore the resultant signal cannot be reconstructed perfectly by a full reconstruction filterbank based on short, linear phase finite impulse response filter (FIR) filters. I use results from a simple *Python* simulation to demonstrate the most significant artifacts that are present in the subbands of analysis filterbanks, which are commonly used in pulsar data recording systems. I also review *Matlab* based modelling for the Square Kilometre Array that has shown that the ultimate limit on the accuracy of MSP timing due to an artifact with relative power ~-34 to -37 dB is about ~100 ns.

In Chapter 4 I present an updated pulsar timing solution using data from four of the five European pulsar timing array (EPTA) telescopes for the PSR J2051−0827, which was the second BWP to be discovered. For this project I processed the data and carried out preliminary investigations to ensure the resulting timesof arrival were free from systematic errors. I then performed the analysis using the technique of pulsar timing. This work resulted in the publication – Shaifullah et al. (2016).

In Chapter 5 I measure the spectral indices of 12 MSPs using data collected with the 300-m radio telescope at the Arecibo Observatory. For this project I carried out simulations to estimate the effect of interstellar scintillation. I also performed all the observations, the radio frequency interference removal and subsequent post-processing of the data. I calibrated the data and obtained flux-density measurements using code from my collaborators. I implemented robust statistics to estimate the spectral indices. I rederived spectral indices for an additional 19 sources for which only flux densities were available in literature to increase the sample size of MSP spectral indices to 74, which is almost two-thirds of the known Galactic population of 195.

In Chapter 6 I summarise the main results presented and offer suggestions for future work.

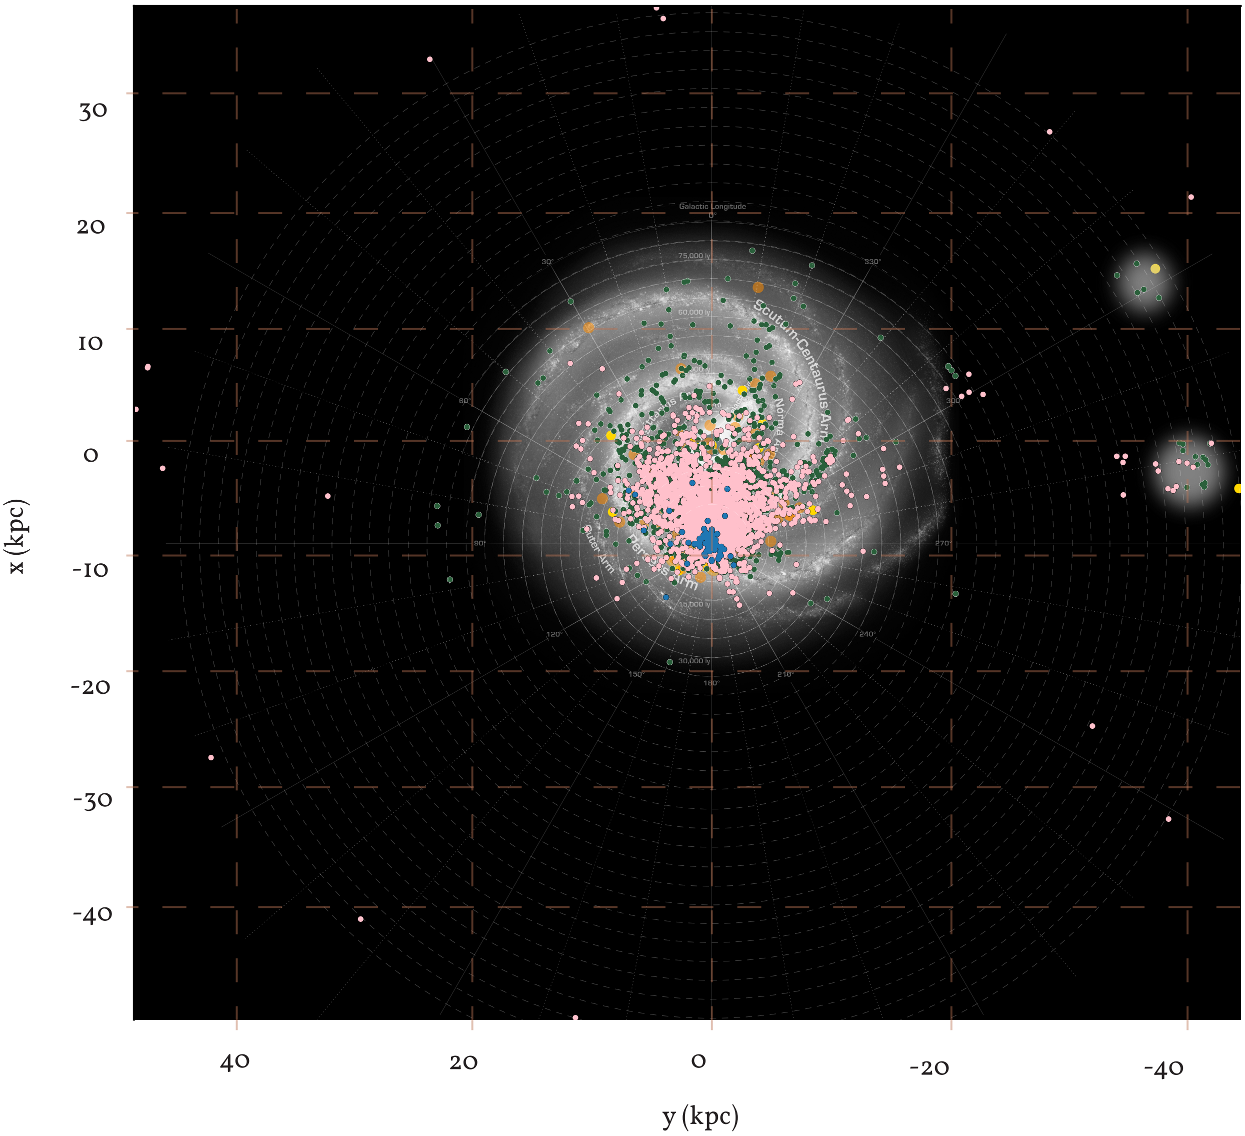

Plate 1: A plot of all pulsars currently known, from the pulsar catalogue (Manchester et al., 2005) projected on the disk of the Milky Way Galaxy. Green dots denote pulsar positions estimated using the YMW17 (Yao et al., 2017) model of the distribution of electrons in the Galaxy, pink dots denote pulsar positions estimated using the NE2001 model (Cordes and Lazio, 2003). The disagreement between the two models is distinct and serves to remind us of the large uncertainty that must be accounted for distances estimated from measured dispersion measure (DM) values.
Blue dots indicate pulsar distances estimated from parallax measurements (Verbiest et al., 2010, 2009) and have been updated with values from Shami Chatterjee's pulsar parallax catalogue. Golden dots represent known magnetar positions retrieved from the McGill Magnetar Catalog(Olausen and Kaspi, 2014) and orange dots indicate known SNRs retrieved from the Chandra SNR Catalog.
The Sun is located 8.3 kpc (Reid et al., 2014) from the galactic centre, at the origin of the Galactic coordinate system shown by gray lines. Dashed gray circles mark Galactic distances in units of 1.5 kpc while the grey lines mark Galactic longitudes. Brown dashed lines show distances in kpc on the disc.
Background Image Credits: Robert Hurt, Spitzer Science Center/NASA, retrieved from the Spitzer mission pages.
The original image has been modified and custom annotations added.

2

Pulsar astronomy

春過ぎて
夏来にけらし
白妙の
衣ほすてふ
天の香具山

The earth, that is sufficient,
I do not want the constellations any nearer,
I know they are very well where they are,
I know they suffice for those who belong to them.

— **Walt Whitman;** *Song of the open road 1*

Pulsars are extremely energetic emitters and are easily observed in the radio regime. However, most often the only visible component of the pulsar is the rotating pulsar beam. The process of deriving pulsar parameters from these beams crossing the line of sight to the pulsar from the earth begins with observations made with some of the most sensitive radio telescopes currently available. The incident electromagnetic radiation must be converted into an electrical signal which is then converted into a digital data file. The data file is finally analysed and the resulting measurements are modelled to convert the beam crossings into meaningful measurements of the pulsar's properties. In this chapter I introduce the technical aspects of pulsar observations and, I present details of the process of recording and calibrating the pulsar data and finally, I explain the process of pulsar timing.

2.1 *Instrumentation for observing pulsars*

The primary instrument for observation in the radio regime is the radio telescope, designs and descriptions of which are disparate and wide-ranging (See e.g., Wilson et al., 2013). A sketch of a generic setup for pulsar observation is shown in plate 2.

The telescope typically consists of a large electromagnetic (EM) reflector arranged in a parabolic shape (or any convenient section through a spheroidal shell) at whose focus a 'horn' antenna is mounted. At low enough frequencies, such parabolic arrangements are substituted for by arrays of simple 'wire' antennas. Depending on the geometry of their cross-sections, these feed horns can be classified as rectangular or cylindrical horns. The horn itself consists of an impedance transfer section (the 'flare' and ridges together in figure 2.1) which increases the efficiency of the coupling of the incident EM waves to the antenna, a waveguide section where the incident wave produces a resonant standing-wave and a single or a pair of orthogonally mounted probes, as shown in the lower part of figure 2.1.

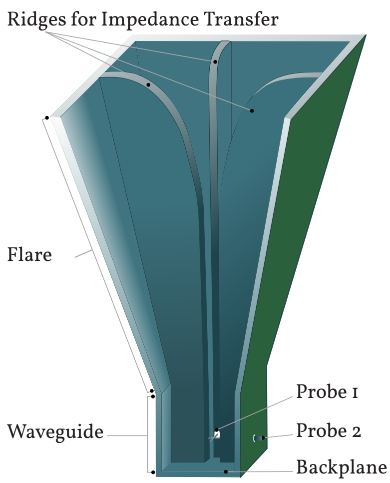

Figure 2.1: Cut-away diagram of a dual-polarisation, quad-ridge rectangular horn antenna for radio astronomy showing the arrangement of the ridges and the orthogonal probe arrangement to recover two polarisations.

The induced electric currents (or their corresponding voltages) are amplified and filtered to remove signals from unwanted frequencies. It is almost standard to then combine this filtered signal with a locally

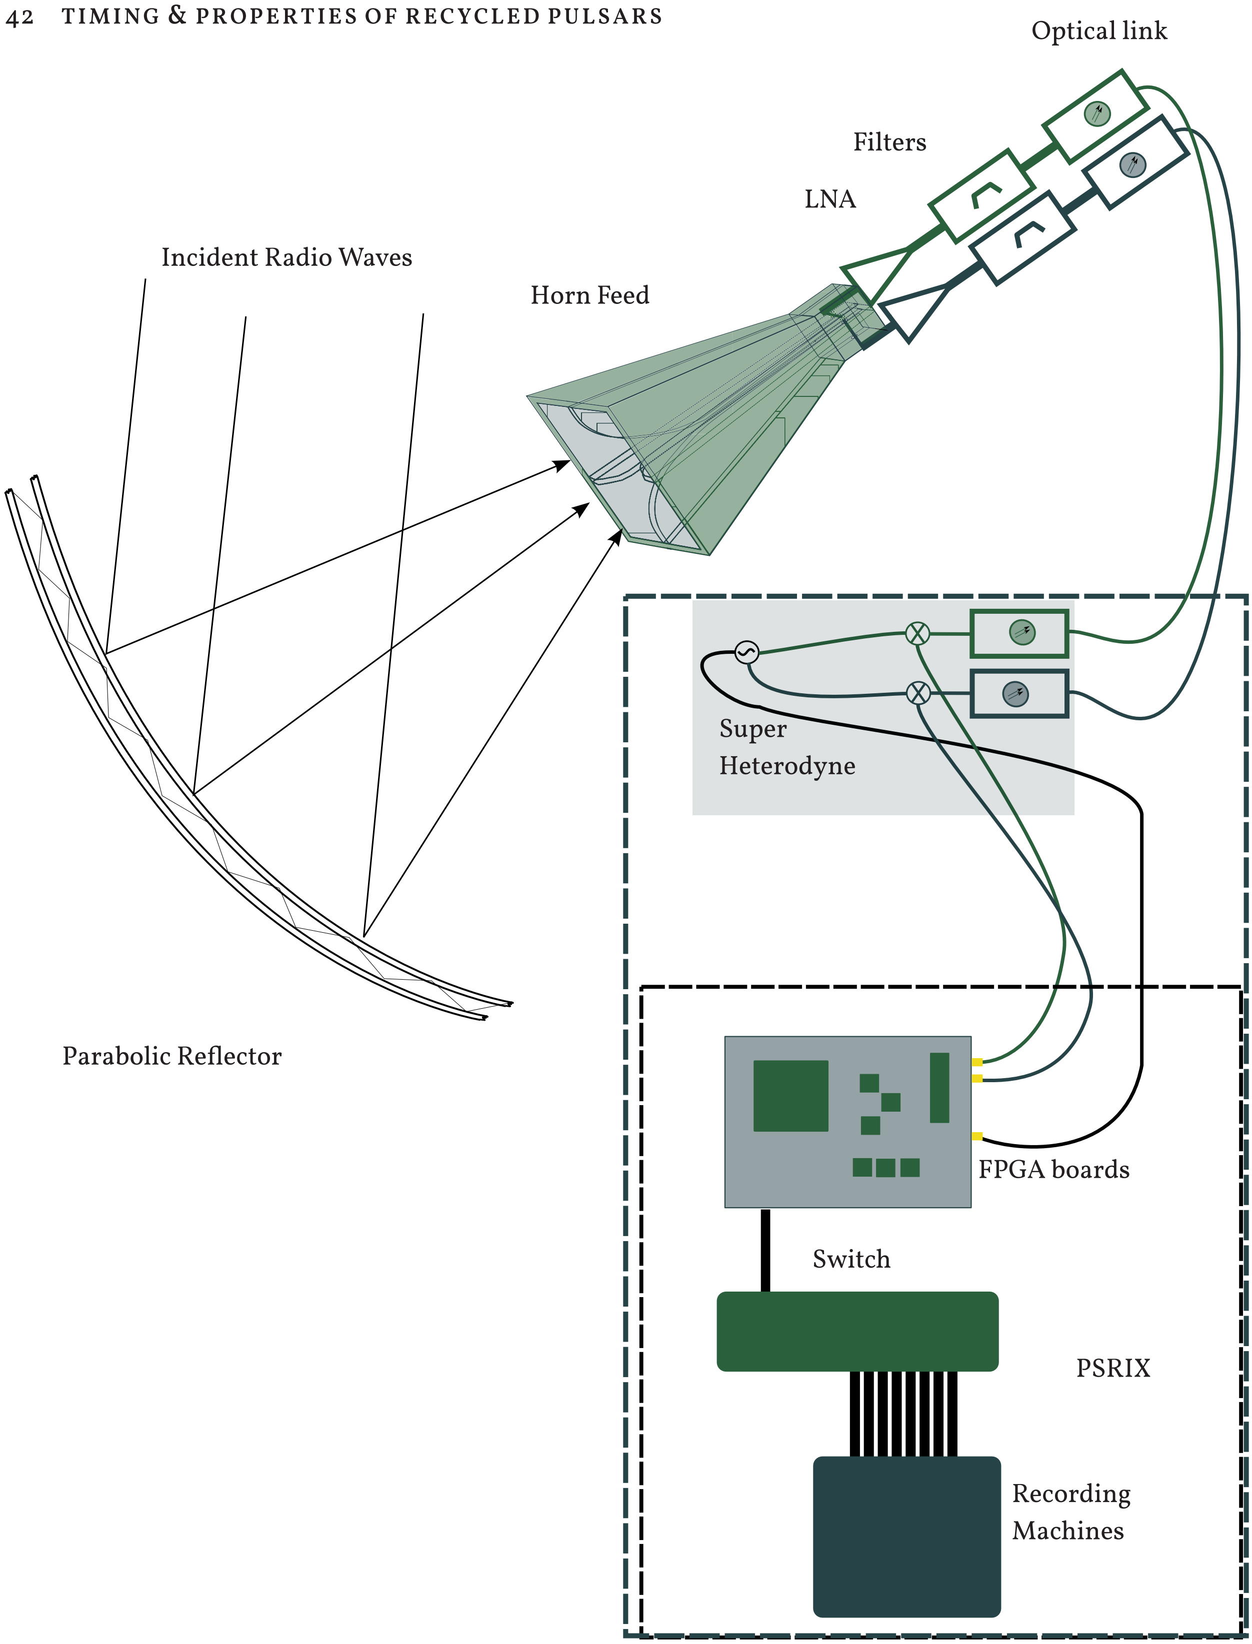

Plate 2: Schematic of a modern radio telescope with a specialised backend for pulsar astronomy. Incident radio waves from a distant pulsar (or any radio source) are collimated by a parabolic reflector, also called a dish, onto a 'feed'. Typically, for large telescopes like the Effelsberg 100-m radio telescope the feed is a horn antenna. The horn antenna converts the incident wavefronts of the radio waves from the pulsar into a standing wave which excites a pair of mutually orthogonal probes (see figure 2.1). Each probe corresponds to a single polarisation. The current induced on these probes is then amplified by a low noise amplifier (LNA) before being filtered and further amplified. It is then converted and transmitted over an optical link to receiver boards at the ground level, where the signal is mixed with a local oscillator signal to convert the high frequency signals into lower frequencies in what is called the super heterodyne method. After some subsequent processing the signal is digitised and transferred to field programmable gate array (FPGA) boards which split the received bandwidth into smaller parts and transmit the split subbands over a high speed network link to a number of computers where the data are further processed and recorded onto hard drives.

generated, stable, monotonic frequency such that the phases of the two signals are added coherently. This is called frequency mixing.

The resulting signal consists of two components with a frequency that is either the sum or the difference of the input signal and local signal. One of these two components, typically the lower frequency, is retained and the other is filtered out before further signal-conditioning and digitisation via an analog-to-digital convertor (ADC). This is known as the super-heterodyne architecture which allows us to observe at very high frequencies using devices which natively work at much lower frequencies. The entire arrangement from the horn up to the ADC in plate 2 is often called the 'front-end'. The frontend, which is usually located very close to the horn is followed by a backend housed away from the telescope, where the signal is passed through further processing stages before being formatted, standardised and then recorded onto some kind of storage media for further analysis.

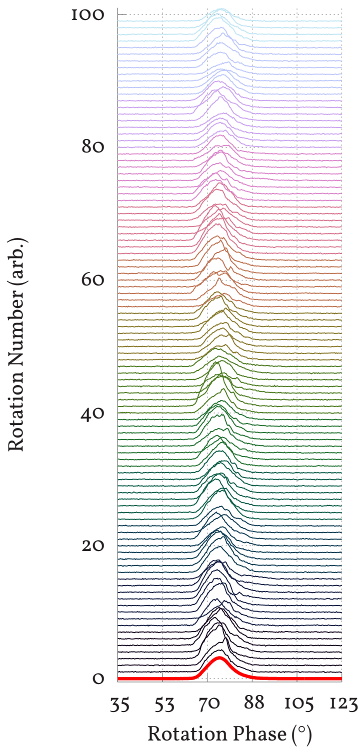

Figure 2.2: Plot of single pulses of PSR J0835-4510 (B0833-45), showing sequential *pulses* plotted on the y-axis as a function of the rotational phase of the pulsar. Note that the plot is restricted to only a small portion of the total rotation, showing the regularity with which the pulsar rotates. The individual pulses offset manually along the *y*-axis for distinction. PSR J0835-4510 is the pulsar associated with the SN G263.6-02.8 in the Vela constellation. The thick red line at the bottom of the plot shows the integrated pulse profile.

Almost all of what is described in the previous paragraphs is common for all radio astronomy. Pulsar observations differ in a number of ways from other fields of radio astronomy because of the unique nature of the sources being observed and often employ special hardware and techniques. Some aspects of pulsars that are immediately distinguishable are:

- They are relatively faint point sources moving with considerable transverse velocities.
- Their radio beams are highly collimated, polarised and show frequency-dependent effects.
- Starting from an arbitrary pulse, it is always possible to predict the time of arrival of the n^{th} pulse, if the rotational period of the pulsar and the rate at which this period increases is known from previous measurements.

Given these properties, pulsar observations benefit greatly from the use of dual-polarisation receivers that either record only raw voltages or store the incoming pulse train by truncating it at every rotational period of the pulsar being observed and averaging the resulting segments. This second mode is often called the 'fold-and-add' or folding method, which exploits the rotational stability of pulsars to add a number of individual pulses to form a bright, well defined pulse-*profile*. This is possible in spite of the random, stochastic nature of individual pulses, shown for example in figure 2.2, because the ensemble average of groups of pulses (more commonly referenced as the 'integrated' pulse profile) observed at discrete, widely separated intervals is exceedingly well-matched, as can be seen from the plot of integrated pulse profiles of PSR J2235+1506 created from four separate observations, separated approximatel by 3 months each shown in figure 2.3. The majority of the visible distinctions in the pulse profiles shown is the result of ambient RFI, which results in the 'spiky' profiles and a relatively smaller contribution to these variations is the result of changes in the column density of electrons along the line of sight to PSR J2235+1506.

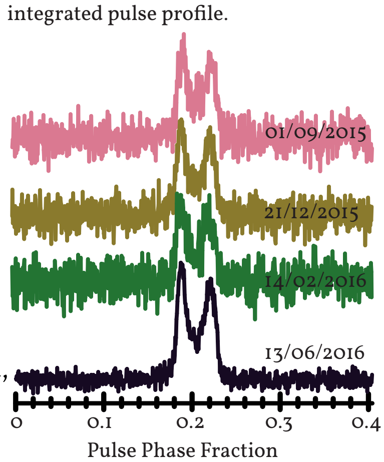

Figure 2.3: Plot of integrated pulse profiles obtained for PSR J2235+1506 on four days separated by 3 months each. While the individual pulses differ from each other mainly due to the changes in the ambient RFI and column density of electrons along the line of sight to the pulsar, the overall shape of these individual pulses remains remarkably similar.

The probes in the horn are designed to be sensitive to orthogonal polarisations of the induced EM field in the waveguide and a complex voltage is received for each polarisation, as sketched by the blue and green . If these orthogonal polarisations are linear (X and Y) as is the case for crossed-dipole (wire) arrangements or rectangular horns, the full Stokes matrix can be recovered as:

$$\begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} |X|^2 + |Y|^2 \\ |X|^2 - |Y|^2 \\ 2\,Re\,(X * Y) \\ 2\,Im\,(X * Y) \end{bmatrix}. \tag{2.1}$$

If the feed-horn is cylindrical, they are sensitive to the circular polarisations, L and R. Then the equation becomes:

$$\begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} |L|^2 + |R|^2 \\ 2\,Re\,(L * R) \\ 2\,Im\,(L * R) \\ |L|^2 - |R|^2 \end{bmatrix}, \tag{2.2}$$

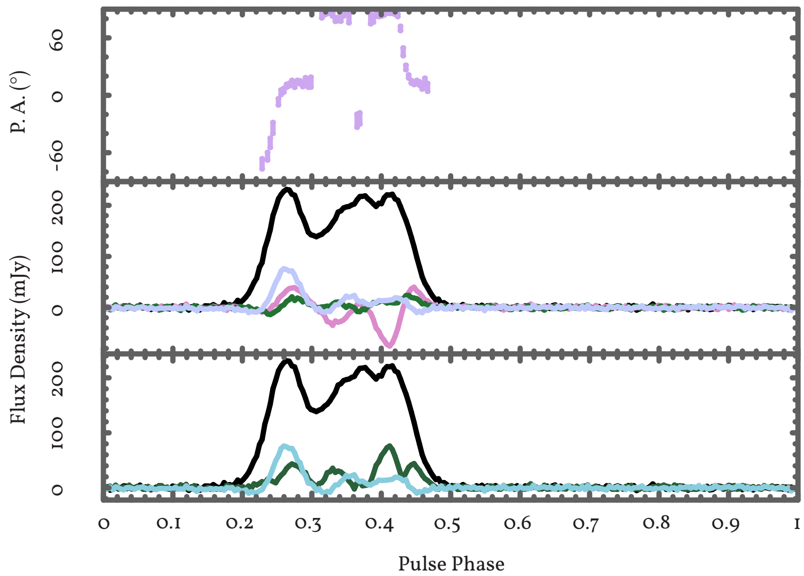

Figure 2.4: Pulse profile of PSR J0407+1607 at 327 MHz, showing the polarised components. The middle plot shows the Stokes components; $\hat{Q}$ in pink, $\hat{U}$ in blue and $\hat{V}$ in green. The bottom plot shows the linearly polarised component in green and the circularly polarised component in blue. The total intensity (Stokes $\hat{I}$) is plotted in black in both, the middle and bottom plots. The top plot shows the position angle (P.A.) for the same observation.

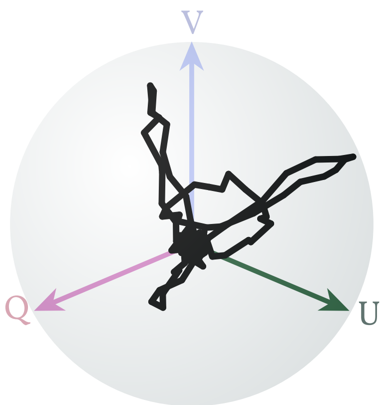

Figure 2.5: Projection of the Stoker's vector on the Poincaré Sphere (shown by the black trace), for the observation plotted in figure 2.4.

The receiver records a complex voltage for each polarisation channel, in what is commonly called the raw mode. Using frequency mixing in hardware or fast Fourier Transforms (FFTs) (Cooley and Tukey, 1965) in software (see e.g. Press et al., 1992, pp. 496 - 536), one can then obtain the full polarisation properties of the incident signal in the form of either the magnitudes and cross-multiplication terms or the magnitudes of the components. The ability to recover the Stokes components is especially important in the context of reconstructing the shape of the pulsar beam, which is important for studying the properties of the emission itself as well as, for the technique of pulsar timing discussed in section 2.5.1.

2.2 *Effects of the ionised inter-stellar medium*

The intervening medium between the pulsar and the earth consists of diffuse gas and dust and is called the interstellar medium (ISM). Although the ISM consists mainly of neutral gas and dust, ionising background radiation and heating due to the expanding shock-fronts of SNRs produces free electrons with a number density, n_e ranging from $\simeq 15 \times 10^{-4}$ cm^{-3} in the hot inter-cloud phase to about 0.4 cm^{-3} in the cold neutral phase of the ISM Tielens (2009)[1]. The group velocity, v_g, of the EM waves propagating through this refractive medium is given by:

$$v_g = c\left(1 - \frac{f_p{}^2}{f^2}\right)^{1/2}, \tag{2.3}$$

where f is the frequency of the EM wave and given e and m_e are the charge and mass of the electron respectively, $f_p \equiv n_e e^2/\pi m_e$. This implies that for two monochromatic frequencies, f_1 and f_2 the propagation times will be slightly different, as shown in figure 2.6. Since pulsar receivers are typically designed to have bandwidths (BWs) of the order of MHz to GHz and the plasma frequency for the estimated number density is $f_p \simeq 2$ kHz, this delay can be expressed to first order in f_p^2/f^2 as

$$t_2 - t_1 = \Delta t \equiv e^2 2\pi m_e c^2 \left(\frac{1}{f_1^2} - \frac{1}{f_2^2}\right) \mathrm{DM} \tag{2.4}$$

DM here represents the column density of free electrons along the line of sight to the pulsar and is called the dispersion measure. It is measured in units of pc cm^{-3} (Lorimer and Kramer, 2005, p. 86).

$$\mathrm{DM} \equiv \int_0^d n_e dz \tag{2.5}$$

Following Lorimer and Kramer (2005) we can simplify this expression for the delay at any signal frequency, f in MHz as

$$\Delta t \simeq 4.15 \times 10^6 \mathrm{DM} f^{-2} \tag{2.6}$$

where Δt is the signal delay in ms. Thus the different frequency components of the beam across the observational bandwidth appear at the detector in a quadratically delayed manner. Dedispersion is the process of removing these delays, relative to an arbitrarily choosen frequency within the observed range, so that integration of the data with respect to frequency would represent the summation of photons which left the pulsar at the same time.

2.2.1 *Incoherent and coherent dedispersion*

Dedispersion of the incoming data, prior to final storage, is the most computationally intensive of the backend processes. To reduce the computational complexity a convenient practice is to split the received bandwidth into a number of narrow band channelsto run filtering and other pre-processing operations, followed by the dedispersion of the translated signal.

[1] The ISM itself is however a complex system and interested readers are referred to table 11.2 of Tielens (2009) for a summary of the electron densities in the different phases.

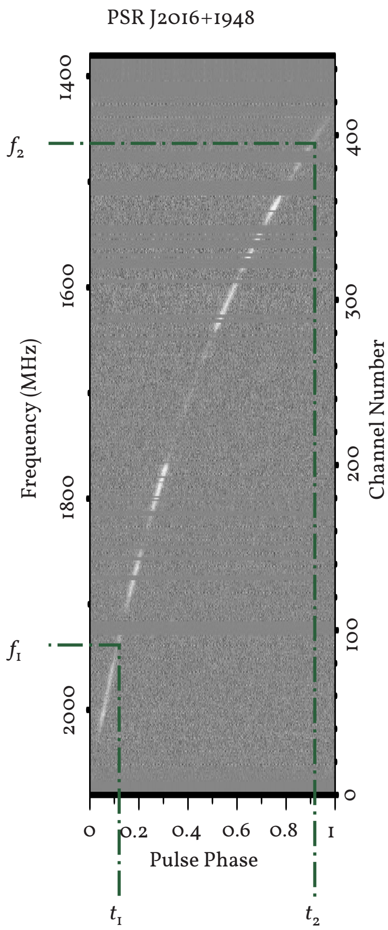

Figure 2.6: Dispersion of the pulsar signal for PSR J2016+1948, shown by the quadratically increasing delay in the arrival time of photons as function of their associated frequency. The plot shows the frequencies starting from ~2100 MHz and decreasing up to ~1300 MHz, plotted as a function of the rotational or pulse phase. At some arbitrary time t_1 the photon of frequency f_1 arrives. Photons associated with higher frequencies (i.e., the frequencies lower on the y-axis) arrive earlier while those associated with the lower frequencies at the top of the plot arrive later. If f_2 is the frequency associated with a photon which arrives at time t_2, the associated DM can be calculated from eqn. (2.4).

Regardless of whether this splitting up is carried out in hardware or software, the array of filters is called a filterbank. The dispersed pulsar signal appears as a drifting pulse as one moves from channel to channel of the filterbank. In the analog case, it is almost self-suggestive then that either a different set of delays introduced per channel or convolution with a local tone that drifts with the opposite phase but an exactly equal rate as the pulsar can be used to introduce offsets such that the pulses in the individual channels align. However, the analog case is also often difficult to construct for a general-purpose receiver.

With the advent of the FPGAs, digital backends have now become the de-facto standard for radio astronomy. In the case of the software backend, the analog filterbank is replaced by a polyphase filterbank (PFB), which is an optimal architecture for filtering. The PFB consists of an analysis stage (the part where the input signal is decomposed into a number of channels), followed by intermediate processing before the channels are recombined by the synthesis stage.

It is quite straightforward then to introduce after the analysis stage a simple DM-dependent delay stage into each channel, which will align the pulse in the individual channels. If this delay per channel is an a-priori calculated constant, then the process is called incoherent dedispersion.

Coherent dedispersion, on the other hand, uses all the available frequency information to dedisperse the signal in the frequency domain (Hankins and Rickett, 1975). In this method, the signal must be sampled at its Nyquist frequency (Nyquist, 1928b). The resulting stream is then multiplied by a frequency-dependent chirp function. This alters the complex phase of the FFT values, removing the dispersive delays. Simultaneous windowing[2] of the sampled block of the complex amplitudes of the FFTs prevents aliasing[3].

[2] Which could be thought of as shaping or tapering of the edges of the sampled block of the signal.

[3] The effect of images of the sampled signal being produced at a different frequency. See chapter 3.

As discussed in chapter 3, certain assumptions on the nature of the properties of the Fourier space lead to drastic departures from the ideal case, causing 'images' of, and aberrations in the pulsar signal to appear at different phase offsets relative to the main signal. This happens at both the digitisation stage and the PFB stage. An alternative that is currently being deployed at the Effelsberg 100-m radio telescope is to avoid filterbanks altogether by performing full-band acquisition and coherent dedispersion. This requires high-speed ADCs and FPGAs. Given that the required FFT lengths increase dramatically for full-band processing, the demands on the computing and memory resources are also significantly higher.

2.3 Polarisation Calibration

Once dedispersion has been performed, the data must be recorded and stored for further analysis. Since it is already in the Fourier domain, it is straightforward to construct the four Stokes components. The four Stokes components are then compressed into a single four-dimensional (phase, frequency, time and intensity) data object, commonly referred to as an 'archive'. The availability of the four Stokes components is crucial in recovering the correct profile of the projected beam of the pulsar, given the high degree of polarisation these beams are known to possess.

The expressions for the Stokes parameters shown in eqn. (2.1) and eqn. (2.2) implicitly assume that the individual polarisations are orthogonal and free from any cross-coupling. However, the mutual isolation of the two polarisations of the receiving horn is never perfect and instead, depends on the electrical properties[4] of the probe arrangement and the horn/feed design, which implies frequency-dependent variations. This cross-contamination which originates from subtle differences in the electrical and geometrical properties of the telescope and the front-end can be modelled and removed as well [5] using multiple injections of the noise diode.

[4] Which in turn depend on the geometry and materials used.

[5] This is done at the Parkes radio telescope in Australia (van Straten, 2004, 2006) and is also suggested for observations at Arecibo which follow sources over long tracks on the sky.

This can be measured by injecting a signal of known polarisation (using a noise diode) periodically and comparing the incident and injected signals. For newer telescopes, these changes are often small enough that a single calibration observation with the noise diode, per epoch for every pulsar observed, can suffice. If the archives produced earlier have had this calibration performed on them, they are said have been polarisation calibrated.

2.4 *Flux Calibration*

Pulsar astronomy backends typically record the spectral density function, S in arbitrary units such that the archive needs to calibrated both in terms of the flux density value and the frequency dependent gain of the receiver, G_r. There also a need to individually calibrate the effect of the frequency and elevation dependent gain of the telescope, G_T.

$$S^{[cal]}(f) = G_T(f)G_r(f)\left(T_{\text{source}}(f) + T_{\text{sky}}(f) + T^{[cal]}_{\text{sys}}(f)\right) \tag{2.7}$$

If G_T and G_r can be determined independently along with the noise power generated by the receiving system T^{cal}_{sys}, then it is possible to recover the flux density of the source, T_{source} for an observing frequency f. In practice, a number of parameters change in the telescope from observation to observation and these can make it quite difficult to measure the different gains and flux densities. Using a noise diode to inject a constant, unknown amount of power in the signal chain and a celestial radio source with a previously measured flux density at the frequency of observation, we can recover the different components of the measured flux density quite easily.

By pointing the telescope at a strong continuum source with a well known flux density (referred to as a standard candle), we can obtain the total power for the On and Off states of the noise diode, H and L respectively, in terms of the relative gain $g \equiv G_T \cdot G_r$, the unknown system temperature T_{sys}, the flux density of the standard candle T_{std} and the unknown flux density of the noise diode, T_{diode}. To simplify the expressions, the notation for frequency dependance has been dropped and the subscript 'on' refers to measurements made with the telescope being pointed at the flux calibrator source and the subscript 'off' indicates

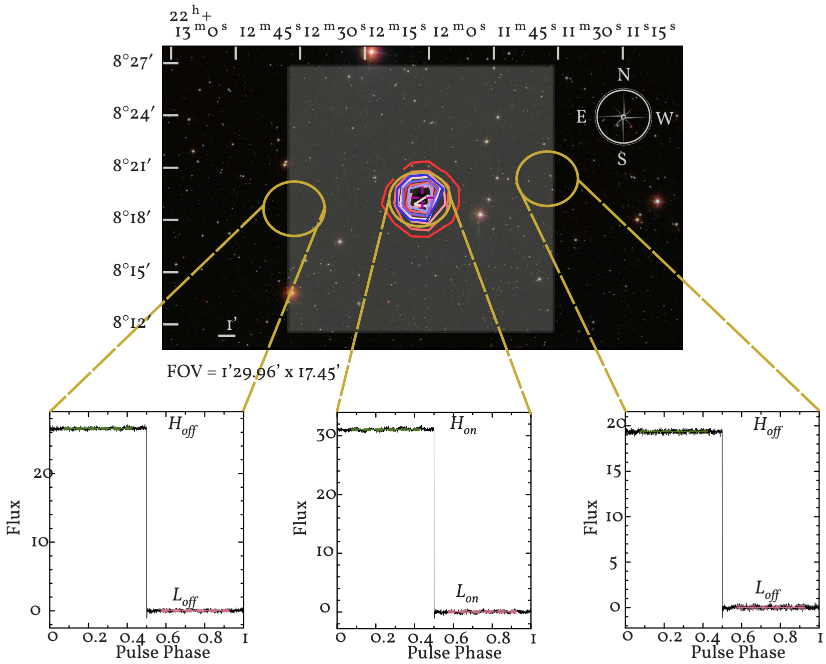

Figure 2.7: Sketch of the method by which the absolute flux received is measured, using a quasar in the sky. In the example shown alongside, the standard candle used is the the quasar 4C 08.64 (B2209+080), the optical image for which is shown from the Sloan Digital Sky Survey (SDSS) archival data, plotted using the Aladdin desktop software (Bonnarel et al., 2000). Overplotted in colours from red to white are the radio flux contours from the Jansky Very Large Array (VLA) Low-Frequency Sky Survey Redux (VLSSr; Lane et al., 2012) 74 MHz and from blue to white, radio flux contours from the National Radio Astronomy Observatory (NRAO) VLA Sky Survey (NVSS; Condon et al., 1998) 1.4 GHz continuum survey. The three brown ellipses denote the beam-width of the L-wide receiver at the Arecibo Observatory, with which these data were collected. A minimum of three pointings (two off source and one on source) are used in practice, to remove an implicit assumption of source symmetry in the equations used in section 2.4. The translucent white squares denote the field of view of the VLSSr and NRAO VLA Sky Survey (NVSS) survey plots, identifying the absence of radio bright sources apart from 4C 08.64.

those with the telescope pointed to a nearby part of the sky as shown in figure 2.7.

$$\begin{aligned} H_{on} &= g_{on}(T_{\text{sys}} + T_{\text{std}} + T_{\text{diode}}) \\ L_{on} &= g_{on}(T_{\text{sys}} + T_{\text{std}}) \end{aligned} \tag{2.8}$$

By pointing the telescope at a nearby location of the sky, sufficiently spaced to exclude the standard candle from the beam of the telescope, we obtain another set of measurements for the On and Off states of the noise diodes.

$$\begin{aligned} H_{off} &= g_{off}(T_{\text{sys}} + T_{\text{diode}}) \\ L_{off} &= g_{off}T_{\text{sys}} \end{aligned} \tag{2.9}$$

From these four measurement, we can now obtain the relation between the flux densities of the standard candle and the noise diode.

$$\frac{L_{on}}{H_{on} - L_{on}} - \frac{L_{off}}{H_{off} - L_{off}} = \frac{T_{\text{std}}}{T_{\text{diode}}} \tag{2.10}$$

Finally, by substituting the value of T_{diode} obtained in eqn. (2.10), we can recover the system temperature using the relation;

$$\frac{H_{off}}{L_{off}} - 1 = \frac{T_{\text{diode}}}{T_{\text{sys}}} \tag{2.11}$$

from which we obtain the system temperature of the receiver.

By using this T_{sys} and measuring the power when the pulsar is just outside the beam of the telescope to obtain $G_T(f)G_r(f)\left(T_{\text{sky}}(f) + T_{\text{sys}}^{[cal]}(f)\right)$, we can measure T_{source}. Finally, comparing T_{source} with T_{std} gives the equivalent flux-density of the source.

2.5 *Pulsar timing*

Since pulsars have been shown to possess rotational stability comparable to some terrestrial time standards, one can treat each sweep of the pulsar beam as a single *tick* from a clock. By comparing these pulses with a terrestrial standard, we can measure offsets in the pulsar's rotation at precisions of microseconds or less, which can then be modelled to recover the physical phenomena that produce these offsets. This is known as the pulsar timing technique.

2.5.1 *Measurement of pulse times of arrival*

The fundamental datum of the pulsar timing technique is the time of arrival (ToA) of the pulse, typically referred to Solar system barycentre (SSB). These ToAs embed information about the telescopes, backends and algorithms used to measure them; the behaviour of the pulsars themselves and the systems hosting them, if any; the effects of any large objects near the line of sight, the Solar System and any other physical phenomena that would affect the propagation of the pulsar signals. Timing models incorporating the spin and astrometric parameters, as well as parameters for the DM and when applicable, orbital parameters are fitted to these ToAs using timing software like TEMPO2 (Hobbs et al., 2006).

Often, however, the individual pulses are too weak to be distinguished from noise and it is preferable to integrate the archive along the time and/or frequency domain(s). This decreases the amplitude of the Gaussian noise and increases that of the pulse and any other non-Gaussian components of the noise. The *scrunched* or integrated profile then represents an ensemble average of the individual pulses that are contained in the archive. Within the limits of the sensitivity of the receiver and the telescope, the ensemble average at one epoch can approximate the ensemble average at a different epoch exceedingly well. In fact, we can express the measured pulse profile $P(t)$ as a function of a standard *template* profile $S(t)$, via the relation:

$$P(t) = a + b \times S(t - \Delta\tau) + n(t) \tag{2.12}$$

where a is the offset between the baselines of the standard and measured profiles, b is a scaling factor, $\Delta\tau$ is the phase offset between the profile and the template and, $n(t)$ is the noise component.

If the discrete Fourier transforms of the profile and template are

$$P_k \mathrm{e}^{\mathrm{i}\theta_k} = \sum_{j=0}^{N-1} p_j e^{\mathrm{i}2\pi jk/N} \tag{2.13}$$

and

$$S_k \mathrm{e}^{\mathrm{i}\phi_k} = \sum_{j=0}^{N-1} s_j e^{\mathrm{i}2\pi jk/N} \tag{2.14}$$

respectively, where k represents a frequency index, P_k and S_k the amplitudes of the complex Fourier coefficients, θ and ϕ are the respective phases and N is the number of frequency channels.

Since the transformation preserves linearity, eqn. (2.12) can be rewritten as

$$P_k e^{i\theta_k} = aN + bS_k e^{i\phi_k} + G_k \tag{2.15}$$

where the index k runs from zero to N-1.

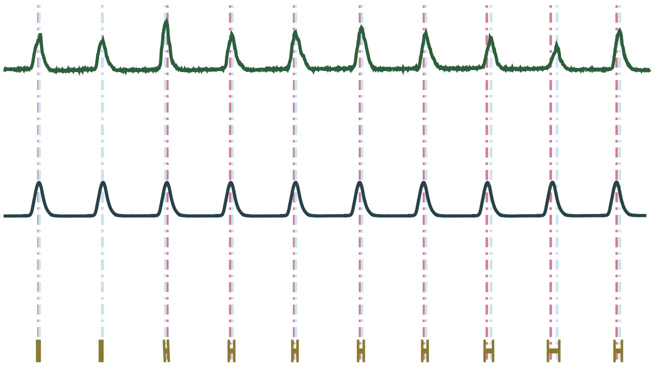

Figure 2.8: Schematic of the method of calculation of ToAs. In practice, the template matching or comparision between the predicted arrival time and the actual arrival time is carried out in the Fourier domain, as explained in section 2.5.1.

Once the individual transforms have been computed, the baseline offset can be measured from the zeroth components of the two amplitudes,

$$a = (P_0 - bS_0)/N. \tag{2.16}$$

The ToA can now be derived along with the scaling factor b by minimising the goodness-of-fit statistic

$$\chi^2 = \sum_{k=1}^{N/2} \left| \frac{P_k - bS_k e^{i(\phi_k - \theta_k + k\tau)}}{\sigma_k} \right|^2 \tag{2.17}$$

where σ_k represents the root-mean-square intensity of the noise at frequency indexed by k. A more detailed derivation can be found in (Taylor, 1992).

While the derivation above shows the simplest case, alternative methods to recover ToAs from low signal-to-noise ratio (S/N) signals using Gaussian interpolation can be found in Hotan et al. (2005) or van Straten (2006) which utilises the full Stokes information to produce ToAs. An equivalent frequency domain formulation of the (Taylor, 1992) method can also be found in Demorest (2007). Updated techniques for ToA generation for wideband backends, where frequency dependent evolution of the pulse profiles must be accounted for can be found in Liu et al. (2014); Pennucci et al. (2014)

2.5.2 *Pulsar timing with* TEMPO2

First, the ToA is time-stamped using a clock at the observatory. This local clock is referenced to a hydrogen maser, which itself is ultimately not stable over extended periods of time. This time must then be translated into the Universal Coordinated Time (UTC) scale from which one can finally derive the corresponding value in the Temps Atomique International (TAI) timescale. However, the atomic clocks used to derive the TAI timescale do not measure the SI second exactly and offsets must be accounted for. The clocks used in the TAI timescale are used to define what is called the Terrestrial Time (TT) scale. The Bureau International des Poids et Mesures (BIPM) publishes the transformation between pairs of timescales and these must be used to derive the correct reference ToA in SI seconds referred to the Geocentric Celestial Reference System (GCRS) based timescale, denoted by Geocentric Coordinate Time (TCG). The current standard for translating from TT to any other timescale is the TT(BIPM15), which applies to measurements extended beyond modified Julian date (MJD) 57379.

"TT(BIPM) is a realization of Terrestrial Time as defined by the International Astronomical Union (IAU). It is computed annually by the BIPM based on a weighted average of the evaluations of the frequency of TAI by the primary and secondary frequency standards."

– BIPM.org

The latest value for the correction is TT(BIPM15) = TAI + 32.184 s + 27702.0 ns

To extract meaningful information from the measured ToAs they must be translated into a proper time of emission at the pulsar. The steps involved in this translation from a ToA on the earth to the time of emission at the pulsar are as follows.

Having corrected the ToA to the GCRS we then translate it to the SSB. This involves, apart from translating the reference frame from the GCRS to the barycentric celestial reference system (BCRS), calculating the delays the photon must have encountered during its propagation through the Earth's atmosphere (Δ_{Atm}), the vacuum retardation due to the motion of the observatory ($\Delta_{\mathrm{R}_\odot}$ and Δ_{p}), that due to dispersion by the ionised solar wind (Δ_{SW}), that due to the relativistic frame transformations due to the co-moving SSB and observatory, also called the Einstein delay (i.e., the gravitational redshift, $\Delta_{\mathrm{E}_\odot}$) and finally that due to the excess path it has to travel through the gravitational potential of the Solar System, called the Shapiro delay ($\Delta_{\mathrm{S}_\odot}$).

$$\Delta_\odot = \Delta_{\mathrm{Atm}} + \Delta_{\mathrm{R}_\odot} + \Delta_{\mathrm{p}} + \Delta_{\mathrm{SW}} + \Delta_{\mathrm{E}_\odot} + \Delta_{\mathrm{S}_\odot} \tag{2.18}$$

The barycentred arrival time (BAT), or the ToA translated to the SSB must be corrected for the effects of propagation through the ISM which include the vacuum propagation delay (Δ_{VP}), the dispersion due to the ISM (Δ_{ISD}) and other frequency dependent effects (Δ_{FDD}) and finally the Einstein delay due to the relativistic motion of the SSB and the binary barycentre[6] ($\Delta_{\mathrm{E}_{\mathrm{SSB,BB}}}$).

$$\Delta_{\mathrm{IS}} = \Delta_{\mathrm{VP}} + \Delta_{\mathrm{ISD}} + \Delta_{\mathrm{FDD}} + \Delta_{\mathrm{E}_{\mathrm{SSB,BB}}} \tag{2.19}$$

If the pulsar is in a binary, we must now correct for the effects of binary motion which include the Römer delay due to the binary companion ($\Delta_{\mathrm{R}_\mathrm{B}}$), the aberration that is introduced due to the proper motion ($\Delta_{\mathrm{A}_\mathrm{B}}$), the Einstein delay due to the companion ($\Delta_{\mathrm{E}_\mathrm{B}}$) and the Shapiro delay due to the companion ($\Delta_{\mathrm{S}_\mathrm{B}}$).

$$\Delta_{\mathrm{BB}} = \Delta_{\mathrm{R}_\mathrm{B}} + \Delta_{\mathrm{A}_\mathrm{B}} + \Delta_{\mathrm{E}_\mathrm{B}} + \Delta_{\mathrm{S}_\mathrm{B}} \tag{2.20}$$

A more detailed treatment of these terms can be found in Edwards

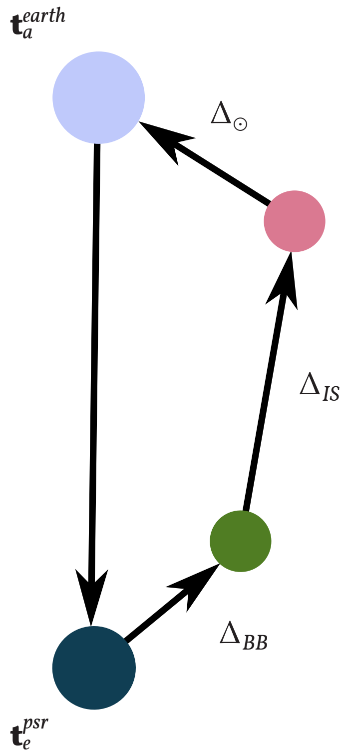

Figure 2.9: Vector representation of the translation repsresented by eqn. (2.21), showing the change from the ToA at Earth, $t_{\mathrm{a}}^{\mathrm{earth}}$ to the time of emission at the pulsar, $t_{\mathrm{e}}^{\mathrm{psr}}$ via the removal of the effect of the binary orbit Δ_{BB}, the effects due to propagation through a non-neutral, turbulent interstellar medium Δ_{IS} and the effects Solar System, $\Delta_\odot$.

[6] This is true for a pulsar in a binary, however for solitary pulsars this corresponds to centre of mass (effectively, the centre of the pulsar).

et al. (2006), which discusses these terms in context of the most commonly used software package for pulsar timing nowadays, TEMPO2.

Using the equations (2.18) to (2.20) we can now derive the time of emission at the pulsar as

$$t_e^{psr} = t_a^{Earth} - \Delta_\odot - \Delta_{IS} - \Delta_{BB}, \quad (2.21)$$

represented by the vector diagram figure 2.9. Having obtained the time at which the photon was emitted at the pulsar, we can now model the rotational phase of the pulsar at this time as an integer number of cycles since the epoch, t_P, at which the rate-change of the phase $\dot{\phi}$, equals the frequency of rotation ν using the following expression (Taylor, 1992)

$$\phi(t) = \sum_{n \geqslant 1} \frac{\nu^{n-1}}{n!} \left(t_e^{psr} - t_P \right) \quad (2.22)$$

where ν^n are the frequency derivatives.

It is evident that many of the parameters listed earlier are not always known apriori. Instead, starting with a minimal set of parameters, a least squares minimisation must be carried out over the expression:

$$\chi^2 = \sum_{i=1}^{M} \left(\frac{\phi(T_i) - n_i}{\sigma_i / P} \right)^2 \quad (2.23)$$

where n_i is the closest integer to the phase $\phi(T_i)$ and σ_i is the uncertainty of the ith ToA.

2.5.3 *Pulsar timing models*

Typically, the first set of parameters that are available after a few epochs of observations are the position, often expressed in terms of the *R.A.* and *DEC* as well as the spin period, P and the spin down rate, $\dot{P}$. From the earlier discussions, it is evident we can also extract the DM along the line of sight to the system and any time derivatives of the DM, provided the data are sensitive to such changes. If the observations last for more than a year, we can extract the proper motion terms (μ_α and μ_δ) as well. In some cases one can even measure the parallax of the system.

If the pulsar is in a binary we can extract the binary parameters as well. These include the five Keplerian parameters, the projected semi-major axis $x \equiv a_1 sin(i)/c$, eccentricity e, binary period P_b, longitude of ascending node, Ω and the epoch of periastron, T_0, all shown in figure 2.10. In addition to these parameters, we can apply orbital dynamics to recover the mass function and relative orientations of the pulsar and the companion. In the special case where the pulsar is orbited by a massive companion we can estimate the mass of the pulsar using the Shapiro delay (Shapiro, 1964). If the companion is visible as an optical source, we may be able to independently derive the companion's mass and then derive the pulsar's mass.

If the mass ratio of the system is suitably high and the orbital parameters are particularly favourable, especially in the case of double NS binaries, then it may even be possible to extract up to eight separately measureable post-Keplerian parameters; the derivatives $\dot{\omega}$, $\dot{P}_b$, $\dot{x}$ and $\dot{e}$,

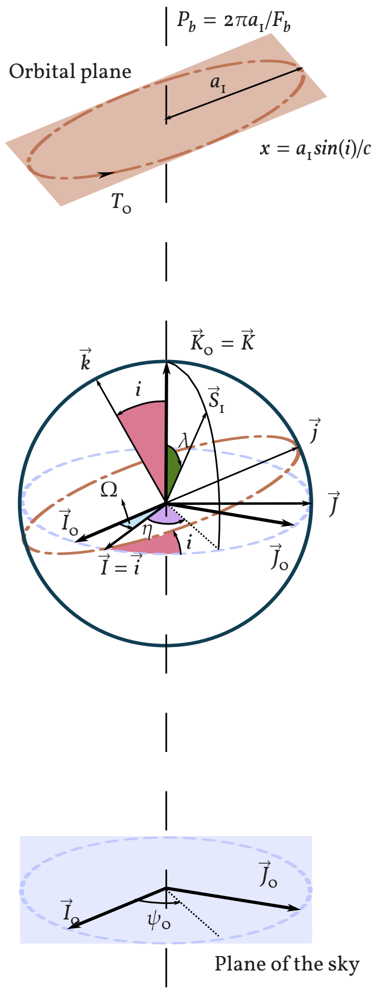

Figure 2.10: The angles and conventions used in some of the most commonly used binary pulsar timing models. The co-ordinate axes $\vec{I}$, $\vec{J}$ and $\vec{K}$ are co-aligned with the centre of mass of the binary, while $\vec{S}$ denotes the spin angular momentum of the binary. $\vec{K}_0$ coincides with the line of sight, shown by the dashed black line. Also shown are the angle of inclination, i ($0 < i < \pi$) the longitude of the ascending node Ω ($0 < \Omega < 2\pi$). Figure following Damour and Taylor (1992).

the Einstein parameter $\gamma_{Einstein}$, the range and shape of the orbital Shapiro delay, r and s and, an orbital shape correction, δ_θ (see Damour and Taylor, 1992).

Extensive treatments of pulsar timing can be found in Hobbs et al. (2006); Taylor (1992); Lorimer and Kramer (2005) while a more complete treatment of the binary models can be found in Blandford and Teukolsky (1976); Damour and Deruelle (1986); Taylor and Weisberg (1989).

2.5.4 *Pulsar timing arrays*

Gravitational waves (GWs) modify space-time as they propagate through it. The first direct detection of such waves by the Laser Interferometer Gravitational-Wave Observatory(LIGO, Abbott et al., 2016a) marks the advent of a new age of astronomy. However, the black hole (BH) binary systems whose mergers Laser Interferometer Gravitational-Wave Observatory (LIGO) probes are limited to tens of $M_\odot$. To probe the GW emission from super massive black hole (SMBH) binary mergers whose masses are $\sim 10^6\ M_\odot$, we must turn to pulsars. Pulsars, as rapidly rotating systems and massive objects in binaries are capable of emitting GWs themselves. They are however, also excellent Einstein clocks with which we can measure the perturbation of the space-time enclosing the Galaxy as GWs with periods $\sim 10^2$ d pass through it.

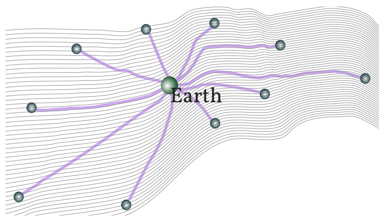

Figure 2.11: Schematic showing a plane section of space-time that is perturbed by the superposition of a number of GWs. The green sphere represents the Earth, while the blue spheres mark a number of pulsars that lie on the plane section. As photons travel through this space-time, their time of flight is affected by the GWs, imprinting them with a specific signature.

An ensemble of luminous MSPs spread across the Galaxy would form a 'pulsar timing array' (PTA, see e.g., Kramer and Champion, 2013). Figure 2.11 sketches the fundamental principle of pulsar timing arrays. As GWs propagate they introduce variations in the time at which a pulse from each pulsar is detected. By carefully monitoring the MSPs in the pulsar timing array (PTA) and correlating the signals from all the MSPs we can detect GWs. The success of PTAs however is critically dependent on the number of bright, stable sources in the PTA. This is especially true in the intermediate to high S/R regime of the GW background, which results from the superposition of the GWs of multiple SMBH mergers.

To a large extent, the work presented here is closely related to PTA research. In fact, much of the work presented in this thesis was carried out within the EPTA (Desvignes et al., 2016). Chapter 4 and chapter 5 seek to address the problem of adding more sources to PTAs by investigating the stability of BW pulsars for high-precision timing and analysing the spectral properties of the MSP population while in chapter 3 I review the limitations introduced on PTAs by artefacts due to signal processing.

3

Artefacts in Polyphase filterbanks

住の江の
岸による波
よるさへや
夢の通ひ路
人目よくらむ

I am silver and exact

— **Sylvia Plath;** *The Mirror*

Millisecond pulsars are by definition rapidly rotating stars which are often weak radio sources. Timing analyses of such sources relies on the accurate measurement of the true pulse profile since this vastly improves the precision with which we can measure the time of arrival of a pulse at an Earth-bound observatory. Digital signal processing has greatly improved the limits on such reconstructions by using many new techniques and reprogrammable devices like field-programmable gate arrays. One such technique, polyphase filterbanking, is extensively used in all domains of communication. The reduced computational complexity and flexible scope of the method has led to widespread application in improving the performance of data recorders for pulsar astronomy. However, in the resource-limited systems that are typically available for use in the digital backends, it is practically impossible to designs systems that do not produce artefacts. As the receiving systems achieve ever lower system temperatures, these artefacts become more pronounced. We present here a short description of the many sources of errors that affect such filterbank schemes which can be found in literature and investigate the effect of coherent dedispersion in such filterbanks. We show that in the case of wide-bandwidth receivers with a limited number of channels, coherent dedispersion breaks the constant phase offset limit and therefore filterbanks based on finite impulse response filters cannot produce perfect reconstruction.

3.1 Introduction

Modern radio astronomy backends make heavy use of digital signal processing (DSP) techniques. One of the most commonly used components in digital radio astronomy receivers is the filterbank (FB), which is essentially a contiguous arrangement of filters. FBs are used to transform time-domain signals into their frequency-domain representations (Vaidyanathan, 1993), for the spectral analysis of signals (Boashash, 2003; Stoica and Moses, 2005) or as transmultiplexers (Fliege, 1993). FBs are constructed in many different schemes and are usually tailored to the specific application. The broadest classification of FBs depending on their output lists three types: analysis, synthesis and full. Essentially, the analysis FB is where the signal is broken down into smaller parts and then filtered while the synthesis FB combines smaller pieces to reproduce the input signal. A full FB consists of both, the analysis and synthesis FBs. However, there exist many other classification schemes for FBs which are too numerous to list here. FBs are especially relevant for pulsar astronomy given the high data rates and the need to perform efficient computations on these data, especially to remove the effect of the ISM (i.e. to perform dedispersion).

The signals we receive from pulsars are continuous and time variable. These must be translated into discrete signals which can then be processed within filterbanks. As a result of this conversion from a continuous function to its discrete version, as well as the fundamental nature of discrete algebra (and therefore DSP), the recorded signal suffers from many artefacts. We present below an overview of the most important of these artefacts, along with their origins and mathematical quantification. While a majority of the discussion deals with full reconstruction FBs implemented within FPGAs, most present-day pulsar data recording systems implement only analysis FBs after which the data are streamed over high-speed network links to recording computers where the data are dedispersed and stored. However, implementing optimised schemes for full-reconstruction FBs requires effectively the same mathematical treatment. Hence we discuss the implementation of a full-reconstruction FB and only address the question of applicability to present-day pulsar data recording systems in Section 3.6.

The most common challenge of digital filter design is that of fitting complex designs into resource-limited devices like FPGAs or microcontrollers. Apart from the limits on the available physical memory, these designs need to have minimum computational complexity to improve processing times and reduce power consumption. A commonly used scheme to this end is the PFB. This relies on the polyphase decomposition technique, introduced by Bellanger et al. (1976) and Vary (1979). It remains extremely popular due to the fact that it allows designers to perform the necessary computations at the "lowest rate permissible within the given context" (Vaidyanathan, 1998).

Thanks to the polyphase structure and the existence of certain 'Noble identities' (Vaidyanathan, 1993, see also Figure 3.1), it is possible to decompose the discrete (or sampled) version of the continuous time-series of interest into small parts which can then be processed using limited resources. This allows designers to construct fast, compact DSP algorithms which form the backbone of most astronomical receivers.

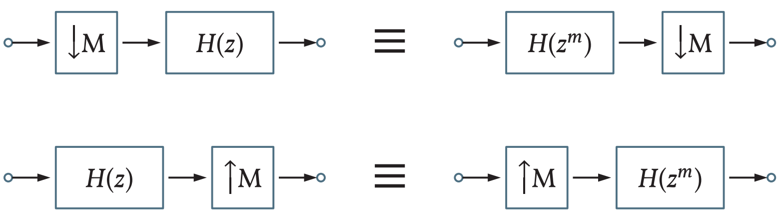

Figure 3.1: Noble identities are commutative relations which allow us to change the order of the up/down-sampling operations and the processing, without affecting the signal. $H(z)$ is the filter transfer function and ↑↓M represents an interpolation or decimation operation with factor M. Here the superscript m denotes that the filters (shown only as an example) on the right hand side of the equivalence symbol are special versions of those on the left hand side. The top-plot shows the equivalence relation for the action of down sampling, while the bottom plot shows that for up sampling. Derivations showing the validity of these relations can be found in Vaidyanathan (1993).

The flow diagram of a simple three-channel filterbank is shown in Figure 3.2. A digital signal $x[n]$ is analysed, sub-processed and synthesised to produce a reconstructed signal, $x'[n]$. This can be explained in the following manner; the sequence $x[n]$ is first split into three parts by applying a decimate-by-three operation represented by $\downarrow 3$. The output of the decimator in each branch can be written as

$$x[n] \downarrow 3 = x_d[3n] \tag{3.1}$$

where $d = 0, 1, 2$. This implies that of the full sequence $x[n]$ the decimated sequence $x_d[n]$ retains only every 3rd sample starting from the sample number corresponding to the channel number. Translated to the frequency domain, this produces an infinite number of replicas of the input signal to appear at integer multiples of the input frequency. To be precise, in terms of the z-transform[1], this can be written as (see Equations (3.17) and (3.19) for hints on the derivation):

$$X_M(z) = \frac{1}{3}\sum_{m=0}^{2} X(z^{1/3}W^m) \quad m = 0, 1, 2. \tag{3.2}$$

[1] Which is defined for a discrete sequence $x[n]$ of finite length k as

$$X(z) = \sum_{n=0}^{k} x[n]z^{-n}$$

where $z = e^{i\omega}$ is any complex number. See e.g., Jury (1964)

In this case, these three terms make up a function which is periodic in ω, the associated frequency for discrete-time signal $x[n]$. This is a basic property for any sequence which has been Fourier transformed (Oppenheim and Willsky, 2013). For the first or 0th channel the terms with $m = 1, 2$ are called *aliases* and are removed by applying a low-pass filter (for real valued signals, while for complex valued signals this becomes a band-pass filter), represented here by $H(z)$. This is often called the anti-aliasing filter. The order of this operation is also shifted as shown in Figure 3.2, using the first Noble identity of Figure 3.1. This can be followed by a number of signal processing steps with the condition that any operation must either preserve the phase of the input signal or alter it only by a constant value. This is a fundamental requirement of the entire PFB as well.

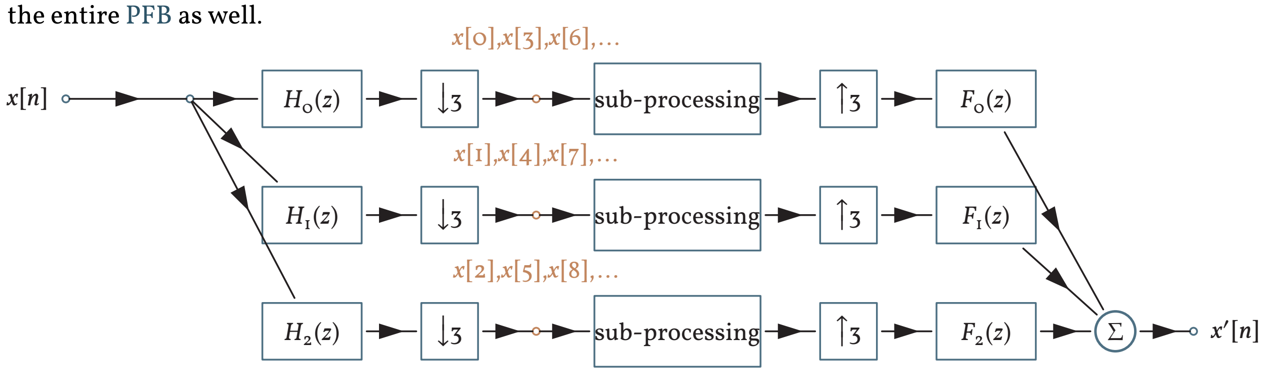

Figure 3.2: Schematic of a polyphase filterbank. A time domain signal $x[n]$, where the index represents individual sample points, is filtered by the analysis filter and split into three parts by the decimators such that every third sample is retained by the decimator. The reduced representations of the input signal are then passed through the subprocessing blocks. Finally, the interpolators pad the branch signals with zeros and the synthesis filters remove imaging artefacts, after which the branch signals are recombined the branch signals in order, to produce the reconstructed output $x'[n]$

After the sub-processing is completed, the signal is now interpolated by the interpolator ↑ 3, which inserts zero-valued samples into the signal from each branch such that:

$$\begin{aligned} x_d[n] \uparrow 3 &= x[n] \ \textit{for}\ n = d, d+3, d+6, \ldots, \\ &= 0 \ \textit{otherwise}. \end{aligned} \tag{3.3}$$

The three streams are then added together to produce the reconstructed signal, $x'[n]$ In the most general case, the z-transform of the output of a polyphase analysis-synthesis FB can be represented by (Vaidyanathan, 1998):

$$X'(z) \ = \ T(z)X(z) \ + \ \textit{terms due to aliasing} \tag{3.4}$$

where $X'(n)$ is the reconstructed signal represented in the time-domain, $X(z)$ is the z-transform of the digitised input signal and $T(z)$ is the matrix operator representing the effect of the PFB. For example, in Figure 3.2 $T(z)$ contains the action of the filters $H_m(z)$, $F_m(z)$ and the sub-processing blocks. It can be shown that aliasing can be completely removed in specific cases (see e.g., Crochiere and Rabiner, 1976), by making a proper choice of synthesis filters $F_0(z)$, $F_1(z)$, etc (see also, Section 3.4). Given a set of specifications, a further simplification can be made by using *modulated* filters, wherein a single prototype filter is modulated by a real or complex function to obtain the analysis and synthesis filters.

Further, when $T(z)$ is equal to a pure delay[2], $T(z) = cz^{-k}$, the output of the PFB is said to be 'perfectly reconstructed'. The simplest design case is that of a perfect reconstruction (PR), two-channel quadrature modulated filterbank (QMF)[3] as demonstrated in Smith and Barnwell (1986) and Mintzer (1985). However, in the most general case, it is non-trivial to satisfy the PR conditions, although in many cases it is possible to obtain a very close approximation or near-perfect reconstruction (NPR).

Infinite impulse response filter (IIR)[4] filters can also be used since they can be implemented using recursive methods but it can be difficult to design universally stable IIR filters. Hence, most PFBs are implemented using FIR filters, since these filters are easy to stabilise across wide bandwidths and can also satisfy the phase linearity requirements which are necessary for PR as discussed in Section 3.4.

In the following sections we first introduce the various sources of errors or artefacts that affect PFBs starting with the artefacts that appear due to the digitisation itself in Section 3.2. The polyphase decomposition is carried out for the two channel QMF and the primary sources of error are demonstrated in Section 3.4.

[2] The frequency response of $T(z)$ is $T(e^{i\omega}) = |T(e^{i\omega})|\, e^{-i\omega}$ where $|T(e^{i\omega})| = c$ is the amplitude response and $e^{-i\omega}$ is the phase response. If the phase response is modified by a constant then the output resembles the input except for a constant delay in time.

[3] The original definition of the QMF is that of a two-channel FB. It is worth noting that the math remains so similar even in the case of multi-channel FBs that they are often also called QMFs. However, in the discussion here, we specify the number of channels to avoid ambiguity.

[4] Digital filters can be classified into two types, IIR and FIR. FIR filters are those filters whose response to an input impulse is a finite duration signal (in other words, the set of filter coefficients is finite) while for IIR filters the response does not go to zero even after the input impulse has disappeared. In the IIR filter, a fraction of the input power is looped back into the filter by design, preventing its response from going to zero after a signal has appeared at its input. In this case, the length of the set of filter coefficients is theoretically infinite but in practice this can be limited to some large, finite value. Proper definitions may be found in Rabiner and Schafer (1978).

3.2 *Digitisation artefacts*

During a pulsar observation, an analog, continuously varying time-domain signal is sampled by a digitiser before it can be passed through the PFB. The action of digitisation itself introduces several errors. While some of these are often easily overcome by using a larger number of bits to represent each sample, others are more deeply linked to the nature of sampling and approximation and require careful treatment. We describe the most important of these below. It should be noted that apart from the first artefact listed below, the others are relevant for digital filters and other components as well, i.e., they introduce non-ideal behaviour in all parts of a digital system.

3.2.1 *Sampling Artefacts*

Consider a continuous-in-time function f(t) which is sampled with a period T_0. This can be represented as a multiplication by a Shah function or a Dirac comb (a train of impulses where each pulse has width tending to zero and has finite amplitude, as shown by blue lines in Plate 3), $\text{III}_{T_0}(t)$.

The sampled function in the time domain is then:

$$g(nT_0) = f(t)_{T_0}(t) \qquad (3.5)$$

where $g(nT_0)$ consists of n discrete values separated by T_0 units of time.

Since the Fourier transform of a Shah function is :

$$\mathscr{F}\left[Ш_{T_0}(t)\right] = \frac{2\pi}{T_0}\sum_{k=-\infty}^{\infty}\delta(\omega - k\omega_0) = \frac{2\pi}{T_0}Ш_{\omega_0}(\omega), \tag{3.6}$$

we can write the Fourier equivalent of Eqn. (3.5) as:

$$\begin{aligned} G(\omega) &= \frac{1}{2\pi}F(\omega) * \frac{2\pi}{T_0}Ш_{\omega_0}(\omega) \\ &= \frac{1}{T_0}F(\omega) * \sum_{k=-\infty}^{\infty}\delta(\omega - k\omega_0) \end{aligned} \tag{3.7}$$

where uppercase letters represent the Fourier transforms of the respective functions and $\delta(\omega)$ is a Dirac delta function.

Eqn. (3.7) now becomes:

$$G(\omega) = \frac{1}{T_0}\sum_{k=-\infty}^{\infty}F(\omega - k\omega_0) \tag{3.8}$$

i.e., the Fourier transform of a sampled function $g(nT_0)$ consists of periodic repetitions of the Fourier transform of the original function, $f(t)$. Each copy is separated from the next by a 'period' $\omega_0 = \frac{1}{T_0}$ where T_0 is the sampling period. If the sampling period is chosen such that it satisfies the Whittaker–Nyquist–Kotelnikov–Shannon (or Nyquist) sampling theorem (see e.g., Nyquist, 1928a; Shannon, 1949) then the copies do not overlap and simply applying a low-pass or band-pass filter leads to a copy-free signal. However, in practice it is far more useful to sample the continuous-time (CT) signal with a frequency greater than the Nyquist frequency since this reduces the restrictions on the filter design and leads to a better approximation (cf. Section 3.2.2).

3.2.2 *Truncation error*

A practical difficulty that has been ignored so far in this discussion is the effect of truncation and we comment on it here, following Strohmer and Tanner (2006). The Nyquist sampling theorem assumes an infinite number of samples but all practical signals are finite. If this finite length of (the set of) samples is L then this introduces an error of the order of $1/\sqrt{L}$, i.e., if the signal (i.e. the Fourier tansform of the voltage stream) has a peak amplitude a, the uncertainty in the sampled peak is $1/\sqrt{L} \times a$. This error (called the truncation error) is dependent on the localisation (in Fourier space) of the *atom*[5] of the signal and can be expressed as (Strohmer and Tanner, 2006):

$$\begin{aligned} \epsilon\,(t, L, T) &\equiv \left| f(t) - \frac{\sqrt{2\pi}}{2\sigma}\sum_{|k|\le L} f\left(\frac{k}{2\sigma}\right)\psi\left(t - \frac{k}{2\sigma}\right)\right| \\ &\le \frac{\sqrt{2\pi}}{2\sigma}\cdot \|f\|_{L_\infty}\sum_{|k|<L}\psi\left(t - \frac{k}{2\sigma}\right) \end{aligned} \tag{3.9}$$

The atom $\psi\,(t)$ in the classical Nyquist sampling theorem is the *sinc* function. However, the *sinc* decays very slowly; at the rate $\lim_{\tau\to\infty}\psi(\tau) \sim$

[5] An atom here denotes a function whose Fourier dual has compact support over the same basis as the Fourier equivalent of the signal or sample (In other words, the Fourier dual has a finite value in the region in which the Fourier dual of the signal is defined and is zero everywhere else). The simplest Fourier dual is a rectangular function, $\chi_{[-\sigma,\sigma]}$, whose time-domain dual is the *sinc*. The terminology used here relates to the theory of wavelet transforms (see e.g., Mallat, 2008)

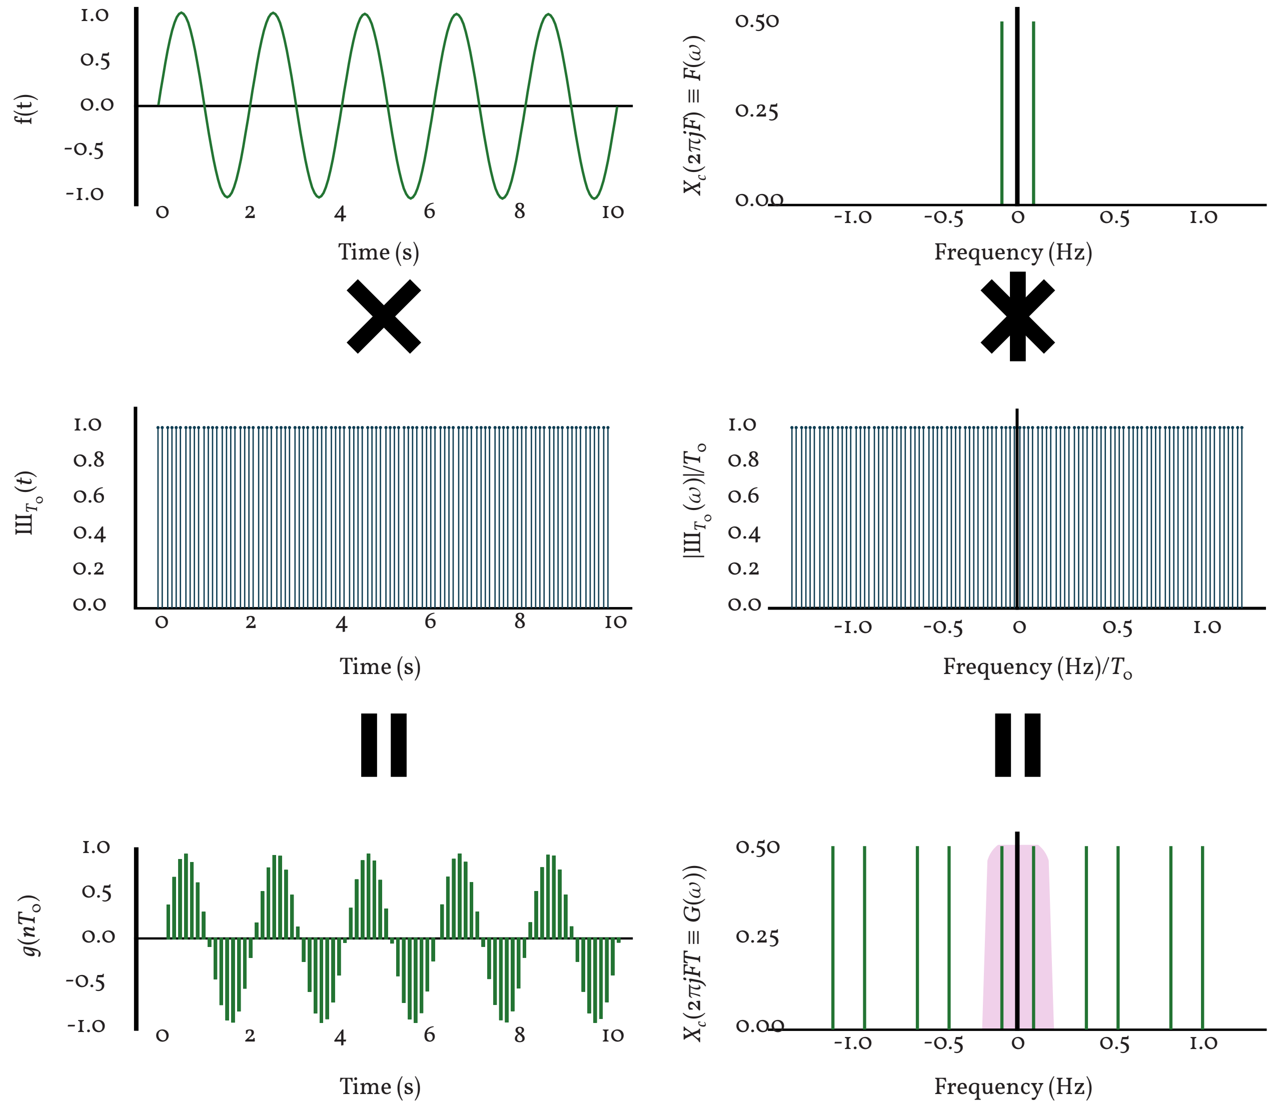

Plate 3: Sketch of the sampling process, where a continuous function $f(t)$ is sampled by convolution with a Shah function (or Dirac comb) $III_{T_o}(t)$ resulting in the sampled equivalent, $g(nT_o)$. The left hand side shows the sampling operation in the time-domain while the right hand side of the plot shows the frequency-domain equivalent of the operations. The frequency domain plots have been recentred to the frequency of interest. The green bars in the bottom-left panel show discrete amplitudes recovered by the sampling process and as such a number of Fourier components can pass through them. These Fourier components will produce infinite multiples (or aliases) of the input frequency to appear in the sampled signal, as shown by the green lines in the bottom-right plot. To remove the aliases, a bandpass filter (in the case of complex valued sampling) is applied, shown by the pink shaded region.

$1/\tau$. This implies that the rate at which the reconstructed function converges to the original function is of the first order which can produce a 'rippled' output at the end of reconstruction. This is the well-known Gibbs phenomenon, shown in Figure 3.3.

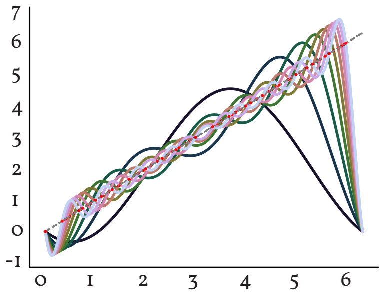

Figure 3.3: Gibbs phenomenon showing the incomplete sampling of an infinite, continuous function sampled via a number of increasing Fourier interpolations. As the line is sampled using more interpolations, the reconstruction moves towards the original signal, as shown by the lines moving from the darkest shade to the brightest. However, with increased interpolations, there is a sharp peak at the edges of (or any sharp transition inside) the finite region over which the interpolation is carried out. This is the Gibbs overshoot.

3.2.3 *Gibbs phenomenon*

Gibbs phenomenon, in one of its many interpretations, deals with the issue of recovering point values of a function from its expansion coefficients. A detailed analysis of Gibbs phenomenon is made in Gottlieb and Shu (1997), from which the following introductory discussion is summarised.

Consider the following problem:

Given 2N+1 Fourier coefficients $\hat{f}_k$, for $-N \geq k \geq N$, of an unknown function $f(x)$ defined everywhere in $-1 \geq x \geq 1$, construct accurate point values of the function.

The simplest method to solve this is to construct the classical Fourier sum:

$$f_N(x) = \sum_{k=-N}^{N} \hat{f}_k e^{ik\pi x}. \tag{3.10}$$

If $f(x)$ is smooth and periodic, this leads to a very good reconstruction of the point values of $f(x)$. In fact, if $f(x)$ is analytic and periodic, the Fourier series converges exponentially fast and the reconstructed signal is a near-perfect representation of the original signal if the sampling is sufficiently dense, i.e.:

$$\max_{-1 \leq x \leq 1} |f(x) - fN(x)| \leq e^{-\alpha N}, \quad \alpha > 0 \tag{3.11}$$

However, if $f(x)$ is either discontinuous or non-periodic, then $f_N(x)$ is no longer a good approximation to $f(x)$.

The Gibbs phenomenon leads to two distinct features in the approximation, which are discussed below for the specific case of a discontinuous function, which is more general and relevant to the case of filtering or channelisation.

- *Ringing artefacts*

  The convergence of Eqn. (3.10) is rather slow even at a finite distance from the discontinuity. If x_0 is a fixed point in $(-1;1)$

  $$|f(x_0) - f_N(x_0)| \sim O\left(\frac{1}{N}\right) \tag{3.12}$$

  This phenomenon produces artefacts known as ringing artefacts. Ringing artefacts are additional frequencies that appear to be introduced into the spectrum due to the finite sampling points. Specifically, ringing artefacts appears inside the band-pass and are visible as low-level fluctuations in it, which do not disappear even if the number of Fourier coefficients is very high, as demonstrated by the coloured curves in Figure 3.4 which are sampled representations of the continuous signal represented by the gray dotted line.

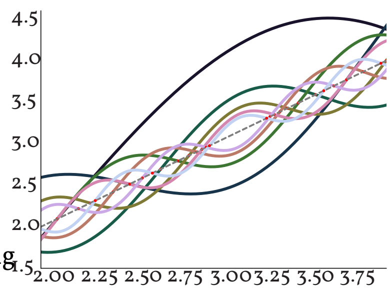

Figure 3.4: As the number of Fourier coefficients is increased, the function is sampled better but very high frequency ripples appear in the sampled signal, although the input signal was unrippled.

- *Gibbs overshoot*
  There is an overshoot, close to the boundary, that does not diminish with increasing N; thus

$$\max_{-1\leq x\leq 1} |f(x) - f_N(x)| \tag{3.13}$$

  does not tend to zero. In terms of the frequencies in the sample, near a discontinuity the Fourier component represents a sinc function whose width is a function of the sampling frequency but whose amplitude remains constant, as shown in Figure 3.5.

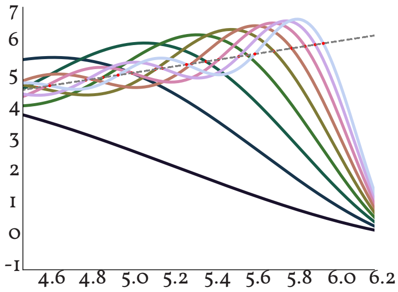

Figure 3.5: As the number of Fourier coefficients is increased, the maximum amplitude of the sampled signal at the edges of any discontinuity rises slowly to a small value above the input signal. This is called the Gibbs overshoot.

The Gibbs phenomenon seems to imply that it is inherently impossible to obtain accurate local information (point values) from the knowledge of global properties (Fourier coefficients) for piecewise smooth functions. However, there exist at least two methods using which it is possible to significantly mitigate this effect (Strohmer and Tanner, 2005).

In most cases, the focus of filter designers has been on reducing the effect of the overshoot. Primarily, these efforts rely on the design of 'windowing' functions, which maximise the local weight of the atom of Eqn. (3.9),

$$W_{\psi,local} = \frac{\int_{-R}^{R} \psi^2(t)dt}{\int_{-\infty}^{\infty} \psi^2(t)dt}. \tag{3.14}$$

In the case of Nyquist sampling, the only possible window is the rectangular window. However, if the sampling is carried out with a sampling rate higher than the Nyquist rate (oversampling), it is possible to construct a much larger family of windowing functions. There are several windowing functions that are commonly used. These include, but are not limited to, the rectangular, triangular, Blackman, Hamming, Hann, Blackman-Harris, Kaiser, Gaussian and the exponential windowing functions. However, the success of these windowing schemes depend greatly on the problem being addressed and there is no general method.

Other methods include addressing the design of the filter itself to make the filter functions 'smoother' . The classical raised cosine and the root raised cosine are examples of such filters.

However, both of these techniques are often difficult to implement with limited computational resources. As such, while it remains possible to significantly mitigate the effect of the Gibbs phenomenon (Gottlieb and Shu, 1997), most practical implementations, especially on FPGAs use only best-approximation approaches.

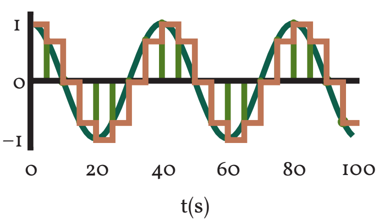

Figure 3.6: A continuous in time signal sampled with an impulse train (green lines) produces the quantised, sampled representation in brown. The sampled signal therefore contains small errors in the positions of the rising and falling edges, as well as inherent variations in the levels of the discrete steps. If we attempt to reconstruct the input signal from the brown curve, it is apparent that the recovered signal may look quite distinct from the input signal.

3.2.4 Quantisation & Rounding errors

In transitioning from an abstract, mathematical filter function to a digital filter implemented inside some computing device, it is necessary to abandon the infinitely well-defined function in favour of a finitely accurate digital approximation, as shown in Section 3.2.4. In practice, this introduces additional errors in the signals that are being processed within the digital system.

Quantisation effects in digital filters can be divided into four main categories:

- quantisation of the filter coefficients,
- errors due to analog-digital conversion,
- errors due to round offs in the arithmetic,
- and a constraint on signal level (i.e., the dynamic range) due to the requirement that overflow be prevented in the computation.

Explicit expressions for these terms can be found in Oppenheim and Weinstein (1972); Oppenheim and Johnson (1972). Statistical estimates of the introduced error can be found in Weinstein (1969a) or Weinstein (1969b).

Thus even before a FB is applied, a digitised signal contains four major sources of error. While aliasing due to sampling is effectively dealt with a simple band-pass filter (for complex data), Gibbs phenomenon is typically dealt with only in an approximate manner and quantisation and rounding errors are completely defined by the system in use. In most cases, system designers define a maximum permissible error limit for each of the separate sources of error, such that the signal is recovered without suffering any serious degradation. For instance, it is very common in radio astronomy to use 3-bit sampling with low ripple filters and 8-bit arithmetic inside the FPGAs.

3.3 *Polyphase Filterbanks*

A common algorithm for the design of polyphase FBs proceeds as follows:

1. Determine the number of branches depending on the maximum acquired bandwidth and the maximum useful filter length. This involves trading the filter quality for speed but in many pulsar applications, the filters are limited to 4 or 8 taps or filter coefficients.
2. Design an analysis FB using filters of the length chosen in the previous step.
3. Define sub-processing sections, if any. For example, for a coherent dedispersion pulsar backend, this would involve the dedispersion sections.
4. Design a synthesis FB which optimally reconstructs the input signal.

This rather general scheme can be optimised by adopting structures where filter design is highly redundant (i.e., the filters in each decimated channel are as similar as possible to the others). The most efficient to compute are analysis and synthesis FBs where all the filters are modulated version of one prototype filter. This automatically places strong constraints on the kind of sub-processing that is possible within the FB. A slightly more flexible approach follows from designing the analysis filters using a prototype, inserting the necessary sub-processing and then designing the synthesis filters using a different prototype filter. Reproduced below is the well-known example of the two-channel QMF.

3.4 *The two-channel QMF and the alias component matrix*

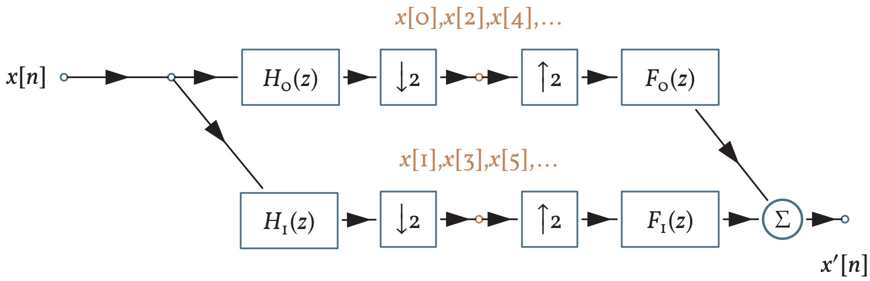

Figure 3.7: Signal flow representation of a two-channel quadrature modulation filterbank.

For the two-channel QMF shown in Figure 3.7, the z-transform of the output of the analysis section is (Vaidyanathan, 1998) :

$$X'_k(z) = H_k(z)X(z) \quad k = 0, 1 \tag{3.15}$$

The frequency domain equivalent of the decimate-by-two operation is (Strang and Nguyen, 1996):

$$v(\omega) = \frac{1}{2}\left[x\left(\frac{\omega}{2}\right) + x\left(\frac{\omega}{2} + \pi\right)\right] \tag{3.16}$$

Since the z-transform is defined such that $z = e^{i\omega}$, setting the frequency to $\omega/2$ implies $e^{i\omega/2} = z^{1/2}$. Similarly, an addition of π sets $e^{i\omega/2+\pi} = -z^{1/2}$. This leads to:

$$V_k(z) = \frac{1}{2}\left[X_k(z^{\frac{1}{2}}) + X_k(-z^{\frac{1}{2}})\right] \quad k = 0, 1 \tag{3.17}$$

which shows that the input signal is repeated every 2π units of frequency, as shown in figure 3.8.

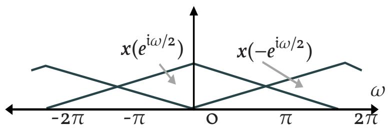

Figure 3.8: Aliasing due to the analysis filterbank of the two-channel QMF.

If $Y_k(z)$ is the result of the subsequent interpolate-by-two operation, shown by the $\uparrow 2$ in figure 3.7 it is equivalent to replacing z by z^2 in the z-transform and we can simplify the expression further.

$$Y_k(z) = V_k(z^2) = \frac{1}{2}\left[X_k(z) + X_k(-z)\right] \tag{3.18}$$

This results in the production of a number of images, as shown in figure 3.9.

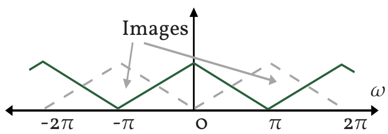

Figure 3.9: Imaging produced due to synthesis filterbank of the two-channel QMF.

From Equations (3.15) and (3.18), we have:

$$Y_k(z) = \frac{1}{2}\left[H_k(z)X(z) + H_k(-z)X(-z)\right] \tag{3.19}$$

If $F_0(z)$ and $F_1(z)$ are the synthesis filters, then the reconstructed output $X'(z)$ is:

$$\begin{aligned} X'(z) &= F_0(z)Y_0(z) + F_1(z)Y_1(z) \\ &= \frac{1}{2}\left[H_0(z)F_0(z) + H_1(z)F_1(z)\right]X(z) \\ &+ \frac{1}{2}\left[H_0(-z)F_0(z) + H_1(-z)F_1(z)\right]X(-z) \end{aligned} \tag{3.20}$$

This can be rewritten in the matrix notation as:

$$2X'(z) = [X(z)\, X(-z)] \underbrace{\begin{bmatrix} H_0(z) & H_1(z) \\ H_0(-z) & H_1(-z) \end{bmatrix}}_{\text{ACM}} \begin{bmatrix} F_0(z) \\ F_1(z) \end{bmatrix} \tag{3.21}$$

The matrix ACM in Eqn. (3.21) is called the *Alias Component Matrix*.
For the linear time variant (LTV)[6] case, we can rewrite Eqn. (3.20) as:

[6] A linear time variant is any system in which the output (or matrix of outputs) $y(t)$ at any time t can be expressed via the action of an operator matrix $\mathscr{O}(t)$ on the input (or matrix of inputs), i.e.: $x(t)$

$$y(t) = \mathscr{O}(t)x(t).$$

$$\begin{aligned} X'(z) &= T(z)X(z) + A(z)X(-z) \\ \Rightarrow X'(z) &= \sum_k \left(t(k)x(k) + (-1)^{k-n}a(k)x(n-k)\right) z^{-k} \end{aligned} \tag{3.22}$$

where t(n) is the impulse response function of T(z), the transfer function of the desired (i.e., mth) channel and a(n) is the impulse response of A(z), the alias transfer function (i.e., the ACM), respectively.

Using Eqn. (3.20), we construct the following even and odd series,

$$\begin{aligned} g_0(k) &= t(k) + (-1)^k a(k) \quad k = 0, 1 \\ g_1(k) &= t(k) - (-1)^k a(k) \end{aligned} \tag{3.23}$$

which can be used to represented the reconstructed equivalent of the sampled function.

$$x'(n) = \begin{cases} \sum_k g_0(k)x(n-k) & \text{for even n} \\ \sum_k g_1(k)x(n-k) & \text{for odd n} \end{cases}$$

We can now proceed to derive the explicit expressions for the distortions contained in the output of a full FB.

3.4.1 *Amplitude and phase distortions*

We want to simplify the expressions derived above such that the undesired effects of the filters in the FB are minimised. Hence, to cancel aliasing we choose

$$F_0(z) = H_1(-z) \,\&\, F_1(z) = -H_0(-z) \tag{3.24}$$

This leads to

$$\frac{F_0(z)}{F_1(z)} = -\frac{H_1(-z)}{H_0(-z)} \tag{3.25}$$

which implies $A(z) = 0$. Thus the expression for the reconstructed output reduces to:

$$X'(z) = T(z)X(z) \tag{3.26}$$

where

$$T(z) = \frac{1}{2}\left[H_0(z)F_0(z) + H_1(z)F_1(z)\right] \tag{3.27}$$

is called the distortion transfer matrix.

Given our choice of synthesis filters, $F_0(z)$ and $F_1(z)$ in Eqn. (3.24), the distortion transfer matrix reduces to:

$$T(z) = \frac{1}{2}\left[H_0(z)H_1(-z) - H_1(z)H_0(-z)\right]. \tag{3.28}$$

To enable us to derive the conditions for which amplitude and phase distortions occur, we set $z = e^{i\omega}$ and writing out $T(z)$ as a product of the amplitude $|T(z)|$ and phase $e^{i\phi(\omega)}$, i.e.,

$$T(e^{i\omega}) = |T(e^{i\omega})|e^{i\phi(\omega)} \tag{3.29}$$

in Eqn. (3.26) we get:

$$X'(e^{i\omega}) = |T(e^{i\omega})|e^{i\phi(\omega)}X(e^{i\omega}) \tag{3.30}$$

This leads to the following conditions:

- If $|T(e^{i\omega})| \neq$ constant, amplitude distortion occurs.
- If $\phi(\omega) \neq a + b\omega$ for constant a & b, phase distortion occurs.

The conditions listed above are the frequency domain measures of PR.

However, in multichannel PFBs and in general, for any PFB with additional processing between the analysis and synthesis sections, it is difficult to simultaneously produce ideal filters and satisfy the PR requirements. The conditions in most practical designs are then relaxed to the NPR where the measure of how 'distant' the design is from PR can be quantified using the following terms:

- Amplitude Distortion

$$e_m(\omega) = 1 - |T_0(e^{i\omega})|^2 \quad \text{for } \omega \in [0, \pi] \tag{3.31}$$

  where T_0 is the expected response of branch processing alone.

- Group delay distortion

$$e_{gd}(\omega) = \tau_T - arg\left[T_0(e^{i\omega})\right] \quad \text{for } \omega \in [0, \pi] \tag{3.32}$$

  where τ_T is the desired group delay of the filter

- Worst case aliasing error

$$e_a(\omega) = \max_{1 \leq l \leq M-1} |T_l(e^{i\omega})| \quad \text{for } \omega \in [0, \pi] \tag{3.33}$$

  where T_l is the undesired aliasing response for the given branch.

The first error term $e_m(\omega)$ involves the amplitude response (i.e., passband behaviour) of the prototype filter and accounts for its deviation from ideal performance, the second $e_{gd}(\omega)$ accounts for the non-ideality of the phase response of the prototype filter and the third $e_a(\omega)$ measures the cumulative effect of aliasing in the neighbouring due to each of the filters in the individual branches. The FB is said to be NPR if $e_m(\omega) \rightarrow 0$; $e_{gd}(\omega) \rightarrow 0$ & $e_a(\omega) \rightarrow 0$.

While it is possible to directly use these equations to construct linear expressions which can then be minimised, these equations can often result in unwieldy solutions, which in turn are computationally expensive to process and therefore not very useful for resource-limited systems like those used in pulsar astronomy. Often designers rely on constructing expressions for the errors introduced due to each part of the FB. These can be used to construct a 'cost function' which, if constructed suitably, can be easily optimised to minimise the various terms of the 'error budget'. This almost always implies that PR has been discarded in favour of NPR. In the following section, the specific example of a critically sampled, modulated FB (which is the most commonly used FB in pulsar astronomy) is investigated and we derive expressions showing that in the case of coherent dedispersion systems, even NPR may be difficult to achieve with simple digital filters.

3.5 *Least-Squares Optimisation*

A very common measure of the distance from PR can be stated in terms of the energy 'lost' to the error terms. The objective then is to minimise the energy removed by the error function for a filter(bank), given that the signal carries finite energy and the filter(bank) itself does not add any energy to the signal. Under such considerations, this is the same as minimising the energy in the stop-band. For an M-channel filter bank, the least squares criterion can be stated as:

$$\min_{h_m, g_m} \int_{\pi}^{-\pi} |E_{h_m, g_m}(\omega)|^2 d\omega, \tag{3.34}$$

where $E_{h_m,g_m}(z)$ is the error function which depends on the impulse responses of the analysis and synthesis filters h_m and g_m, and contains the desired properties of the filter bank. The least squares error in Eqn. (3.34) is, in its most general form, a set of non-linear equations, which must then be solved using non-linear optimisation procedures. However, if the analysis filters are modulated versions of a single prototype filter $H_0(z)$, and the synthesis filters are modulated versions of a prototype filter $G_0(z)$, such that $H_0(z) \prec G_0(z)$[7]; then the design problem can be divided into two sequential quadratic optimisation problems, following de Haan (2001), from which we reproduce the most relevant equations below.

[7] which implies that the analysis filters are defined before the synthesis filters, since the properties of the second depend on the first.

3.5.1 *Uniformly Modulated FB.*

A uniformly modulated FB consists of M branches in each of which a single prototype filter is modulated[8] to produce the individual branch filters. In the simplest case this implies that the prototype is employed to construct both the analysis and synthesis FBs. However, in most practical solutions, the synthesis filters are defined as separate filters which depend on the analysis filter.

[8] I.e., all the filters in the filterbank are obtained by applying a single transformation to the prototype filter. Typically, this should be at most a scaling operation combined with a phase transformation since we are using complex valued filters.

Consider an M-channel modulated filter bank, each branch of which consists of the following elements :

- a decimator with decimation factor D,
- an analysis filter $H_m(z)$,
- a synthesis filter $G_m(z)$ and
- an interpolator with interpolation factor $I = D$.

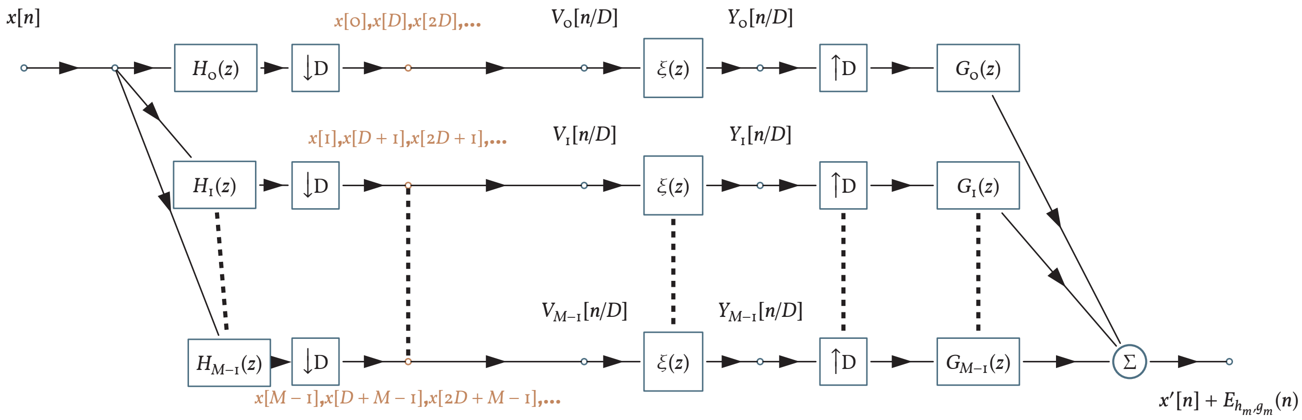

Figure 3.10: Signal flow representation of an M-channel polyphase FB that includes subband processing.

One could simplify the following discussion by forcing the decimation factor to be the same as the number of branches, i.e., follow the critically decimated scheme and set $M = D$, however, since we are only interested in the investigation of whether it is possible to easily construct an NPR FB we do not force this condition.

Let the analysis filters be FIR filters of length L_h. FIR filters are generally preferred since they require shorter filter lengths compared to their IIR counterparts and are linear in phase, leading to ease of construction of an 'analysis + synthesis' FB. In order to construct a uniform filter bank, i.e., with sub-bands of equal width, we define low pass analysis filters, H(z). Since all the analysis filters in the filter bank are modulated versions of the prototype analysis filter we can express the individual filters as:

$$h_m(n) = h(n)W_M^{-mn} = h(n)e^{i2\pi\frac{mn}{M}} \leftrightarrow H_m(z) = H(zW_M^m) \tag{3.35}$$

where $W_M = e^{-i2\frac{\pi}{M}}$. In the trivial case (m = 0) obviously the analysis filters reduce to their respective prototype filters, $H_0(z) = H(z)$. In order to analyse the properties of this (analysis) FB we derive an input-output relation.

$$V_m(z) = H_m(z)X(z) = H(zW_M)X(z). \tag{3.36}$$

Down-sampling (or decimation) is equivalent to expanding the spectra of the signal in each branch, i.e., the same expansion as in Equations (3.16) and (3.17):

$$\begin{aligned} X_m(z) &= \frac{1}{D}\sum_{d=0}^{D-1} V_m\left(z^{1/D}W_D^d\right) \\ &= \frac{1}{D}\sum_{d=0}^{D-1} H\left(z^{1/D}W_M^m W_D^d\right) X\left(z^{1/D}W_D^d\right) \end{aligned} \tag{3.37}$$

where D is the decimation factor and $W_D = e^{-i2\frac{\pi}{D}}$. The summation in Eqn. (3.37) shows that the sub-band signals consist of D aliasing terms. Depending on the sub-band index and the decimation factor, the desired spectral content is present in one or more aliasing terms. In general, therefore an analysis FB will introduce a set of artefacts into each channel that is processed. If a reconstruction is desirable, then these artefacts can be cancelled using a well designed synthesis FB.

In most practical applications, the sub-band signals, $X_m(z)$, are typically further processed before synthesis. In the simplest case, we can consider them to be an additional set of filters, denoted by $\xi_m(z)$. The processed sub-band signals, $Y_m(z)$, are then given by

$$Y_m(z) = \xi_m(z)X_m(z) \tag{3.38}$$

where $X_m(z)$ is the branch input signal.

As discussed in Section 2.2, pulsar signals suffer from a quadratic frequency dependent delay, i.e., dispersion, due to propagation through the ionised ISM (IISM) which must be removed. A computationally inexpensive method to mitigate the effect of dispersion, i.e., to perform *dedispersion*, is to introduce in to each branch of the FB a delay [9] estimated such that the individual branch signals are aligned with respect to the rotational phase of the pulsar. This is known as incoherent dedispersion. Since these are zero-phase delay only blocks and if the desired frequency response of the analysis prototype filter is designed so that the transfer functions of the analysis filters have power-complementary transfer functions, i.e. the sum of the squared filter magnitudes is unity (Vaidyanathan, 1993),

[9] or a delay-only block in the language of DSP.

$$\sum_{m=0}^{M-1} |H_m(e^{i\omega})|^2 = 1, \quad \omega = [-\pi, \pi] \tag{3.39}$$

we can define a distance function as follows

$$\epsilon_h = \frac{1}{2\omega_p}\int_{-\omega_p}^{\omega_p} |H_m(e^{i\omega}) - H_d(e^{i\omega})|^2 d\omega \tag{3.40}$$

where $H_d(z)$ is a desired frequency response of the prototype analysis filter in the pass band region $\Omega_p = [-\omega_p, \omega_p]$. The desired frequency response is then given by

$$H_d(e^{i\omega}) = e^{-i\omega\tau_H}, \text{ such that } \omega \in \Omega_p \tag{3.41}$$

where τ_H is the desired group delay of the prototype analysis filter of Eqn. (3.35). The pass band response error (for the analysis FB alone, since we have already determined that incoherent dedispersion is a delay only operation) is:

$$\alpha_h = \frac{1}{2\omega_p}\int_{\omega_p}^{\omega_p} |H(e^{i\omega}) - H_d(e^{i\omega})|^2 d\omega \tag{3.42}$$

Similarly, we can define the inband-aliasing distortion:

$$\beta_h = \frac{1}{2\pi D^2}\sum_{m=0}^{M-1}\sum_{\substack{d=(M-m)\\ modM}}\int_{-\pi}^{\pi} |H(e^{i\omega/D}W_M^m W_D^d)|^2 d\omega \tag{3.43}$$

where all inband-aliasing terms are included.

In summation in Eqn. (3.43), for the critically sampled ($M = D$) case M equal terms are repeated M times in the summation. Since this is a modulated FB we can drop all the terms apart from those in the first sub-band (m = 0), i.e. the terms for $d = 1, ..., D-1$. Therefore, β_h in Eqn. (3.43) can be rewritten as:

$$\beta_h = \frac{1}{2\pi D^2} \int_{-\pi}^{\pi} \sum_{d=1}^{D-1} |H(e^{i\omega/D} W_D^d)|^2 \tag{3.44}$$

The distance function can now be rewritten as:

$$\epsilon_h = \alpha_h + \beta_h. \tag{3.45}$$

Following de Haan (2001, also see original derivations there), we can rewrite Equations (3.42) and (3.44) in terms of the impulse response functions of the filters:

$$\begin{aligned} \alpha_h &= h^T A h - 2h^T b + 1 \text{ and} \\ \beta_h &= h^T C h \end{aligned} \tag{3.46}$$

where we have expanded Equations (3.42) and (3.44) and substituted using Eqn. (3.35). The quantities A and C are matrices while b represents a vector. This allows us to rewrite Eqn. (3.45) as:

$$\epsilon_h = h^T (A + C) h - 2h^T b + 1. \tag{3.47}$$

and the ideal analysis prototype filter can be found by minimising this function. In proper notation, this becomes:

The notation $\underset{h}{argmin}$ implies we are interested in the minima of argument of the function, where we are modifying the impulse response $h(n)$ to achieve this.

$$h_{opt} = \underset{h}{argmin}\, h^T (A + C) h - 2h^T b + 1, \tag{3.48}$$

and the minimisation is achieved by solving the set of linear equations

$$(A + C) h = b \tag{3.49}$$

The preceding discussion is generally applicable in the case that the action of the sub-processing introduces either zero or constant phase variation in the signal[10]. However, this is not very useful for the most sensitive pulsar backends, for reasons discussed below.

[10] This is applicable to Section 3.4 and typically, for all FIR filter based PFBs.

A much more powerful method of dedispersion than discussed above involves convolving the sampled stream with the inverse of the dispersion action and is called coherent dedispersion (see Section 2.2). While coherent dedispersion was originally introduced by Hankins and Rickett (1975), a reimplementation by Stairs et al. (2000) proceeds by splitting the incoming signal into several (even) branches and followed by Fourier transformation. Following this, a multiplication in the Fourier domain by a specially constructed chirp-like signal is carried out. This function depends on the DM and is essentially the inverse of the ISM transfer function. Finally, an inverse FFT is carried out and the resulting data stream is further processed to obtain the final data products.

This is a description that lends itself very well to the PFB architecture. In this case, the action of multiplication with the inverse of the ISM transfer function is easily represented as the action of a filter with following form:

$$H_{ISM}(\omega_0 + \omega) = e^{-i \cdot \frac{2\pi \mathscr{D}}{(\omega + \omega_c)\omega_c^2} \omega^2}, \tag{3.50}$$

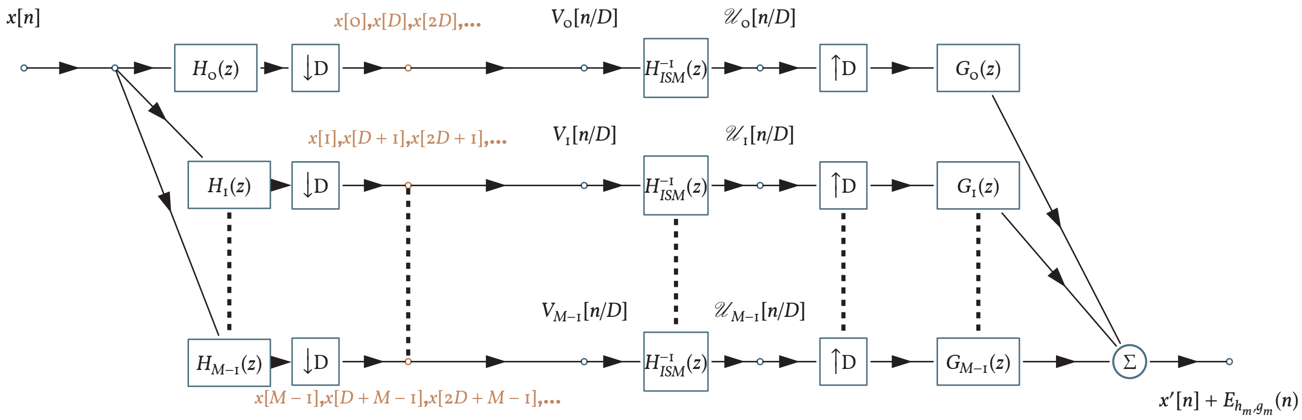

Figure 3.11: Signal flow representation of an M-channel PFB that includes coherent dedispersion.

where $\mathscr{D}$ is given by:

$$\begin{aligned} \mathscr{D} &= \frac{e^2}{2\pi m_e c}\left\{\frac{1}{\omega_L^2}-\frac{1}{\omega_U^2}\right\}\int_0^L n_e dl \\ &= \frac{e^2}{2\pi m_e c}\left\{\frac{1}{\omega_L^2}-\frac{1}{\omega_U^2}\right\}\mathrm{DM}. \end{aligned} \tag{3.51}$$

The z-transform representation of this transfer function is:

$$H_{ISM}\left(\omega_0+\omega\right) \equiv H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) \tag{3.52}$$

The resulting output of each channel of the analysis FB after coherent dedispersion is then given by:

$$\begin{aligned} \mathscr{U}_m(z) &= \frac{1}{D}\sum_{d=0}^{D-1} H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) V_m\left(z^{1/D}W_D^d\right) \\ &= \frac{1}{D}\sum_{d=0}^{D-1} H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) H\left(z^{1/D}W_M^m W_D^d\right) X\left(z^{1/D}W_D^d\right) \end{aligned} \tag{3.53}$$

For the 0th channel the output can be rewritten as two terms; the desired sub-band signal and the error terms.

$$\begin{aligned} \mathscr{U}_0(z) = &\overbrace{\frac{1}{D}H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) H\left(z^{1/D}\right) X\left(z^{1/D}W_D^d\right)}^{\textit{Desired Response}} \\ &\underbrace{+\frac{1}{D}\sum_{d=1}^{D-1} H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) H\left(z^{1/D}W_M^m W_D^d\right) X\left(z^{1/D}W_D^d\right)}_{\textit{Error terms}} \end{aligned} \tag{3.54}$$

The dedispersion transfer function is therefore no longer zero-phase or linear phase. Hence, the synthesis FB must now account for an additional non-linear phase transfer function. This can be exploited to simplify the design of the synthesis prototype by removing the requirement for linear phase FIR using, e.g., the Parks-McClellan scheme (Parks and McClellan, 1972), which is based on the Remez exchange algorithm (Remez, 1934), to obtain the necessary filter. In the ideal case that the number of branches is sufficiently high, the bandwidth of the prototype filter can be made small enough that the non-linear phase change is well-modelled by a linear function of phase and the assumptions on the design fall back to the case of incoherent dedispersion.

Only if we can assume that the subband processing introduces zero or constant phase delays, we can then return to the de Haan (2001) scheme and define an additional set of error functions; which resemble the error functions defined for the analysis FB. Thus we define a total response error for the analysis+synthesis FB :

$$\gamma_g(h) = \frac{1}{2\pi}\int_{-\pi}^{\pi} |\sum_{m=0}^{M-1} A_{m,0}(e^{i\omega}) - H_d(e^{i\omega})|^2 d\omega \tag{3.55}$$

and the residual aliasing distortion:

$$\delta_g(h) = \frac{1}{2\pi}\sum_{d=0}^{D-1}\sum_{m=0}^{M-1} |A_{m,d}(e^{i\omega})|^2, \tag{3.56}$$

The term

$$A_{m,d} = \frac{1}{D}\xi_m\left(z^D\right) H\left(zW_M^m W_D^d\right) G\left(zW_M^m\right) \tag{3.57}$$

which is present in both Equations (3.55) and (3.56) (in Eqn. (3.55) only the $d = 0$ term is retained) is derived by expanding Eqn. (3.38) and rearranging such that it becomes

$$Y(z) = \sum_{d=0}^{D-1}\sum_{m=0}^{M-1} A_{m,d}(z) X\left(zW_D^d\right). \tag{3.58}$$

The optimal prototype synthesis filter, in terms of minimal total response error and minimal energy in the aliasing components, is found by minimising the objective function:

$$\epsilon_g(h) = \gamma_g(h) + \delta_g(h) \tag{3.59}$$

Inserting Equations (3.55) and (3.56) into Eqn. (3.59) yields

$$\epsilon_g(h) = g^T(E+P)g - 2g^T f + 1 \tag{3.60}$$

The solution

$$g_{opt} = \underset{g}{argmin}\, \epsilon_g(h) \tag{3.61}$$

can be found by solving the set of linear equations

$$(E+P)g = f \tag{3.62}$$

At this stage, it is pertinent to mention that we have investigated the application of the de Haan (2001) least-mean-squares (LMS) algorithm to the simplified problem only. Eventually, we find that a linear phase FIR synthesis FB alone is not sufficient to mitigate the artefacts of the analysis FB due to the phase terms in Eqn. (3.54) and some care must be taken to produce a performant synthesis FB. It is also proper to mention that in practical PFB designs, a few more necessary optimisations are commonly made and therefore the treatment above is only an oversimplified approximation. However, this does not impact our interpretations since we are interested in the most general investigation.

3.6 *Artefacts in pulsar backends*

The data recording systems or backends that are currently used for pulsar observations do not use full reconstruction filterbanks. Instead, high-speed FPGA boards, as shown schematically in Plate 2, typically perform only analysis on the incoming data. The analysed streams are then packetised and transmitted via high speed network links to a number of computers[II]. These computers, colloquially called 'recording machines' are then used to process and store the data. In the case of observing known pulsars, the incoming subbands are usually integrated over time by folding the subband data modulo the spin-period of the pulsar. For pulsar searching the data is recorded without any processing, in what is known as the 'raw' mode. In either case, the design of the analysis filters is critical and often specially shaped FIR filters with a small number of taps are applied to the analysis streams. These short length filters are often unable to suppress a number of artefacts, of which typically the most pronounced are the 'band-edge' artefacts.

[II] The technically correct term for such groups of computers is a 'server farm'.

A simple *Python* simulation which demonstrates the artefacts produced in an analysis FB based on short FIR filters was implemented using finite-precision arithmetic. The code takes as input the number of taps in the filter N_{taps}, the sampling frequency F_s in units of the centre frequency of the input signal f_c, the desired arithmetic precision N_{bit}, the S/N of the signal $C_{s/r}$ and the decimation factor D. The results from a run with $N_{taps} = 4$, $F_s = 2f_c$, $N_{bits} = 32$, $C_{s/r} = 10\,000$ and $D = 4$ are shown in Figure 3.12.

The top panel shows a synthetic, high S/N pulsar signal that is generated for exactly (hypothetical) one spin-period. The frequency axis is normalised so that the top of the input band corresponds to π and the bottom to $-\pi$. The bottom plot shows the central channel of the analysis FB after it has been filtered by a low-pass filter which has been shaped using the Kaiser window (Kaiser and Schafer, 1980). The Kaiser window is a windowing scheme that utilises Bessel functions to provide a very high degree of attenuation outside the passband while maintaining a decrease of 6 dB per octave of frequency from the passband to the stop-band (i.e. the filter roll-off).

Even though the Kaiser window is quite efficient at suppressing the majority of the aliasing artefacts, the bottom plot shows two main artefacts, which achieve maximum power near the edges of the band. Typically, the power is distributed in such a fashion as to make only the components near the edges visible in spectrograms (or frequency versus rotational phase plots for folded observations), leading to the name 'edge' or band-edge artefacts. Closer inspection (as well as an increase in the S/N) also reveals the presence of more low-power artefacts in the analysed spectrum.

Real pulsar signals contain RFI and often the S/N is much lower, leading to a portion of the artefacts being absorbed in the noise. However, as the sensitivity of the receivers increase and with increase in collecting area expected for telescopes like the the Square Kilometre Array (SKA), the S/N will increase significantly and even the low-power arte-

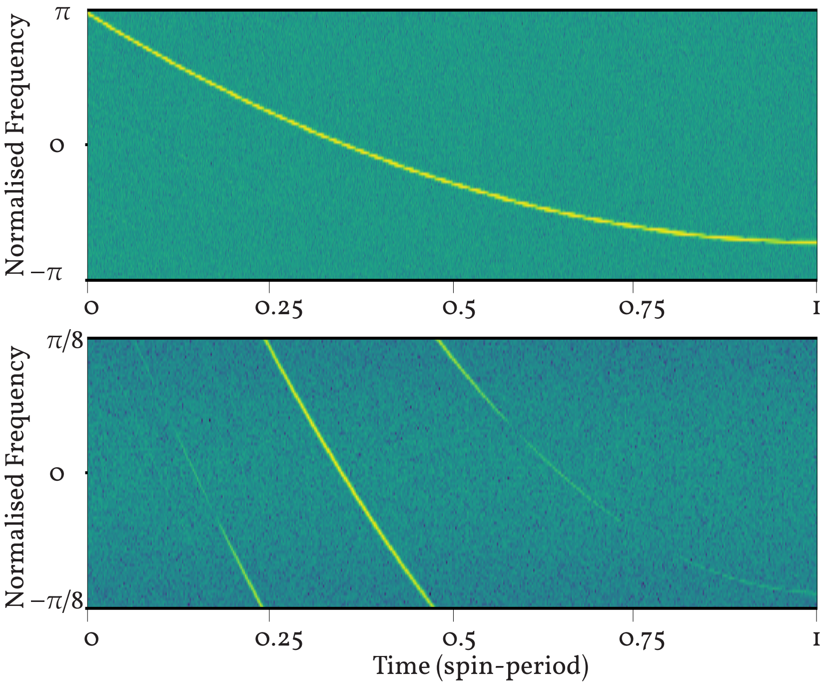

Figure 3.12: Simulation to study the generation of artefacts in an analysis FB based on 4-tap FIR filters. The FIR filters are shaped using the Kaiser window, which provides a smooth taper from the passband to the stopband and minimises ringing artifacts in the filtered signal.

facts should become significant. Of these, only the band-edge artefacts can typically be mitigated by reshaping the passband during post-processing of the data since modifying the passband far from the edges will also remove the most powerful parts of the signal. Thus, reshaping the channels in post-processing simply reduces the total power of the signal and reduces the S/N. For very faint pulsars, this reduction can render them invisible. Well designed analysis FBs are therefore necessary to make sensitive observations. To improve the performance of such analysis FBs we can optimise the filter response using the LMS algorithm of de Haan (2001), by solving Eqn. (3.49) to obtain the optimal analysis filters.

3.7 *Limits on the pulsar timing precision*

From the perspective of high-precision pulsar timing, the presence of artefacts leads to significant limits on the ultimate precision of the timing exercises and we comment below on this.

A very specific investigation that has been recently completed is presented in Morrison et al. (2015) where the authors construct a *Matlab*-based model of a PFB based polyphase filterbank. Using this model the effects of DSP on a simulated pulsar which is constructed to be similar to the brightest MSP known, PSR J0437−4715, are quantified. For the full details of the implementation and the algorithms we refer the reader to the original document and only reproduce selected results here.

From the perspective of high sensitivity timing, the most important aspect is the effect of the PFB artefacts on the timing precision. As mentioned earlier, this relies not only on the level of digitisation, which places a fundamental limit with which a pulse profile can be reconstructed from the recorded data but also the amount of power in the artefacts, which can reduce the precision with which we can measure the time of arrival of a pulse at an Earth-bound observatory.

Morrison et al. (2015) simulate the effect of a single 'ghost component', which they model as an attenuated image of the main pulse, which is allowed to drift along the rotational phase axis, relative to the main pulse. In our general demonstration above, we have shown that in practice the number of such components will actually be D-1, where D is the number of PFB channels, each appearing to lag behind the main pulse in a non-linear manner. The Morrison et al. (2015) treatment which is therefore equivalent to considering the effect of the brightest of the D-1 components, shows that the expected limit on the timing precision is ~100 ns for a ghost component which has a peak amplitude −34 to −37 dB relative to the main pulse (cf. Table 6-B of Morrison et al., 2015). It should be noted that unlike the considerations presented in the rest of this chapter, the filters considered by (Morrison et al., 2015) are typically longer. In an investigation of the required number of taps to adequately suppress stop band leakage and maintain a suitable pass band ripple,~1 dB, they suggest using 14-tap filters while current implementations at the 100-m Effelsberg Radio-Telescope use 8-tap filters.

3.8 Conclusions

We have reviewed currently available literature with the aim of finding the ultimate limits on the timing of MSPs due to artefacts in DSP based pulsar backends which employ PFBs. We find that even without PFBs digitisation introduces well-defined constraints on the accuracy of any real-life signal that has been sampled using a finite number of bits and is processed with precision limited arithmetic, as must be done in an FPGA based backend. We introduced the concept of a PR FB and then define the NPR FB using a distance measure. By using a well-known method where the distance measure is cast as a measure of energy contained in the various components of the PFB we show using only general arguments that for a coherent dedispersion pulsar backend, there is an additional non-linear term in the phase response of the dedispersion filter which is non-trivial to cancel. Similar work, using *Matlab*-based models has shown that the ultimate limit on the accuracy of MSP timing due to an artefact with relative power ~−34 to −37dB is about ~100 ns

4

21-year timing of the black-widow pulsar J2051−0827

月見れば
千々に物こそ
悲しけれ
わが身ひとつの
秋にはあらねど

¡Ay, ay, ay, ay!
Toma este vals que se muere en mis brazos.

— **Federico García Lorca;** *Pequeño vals Vienés*

Recycled pulsars are old pulsars which have been spun-up to high rotational rates by tranfser of angular momentum via infalling matter from a companion star. Presented below is a discussion of the long-term behaviour of the *black-widow* pulsar J2051−0827. This analysis uses a 21-year dataset from four European Pulsar Timing Array telescopes and the Parkes radio telescope. This dataset, which is the longest published to date for a black-widow system, allows for an improved analysis that addresses previously unknown biases. While secular variations, as identified in previous analyses, are recovered, short-term variations are detected for the first time. Concurrently, a significant decrease of $\sim 2.5 \times 10^{-3}$ cm^{-3} pc in the dispersion measure associated with PSR J2051−0827 is measured for the first time and improvements are also made to estimates of the proper motion. Finally, PSR J2051−0827 is shown to have entered a relatively stable state suggesting the possibility of its eventual inclusion in pulsar timing arrays.

– This chapter is an enhanced version of Shaifullah et al. (2016)

4.1 Introduction

Of the ~2600 pulsars known today, roughly 10% appear to have rotation-periods of the order of a few milliseconds and are known as MSP. Within the MSP population there exist a variety of configurations, however, most MSPs are found in binary systems. Among these, about 10% are in tight, eclipsing binaries. Such systems are further classified into the BWP systems, with very light companions of mass ($\dot{m}_c \lesssim 0.05$ M$_\odot$) and redback systems, with heavier companions (0.1 M$_\odot$ $\lesssim \dot{m}_c \lesssim 0.5$ M$_\odot$; Roberts, 2013; Chen et al., 2013). PSR J2051−0827 is the second black-widow system that was discovered (Stappers et al., 1996). Its companion is expected to be a ~0.02 to 0.06 $M_\odot$ star, whose exact nature is yet to be determined (see discussions in Stappers et al., 2001; Lazaridis et al., 2011).

Pulsar timing relies on making highly precise measurements of the time at which the radio-beam from a rotating pulsar crosses a radio telescope. These measured times are then compared to a theoretical prediction of these crossing events to derive various properties of the pulsar. A more extensive discussion on pulsar timing and the benefits of MSPs for pulsar timing can be found in Lorimer and Kramer (2005) and other reviews of pulsar timing.

MSPs are particularly well-suited for this because of their inherent stability and short rotation periods. Even though the pulsars in black-widow systems are MSPs they are typically excluded from high-precision pulsar timing experiments since several of them have been observed to display variability in their orbital parameters, in particular the orbital period. This variability may be due to many reasons like the interaction of the pulsar with the companion, the presence of excess gas around the companion's orbit or the companion's mass loss.

However, only a limited number of studies so far have tried to identify if the variability of such pulsars can be modelled by introducing new parameters into the pre-existing timing models or by defining new timing models for such systems. Given the recent increase in the number of MSPs detected, in large part from surveys of Fermi-LAT sources (Abdo et al., 2013), and the rapid growth in the sensitivity and bandwidth of modern digital receiver systems for pulsar timing making it possible to detect variations in much greater detail, it is pertinent to address this long-standing question.

PSR J2051−0827 has been continuously timed since its discovery in 1995 (Stappers et al., 1996) and therefore the dataset presented in the following analysis represents the longest timing baseline currently published for eclipsing black-widow systems. Given this long time-baseline and other favourable properties discussed in the following sections, this dataset offers an ideal opportunity to attempt such an exercise.

Previous pulsar timing analyses of PSR J2051−0827 have shown that the orbital period, P_b, and projected semi-major axis, x, undergo secular variations (Stappers et al., 1998; Doroshenko et al., 2001; Lazaridis et al., 2011). These variations are possibly linked to the variations of the gravitational quadrupole moment of the companion and induced by variations of the mass quadrupole of the companion as its oblateness varies due to rotational effects (Lazaridis et al., 2011). These variations may arise due to a differential rotation of the outer layers of the companion (Applegate and Shaham, 1994) or due to variations in the activity of the magnetic field of the companion as in the Lanza and Rodonò (2001) model. Similar variations have been measured for a few other pulsars in BWP systems like PSR J1959+2048 (PSR B1957+20; Fruchter et al., 1988), PSRs J0024−7204J and PSRs J0024−7204O (47 Tuc J and O; Freire et al., 2003), PSR J1807−2459A (NGC 6544A; Lynch et al., 2012) and PSR J1731−1847 (Ng et al., 2014).

The binary system containing PSR J2051−0827 has also been recently detected in *Fermi* and *Chandra* data (Wu et al., 2012). The γ-ray luminosity is 7.66×10^{32} erg s^{-1}. The inferred spin-down power, $\dot{E}$, from radio observations is $\sim 5.49 \times 10^{33}$ erg s^{-1}. The γ-ray luminosity, therefore, represents $\sim 15\%$ of the total spin-down power, which is consistent with other MSPs for which such a detection has been made. The γ-ray emission from the system appears to be well fit by a model of emission in the 'outer gap accelerator', as discussed in Takata et al. (2012). Using the new ephemerides presented here, it may be possible to detect the orbital dependence of pulsed emission from PSR J2051−0827.

The X-ray luminosity is 1.01×10^{30} erg s^{-1} (Wu et al., 2012) and the data do not present any evidence for bursts, which suggests that the companion is stable and does not undergo sudden deformations. The flux values fit well for a model with emission from the intra-binary shock, the polar caps and synchrotron emission from the pulsar magnetosphere (Wu et al., 2012).

This work provides an update on the timing of PSR J2051–0827 and presents an improved analysis. Two complementary timing models for PSR J2051–0827 are provided, one capable of handling small eccentricities and another, utilising orbital-frequency derivatives. A new method for measuring the variations in the orbital period, ΔP_{b}, by measuring the change in the epoch of ascending node, T_{asc} is also presented.

4.2 *Observations and Data Analysis*

The bulk of the dataset used for the timing analysis consists of pulse times-of-arrival (henceforth; ToAs) derived from data from four EPTA telescopes[1] and extend from 2003 to 2015. To extend the analysis and to test for consistency with previous analyses, ToAs (obtained from the Lazaridis et al., 2011, dataset) from the EPTA telescopes, in the period 1995 to 2009, and the Parkes radio telescope, extending from 1995 to 1998, were added to the dataset. Wherever possible, these ToAs were replaced with new ToAs derived from data processed as described later in this section.

[1] These are the Effelsberg 100-m radio telescope, the Lovell radio telescope at Jodrell Bank, the Nançay radio telescope and the Westerbork Synthesis Radio Telescope (WSRT). A fifth telescope; the Sardinia Radio Telescope (SRT), has just entered its initial operational phase and therefore no data from the SRT are included here.

As a result of the extended temporal coverage, data-files (henceforth, archives) from a number of pulsar data recording instruments or 'backends' are included in the dataset. These include the Effelsberg-Berkeley Pulsar Processor (EBPP), the Berkeley-Orleans-Nançay (BON) instrument, the Digital Filter-bank system (DFB) and the Pulsar Machine I (PuMa-I) backend, all described in (Desvignes et al., 2016) as well as the Analogue Filter-bank system (AFB) (Shemar and Lyne, 1996) at Jodrell Bank and, the new generation of pulsar timing backends, namely, PuMa-II at the WSRT (Karuppusamy et al., 2008), PSRIX at Efflesberg (Lazarus et al., 2016), ROACH at Jodrell Bank (Bassa et al., 2016) and the Nancay Ultimate Pulsar Processing Instrument (NUPPI) at Nançay (Desvignes et al., 2011). The names of all the backends and their respective telescopes can be found in Table 4.1.

The archives from all the backends were first re-weighted by the sqaure root of the S/N and then grouped into 5-minute integrations using the psradd tool from the PSRCHIVE[2] suite (Hotan et al., 2004; van Straten et al., 2012).

[2] Commit hash - 87357c2; psrchive.sourceforge.net

ToAs were generated via cross-correlations of the time-integrated, frequency-scrunched, total intensity profiles with noise-free analytical templates, constructed by fitting high S/N pulse profiles with a set of von Mises functions using the paas tool of PSRCHIVE. These templates were manually aligned using pas. The pat tool from the same suite was used to generate ToAs, with the Fourier Domain with Markov-chain Monte-Carlo (FDM) algorithm (a re-implementation of Taylor, 1992) and goodness-of-fit (GOF) flags were enabled for the ToAs, as advised by Verbiest et al. (2016). A summary of the data from the different backends and telescopes is provided in Table 4.1. Figure 4.1 shows a plot of the timing residuals for the entire 21-year span, when the ToAs are fitted to the BTX model, as explained below.

Instrumental offsets between the various backends were corrected for by using 'JUMP' statements, which allow correct error-propogation.

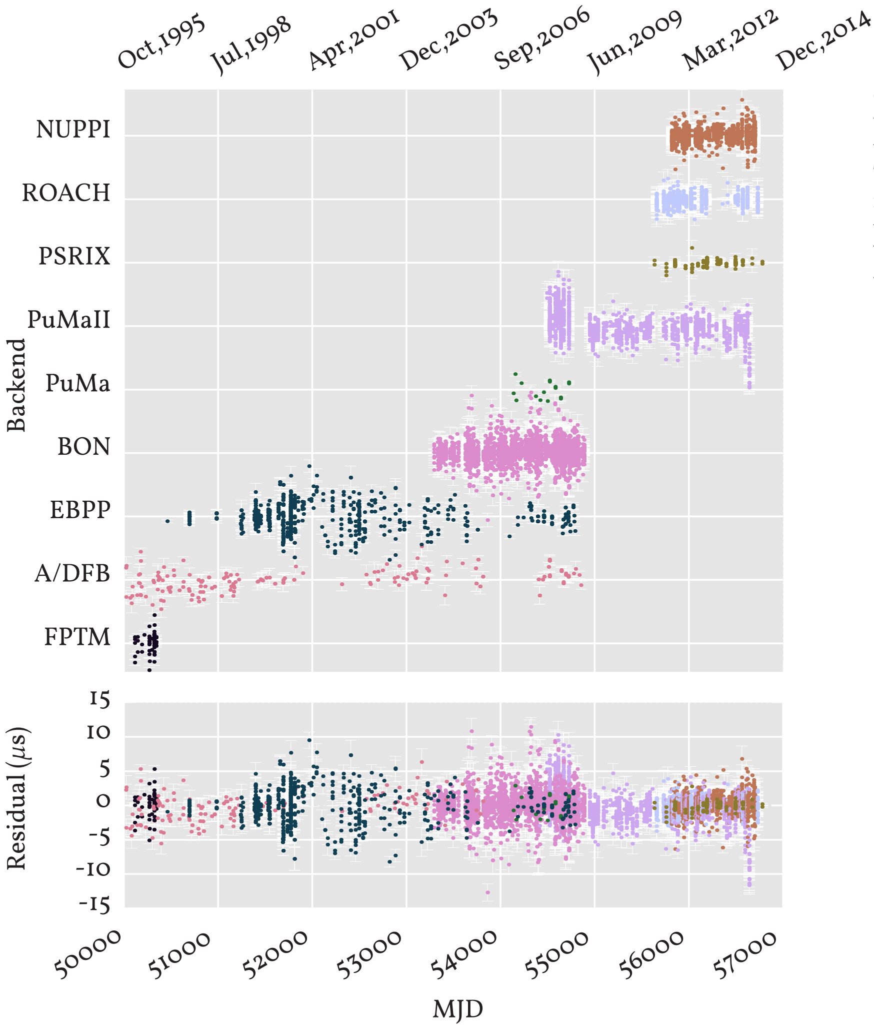

Figure 4.1: Plot of ToAs as a function of MJD. The bottom plot shows the timing residual from fitting the ToAs to the BTX model (see Section 4.2). The top plot shows the same but with manually introduced offsets to show the ToAs grouped by their respective backends. See Table 4.1 for the details of the backends.

For the PSRIX (PSRIX) backend (Lazarus et al., 2016) at the Effelsberg radio-telescope, which has a total bandwidth of 200 MHz at 21-cm wavelength and the archives with the highest S/N (up to ~4000, for a particular observation), archives were tested for frequency evolution of the pulse shapes. The data were split into 25 MHz channels and analytical templates were generated for each band, as explained above. These templates were manually compared using the paas tool. No significant differences were detected and the data were recombined into the full 200 MHz band. For the other backends such an exercise is not possible since either the S/N is typically worse or the bandwidth is too low to detect any obvious frequency evolution in the pulse profile. ToAs were also generated by using templates from different backends to test for pulse shape differences between backends. The timing analysis was then carried out using the TEMPO2[3] package (Hobbs et al., 2006). Observations which were linked to ToAs with unexplained residual offsets $\gtrsim 3\sigma$ were manually investigated. In some cases, manual

[3] version - 2013.9.1 with updated clock files; www.atnf.csiro.au/research/pulsar/tempo2

Telescope +Backend	BW (MHz)	f_c (MHz)	No. of ToAs	MJD range
Effelsberg	28	840	24	51 772–53 159
+EBPP	56	1410	690	50 460–54 791
	112	2639	35	51 247–54 690
+PSRIX	200	1360	120	55 759–56 494
	100	2640	116	55 632–56 779
Lovell	16	410	8	50 157–50 695
+A/DFB	16	610	42	50 021–51 452
(see note alongside)	16	1400	154	49 989–54 853
+ROACH	400	1532	844	55 656–56 729
Nançay+BON	128	1397	4502	53 293–54 888
+NUPPI	512	1484	2324	55 817–56 700
WSRT+PuMa	80	1380	20	54 135–54 372
+PuMaII	80	345	1173	54 490–56 640
	160	1380	536	54 520–56 640
Parkes	128	1400	23	50 116–50 343
+FPTM	128	1700	31	50 116–50 343

Table 4.1: Telescope and receiver-wise description of the dataset, showing the bandwidth (BW), the centre frequency of observations (f_c), the number of ToAs retained after the the selection process described in the text and the MJD ranges over which the ToAs exist. For the older backends (see text), only ToAs were available. For the new backends archives were processed as described in section 4.2. **Note :** The figures for bandwith (BW) and centre frequency (f_c) for the Jodrell Bank A/DFB and Parkes data are indicative only since the observations were made with various configurations. Details for these can be found in Stappers et al. (1998). Similar details for the other telescopes can be found in Desvignes et al. (2016), Bassa et al. (2016) or other specific references listed in section 4.2.

RFI excision was sufficient to remove the offset. A few ToAs were found to be linked to observations with previously determined time offsets, which were corrected for using the TEMPO2 TIME keyword in the relevant sections of the ToA files. In a few cases ToAs were found to have offsets which could not be corrected by either of the two methods. In most cases these ToAs were found to have poor GOF values (≥2) from the template matching and therefore, removed from the dataset. These ToAs are being investigated further to determine their possible association with micro-eclispes of the kind demonstrated by Archibald et al. (2009). However, their exclusion does not affect any of the conclusions in this analysis.

Similar to previous analyses, ToAs corresponding to orbital phases 0.2 to 0.35 (determined using the ephemeris presented in Lazaridis et al. (2011)) were removed as the eclipse region lies within that range. When carrying out a weighted fit, ToAs with large uncertainties contribute only weakly to the timing solutions and can often be discarded without greatly affecting the results. For MJD ranges with dense temporal sampling, a cut-off of 20 μs was applied. For the MJD range ~52000 - 53000, where the number of ToAs was very low even before a cut-off was applied, only ToAs with uncertainties greater than 60 μs were removed.

After the ToA selection procedure described above, the ToAs were split into ~1095 day (or 3 year) long 'aeons' with an overlap of 365 days between successive aeons. For each aeon the ToAs were fitted to the ELL1 (Lange et al., 2001) timing model, while keeping the DM fixed and setting the reference epoch to the centre of the aeon. The timing solutions were derived using the NASA-JPL DE421 planetary ephemeris (Folkner et al., 2009). The reference clock used was the Terrestrial Time standard derived from the 'Temps Atomique International' time standard, denoted by TT(TAI) and the final ToAs were corrected according to the BIPM standards (see e.g. Hobbs et al., 2006, and references therein). The default TEMPO2 assumptions for the Solar-wind model were retained for this analysis.

When using data from multiple instruments, it is necessary to correct the possible mis-estimation of the uncertainty of the ToAs in order to correct for the relative weighting of data from different backends. TEMPO2 error scaling factors (or T2EFACs) were calculated for each backend by applying the timing model derived in the previous step (without re-fitting) and then taking the square root of the reduced χ^2. The corresponding ToA uncertainties were then multiplied by these T2EFAC values.

MJD Range	Weighted RMS Timing Residual (μs)	Reduced χ^2 value	Number of ToAs
49 989–51 062	8.9	1.0	143
50 724–51 812	13.2	1.0	331
51 451–52 538	14.2	0.9	195
52 213–53 258	19.3	1.0	146
52 927–54 004	9.2	1.5	1037
53 643–54 733	9.5	1.0	2518
54 372–55 444	10.8	1.4	1959
55 121–56 189	5.0	1.1	1679
55 836–56 880	6.2	0.9	2110

Table 4.2: Properties of the TOA sets for each individual aeon (~1095 MJD period), determined using the respective ELL1 models. Note: The reduced χ^2 values shown below are derived after applying error scaling or EFACs as described in section 4.2

For the ELL1 model the $\sigma/\sqrt{N}$ statistic, where σ is the timing residual and N is the number of ToAs is used to select the aeon with the most information. From Table 4.2 this is identified as the epoch starting at MJD 55121. The timing parameters for this aeon are presented in Table 4.4 and a comparison with published literature is provided in Table 4.3.

As is obvious from the preceding discussions, the ELL1 model requires updating at regular intervals or aeons. This is a consequence of the orbital variability of this system, as discussed in Section 4.3.3. Therefore, the BTX model was used to construct a single timing model encompassing the entire 21 year period.

The BTX model is a re-implementation of the BT model (Blandford and Teukolsky, 1976) and incorporates higher order derivatives of the orbital-frequency. This model is completely phenomenological and thus, has no predictive power. The model also demands judicious usage since the highest order orbital-frequency derivatives can easily introduce correlations with proper motion components, DM variations and instrumental offsets, particularly in this highly heterogeneous dataset. Eccentricity measurements from the ELL1 models show large variability along with low measurement significance, indicating that these measurements are probably unreliable. Hence, the BTX model was created with eccentricity set to zero.

To limit the number of orbital-frequency derivatives (OFDs) employed

Parameter	Doroshenko et al. (2001)	Lazaridis et al. (2011)	ELL1 Model (Best fit)
MJD range of timing model fit	49 573–51 908	53 293–54 888	55 121.8–56 189.9
Proper motion in R.A., μ_α : (mas yr^{-1})	5.3(10)	6.6(2)	5.63(4)
Proper motion in Decl., μ_δ : (mas yr^{-1})	0.3(30)	3.2(7)	2.34(28)
Dispersion measure, DM (pc cm^{-3})	20.7449(4)	20.7458(2)	20.7299(17)
Epoch of DM measurement, (MJD)	51 000.0	49 530.0	56 387.8
Eccentricity, e	$< 9.6 \times 10^{-5}$	$6(1) \times 10^{-5}$	$5.1(8) \times 10^{-5}$
Reduced χ^2 (with scaled uncertainties)	–	1	1.1
Number of TOAs	584	3126	1679
Solar-system ephemeris model	DE200	DE405	DE421
Timing Software Package	TIMAPR/TEMPO	TEMPO2	TEMPO2

Table 4.3: Comparision of selected parameters of the black-widow pulsar system J2051–0827 with published values. μ_α and μ_δ values for the ELL1 model are obtained from a weighted fit to the position measured at succesive aeons. The epoch of DM determination need not correspond with the epoch of the timing model since the DM value for the ELL1 models are fixed from Kondratiev et al. (2016), as explained in subsection 4.3.2. Similarly, the DM value used in the Lazaridis et al. (2011) analysis is taken from Stappers et al. (1998).

in the BTX model, the reduced χ^2 was used as the primary selection criterion. The reduced χ^2 remains well above ten until the 17$^{\text{th}}$ OFD is introduced. Subsequent OFDs do not affect the reduced χ^2 and are not determined with any significance by TEMPO2. Amongst the timing models with 13 or more OFDs. the Akaike Information Criterion (Akaike, 1974) also favours the model with 17 OFDs. The BTX timing parameters with 17 OFDs for PSR J2051–0827 are presented in Table 4.4.

The timing models and ToAs are available under 'additional online material'[4] at the EPTA web page.

[4] The timing models, ToAs and the standard templates used for timing are accessible via http://www.epta.eu.org/aom.html.

4.3 *Timing Results*

4.3.1 *Proper motion*

PSR J2051–0827 has a low ecliptic latitude of ~8°51′. Typically for such low latitudes, the determination of position is relatively poor (Lorimer and Kramer, 2005). Therefore, the resulting measurement of proper motion in declination or ecliptic latitude (depending on the coordinate system used) is imprecise. This is evident in the published values of proper motion in declination, μ_δ, presented in Table 4.3.

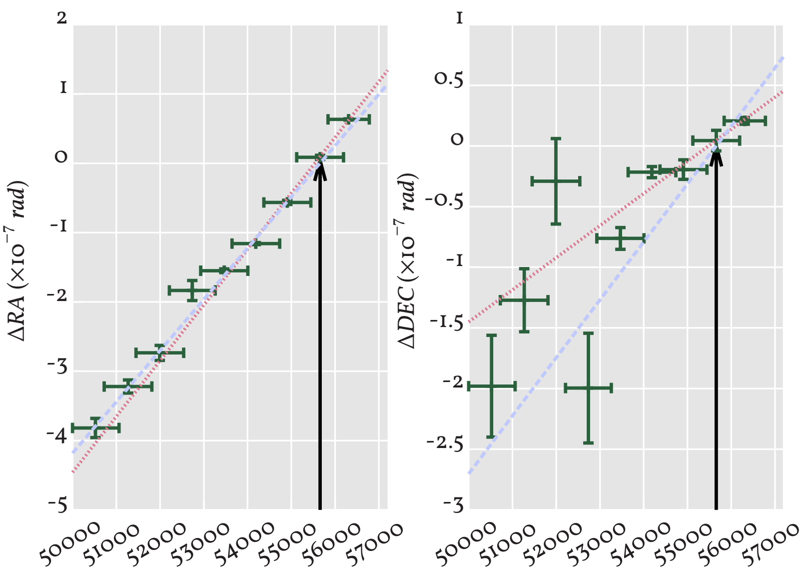

Figure 4.2: Measured values of R.A. (left) and DEC (right) of PSR J2051–0827 for each aeon (green +) and linear fits to those. Black arrows indicate the values at the reference epoch at which the two timing models of Table 4.4 are defined, MJD 55655. The fit to the position at the median MJD of each aeon (the finely-dotted pink line) returns μ_α =5.63(10) mas yr^{-1} and μ_δ =2.34(28) mas yr^{-1} while the dashed, lilac line represents the values obtained from the BTX model, shown in Table 4.4.

Pulsar name	PSR J2051−0827	
Binary model	T2	BTX
MJD range	55 121.8 to 56 189.9	49 989.9 to 56 779.3
NToA	1679	11391
RMS $t_{\rm resid}$ (μs)	5.0	5.2
Red. χ^2	1.1	4.2
R. A., α	$20^{\rm h}51^{\rm m}7.519\,768(18)^{\rm s}$	$20^{\rm h}51^{\rm m}7.519\,763(8)^{\rm s}$
Dec., δ	$-8°27'37.7497(8)''$	$-8°27'37.7505(4)''$
Ec. Long., λ	312.835 726 88(8)°	312.835 727 10(2)°
Ec. Lat., β	8.846 341 8(5)°	8.846 342 30(9)°
μ_α (mas yr^{-1})	5.63(4)	5.57(4)
μ_δ (mas yr^{-1})	2.34(28)	3.60(10)
μ_λ (mas yr^{-1})	7.2(3)	6.34(1)
μ_β (mas yr^{-1})	4.6(23)	1.9(1)
ν (s^{-1})	221.796 283 653 017(5)	221.796 283 653 049 2(10)
$\dot{\nu}$ (s^{-2})	$-6.264(3) \times 10^{-16}$	$-6.265\,32(6) \times 10^{-16}$
P (ms)	4.508 641 820 006 43(8)	4.508 641 820 006 1(5)
$\dot{P}$	$1.2732(4) \times 10^{-20}$	$1.273\,74(3) \times 10^{-20}$
DM (cm^{-3} pc)	20.7299(17)	20.7299(17)
x (lt − s)	0.045 072 0(3)	0.045 070 74(20)
$\dot{x}$	$1.3(148) \times 10^{-16}$	$9.6(12) \times 10^{-15}$
P_b (d)	0.099 110 254 90(4)	N/A
$\dot{P}_b$	$-5.9(3) \times 10^{-12}$	N/A
FB0(Hz)	N/A	$1.167\,797\,940\,6(7) \times 10^{-4}$
FB1(/s^2)	N/A	$8.2(4) \times 10^{-20}$
FB2(/s^3)	N/A	$-7.4(3) \times 10^{-27}$
FB3(/s^4)	N/A	$-6.3(16) \times 10^{-35}$
FB4(/s^5)	N/A	$3.9(8) \times 10^{-42}$
FB5(/s^6)	N/A	$1.8(7) \times 10^{-49}$
FB6(/s^7)	N/A	$6.5(24) \times 10^{-57}$
FB7(/s^8)	N/A	$-5.8(23) \times 10^{-64}$
FB8(/s^9)	N/A	$-4.0(8) \times 10^{-71}$
FB9(/s^{10})	N/A	$1.6(7) \times 10^{-78}$
FB10(/s^{11})	N/A	$1.4(3) \times 10^{-85}$
FB11(/s^{12})	N/A	$-3.2(18) \times 10^{-93}$
FB12(/s^{13})	N/A	$-3.7(8) \times 10^{-100}$
FB13(/s^{14})	N/A	$3.0(30) \times 10^{-108}$
FB14(/s^{15})	N/A	$7.3(19) \times 10^{-115}$
FB15(/s^{16})	N/A	$5.2(20) \times 10^{-123}$
FB16(/s^{17})	N/A	$-7.9(25) \times 10^{-130}$
FB17(/s^{18})	N/A	$-1.8(5) \times 10^{-137}$
Ref. epoch (MJD)	55 655	55 655
ω (deg)	36(10)	0
EPS1	$3.0(10) \times 10^{-5}$	N/A
EPS2	$4.1(9) \times 10^{-5}$	N/A
e	$5.1(8) \times 10^{-5}$	0
TASC (MJD)	54 091.034 307 9(8)	54 091.034 349 36(14)
T0 (MJD)	54 091.044(2)	54 091.034 349 36(14)
$\log_{10} \tau_{\rm char}$ (yr)	*9.75*	*9.75*
$\log_{10} B_{\rm surf}$ (G)	*8.38*	*8.38*

Table 4.4: Timing parameters for PSR J2051−0827 for the ELL1 (implemented via the TEMPO2 hybrid model T2) and the BTX models. The values of derived parameters are italicised while parameters that should be neccesarily excluded from the respective timing models are marked as N/A. Note that the DM values presented here are obtained from Kondratiev et al. (2016). For brevity, the table below uses the following abbreviations: FB0 indicates orbital frequency and higher numbers the resp. derivative, NToA denotes the number of TOAs, RMS $t_{\rm resid}$ denotes the RMS timing residual and Red. χ^2 is the reduced χ^2 value for the weighted TEMPO2 fit. $\tau_{\rm char}$ is the characteristic age associated with the pulsar and $B_{\rm surf}$ is the estimated surface magnetic field strength. The TT(TAI) clock correction procedure and the DE421 Solar Sytem Ephemerides were used for both the models. The units are in TCB (See Hobbs et al., 2006, for details). The figures in parentheses are the nominal 1σ TEMPO2 uncertainties in the least-significant digits quoted. The coordinates refer to J2000.

To improve the measurement and utilise the entire 21-year span of the dataset, the measured value of R.A. and DEC for each aeon were fitted with a linear function to obtain a mean proper motion. This results in a significant measurement of μ_α and μ_δ, as shown in Figure 4.2 and Table 4.3. The fitted values of μ_α and μ_δ are inserted into the ELL1 models for each aeon and those models are refitted for the other parameters.

Using an estimated distance of ≃1040 pc (from the NE2001 model of free-electron distribution in the Galaxy, Cordes and Lazio, 2003) and a total proper motion, $\mu_t = \sqrt{\mu_\alpha^2 + \mu_\delta^2} = 6.1(1)\,\mathrm{mas\,yr^{-1}}$; a 2-D transverse velocity of $v_t = 30(9)\,\mathrm{km\,s^{-1}}$ was calculated. This assumes an uncertainty of 30%[5] in the DM derived distance mentioned above. The measurement is in agreement with the value of 30(20) km s^{-1} measured by Stappers et al. (1998) and represents a two-fold increase in precision, even though the uncertainty of the DM derived distance is assumed to be much greater. It should be noted that this is significantly lower than the average value of 93(13) km s^{-1} reported in Desvignes et al. (2016) for the transverse velocities of binary MSPs. However, it agrees well with the value of 56(3) km s^{-1} reported for the binary MSPs with distance measurements from parallaxes.

[5] See Desvignes et al. (2016) for a discussion on the possible underestimation of uncertainties of the DM derived distances.

The proper motion values obtained from the BTX model appear to be inconsistent with those obtained from fitting to position measurements for every aeon using the respective ELL1 models. This is because the proper motion terms and the orbital-frequency derivatives are strongly covariant and therefore the uncertainties in the values obtained from the BTX model are heavily underestimated, reinforcing the need for cautious usage of this model.

4.3.2 *DM variations*

Since the DM is a measure of the density of the IISM along the line of sight to the pulsar, both the motion of the pulsar and the dynamical evolution of the IISM affect this value. While it is possible to obtain the DM from timing, 'JUMPS' or instrumental offsets introduced to align the ToAs from the different backends are fully covariant with the DM and prevent an accurate measurement directly from the dataset presented above. Therefore, a DM value of 20.7299(17) cm^{-3} pc is adopted from the LOFAR measurements of Kondratiev et al. (2016).

When simultaneous dual (or multi) frequency observations are available, it is, however, possible to accurately estimate the *variation* in the DM. The WSRT Puma-II backend provides observations centred at 345 and 1380 MHz, with a cadence of roughly three weeks. Observations between the two frequencies are sequential, which are separated by, at most, a few days and available for the MJD range ~54 600 to 56 800. Since low-frequency observations are more sensitive to the DM variations, these are utilised to measure them instead of the two frequency-band observations of the PSRIX backend.

To measure DM variations, the PuMa-II ToAs were fitted for DM using the ELL1 model presented in Table 4.4. The ToAs were then split into 100-day long intervals, to ensure enough data were available for a re-

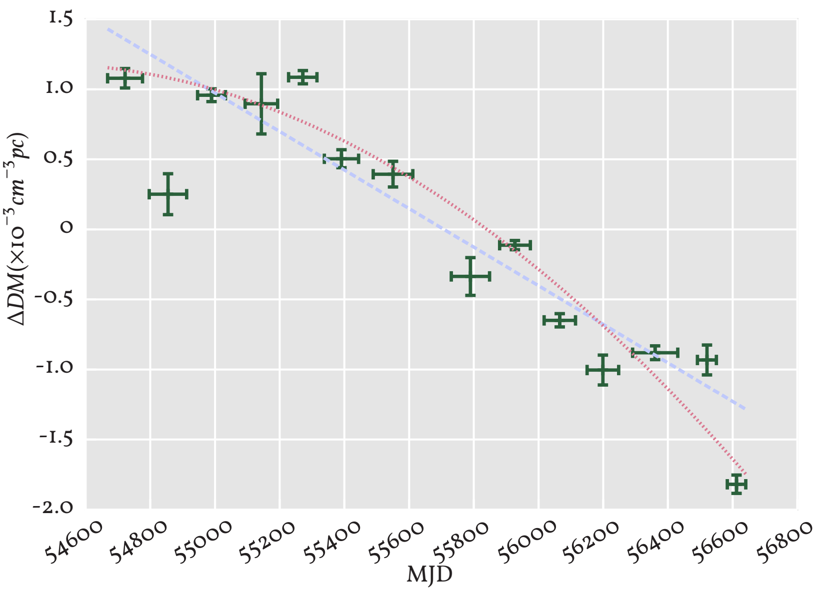

Figure 4.3: DM variation from consecutive 327 MHz and 1380 MHz observations with the WSRT which extend over the period 54 600 to 56 800. A linear fit (lilac, dashed) and a quadratic fit (pink, finely-dotted) are also shown.

liable estimate. Each 100-day interval was then refitted for the DM. P_b and T_0. The fit for P_b is necessary to ensure that orbital-phase dependent effects do not contaminate the DM measurement, since the observations at 345 MHz and 1380 MHz do not necessarily coincide in orbital phase.

This leads to a significant detection of a DM trend after MJD 54600, which is plotted in Figure 4.3. A quadratic fit returns a reduced-χ^2 of 3.5 while a linear fit performs not much worse, with a reduced-χ^2 of 6. The linear trend appears to show a weakly sinusoidal residual, with a 'best-fit' period of ~940 days but this residual becomes insignificant with the quadratic model and therefore, no higher-order model was considered.

While it is quite possible that such variations may be present before MJD 54600, the lack of sensitivity due to sparse and inhomogeneous multi-frequency observations lead to typical DM measurement uncertainties of $\sim 1 \times 10^{-3}$ to about $\sim 3 \times 10^{-4}\,\mathrm{cm}^{-3}\,\mathrm{pc}$. These uncertainties, which may well be severely underestimated, prevent any firm conclusion on the DM evolution. Furthermore, because no combination of two observing systems at different frequencies is continuously present before MJD 54600, any effort to measure DM variations in that MJD range is necessarily corrupted by the arbitrary phase offsets used to align the data from different instruments. The WSRT data which provide continuous data at two frequencies after MJD 54600 provide a DM precision of $\lesssim 3 \times 10^{-4}\,\mathrm{cm}^{-3}\,\mathrm{pc}$ over 100-day intervals, allowing accurate DM modeling over that period.

Traditionally, wherever a DM trend is observed, it is corrected for by introducing DM derivatives.[6] Given the large uncertainties in the earliest eras and to prevent over-fitting or accidentally introducing excess white noise in the timing, only those ToAs belonging to the period over which a clear DM trend is measured are corrected for the DM trend modelled by the quadratic fit shown in Figure 4.3. This is implemented by introducing a TEMPO2 DM offset flag (-dmo) for the ToAs lying in the MJD range 54 600 to 56 800.

[6] For detailed reviews on modern DM correction methods see Verbiest et al. (2016), Demorest et al. (2013) or Lentati et al. (2015).

4.3.3 Secular variations

Following Lazaridis et al. (2011), variations in the binary period (P_b) and the projected semi-major axis (x) were measured by splitting the ToAs into 'eras' of approximately 365 days. The results of reproducing and extending the Lazaridis analysis[7] are presented in Figure 4.4.

[7] In the Lazaridis et al. (2011) analysis, timing models are first derived for the largest MJD range over which a TEMPO2 fit converges, which are analogous to 'aeons' in the present work. Then, the variations in P_b and x are measured by fitting for P_b, x and T_{asc} simultaneously for 300-day periods with an overlap of 30 days.

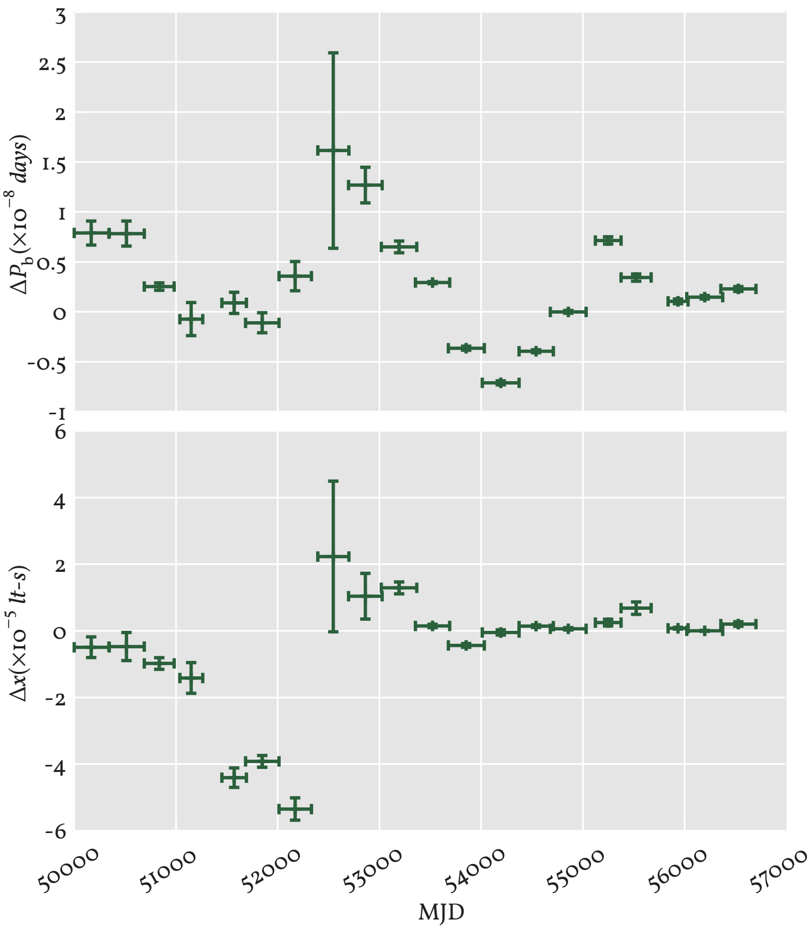

Figure 4.4: Change in P_b and x measured by fitting for P_b, x and T_{asc} for eras of length 365 days with an overlap of 30 days, where possible.

The simultaneous fitting of P_b, x and T_{asc}, as in Lazaridis et al. (2011), is undesirable since P_b and T_{asc} are fully covariant parameters. In practice, wherever good orbital phase coverage ($\geq 60\%$) is available, the measurement of T_{asc} is far more accurate and reliable since it measures the orbital phase and requires less information for its calculation than P_b. Due to the high cadence and long durations of the Nançay, Jodrell Bank and WSRT observations and full orbital observations at Effelsberg, especially in the latest years, it is possible to carry out such a measurement with much greater precision than was previously attempted.

By keeping P_b constant for all eras, and fitting for T_{asc}, x and the Laplace-Lagrange parameters, $\eta = e \cdot sin\omega$ and $\kappa = e \cdot cos\,\omega$ simultaneously, the change in T_{asc} is measured. The change in P_b measured at time t_1, $\Delta P_{b,t_1}$, is then calculated using the equation

$$\Delta P_{b,t_1} = \frac{T_{asc,t_1} - T_{asc,t_0}}{t_1 - t_0} \times P_{b,ref} \tag{4.1}$$

where T_{asc,t_0} and T_{asc,t_1} are the values of T_{asc} at two neighbouring eras t_0

and t_1. $P_{b,ref}$ is a constant P_b value chosen from the P_b values for each epoch, such that the measured ΔT_{asc} values do not show any obvious slope. The resulting ΔP_b variations and the ΔT_{asc} from which they are derived are plotted in Figure 4.5, along with the simultaneous Δx measurements. The measured values for all three parameters are over-plotted with the interpolation of the change in ΔT_{asc} as obtained from the BTX model (see, e.g., Ng et al., 2014). The excellent agreement serves to further confirm the applicability of the BTX model.

Comparing the P_b variations derived from the T_{asc} variations in Figure 4.5 and those in Figure 4.4, derived from the Lazaridis et al. (2011) method, it is apparent that fitting for all three parameters introduces a 'smoothing' effect. This is likely due to the covariance of P_b and T_{asc} and thus demonstrates the importance of estimating ΔP_b from fitting for T_{asc} and x simultaneously. It should be noted that for all the eras that were analysed, P_b and T_{asc} were found to be either strongly correlated or anti-correlated ($|\text{corr.}| \geq 0.9$), with a somewhat alternating behaviour, while the P_b and x are always weakly correlated ($|\text{corr.}| \leq 0.3$). Finally, T_{asc} and x are always very weakly correlated ($|\text{corr.}| \ll 0.3$).

As can be seen from Figure 4.5 the new analysis is in qualitative agreement with the measurements presented in Lazaridis et al. (2011) and the system appears to have entered a 'quieter' phase. For brevity, only a summary of the maximum possible contribution to the secular variations from the various possible sources is presented in Table 4.5 [8] [9]. For a full discussion of these, see Lazaridis et al. (2011).

Variations in the orbital period can be attributed to contributions due to gravitational-wave emission ($\dot{P}_b^{GW}$), changing Doppler shift ($\dot{P}_b^{\dot{D}}$), mass loss from the companion ($\dot{P}_b^{\dot{m}}$), tidal interactions between the companion and the pulsar ($\dot{P}_b^{T}$) and variations of the gravitational quadrupole moment of the companion star ($\dot{P}_b^{Q}$) (see, for instance Lorimer and Kramer, 2005)[10]:

$$\dot{P}_b^{obs} = \dot{P}_b^{GW} + \dot{P}_b^{\dot{D}} + \dot{P}_b^{\dot{m}} + \dot{P}_b^{T} + \dot{P}_b^{Q} \tag{4.2}$$

Similarly, the secular variations in the projected semi-major axis can be split into contributions due to radiation of gravitational waves ($\dot{x}^{GW}$), the proper motion of the pulsar ($\dot{x}^{PM}$), varying aberrations ($\frac{d\epsilon_A}{dt}$), changing Doppler shift ($\dot{x}^{\dot{D}}$), mass loss in the binary system ($\dot{x}^{\dot{m}}$), variations of the gravitational quadrupole moment of the companion star ($\dot{x}^{Q}$), spin-orbit coupling of the companion ($\dot{x}^{SO}$) and a second, or planetary, companion ($\dot{x}^{P}$).

$$\dot{x}^{obs} = \dot{x}^{GW} + \dot{x}^{PM} + \frac{d\epsilon_A}{dt} + \dot{x}^{\dot{D}} + \dot{x}^{\dot{m}} + \dot{x}^{Q} + \dot{x}^{SO} + \dot{x}^{P} \tag{4.3}$$

For the observed 21-year baseline, the maximum $\dot{P}_b$ is $\sim 1.41 \times 10^{-11}$ and the minimum is $\sim -2.03 \times 10^{-11}$. From Table 4.5 it is evident that the first four terms of Eqn. (4.2) cannot drive the observed ΔP_b variations independently. Therefore, the hypothesis of Lazaridis et al. (2011) that the mass quadrupole variations in the companion are the most likely drivers of the observed ΔP_b variations is recovered.

Similarly, from Figure 4.5, the variation of the projected semi-major

[8] $\dot{m}_c$ refers to the rate at which mass is lost by the companion.

[9] The contributions from the gravitational quadrupole (GQ) and the classical spin-orbit coupling (SOC) variations require assumptions based on Lazaridis et al. (2011). Since the derived values are then identical to those presented there, readers are referred to the original source instead.

[10] The sign on the $\dot{P}_b^{\dot{D}}$ and $\dot{x}^{\dot{D}}$ terms are made positive for the sake of uniformity here.

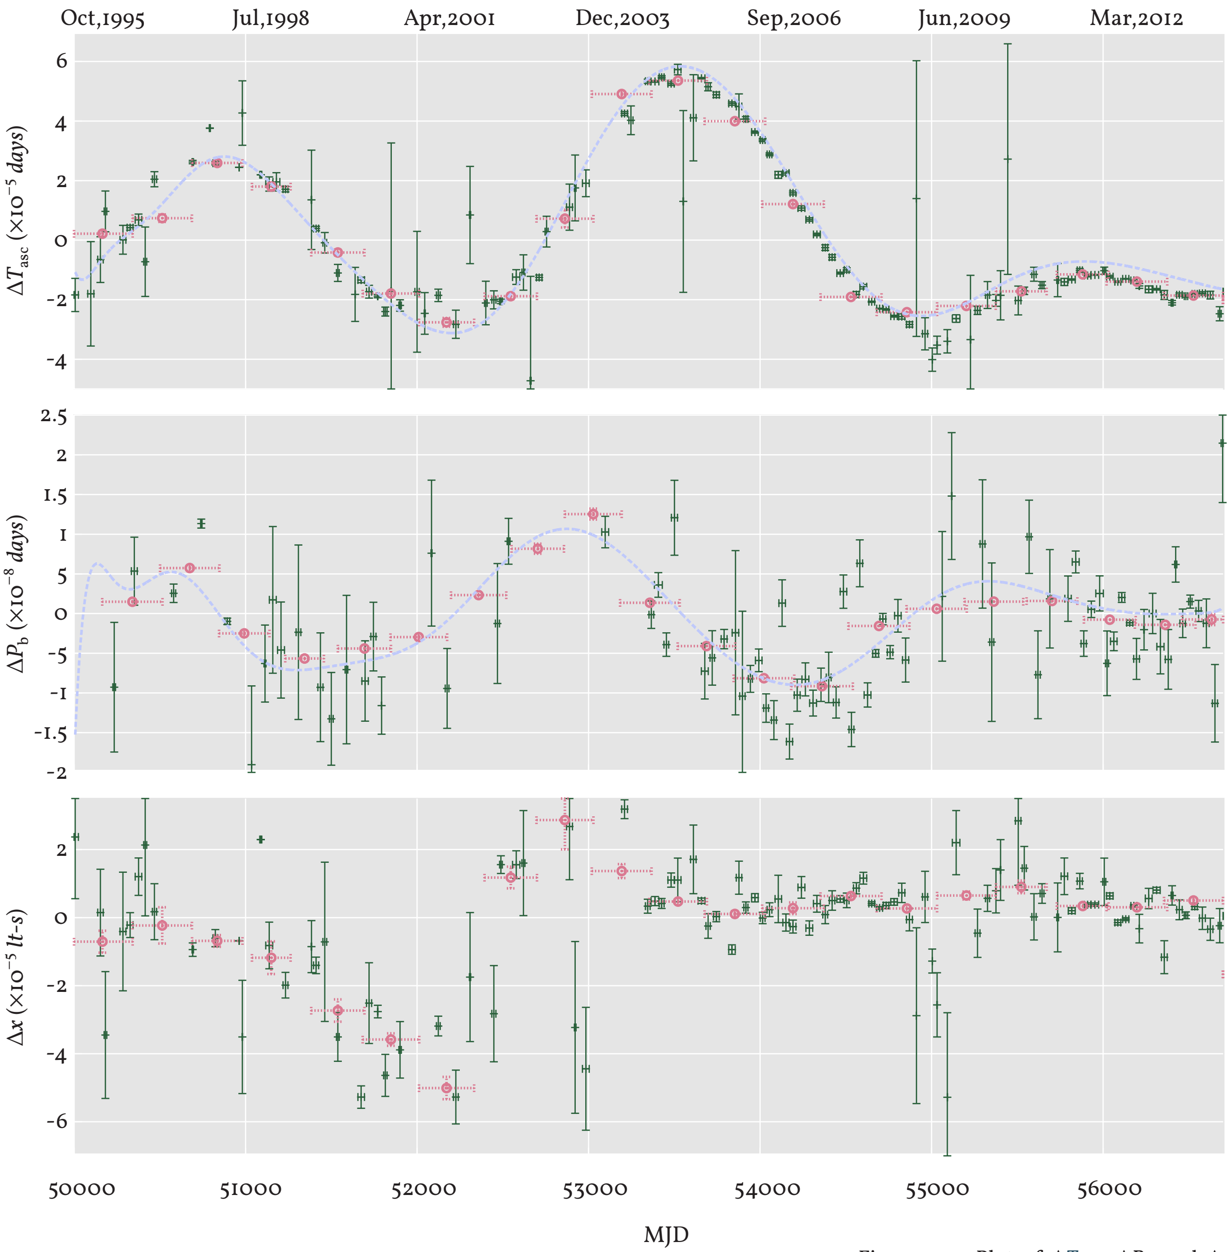

Figure 4.5: Plot of $\Delta T_{\rm asc}$, $\Delta P_{\rm b}$ and Δx measured from fitting for x and $T_{\rm asc}$ only for epochs with a length of 45 (green +) and 365 (pink ⊙) days, along with the variations described by the BTX model (lilac, dashed). To improve the readability of the graphs for $\Delta P_{\rm b}$ and Δx, points with uncertainties comparable to the y-range of the graph (typically in the earliest epochs) are removed. The prominent fluctuations for the BTX prediction of $\Delta P_{\rm b}$ at ~ MJD 50 100 to 50 600 agree with the measured (but unplotted) values, as can be discerned from the $\Delta T_{\rm asc}$ plot.

Source	$\dot{P}_b$	$\dot{x}$
GW emission	-7.61×10^{-14}	-2.67×10^{-19}
Doppler Correction	-4.06×10^{-21}	-4.06×10^{-21}
Proper Motion Corr.	N/A	4.99×10^{-17}
Varying Aberration	N/A	-4.41×10^{-17}
Mass loss	Requires unphysical $\dot{m}_c \sim 1 \times 10^{-7}\,M_\odot\,{\rm yr}^{-1}$	
Mass/GQ Variations	See Lazaridis et al. (2011)	
Spin-Orbit Coupling	See Lazaridis et al. (2011)	
Max. Measured	1.41×10^{-11}	2.29×10^{-13}
Min. Measured	-2.03×10^{-11}	-5.08×10^{-13}

Table 4.5: Maximum contributions from the various sources of secular variations in $P_{\rm b}$ and x as presented in Equation 4.2 and 4.3.

axis shows a strong 'feature' in the MJD-range ~51000 to 53000, which is not present in the remaining data. Since the correlation between x and T_{asc} or x and P_{b} is very weak, the differences between the bottom panels of Figure 4.4 and Figure 4.5 are marginal, although the uncertainties in the second case are typically smaller for the 365-day epochs.

As in the case of the ΔP_{b} variations, the terms of Eqn. (4.3) for which values are presented in Table 4.5 are not likely to be independent drivers of the variations in Δx. This implies that the Lazaridis et al. (2011) conjecture that the classical spin-orbit coupling term combined with the GQ term is the most likely driver for the Δx variations is also recovered.

In addition to recovering the long-term fluctuations, the derivation of ΔP_{b} from ΔT_{asc} reveals small-scale variations, as indicated with black arrows in Plate 4. These points lie $\gtrsim 4\sigma$ away from their local means and do have corresponding values with negative offsets. Given the results from Wu et al. (2012) presented in Section 4.1, it remains unclear what processes could lead to such deviations.

It is evident that continued multi-band monitoring of PSR J2051−0827 is necessary to reveal the origin of these sudden, sharp increases in the orbital period. If these changes are a result of activity of the companion, a greater understanding of the origin of these changes might help to understand the processes which drive state changes in the 'transitioning' MSP systems, i.e., binaries where the MSP alternates between accreting and radio-pulsar states (see, e.g., Stappers et al., 2014).

Given the high cadence and regular sampling in the later aeons, a test for the presence of a second companion, possibly of planetary dimensions, is carried out as well. This involves testing for the presence of higher-order derivatives of pulse frequency in the timing solution (Joshi and Rasio, 1997). The extrema of the second and third order frequency derivatives from TEMPO2 fits to the aeons are

- $-4.1(8) \times 10^{-24}\,\text{s}^{-3} \leqslant f^{(2)}_{\text{max}} \leqslant 3.0(19) \times 10^{-24}\,\text{s}^{-3}$ and
- $1.1(6) \times 10^{-30}\,\text{s}^{-4} \leqslant f^{(3)}_{\text{max}} \leqslant 2.1(9) \times 10^{-30}\,\text{s}^{-4}$.

Since these values are at best marginally significant and in the absence of any supporting evidence from optical observations, the hypothesis of a second companion to PSR J2051−0827 remains unjustified.

4.4 High-precision timing prospects

Due to the complicated and somewhat arbitrary orbital variability that some pulsars in BWP systems have been shown to exhibit (e.g., Nice et al., 2000; Freire et al., 2003; Lynch et al., 2012; Ng et al., 2014, etc.), these sources have been traditionally left out of high-precision pulsar-timing campaigns. With the recent increase in the number of BWP systems discovered among the *Fermi* Large Area Telescope (*Fermi*-LAT) sources (Abdo et al., 2013), it will soon be possible to quantify these instabilities for a larger sample. As a counter-example to the current practice, the pulsar of the BWP system J0610−2100 has recently been added to the list of sources for the EPTA (Desvignes et al., 2016) and has, so far, provided stable timing.

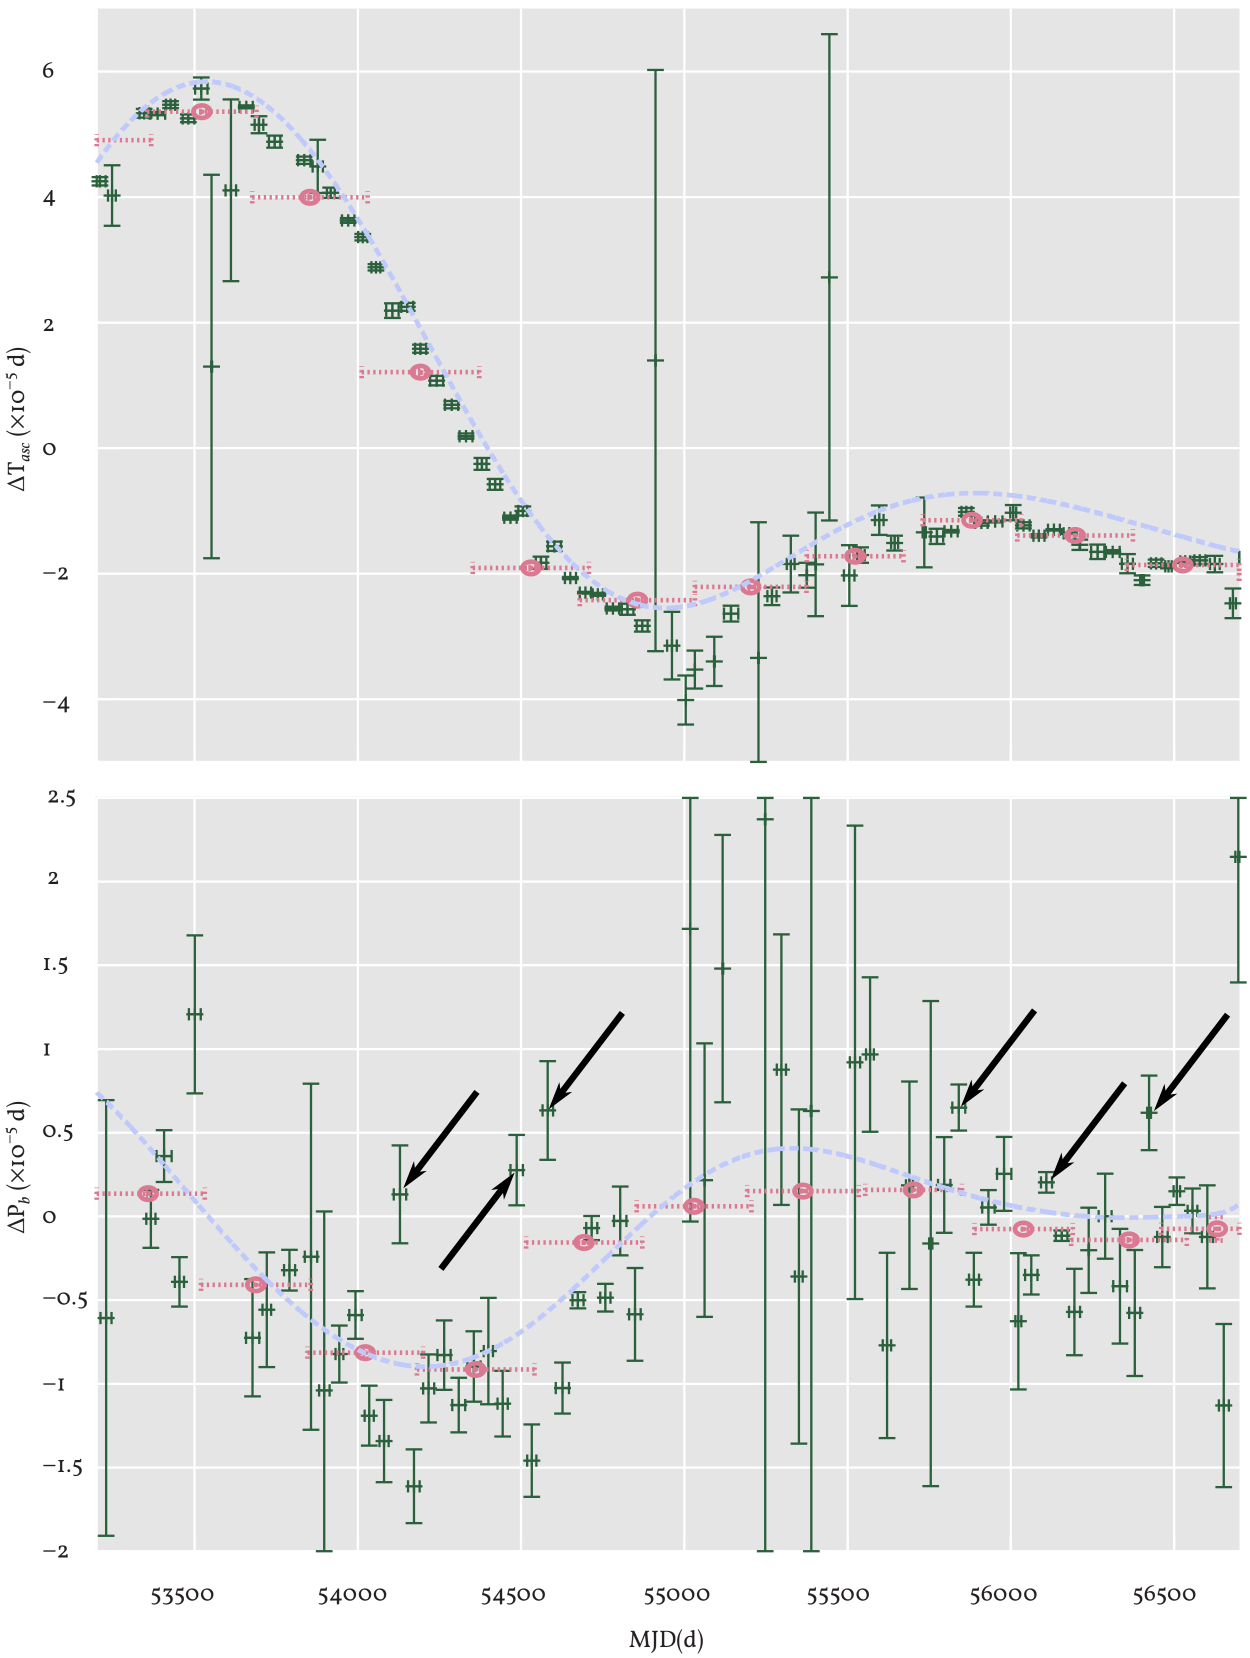

Plate 4: Zoomed in plot of ΔT_{asc}, ΔP_{b} and Δx measured from fitting for x and T_{asc} only for 45 (green +) and 365 (pink ⊙) day long epochs, along with the predicted variations from the BTX model (lilac, dashed). Solid black arrows indicate ΔP_{b} measurements corresponding to epochs where the derivative of T_{asc} abruptly changes sign.

Simulations for pulsars timed using the BTX model by Bochenek et al. (2015) show that only a small percentage of the power from gravitational waves is likely to be absorbed into the higher-order orbital-frequency derivatives and again, appear to favour the inclusion of such pulsars in PTAs. However, Bochenek et al. (2015) do not take into consideration variations of x, as identified for the BWP system J2051−0827.

The timing analysis presented here demonstrates the practical usability of the BTX model for such systems. However, it should also be noted that the GOF for the BTX model is still rather low as some variations remain unaccounted for.

It is probably an opportune coincidence that the BWP system J2051−0827 has entered a relatively stable phase, suggesting greater usability for a PTA. Even without addressing some of the ambiguities in the fundamental properties of this system, for both the ELL1 and BTX models, the present analysis shows it is possible to obtain timing residuals of the order of $\sim$5.0 μs, quite comparable to the timing precision of several sources already in the PTAs (Verbiest et al., 2016). In the intermediate-to-high S/N regime of gravitational wave background observations, where the number of pulsars becomes more important than very high timing precision (Siemens et al., 2013) timing residuals of the order of 1 μs could be sufficient. With the advent of the new 'ultra-broadband' backends (Karuppusamy, private communication) and rapid increases in sensitivity, this does not appear to be an unrealistic goal.

4.5 *Summary*

A timing update on PSR J2051−0827 is presented, along with timing models for the BTX and ELL1 models of TEMPO2. An improved estimate of the mean proper motion is also made, giving a value of 30(9) km s^{-1}. A significant decrease in the DM of $\sim 2.5 \times 10^{-3}$ cm^{-3} pc is detected for the MJD range 54 600 to 56 800 and corrections are incorporated in the ToA file.

A more robust analysis is performed by reducing covariant terms and it is shown that the resulting measurements are more precise and consistent with earlier analyses. The variations of the orbital period are detected over more than a full 'period', supporting earlier analyses that suggested that these variations arise from cyclic variations in the companion, instead of a tertiary star or planet. In addition, small-scale fluctuations in the P_b variations are detected.

The continued timing of PSR J2051−0827 shows that the variation of the projected semi-major axis appears to have decreased and does not show the extreme behaviour observed at an earlier epoch, lending hope that the black widow system containing PSR J2051−0827 may be included in PTAs in the near future.

5

Spectral Indices of Millisecond Pulsars

風をいたみ
岩うつ波の
おのれのみ
くだけて物を
おもふ頃かな

> Non ha l'ottimo artista alcun concetto
> c'un marmo solo in sé non circonscriva
> col suo superchio, e solo a quello arriva
> la man che ubbidisce all'intelletto
>
> — **Michelangelo;** *Sonnet, circa. 1538*

We present the spectral indices of 12 millisecond pulsars, measured via dense monitoring of their flux densities at three frequency bands with the 500-m radio telescope at Arecibo Observatory. We compare these spectral indices against literature values and find that our estimates are able to predict flux density values from literature at other frequencies quite well. We have also have rederived the spectral indices of an additional 62 using flux density values from literature alone, increasing the total number of millisecond pulsar spectral indices to 74. We find the median spectral index for the combined population to be −1.74(4). A population analysis shows that the distribution of measured spectral indices of millisecond pulsars and classical pulsars is largely identical, except for a few steep spectrum sources in the former class. We find a similar agreement between the populations of isolated millisecond pulsars and those in binaries. Our results also suggest that *Fermi* sources are typically steep spectrum sources, explaining why ∼1400 MHz surveys were unable to detect a larger number of those sources. Finally, we find that the spectral indices of millisecond pulsars are weakly correlated with their spin-periods and weakly anti-correlated with the associated spin-down energy.

5.1 Introduction

"...caution must be taken so as to avoid being deceived, and also to refer the phenomena to the simple laws." –

von Fraunhofer, J. Neue modifikation des lichtes. *Denkschriften der Königlichen Akademie der Wissenschaften zu München für das Jahre 1821 und 1822*, 8, 1 (1822)

Spectra are fundamental observables of astronomical objects. They offer the most direct tools for analysing the physical processes driving the various emission and absorption processes taking place on or near the source and along the line of sight to it.

In the simplest form, the flux density of pulsar emission is modelled as an exponentially decreasing function of the observing frequency, i.e.:

$$S_\nu \propto \nu^\alpha \tag{5.1}$$

where α is called the 'spectral index' of the source. The shape of the pulsar's spectrum depends directly on the pulsar emission process and the properties of the plasma surrounding the pulsar (Malofeev and Malov, 1980). Thus, a precise measurement of its spectrum can provide observational constraints on pulsar magnetosphere models[1]. The spectral index is a necessary input parameter for pulsar population synthesis and survey yield projections.

[1] which in itself does not amount to *solving* the pulsar emission 'problem'.

By the time of the discovery of pulsars, radio spectra had already been established to be invaluable tools for distinguishing the physical drivers of emission (See e.g., Conway et al., 1963). Initial measurements, some made even before the official announcement of the discovery of pulsars (Hewish et al., 1968; Ryle and Bailey, 1968) and subsequent efforts in the months following it (Eg., Robinson et al., 1968; Lovelace and Craft, 1968) also made it clear that pulsars show an extreme degree of temporal variability in their relative fluxes (See E.g., Scheuer, 1968). A brief search through historical literature on pulsar fluxes shows the continued disagreement between published flux density values from different groups, except for some of the strongest sources. This situation is far from resolved even today. For example, Levin et al. (2013) demonstrate that an analysis of archival data from the Parkes radio telescope provides flux density measurements that are more often in disagreement with the values from literature than not.

That these flux density values rarely coincide is well known and believed to be the result of a combination of factors. These include the dominant effects of propagation through a turbulent ISM and the changes or instabilities in the intrinsic properties of the pulsars themselves which are observed less often. The characteristic dispersive sweep of pulsar signals resulting from propagation through the IISM that led Pilkington to identify these sources to be located outside the Solar System but inside the Galaxy (Hewish et al., 1968) also leads naturally to the understanding that the turbulence of the IISM affects the pencil-beams of pulsar radiation quite strongly as well.

While this fluctuation of the observed flux density is a consequence of the nature of the pulsars' beamed radiation and its propagation, Rickett et al. (1984) were the first to demonstrate that scintillation due to the IISM affected flux measurements over long enough timescales to account for much of the disagreements at different epochs. A comprehensive study of the spectra of classical pulsars[2] accounting for these effects was carried out by Lorimer et al. (1995), measuring the spectra of 280 classical pulsars over a period of ~4 years. This survey obtained a mean spectral index of −1.6.

[2] Specifically, slowly rotating pulsars. See Chapter 1 for clarifications.

In contrast, MSP and pulsars in binaries initially had their spectra measured via much more specific campaigns (See e.g., Foster et al., 1991). Over the last few decades as MSPs have become increasingly important as extreme objects in their own right and as probes of fundamental physics, a few dedicated campaigns (Kramer et al., 1998; Toscano et al., 1998) have been carried out to measure the spectra of MSPs. The Kramer et al. (1998) survey used observations made over a few years at the Effelsberg 100-m radio telescope, with a centre frequency of either 1.4 GHz or 1.7 GHz with a total bandwidth of 300 MHz. The Toscano et al. (1998) campaign surveyed MSPs in the Southern sky using the Parkes 64-m radio telescope, centred at 1.4 GHz. Both Kramer et al. (1998) and Toscano et al. (1998) determine the mean spectral index of MSPs to be about −1.8.

More recently Kuniyoshi et al. (2015) and Frail et al. (2016) have measured a number of spectral indices using data from imaging surveys, the

VLA 74-GHz survey (Condon et al., 1998) and the the GMRT Southern Sky (TGSS) survey, respectively.

Compared to the spectral index of classical pulsars, these surveys all report indices that appear to be slightly steeper. Although MSPs are by definition smaller in terms of their physical extent and believed to have weaker surface magnetic fields, there is no well-founded reason to expect that the emission processes are fundamentally different for MSPs and classical pulsars. However, the first MSPs with published spectra and pulse profiles also seemed to show, in some sense, complicated profiles. Kramer et al. (1998, and subsequent papers in their series) argue that this notion is the result of a selection effect. By defining the number of Gaussian components required to model the pulse profile as a measure of complexity, they show that most MSP profiles in their survey require between two and four components, which is the same as that for classical pulsars.

If this is a statistical truth and MSPs are indeed similar to classical pulsars, then the selection effects which affect population synthesis and survey yield projections for those pulsars must have similar effects on MSPs too. For instance, Bates et al. (2013) show that the generally accepted median spectral index for classical pulsars of ~ -1.6 is due to selection effects and the true spectral index distribution is more likely to be centred on -1.4. That the pulsar spectral index is not much flatter than this is quite well borne out by the limited number of detections of pulsars at frequencies as high as ~10 GHz–100 GHz for classical pulsars (see e.g., Morris et al., 1997; Maron et al., 2004). Wether such a bias affects the measured spectral index distribution of MSPs remains an open question. Similarly, while a significant fraction of the classical pulsars were discovered at low (≃400 MHz–800 MHz) frequencies, only a limited number of the currently known MSPs were discovered by such surveys while the larger fraction were found in relatively deeper follow-up observations of radio-bright *Fermi* sources. One of the possible reasons why the number of MSPs discovered at low-frequencies is low could be the existence of broken-power-law spectra or GHz peaked spectra. However, the number of MSPs known and expected to demonstrate such spectra remains a fairly small number and the spectral breaks typically occur at frequencies $\lesssim$400 MHz(see e.g., Kuniyoshi et al., 2015).

In this chapter we present new spectral indices for 12 MSPs measured using flux densities measured at 327 MHz, 1700 MHz and 2400 MHz with the 500-m radio telescope at Arecibo Observatory . These are complemented with spectral indices for 62 MSPs measured using flux densities collected from literature. We present details of the observations using one of the most wide bandwidth radio-receivers currently available for pulsars and initial inferences from these data in the following sections.

5.2 *Observational Setup and Data Analysis*

We used the 327-MHz, L-wide and S-low frontends with the Puerto Rico Ultimate Pulsar Processing Instrument (PUPPI) backend at the Arecibo Observatory for the sources listed in Table 5.2. The PUPPI instrument is a clone of the Green Bank Ultimate Pulsar Processing Instrument (GUPPI DuPlain et al., 2008). PUPPI is capable of simultaneously observing up to 700 MHz[3] of BW for the L-wide and S-low frontends, with centre frequencies of 1730 and 2380 MHz respectively and 87.5 MHz of BW for the 327-MHz frontend.

[3] This effective bandwidth was due to one of the observing machines being unavailable. As of late 2016, PUPPI can acquire the full bandwidth of 800 MHz again.

Integration times required for a 10-σ detection at 1400 MHz with a nominal bandwidth of 800 MHz were calculated for each of the 18 sources, assuming a spectral index of $\alpha = -1.4$ (following Bates et al., 2013) and using any flux measurements previously published, scaled appropriately.

5.2.1 *Estimation of the required number of epochs*

Interstellar scintillation is the most dominant factor affecting the observed flux density of a given pulsar at a given epoch (e.g., Rickett, 1977). As a result, observational strategies must include this effect to maximise the number of detections.

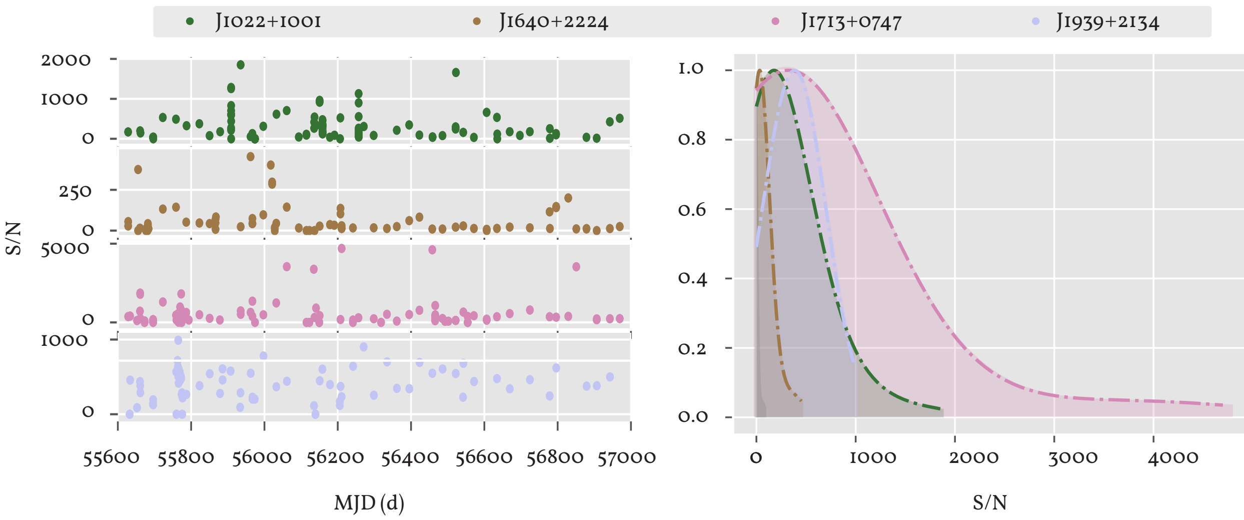

Figure 5.1: The figure shows the S/N's for selected pulsars observed regularly at the Effelsberg radio telescope. The left panel shows the scatter of the S/N per pulsar as a function of observing epoch. The right panel shows the distribution of the S/N.

The Effelsberg radio telescope has been regularly observing a large number of MSPs. We measured the S/Ns of several MSPs for observations made over several years. We verified that scintillation affected these data randomly, as is apparent in the left hand panel of Figure 5.1. This implies that the *number of observing epochs*, and not the total campaign length, determines how well scintillation effects can be characterised. Specifically, since we have not selected our samples to be limited to by their DM values, it is apparent that the timescales over which scintillation affects these data are either on the order of less than a month or larger than several years. From this set of S/Ns we estimated the standard deviation of S/N at 11 cm and 21 cm wavelengths. This was done independently for low ($\lesssim$70 cm^{-3} pc) and high-DM pulsars since their scintillation properties differ significantly. In addition, we also assume that the S/N values are Poisson distributed, similar to the right hand panel of Figure 5.1.

Assuming that the flux density measurement distributions of our samples are not significantly different from those for the sources observed regularly at Effelsberg, we then ran a Monte-Carlo simulation to obtain the expected standard deviation in the measured spectral index distribution as a function of the number of observing epochs. For this, we generate random flux density values at 327 MHz, 1440 MHz and 2450 MHz assuming a 'true' spectral index that is varied discretely from −3.1 to 1.0. We then estimate a spectral index from these generated flux density values and repeat the entire computation over 1000 cycles. The standard deviation of the resulting set of measured spectral indices for each input spectral index value is plotted in figure Figure 5.2 as a function of the number of observing epochs. Thus it is evident that for pulsars with a high DM two or three epochs are enough to obtain spectral indices with an uncertainty comparable to those found in literature ($\sigma_{\alpha_{meas}} \lesssim 0.3$) while for low-DM pulsars we would require at least six epochs.

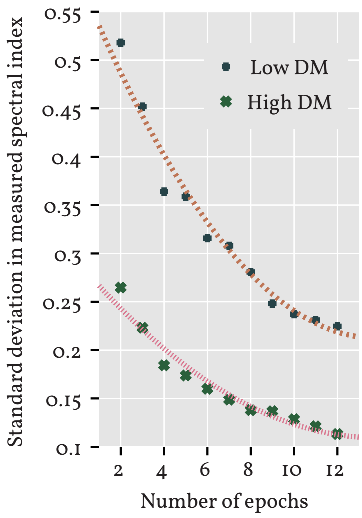

Figure 5.2: Plot of the standard deviation of the measured spectral index as a function of the number of epochs of observation. We can see that the standard deviation in the measured spectral index for low DM pulsars falls to ~0.3 only after six epochs while for high DM pulsars after three epochs of observation the standard deviation is less than 0.25.

5.2.2 Observations and RFI excision

Having accounted for the effects of scintillation we observed the listed sources for six epochs.The observations were folded 'online' using pulsar ephemerides from the ATNF Pulsar Catalogue (Manchester et al., 2005) [4]. However, in most cases these ephemerides were found to have significant errors in either the DM values or other pulsar timing parameters. These were corrected by re-deriving the pulsar timing models for these sources using the TEMPO2 pulsar timing package[5] (Hobbs et al., 2006) as described in Chapter 2. While this does not necessarily affect the flux measurements themselves, it is necessary to make precise estimates of the uncertainties in the flux density measurements.

The data were processed using the PSRCHIVE suite [6] (Hotan et al., 2004; van Straten et al., 2012). Due to the extremely wide-bandwidths the observations often suffered from excessive RFI particularly the S-band receiver. RFI excision was performed using the *zap* tool from PSRCHIVE along with a custom python script[7].

[4] www.atnf.csiro.au/research/pulsar/psrcat

[5] www.atnf.csiro.au/research/pulsar/tempo2

[6] Commit hash - fc8f777; psrchive.sourceforge.net

[7] Modified version of clean.py from the CoastGuard package; https://github.com/plazar/coast_guard

5.2.3 Polarisation and Flux calibration

Following the cleaning and updating of the timing ephemerides, the data were polarisation and flux calibrated using the *pac* and *fluxcal* tools from the the PSRCHIVE suite in the manner described in Section 2.3 and Section 2.4, respectively.

The *fluxcal* program relies on the spectra of the continuum source being supplied in the the form of either a power law or a log-polynomial of the form specified by Baars et al. (1977). While the power law approximation is often broadly applicable when limited flux density measurements are available, the log-polynomial fit is far more precise for well determined measurements, as is the case with most regularly observed

Source	RA (J2000)	DEC
B0038+328	00:40:55.06	+33:10:08.18
B0428+205	04:31:03.73	+20:41:04.30
B1040+123	10:42:44.61	+12:03:31.47
B1442+101	14:45:16.46	+09:58:35.89
B2209+080	22:12:01.41	+08:19:15.98

Table 5.1: List of continuum sources used as flux calibrators. Log-polynomial fits to their spectra are provided in the appendices.

continuum sources. Therefore we derive the spectra for each of the continuum sources observed using data from the NASA/IPAC Extragalactic Database[8]. As an example, the spectrum of is plotted in Figure 5.3. Similar log-polynomials were constructed for the flux-calibrators selected from the flux calibrator catalogue for Arecibo, which are listed in Table 5.1.

[8] https://ned.ipac.caltech.edu

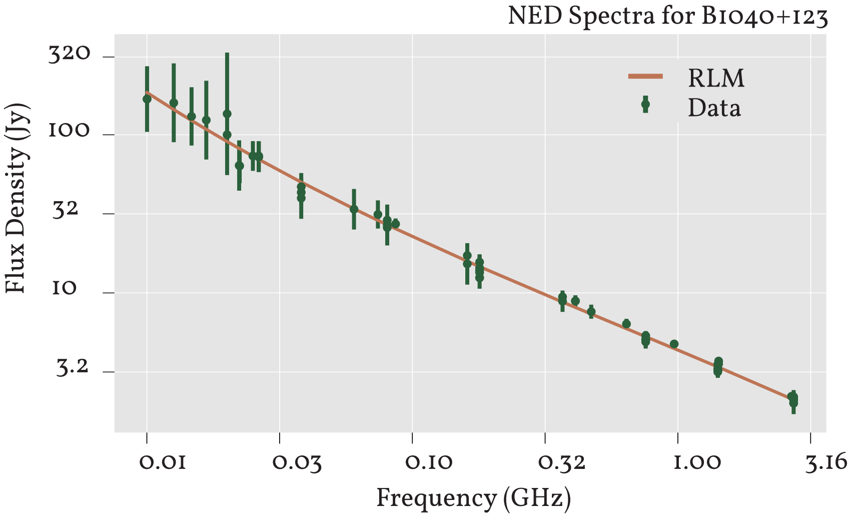

Figure 5.3: Spectrum of B1040+123 derived from recalibrated NASA/IPAC Extragalactic Database (NED) data (green points), using a robust linear models fit and fitted to a 3rd order log-polynomial following (Baars et al., 1977) (brown line).

5.3 Observed sources, flux density measurements and spectral indices

The complete list of sources observed at Arecibo is shown in Table 5.2. While all 19 sources were detected in the final observations in at least one band, spectral indices were obtained for only 12. The number of detections for each source per epoch is low. In most cases RFI combined with fluctuations due to scintillation prevents the detection of the source in some of the bands[9]. Typically the S-band receiver, which is also the most susceptible to RFI showed the smallest number of detections.

[9] correctly speaking, the elimination of a large number of channels due to the RFI mitigation code causes the remaining signal to be too faint.

5.3.1 Measurement of flux densities

The uncertainty on the flux density measurement depends strongly on the proper determination of the baseline root-mean-square (rms) error or the noise level, which depends on correctly identifying the off-pulse region. In a comparison of the effects of scaling the number of phase bins, we found that the most commonly available tools from PSRCHIVE, *psrstat* and *pdv* report inconsistent flux densities, uncertainties or both. Hence, flux densities were estimated using custom code, which allows manual identification of on- and off-pulse regions. While manual identification of the on- and off-pulse regions inherently suffers from small variations in the identified phase bins, the stability of the reported flux values and their uncertainties as a function of the total number of phase bins was found to be much greater than for any of the automatic tools listed above.

Pulsar J2000	G.Long. l (deg)	G.Lat. b (deg)	Dist. (kpc)	DM (pc cm^{-3})	$\nu_{\rm discov}$ (MHz)	$\nu_{\rm discov}$ reference	Period, $P_{\rm o}$ (ms)	$P_{\rm o}$ reference	BinComp
J0030+0451	113.14	−57.61	0.36	4.332 52(4)	430	Lommen et al. (2000)	4.865 453 210 909 7(8)	Arzoumanian et al. (2015)	Isolated
J0337+1715	169.99	−30.04	1.30	21.3162(3)	350	Boyles et al. (2013)	2.732 588 632 44(9)	Ransom et al. (2014)	He-WD, Triple
J0407+1607	176.62	−25.66	1.34	35.65(2)	430	Lorimer et al. (2005)	25.701 739 194 63(2)	Lorimer et al. (2005)	He-WD
J1038+0032	247.15	48.47	4.25	26.59(20)	1400	Burgay et al. (2006)	28.851 557 951 31(8)	Burgay et al. (2006)	Isolated
J1453+1902	23.39	60.81	1.27	14.049(4)	430	Lorimer et al. (2007)	5.792 302 734 966 4(4)	Lorimer et al. (2007)	Isolated
J1640+2224	41.05	38.27	1.50	18.4260(8)	430	Wolszczan (1995)	3.163 315 867 760 34(5)	Desvignes et al. (2016)	He-WD
J1910+1256	46.56	1.80	1.50	38.0701(8)	1400	Faulkner et al. (2004)	4.983 583 941 388 65(15)	Arzoumanian et al. (2015)	He-WD
J1913+0617	40.98	−1.90	5.76	155.8(0)	1400	Lazarus et al. (2015)	5.026 861 113 83(9)	Lazarus et al. (2015)	*
J1923+2515	58.95	4.75	1.20	18.857 66(19)	350	Lynch et al. (2013)	3.788 155 519 613 03(52)	Lynch et al. (2013)	Isolated
J1944+0907	47.16	−7.36	1.36	24.34(2)	430	Champion et al. (2005)	5.185 201 904 126 0(3)	Champion et al. (2005)	Isolated
J1949+3106	66.86	2.55	7.47	164.1264(5)	1400	Deneva et al. (2012)	13.138 183 343 704 0(6)	Deneva et al. (2012)	CO-WD
J1955+2527	62.74	−1.57	8.18	209.971(3)	1400	Deneva et al. (2012)	4.872 765 862 979 3(17)	Deneva et al. (2012)	Isolated
J2016+1948	60.52	−8.67	2.16	33.8148(16)	430	Navarro et al. (2003)	64.940 388 241 514(6)	Gonzalez et al. (2011)	He-WD
J2033+1734	60.86	−13.15	1.74	25.0791(3)	430	Ray et al. (1996)	5.948 957 534 850 2(6)	Splaver (2004)	He-WD
J2043+1711	61.92	−15.31	1.25	20.709 87(3)	350	Hessels et al. (2011)	2.379 878 923 709 447(11)	Arzoumanian et al. (2015)	He-WD
J2214+3000	86.86	−21.67	0.60	22.5511(8)	820	Ransom et al. (2011)	3.119 226 580 479 40(9)	Arzoumanian et al. (2015)	UL
J2234+0611	72.99	−43.00	0.69	10.72(5)	327	Deneva et al. (2013)	3.576 581 686(2)	Antoniadis et al. (2016)	He-WD
J2235+1506	80.88	−36.44	1.54	18.09(5)	429	Camilo et al. (1993)	59.767 357 984 547(9)	Camilo et al. (1996)	Isolated
J2322+2057	96.51	−37.31	1.01	13.36(4)	430	Nice et al. (1993)	4.808 428 289 464 1(17)	Desvignes et al. (2016)	Isolated

Table 5.2: List of sources observed at Arecibo with their galactic coordinates (G.Long. and G.Lat.) and distances estimated from the Yao et al. (2017) model of the distribution of electrons in the Galaxy. Also shown are the frequencies at which the sources were discovered and the respective spin-periods at discovery. The final column shows the nature of the companion if it is known and * indicates that a companion is expected but its nature is not known.

The errors in the PSRCHIVE tools were found to be linked to the algorithms used to identify the on- and off-pulse regions. However, in our investigations of the results from the PSRCHIVE tools, we found that even using predefined phase bin ranges using *psrstat* with the *set* method,[10] did not lead to stable flux density values being reported by the tools.

[10] From the psrstat manual available at http://psrchive.sourceforge.net/manuals/psrstat/algorithms.shtml: "**set** A static selection of regions of pulse phase."

5.3.2 *Spectral Indices*

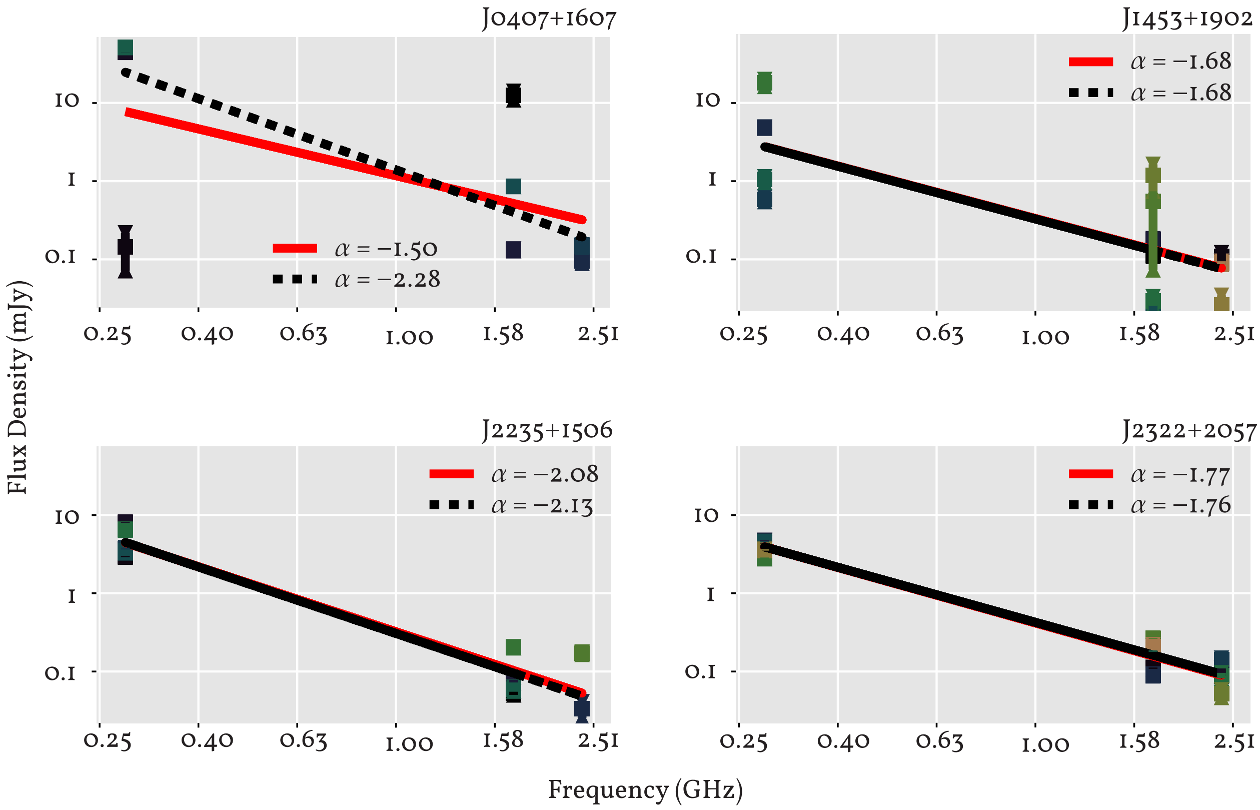

Figure 5.4: Spectral indices of four of 12 pulsars observed using the Arecibo Observatory. A total of 19 pulsars were observed, of which 18 were detected and spectral indices were obtained for 12. The flux density measurements represented by filled squares, are mapped to colour-intensity as a function of the epoch of observation, with darker colours representing the earlier epochs and the lighter colours, the later epochs. The solid red lines are the simple weighted linear regression fits while the dotted black lines show the robust linear regression fits.

Having established that the measured flux densities and uncertainties were reliable, we fitted for the spectral indices using a simple power-law fit. To minimise the effect of outliers, we used the median absolute deviation as a discriminant and performed robust linear fits using the *statsmodels*[11] package. However, given that the total number of observations is rather low and often individual epochs are not sufficient to measure a spectral index, the robust linear fits typically converge to the simple linear fits, as demonstrated in the selection of sources shown in Figure 5.4, except in the case of sources like PSR J0407+1607 where the robust fit performs better at rejecting the unexpectedly faint detection at 327 MHz and PSR J1453+1902, where the faint detections at 1730 MHz and 2380 MHz are correctly accounted for in the fit. The confirmation for the robustness of the fitting algorithms and the flux calibration is provided by the good prediction of values at other frequencies (where available) obtained from literature and plotted in Figure 5.5.

[11] http://statsmodels.sourceforge.net/

As can be seen in Figure 5.4 and Figure 5.5, scintillation induces a

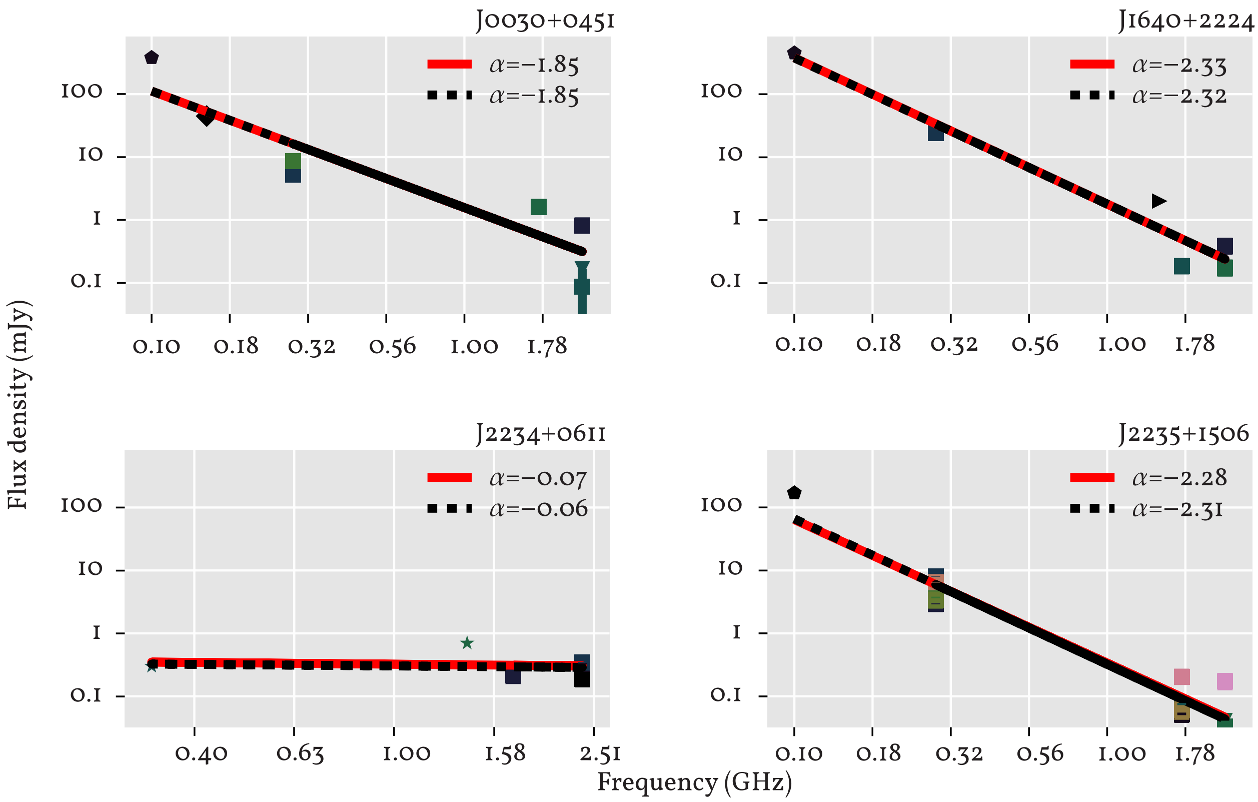

Figure 5.5: Spectral indices of some more of the 12 pulsars observed using the Arecibo Observatory, compared with flux-density values from literature, shown with symbols other than the square markers. Other symbols and colours follow the same scheme as Figure 5.4. Except in the case of PSR J0030+0451, where the spectral index measurement is only made more precise, the addition of flux-density values from literature does not significantly alter the spectral indices measured from our observations alone.

fairly large amount of variation in the estimated flux density per epoch for all the pulsars in our sample. However, in none of the sources observed do we find any regular or periodic trends, justifying the assumptions made for our simulations in Section 5.2.1.

Table 5.3 shows the measured spectral indices for the 12 sources for which observations at more than one band per epoch is available. Spectral indices measured using Arecibo data alone and those obtained by combining flux densities gathered from literature are shown in columns two to five and six to eight, respectively. Six of the sources were detected at one band only and we do not attempt to measure their spectral indices while one source, PSR J1913+0617, was not detected at all.

For five sources listed at the top of Table 5.3, however, we have a large number of detections and these display the resilience of our fits against outliers, e.g., for the case of the scintillation brightened observations of PSR J2235+1506 at L- and S-bands or the scintillation dimmed observations of PSR J1453+1902 at the same bands. The fits are also able to reject outliers at the edges of the fit, e.g., the faint observation of PSR J0407+1607 at 327 MHz.

Figure 5.6 shows the Gaussian kernel density estimates (KDE; Rosenblatt, 1956) of the 12 measured spectral indices [12] compared with published MSP spectral indices from Frail et al. (2016); Kramer et al. (1998) and Toscano et al. (1998). Of these, Frail et al. (2016) uses flux densities from the TGSS (TGSS Frail et al., 2016) survey at 150 MHz and use any flux densities available at other frequencies from literature to measure two-point spectral indices while Kramer et al. (1998) and Toscano

[12] column 6 of Table 5.3

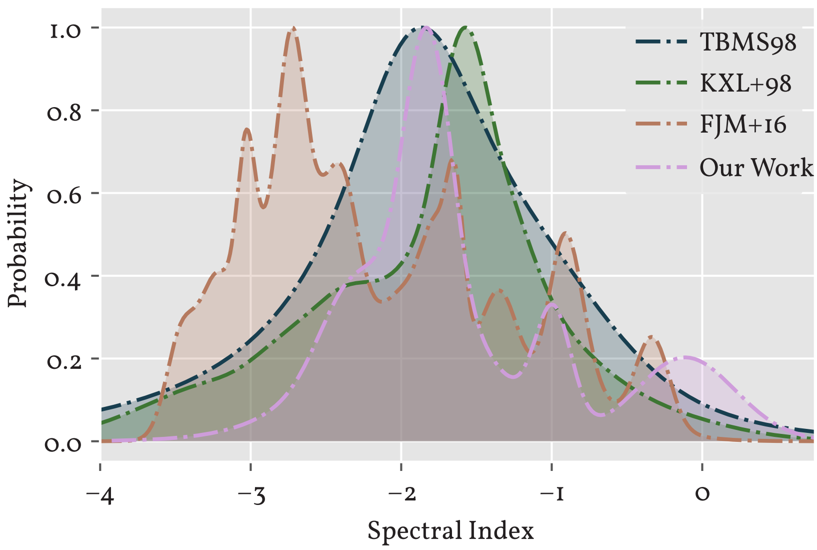

Figure 5.6: Comparison of spectral indices measured with those available in literature. The key codes correspond to the following publications, in order; Toscano et al. (1998); Kramer et al. (1998); Frail et al. (2016) and finaly, spectral indices measured from our observations and combined with available flux densities in literature. Data for the Toscano et al. (1998) values were obtained from the ATNF pulsar catalogue (Manchester et al., 2005).

et al. (1998) report spectral indices using observations at ~1.4 GHz. The rather broad spread of the Frail et al. (2016) results reflects the larger sample of MSPs while the skew towards steeper spectral indices is expected due to the sample bias caused by the low frequency of the TGSS. The results from Toscano et al. (1998) show excellent agreement with our measurements, while Kramer et al. (1998) measures a slightly steeper median spectral index. The broad 'bump' towards the flat-spectrum tail for our measurements is due to three sources in Table 5.3, PSRs J0337+1715, J1944+0907 and J2234+0611. Perhaps the most distinctive feature is the rather precise determination of spectral indices for our sample, as evinced by the narrow peaks.

Pulsar	Arecibo data only				Arecibo data + Lit. values			
(J2000)	α_{rlm}	σ_{rlm}	95 C. L.		α_{rlm}	σ_{rlm}	95 C. L.	
J1453+1902	−1.7	0.4	−2.6	−0.9	–	–	–	–
J1955+2527	−1.0	0.1	−1.2	−0.8	–	–	–	–
J2016+1948	−1.9	0.2	−2.2	−1.6	–	–	–	–
J2235+1506	−2.2	0.2	−2.6	−1.8	−2.4	0.2	−2.7	−2.0
J2322+2057	−1.7	0.1	−1.9	−1.5	–	–	–	–
J2033+1734	−1.7	0.2	−2.0	−1.4	−1.9	0.1	−2.1	−1.7
J0030+0451	−1.0	0.2	−1.4	−0.6	−1.9	0.4	−2.5	−1.2
J0337+1715	0.3	0.6	−0.9	1.4	−0.1	0.3	−0.8	0.5
J0407+1607	−2.3	0.6	−3.5	−1.1	–	–	–	–
J1640+2224	−2.2	0.4	−3.0	−1.4	−2.3	0.4	−3.1	−1.6
J1944+0907	−0.0	0.0	−0.1	0.1	−1.7	0.4	−2.4	−1.0
J2234+0611	0.7	1.7	−2.7	4.0	−0.1	0.3	−0.7	0.6

Table 5.3: Spectral indices measured using only our observations at Arecibo (columns 2-5) and after including any avaiable values from literature (columns 6-9). For the first five sources, we have a large number of observations per epoch per band and the measured spectral indices do not show any significant changes after inclusion of available values from literature. The second set typically have less than three epochs of detections (see text). The third set show two pulsars for which the spectral indices measured are unsually flat.

We note that compared to Kramer et al. (1998), who used a centre frequency tunable from 1300 MHz to 1700 MHz with a system bandwidth of 40 MHz and Toscano et al. (1998) who used four bands centred at 436 MHz, 660 MHz, 1400 MHz and 1660 MHz with system bandwidths of 32 MHz and 128 MHz for the former and latter pairs of bands, respectively, our survey uses data from system with effective bandwidths of 87.5 MHz and 700 MHz for the 327 MHz, and the 1700 MHz and 2400 MHz bands, respectively. We also note that for the Frail et al. (2016) values, the plot would appear smoother if we had used inflated errors as recommended by those authors.

5.3.3 Sub-populations

The comparison of the spectral index distributions for the various sub-populations reveal some potentially interesting features. The data used for these plots consist of the weighted means of published spectral indices, along with those measured via our observations.

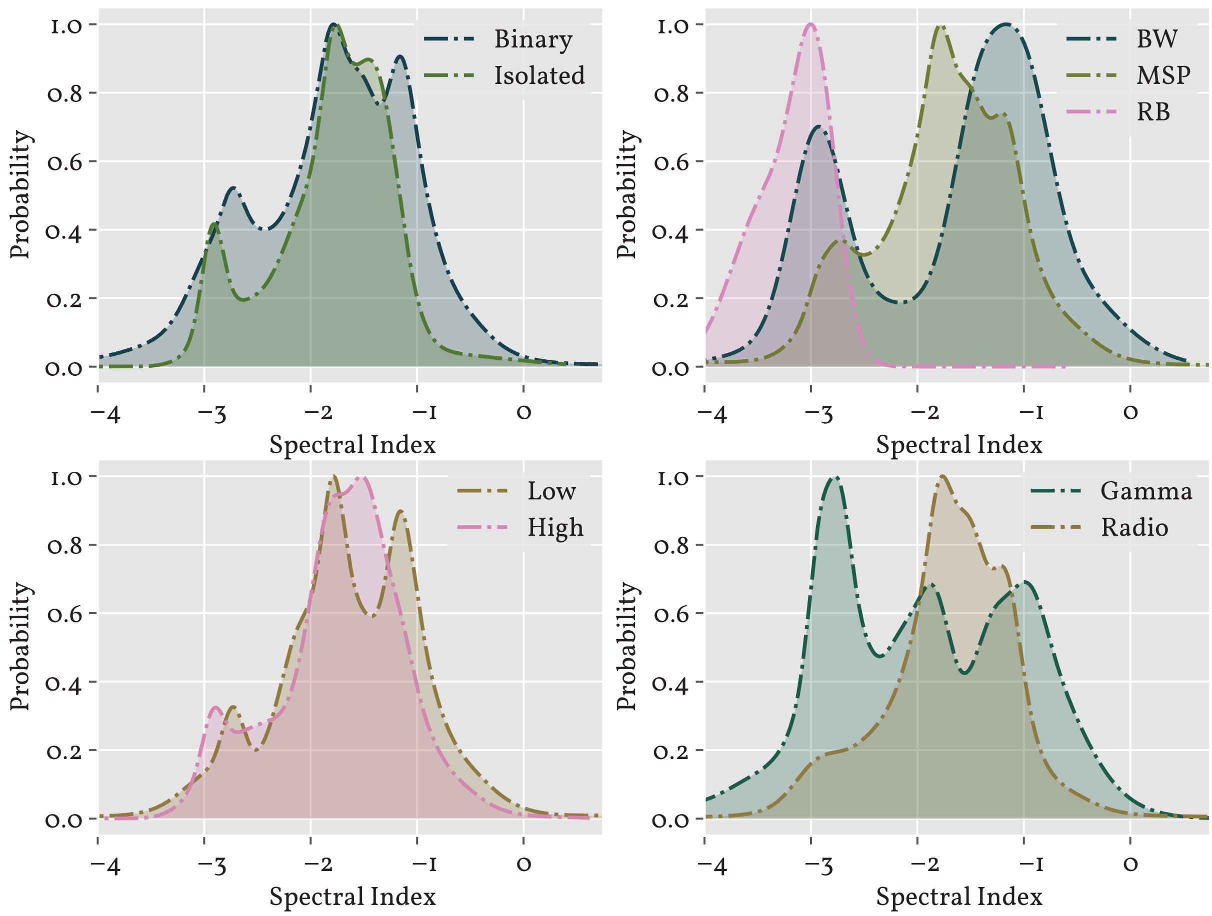

There have been claims of a differentiation in the spectral indices of isolated MSPs and those in binaries. Our sample does not offer much evidence to support this claim, even though it includes 74 Galactic MSPs of which 16 are isolated sources while 56 are in binary systems. Instead we find that the median spectral indices for the two groups coincide, as shown in the top left panel of Figure 5.7, although the isolated sources appear to have marginally smaller bounds on their spectral index distribution.

The comparison of the BWP and RBP (See e.g., Chen et al., 2013) systems in the top right panel of Figure 5.7, with the remaining MSP population shows that the seven such sources (five BWP and two RBP sys-

Figure 5.7: Comparison of spectral indices by sub-populations. The top left panel compares the spectral indices of isolated MSPs against those in binaries. The top right panel compares MSPs with RBP and BWP systems (see text). The bottom left panel shows the comparison of the spectral indices of sources classified by discovery frequency. The large peak at the left hand side of the high frequency group is the result of *Fermi* follow-up and targeted searches. The bottom right panel compares the spectral indices of known γ−ray sources with those detected at radio frequencies alone.

Pulsar J2000	G.Long. l (deg)	G.Lat. b (deg)	Spectral Index α	σ	Pulsar J2000	G.Long. l (deg)	G.Lat. b (deg)	Spectral Index α	σ
J0034-0534	111.49	−68.07	−2.7	0.1	J1747-4036	350.21	−6.41	−2.8	0.2
J0214+5222	135.63	−8.42	−2.9	0.4	J1802-2124	8.38	0.61	−1.9	0.2
J0218+4232	139.51	−17.53	−2.7	0.1	J1804-2717	3.51	−2.74	−2.9	0.2
J0437-4715	253.39	−41.96	−1.0	0.2	J1810+1744	44.64	16.81	−2.4	0.8
J0613-0200	210.41	−9.30	−1.9	0.1	J1816+4510	72.83	24.74	−3.4	0.3
J0621+1002	200.57	−2.01	−1.6	0.2	J1832-0836	23.11	0.26	−1.5	0.2
J0636+5129	163.91	18.64	−1.1	0.1	J1843-1113	22.05	−3.40	−2.9	0.1
J0645+5158	163.96	20.25	−1.4	0.3	J1857+0943	42.29	3.06	−1.3	0.2
J0711-6830	279.53	−23.28	−1.5	0.2	J1902-5105	345.65	−22.38	−2.5	0.2
J0751+1807	202.73	21.09	−0.7	0.3	J1903+0327	37.34	−1.01	−2.0	0.2
J1012+5307	160.35	50.86	−1.2	0.2	J1905+0400	38.09	−1.29	−2.8	0.3
J1022+1001	231.79	51.10	−1.1	0.1	J1910+1256	46.56	1.80	−1.4	0.2
J1024-0719	251.70	40.52	−1.3	0.2	J1911-1114	25.14	−9.58	−2.1	0.3
J1038+0032	247.15	48.47	−0.9	0.7	J1915+1606	49.97	2.12	−1.2	0.1
J1045-4509	280.85	12.25	−1.4	0.1	J1918-0642	30.24	−9.12	−2.0	0.1
J1231-1411	295.53	48.39	−2.8	0.2	J1923+2515	58.95	4.75	−1.9	0.2
J1300+1240	311.31	75.41	−1.9	0.2	J1939+2134	57.51	−0.29	−2.4	0.1
J1455-3330	330.72	22.56	−1.8	0.1	J1949+3106	66.86	2.55	−1.8	0.3
J1518+4904	80.81	54.28	−0.7	0.3	J1955+2908	65.84	0.44	−1.8	0.2
J1537+1155	19.85	48.34	−1.8	0.3	J1959+2048	59.20	−4.70	−2.9	0.1
J1543-5149	327.92	2.48	−2.5	0.2	J2019+2425	64.75	−6.62	−5.8	2.2
J1544+4937	79.17	50.17	−1.0	0.2	J2043+1711	61.92	−15.31	−4.1	0.9
J1603-7202	316.63	−14.50	−2.0	0.1	J2051-0827	39.19	−30.41	−1.4	0.2
J1643-1224	5.67	21.22	−1.7	0.1	J2124-3358	10.93	−45.44	−1.8	0.2
J1713+0747	28.75	25.22	−0.6	0.4	J2129-5721	338.01	−43.57	−2.4	0.2
J1721-2457	0.39	6.75	−1.5	0.2	J2145-0750	47.78	−42.08	−1.5	0.1
J1730-2304	3.14	6.02	−1.6	0.1	J2214+3000	86.86	−21.67	−0.7	0.5
J1737-0811	16.93	12.32	−1.8	0.2	J2215+5135	99.87	−4.16	−3.0	0.2
J1738+0333	27.72	17.74	−1.5	0.1	J2229+2643	87.69	−26.28	−1.5	0.2
J1744-1134	14.79	9.18	−1.3	0.2	J2302+4442	103.40	−14.00	−0.9	0.3
J1745-0952	16.37	9.90	−1.1	0.1	J2317+1439	91.36	−42.36	−0.9	0.2

Table 5.4: Spectral Indices of MSPs derived from published flux values. Fluxes were collected from the ATNF pulsar catalog (v 1.54), Kondratiev et al. (2016), Kuniyoshi et al. (2015) Kramer et al. (1998) and Frail et al. (2016).

Type	Number	α_{mean}	σ_{mean}	α_{median}	σ_{median}
Binary	56	−1.77	0.19	−1.85	0.07
Isolated	16	−1.69	0.23	−1.80	0.07
BW	5	−1.43	0.24	−1.69	0.21
RB	2	−3.19	0.19	−3.19	0.19
GRS	23	−2.00	0.25	−2.00	0.07
Radio	51	−1.70	0.40	−1.74	0.07
Low freq. discov.	32	−1.69	0.23	−1.70	0.10
High freq. discov.	34	−1.66	0.17	−1.73	0.04
All MSPs	74	−1.74	0.18	−1.74	0.04
Classical PSRs	367	−1.60	0.20	−1.57	0.02

Table 5.5: Shown on the left are the rederived mean and median spectral indices of MSPs, grouped by sub-populations analysed in Figure 5.7 and those of classical pulsars, collected from the ATNF pulsar catalogue (Manchester et al., 2005). The comparision of MSPs against classical pulsars is shown in Figure 5.8.

tems, respectively) in our sample do show a marginal bias towards steeper spectral indices. The apparent split in the BWP distribution is more likely the result of the small sample as there are no obvious correlations in the measured properties of the systems or even the companions to these pulsars.

For the comparison of spectral indices by discovery frequency (i.e., bottom left panel of Figure 5.7), we define all surveys with central frequencies less than 1 GHz as 'low'-frequency surveys and the rest, excluding high energy surveys, as high-frequency surveys. The resulting similarity of the two distributions, which appears surprising initially, is due to the fact that in the post-*Fermi* era, many of the MSPs were discovered through deep targeted surveys.

The bottom right panel of Figure 5.7 shows the strong bias of the 23 pulsars that are detectable only in the γ-ray regime towards steep spectral indices, while the 51 pulsars that are visible as radio sources seem to be have spectral indices that are narrowly distributed, tending towards the median spectral index measured for the combined sample.

Domain	Corr. coeff	p-value
$\alpha - P_0$	0.27	0.03
$\alpha - P_1$	−0.17	0.16
$\alpha - \dot{E}$	−0.31	0.01
$\alpha - B_{surf}$	0.00	0.98
$\alpha - P_b$	0.03	0.78
α − Ecc	0.04	0.75

Table 5.6: Spearman rank correlations for the spin-period (P_0), the spin-down rate (P_1), spin-down energy loss for a dipolar magnetic field ($\dot{E}$), the surface magnetic field strength (B_{surf}), the orbital period for binary systems (P_b) and eccentricity of the binary orbit (Ecc) with the spectral index (α).
The p-value

5.3.4 *Correlations between MSP properties and spectral index*

Using all 74 MSPs we also investigate the correlations between selected pulsar and binary parameters such as the spin-period (P) and the spin-down rate ($\dot{P}$) in the top panels of Plate 5, the estimated spin-down energy ($\dot{E}$) and surface magnetic field (B_{surf}) assuming a median pulsar mass of 1.8(10) $M_\odot$ and a spherical mass distribution, in the two middle panels and finally, for the 56 systems with one (known) companion, we plot the spectral index as a function of the orbital period (P_b) and the orbital eccentricity (e) in the bottom two panels. Of these we find that the spin-period is weakly correlated with the spectral index while the spin-down energy is weakly anti-correlated, as evinced by the small coefficients and p-values (see Table 5.6) returned by the Spearman rank correlation test. The spin-down rate on the other hand shows a very weak anti-correlation with the spectral index while the associated surface magnetic, the orbital period and the eccentricity are uncorrelated with the spectral index.

Following Lorimer et al. (1995), we construct an expression for the spectral index in terms of the spin-period and the spin-down rate,

$$\alpha = k_1 + k_2 log\left(\frac{P}{\dot{P}}\right) \tag{5.2}$$

which is solved using linear regression to obtain

$$\alpha = -15(4) + 0.34(9) log\left(\frac{P}{\dot{P}}\right) \tag{5.3}$$

5.3.5 *A note on the Fermi sources in our data*

Of the 19 sources we have observed, 5 are known *Fermi* sources (see updated version of Abdo et al., 2013, available at the *Fermi* website[13]). The steepest spectral index in the *Fermi* sources is measured for PSR J1640+2224, which may host a pulsar of mass ~1.4(4) $M_\odot$ (Löhmer et al., 2005) while

[13] https://fermi.gsfc.nasa.gov/ssc/data/access/lat/4yr_catalog/

PSR J0337+1715, which is in a canonical triple-system (Ransom et al., 2014) and hosts a pulsar of mass 1.4378(13) $M_\odot$, appears to have a rather flat spectral index. PSR J2234+0611, which also hosts a 1.393(13) $M_\odot$ pulsar and harbours a HeWD companion of mass 0.275(8) $M_\odot$ (Antoniadis et al., 2016), shows a rather flat spectrum when flux density from literature (Deneva et al., 2013) are combined with our measurements. The low number of detections for this pulsar appear to support claims of extreme scintillation induced fluctuations due to the low DM of 11 cm^{-3} pc. Both the spectral index and the scintillation seem to agree with the expectations laid out in Deneva et al. (2013). PSR J0030+0451 is an isolated γ-ray bright pulsar (Abdo et al., 2009) with a spectral index −1.9(4) while PSR J2016+1948 (Navarro et al., 2003) is in an eccentric binary with a WD companion and has a spectral index of −1.9(2). The most likely conclusion from this small sample simply confirms that γ-ray production is an ubiquitous phenomenon in recycled pulsars, as expected.

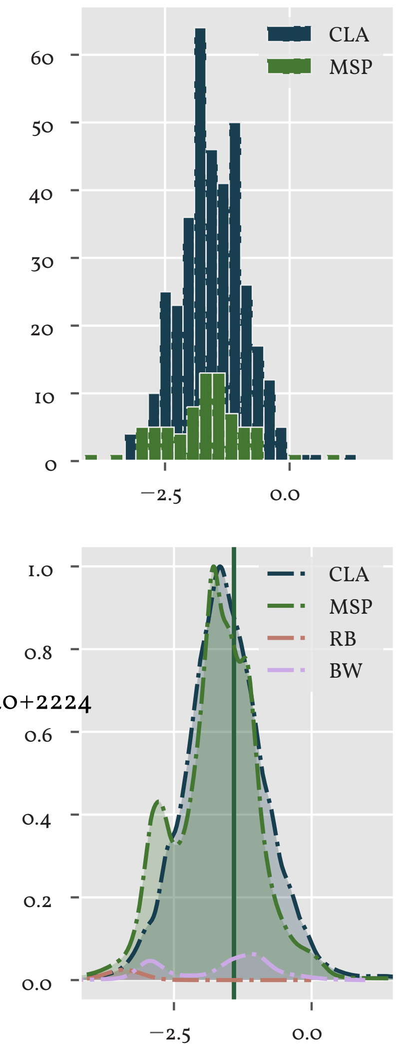

Figure 5.8: Comparison of the spectral indices of classical pulsars (denoted by CLA) with MSPs. The top plot shows simple histograms of the measured spectral indices, while the lower plot shows the kernel density estimator (KDE) for the same. The vertical line shows the Bates et al. (2013) corrected median spectral index for classical pulsars at −1.49(96).

5.4 *Summary and conclusions*

We present the results of a dedicated campaign at three frequency bands, to measure the spectral indices of galactic MSPs. We have confirmed that our initial assumptions are reasonable and show that our choice of calibration techniques and the final spectral index measurements are robust and able to cope with outliers well, even though the methods we use are generic. We demonstrate that it is the number of epochs of successful observations that determines the accuracy of the measured spectral indices and spectral indices obtained from our measurements alone are able to predict fairly well the flux densities at lower frequencies.

We have measured the spectral indices of 12 MSPs of which only PSR J1640+2224 had a previously published value. In combination with values from literature, we now have spectral indices for 74 MSPs. This is only slightly less than half the total number of currently known Galactic MSPs which stands at 195 (Manchester et al., 2005). Sub-population analyses show some interesting features, including suggestions for a uniform distribution of spectral indices for isolated and binary MSPs. The RBP population, although poorly sampled, appears to consist of steep spectrum sources exclusively. In contrast the BWP population shows a broader distribution but with wider limits on the determined indices.

Perhaps the most significant result of the sub-population analysis is the result that *Fermi* sources tend to have steeper spectral indices. This suggests that radio frequency surveys in the era just before *Fermi* , which had typically been centred around ~1.4 GHz were less sensitive to such sources, explaining the large number that were missed by those surveys. This is perhaps also the reason why *Fermi* follow-up searches with the Effelsberg 100-m radio telescope at ~1.4 GHz was far less successful than low frequency (~800 MHz) searches at the Green Bank telescope (Boyles et al., 2013; Stovall et al., 2014). While we see only a suggestive trend by the small number of BWP and RBP in our sample (5 and 2, respectively) to favour steep spectral indices, this would also explain the large number of such objects that *Fermi* has detected.

We find that the spectral index distribution of classical pulsars and that of MSPs are very similar as shown in Figure 5.8. The two-tailed Kolmogorov-Smirnov test (see e.g., Peacock, 1983) suggests that there is a strong likelihood (p-value 0.28) that the spectral index distributions of the two populations are identical. Welch's t-test (Welch, 1947) weakly rejects the hypothesis that the median spectral indices are similar with a p-value of 0.08. We note however the suggested median spectral index (shown by the green vertical line in Figure 5.8) for classical pulsars from Bates et al. (2013) lies at −1.41(96), a value that is quite flatter than the median value we have measured for these pulsars using published spectral indices.

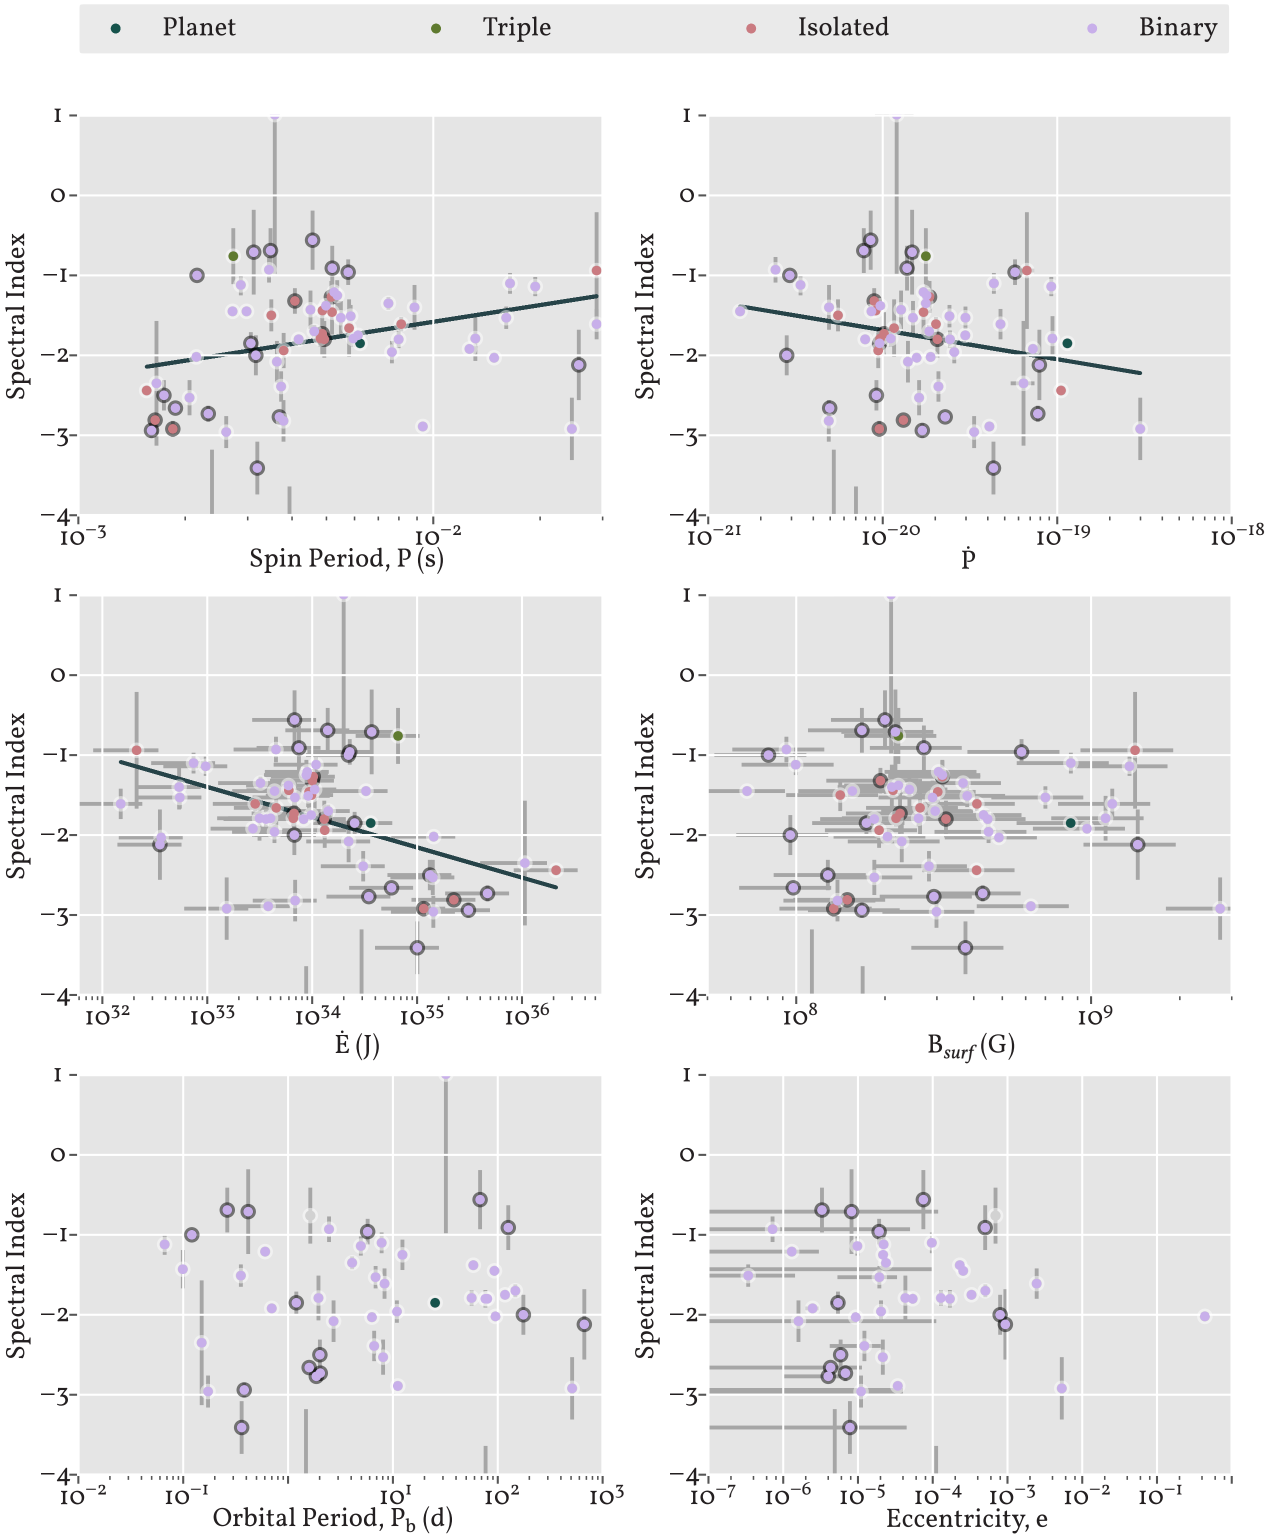

Plate 5: Comparisons of the spectral index of the larger sample of MSPs as a function of their spin period (top left), the orbital period (top right), the surface magnetic field (bottom left) and the eccentricity (bottom right). Points overdrawn with a gray circle represent γ–ray sources and the keys 'Planet' and 'Triple' refer to the pulsar with planetary companions, PSR J1300+1240 (B1257+12) and the triple system PSR J0337+1715, respectively.

6 Conclusions

こぬ人を
まつほの浦の
夕なぎに
やくやもしほの
身もこがれつつ

সখী, ভাবনা কাহারে বলে।
সখী, যাতনা কাহারে বলে।
সে কি কেবলই যাতনাময়।
সে কি কেবলই চোখের জল ?

–**রবীন্দ্রনাথ ঠাকুর;** 1881, গীতবিতান

In this thesis I have studied three aspects of millisecond pulsar studies. The first is linked to the artefacts that result from the digital signal processing within data-recording systems (or backends) used for pulsar astronomy. The second aspect involves the study of a black-widow pulsar, PSR J2051–0827 using pulsar timing data derived from 21 years of continuous observation. The final aspect investigated here is the measurement of the spectral indices of millisecond pulsars. In this final chapter I summarise the main results from the various chapters and comment on how the work presented here may be improved and expanded.

6.1 Introduction

MSPs are some of the most fascinating objects known to man. Their study has led to remarkable insights into the nature of matter at densities greater than that of nuclear matter and their extreme timing precision allow us to test fundamental physics. However recycled pulsars, which account for only ~10% of the currently known pulsars, form an extremely varied population. While they have been observed for more than a quarter of a century now, many claims regarding the properties of these objects remain hotly contested. These range from the more exotic ideas regarding their internal structure to the more general properties linked to their emission and rotational stability. The recent discovery of transitioning MSPs (Stappers et al., 2014), which switch between accreting X-ray and radio-pulsar states has shown the validity of the long-held belief that recycling is the general process by which MSPs are produced. However, detections of glitches in MSPs by Cognard and Backer (2004) and McKee et al. (2016) have shown that MSPs are not always as stable rotators as was previously believed. Mass determinations by Antoniadis et al. (2013); Demorest et al. (2010) also suggest that these are some of the most massive pulsars. Finally, MSPs are crucially important for PTAs which are searching for nHz GWs that are emitted during the early inspiral of SMBHs. In summary, the study of MSPs has deep impacts on our understanding of fundamental physics.

6.1.1 *Summary of work presented in this thesis*

In the course of this thesis, I have presented three independent studies related to the observation of MSPs. The first study is an attempt to investigate the current limits on the observational accuracy with which we can reproduce MSP profiles. The second investigation then attempts to answer the question, "Can RBP and BWP be included in pulsar timing arrays?" and the final study seeks to find the spectral behaviour of the broader population. Admittedly, these studies represent only a small fraction of what we wish to understand about recycled pulsars. However, together they serve to address important questions on the possibility of improving the sensitivity of PTAs the broader framework within which a significant portion of this work has been carried out. In the following sections, I summarise the main findings of this work.

In Chapter 3 I have provided an overview of the different artefacts that are generated in PFBs (Bellanger et al., 1976). PFBs are now the most commonly used architecture for digital receivers, which allow one to construct computationally efficient systems that are currently capable of processing several hundred MHz of bandwidth instantaneously. While theoretical results have shown that for simple PFBs (Vaidyanathan, 1993; Crochiere and Rabiner, 1976, etc) it is possible to construct a PR FB it is also well known that solutions to the general problem may not converge quickly enough to be implemented efficiently on reprogrammable devices like FPGAs. By relaxing the design constraints to allow the power contained in the artefacts that arise from from process of filterbanking to be low enough that they do not significantly impact the signal of interest, PR FBs can be made into NPR FBs.

Typically, the measure of NPR can defined in terms of a 'distance' from a PR FB which depends on the number of branches in the analysis section the PFB M, the decimation factor; D, the interpolation factor; I, and the number of branches in the synthesis section; N [1]. The problem of finding the optimal NPR PFB can be further simplified by using modulated FBs where a single prototype filter is modulated to construct the entire FB. General relations for this distance are presented in Section 3.5.

[1] Proper definitions may be found in Chapter 3

By defining the distance from the PR conditions in terms of the energy content of the alias components in the polyphase filterbanks as

$$\min_{h_m, g_m} \int_{\pi}^{-\pi} |E_{h_m, g_m}(\omega)|^2 d\omega, \tag{6.1}$$

we have recovered the result that any NPR filterbank with sub-processing will have analysis artefacts with substantial energy content. We have also investigated a popular method of optimisation for NPR FBs including the action of dedispersion of pulsar signal. We find that in the case of coherent dedispersion, the sub-processing introduces a non-linear phase term in the filter response

$$E_{dedisp}(z) = \frac{1}{D} \sum_{d=1}^{D-1} H_{ISM}\left(z^{\frac{2\pi\mathscr{D}}{\omega_c^2}}\right) H\left(z^{1/D} W_M^m W_D^d\right) X\left(z^{1/D} W_D^d\right), \tag{6.2}$$

which is typically not accounted for in NPR PFBs based on FIR filters.

These can introduce low-level aliasing artefacts in pulsar data. Even though present-day pulsar backends do not employ full reconstruction FBs on the FPGA based backends, the analysis FBs that they do employ can be improved by applying the LMS optimisation introduced in Section 3.5.1, specifically in Eqn. (3.45).

Results from a simple *Python* simulation show that the major artefacts in pulsar data recorded using FBs based on short FIR filters can become very significant if the S/N of the pulsar is very high. Literature suggests that these artifacts can limit the precision of PTA measurements to ~100 ns (Morrison et al., 2015).

In Chapter 4 I studied the long-term timing of PSR J2051−0827. This was the second BWP to be discovered (Stappers et al., 1996), orbiting a semi-degenerate companion every ~2.4 h (Shaifullah et al., 2016). Previous timing analyses of PSR J2051−0827 by Lazaridis et al. (2011) and Doroshenko et al. (2001) have shown that there is a sign-changing variation of the change in the orbital period due to variations in the GQ of the companion. The GQ variations are believed to arise due to coupled variations in the structure and magnetic field of the companion, which may be linked to differential rotation of the various layers of the star, as expected from the Applegate and Shaham (1994) model or due to dynamical variations within the star as in the Lanza and Rodonò (1999) model.

We have derived an update on the timing of PSR J2051−0827, presented in Shaifullah et al. (2016, and Table 4.4 of Chapter 4), along with timing models for the BTX (an interpolative model based on the work by Blandford and Teukolsky, 1976) and ELL1 (Lange et al., 2001) models of TEMPO2 (Edwards et al., 2006). An improved estimate of the mean transverse velocity is also made, giving a value of 30(9) km s^{-1}. By fitting measurements made over three year epochs, we are able to obtain for the first time, significant measurements of the proper motion in both R.A. and DEC, $\mu_\alpha = 5.63(10)$ mas yr^{-1} and $\mu_\delta = 2.34(28)$ mas yr^{-1} respectively. A significant decrease in the DM of $\sim 2.5 \times 10^{-3}$ cm^{-3} pc is detected over the MJD range 54 600–56 800 and corrections are incorporated in the ToA file.

A more robust analysis than that of earlier authors is performed by reducing covariances between the model parameters and it is shown that the resulting measurements are more precise and consistent with those from earlier analyses. The variations of the orbital period are detected over more than a full 'period' of ~8 yr, supporting earlier analyses that suggested these variations arise from cyclic variations in the companion, instead of a tertiary star or planet. In addition, short-term, seemingly random fluctuations in the P_b variations are detected.

The continued timing of PSR J2051−0827 shows that the variation of the projected semi-major axis appears to have decreased and does not show the extreme behaviour observed at an earlier epoch. If the behaviour of PSR J2051−0827 is indeed becoming more stable and given that another BWP system, PSR J0621+1002 is currently included in the EPTA source list, these results lend hope that more BWP systems may be included in PTAs in the near future.

In Chapter 5 I have presented the results of a dedicated campaign to measure the spectral indices of MSPs at three frequency bands was carried out at the Arecibo Observatory. For this work we observed 19 MSPs for which no published spectral index was available. This resulted in 12 new MSP spectral indices measured from flux density measurements made at three frequency bands over six epochs[2].

[2] Note: Not all sources were observed or detected at every epoch and band

We complemented this sample with 19 new MSP spectral indices derived using flux density values from literature, while for an additional 43 the spectral indices were rederived using all available flux density measurements available. In this way, we are able to increase the sample of MSPs with known spectral indices to 74, which represents almost half of the known MSPs in the Galactic disk (ATNF pulsar catalogue Manchester et al., 2005)[3].

[3] www.atnf.csiro.au/research/pulsar/psrcat, ver. 1.54

We find that the spectral index distributions of classical pulsars and MSPs agree quite well although their median values lie at 1.57(2) and 1.74(4) respectively. Sub-population analyses show some interesting features, including suggestions for a uniform distribution of spectral indices for isolated and binary MSPs and MSPs discovered at high and low frequencies. The RBP population, although poorly sampled, appears to consist of steep spectrum sources exclusively. In contrast the BWP population shows a broader distribution but with lower precision on the determined indices. *Fermi* sources show a strong bias towards steeper spectra which may partly explain the lack of MSP discoveries at high-frequency *Fermi* follow-up searches.

We also investigated the correlations of pulsar parameters with the spectral index and find that the spin-period P is weakly correlated with the spectral index while the spin-down energy $\dot{E}$ is weakly anti-correlated. The spin-down rate $\dot{P}$ is only very weakly correlated to the spectral index. We find the spin-period, spin-down rate and the spectral index are related as:

$$\alpha = -15(4) + 0.34(9) log\left(\frac{P}{\dot{P}}\right). \quad (6.3)$$

6.2 Scope for future work

As already mentioned, while we present some significant new observations and deductions, there is much scope to improve upon the work presented here. We list some of the most interesting avenues for future work below.

6.2.1 Artefacts in polyphase filterbanks

The studies carried out here are only preliminary exercises investigating the possibility of implementing artefact-free digital data recording systems using full reconstruction PFBs implemented on resource-limited devices like FPGAs. The first step beyond this work would be to utilise the least mean square algorithm of de Haan (2001) to design artefact-free analysis filterbanks, which are the kinds of filterbanks implemented currently of the FPGA boards used for recording pulsar observations. This should be followed up with a test for the effects of artefacts in real data, using some of the brightest MSPs as candidates. Implementing a full reconstruction PFB in itself may be avoided altogether by using time-domain multiplexing and full bandwidth acquisition, although those systems come with their own set of challenges.

6.2.2 *The long term timing of PSR J2051–0827*

In the study presented earlier and those which have preceded it (Lazaridis et al., 2011; Doroshenko et al., 2001) the nature of the companion has remained unknown. Even though the companion has been detected in optical observations, the modelling has only led to measurements of the sharp temperature differences between the illuminated and the non-illuminated faces that are expected for a companion that is impinged upon by the pulsar wind (Stappers et al., 2001).

Surprisingly, there is a convincing suggestion for the presence of short-term variations, on timescales of the order of ∼150 days. Follow-up observations that are currently underway are expected to allow us to model these variations far more precisely and search for the origin of these sudden variations.

The Applegate and Shaham (1994) model was originally applied to PSR J1959+2048, the first BWP to be discovered. Recently, it has been shown that the red-back pulsar PSR J2339–0533 also shows secular variations of a similar nature. In contrast, PSR J0621+1002 (Desvignes et al., 2016) shows a high degree of stability and is included in the EPTA source list. Hence, it is important to explore the long-term timing behaviour of known BWP and RBP systems to test how common the tendency for sign-changing variations in the orbital period are. While the Applegate and Shaham (1994) model is very successful at modelling the secular variations of the orbital period, it is also not understood what physical process is the originator of this model. Lanza and Rodonò (1999) propose a similar model where the dynamic variations in the magnetic field of the companion drive the secular variations observed. However, this model has not been applied to MSPs in tight binaries, primarily due to the lack of knowledge about the companion's composition and hence, their magnetic fields. The full orbit observations currently being carried out might allow us to address this question if we are able to obtain rotation measures from bright observations when the pulsar is in eclipse region.

6.2.3 *Spectral indices of MSPs*

The data presented here represent only half of our observing campaign, incorporating only the faint MSPs that were observed at the Arecibo Observatory. The second half of this campaign includes 31 MSPs which will improve the sample-size of the MSP population as a whole, and of the various sub-populations. Including the re-derived spectral indices of MSPs using flux densities from literature, we will be able to sample well over two-thirds of the 195 Galactic MSPs known currently.

Studies of the luminosity distribution[4] of the 74 MSPs in our combined sample, like those of Kramer et al. (1998), would allow us to im-

[4] Which will only be possible if precise distance measurements to these MSPs can be made.

prove the input to population synthesis codes, thereby improving the predicted results. Although, our sample of 12 sources observed at Arecibo offers only a limited dataset studies of the polarimetric profiles and their pulse widths using the brightest observations could place improved limits on the emission geometries of these sources. For sources which are bright at 300 MHz, rotation measures can also be estimated, allowing us to probe the Galactic magnetic field along the line of sight and for some of the brightest sources, analysis of their dynamic spectra could reveal interesting features in the ISM along the those line of sight as well.

A rather pertinent test would be to extend the observation campaign and reduce the impact of scintillation and finally, increasing the integration times for some of the bands where some of the sources are not detected. All of this would allow us to measure the spectral indices of those MSPs with greater precision and reliability.

6.2.4 *The question of improving PTA sensitivity*

One of the most important goals of the pulsar timing community is to detect the nHz GW that are expected to be emitted by SMBH mergers, specifically in the early stages of their in-spiral.

The recent detections of GWs from the merger of two $\sim 30\,M_\odot$ and two $\sim 20\,M_\odot$ black-holes by the Laser Interferometer Gravitational-Wave Observatory (LIGO; the events were named GW150914 and GW151226, respectively; Abbott et al., 2016b,c) has led to a positive revision of the expectations of black hole binary merger rates (Sesana, 2016) in the Universe. This has led to further expectations that pulsar timing arrays (PTAs; see Section 2.5.4) which use lines of sight to MSPs as their primary detectors, are on the cusp of making significant detections of GWs from such mergers. However, recent analyses of the four PTA data sets have only lead to upper limits on the strain amplitude being placed (see Verbiest et al., 2016; Babak et al., 2016; Lentati et al., 2015; Reardon et al., 2016; Arzoumanian et al., 2016; Taylor et al., 2015, etc). While some of the pulsars in the PTAs have been timed continuously for more than 20 years, we are yet to achieve the required sensitivity to detect nHz GWs. In this context, the easiest way to improve our sensitivity is to add new sources as suggested by scaling laws from Siemens et al. (2013), especially in the weak-to-intermediate S/N range of GWs where the signal amplitude is comparable to or slightly greater than the noise.

A simpler quantification, from Jenet et al. (2005) shows that the detection significance of a PTA depends on the number of pulsars in the PTA as:

$$S = \sqrt{\frac{M(M-1)/2}{1 + \left[\chi(1+\bar{\zeta}) + 2(\sigma_n/\sigma_g)^2 + (\sigma_n/\sigma_g)^4\right]/N\sigma_\zeta^2}} \tag{6.4}$$

if the characteristic strain amplitude h_c induced by the stochastic GW background is related to the GW frequency f by a power law, i.e. $h_c(f) = Af^\alpha$ (Jenet et al., 2005).

Here M is the number of pulsars, N is the number of pulsar timing residuals and σ_g is the rms fluctuation in the measured ToA due to the passage of a GW through the space-time containing the Earth and the pulsars, i.e.,

$$\sigma_g = \frac{A^2}{4\pi^2(1-\alpha)}\left(f_1^{2(1-\alpha)} - f_2^{2(1-\alpha)}\right). \quad (6.5)$$

σ_n is the rms fluctuation due to all non-GW sources while $\chi = (1/\sigma_g^4 N)\sum_{i=0}^{N-1}\sum_{j=0}^{N-1} c_{g_{ij}}^2$ estimates the Gaussian behaviour of the timing residuals. σ_ζ is the variance of the ζ given by

$$\zeta = \frac{3}{2}\frac{1-cos\theta}{2} log\left(\frac{1-cos\theta}{2}\right) - \frac{1-cos\theta}{8} + \frac{1}{2} + \frac{1}{2}\delta\left(\frac{1-cos\theta}{2}\right) \quad (6.6)$$

where θ is the angular separation between the pairs of pulsars in the PTA.

Increasing the number of pulsars in the PTA therefore, is the easiest way to improve our sensitivity to GWs. Even though the expected population of MSPs in the Galaxy is expected to be about ~40 000–120 000 (see e.g., Levin et al., 2013), radio frequency surveys have met with limited success at discovering new MSPs. Surveys of ever increasing sensitivity project discovery rates of tens of sources and almost always fall short of such projections by large margins. For example, for the High Time Resolution Universe (HTRU) survey for pulsars and fast transients (Barr et al., 2013) which has been one of the more successful campaigns in recent times, the predicted number of 78 new MSPs (Levin et al., 2013) and actual discovery number of 27 in Ng et al. (2014). While a great part of the challenge of detecting new MSPs lies in the computational aspects of pulsar searching, at least some part of the problem also lies in how poorly we understand the overall population of MSPs.

6.3 *Concluding remarks*

Recycled pulsars offer a veritable zoo of possible science, from studies of their masses to the processes through which they evolve. If pulsars are 'the gift that keeps on giving' then MSPs are the prime of those gifts. While a significant thrust of this thesis lies in the associated goals of improving survey detection rates or the sensitivity of PTAs these are fascinating objects themselves, and the challenges of observing and classifying these objects in itself is hardly a lesser goal. The advent of the SKA will open the path to discovering great numbers of MSPs and thanks to the expected sensitivity of the SKA we will finally be able to address long-standing questions about the formation and evolution of these objects. Real-time, parallelised data processing on FPGA based backends will be of utmost importance in increasing the efficiency and throughput in the era of data-intensive astronomy, the designs of which must meet the challenges of the extremely low system temperatures expected. Observations with the SKA will probably reveal MSPs with a plethora of companions which will require newer and more powerful timing models, as we race to measure masses, constrain equations of state and hunt for gravitational waves. For surveys to be efficient with a telescope as sensitive as the SKA the bottleneck will be primarily our understanding of the spectral characteristics of the MSPs we search for.

A

Spectra of standard calibrators

The spectra for the absolute flux-density calibrations used in Chapter 5 are shown below. The indicated flux values are the recalibrated flux-density values obtained from the NED server. The brown lines are 3rd-order polynomial fits following Baars et al. (1977). The polynomial coefficients shown in Table A.1 were supplied to the *fluxcal* program of the **PSRCHIVE** suite (van Straten et al., 2012; Hotan et al., 2004).

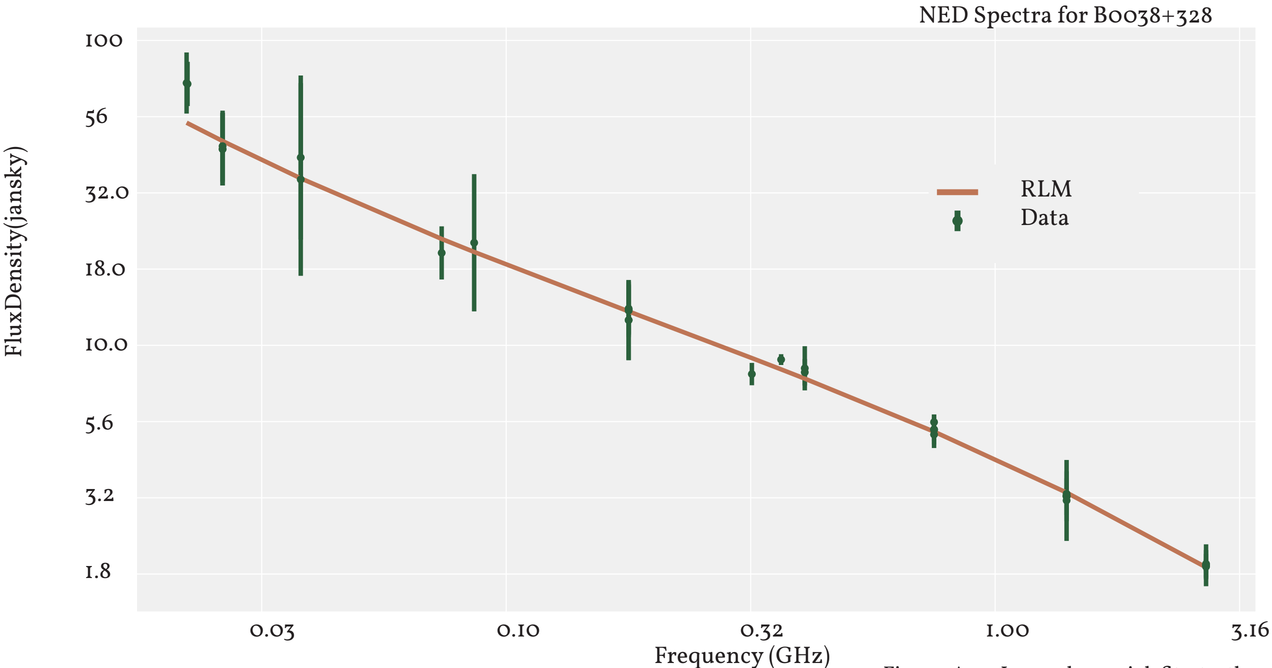

Figure A.1: Log-polynomial fits to the spectrum of B0038+328

Table A.1: Polynomial coefficients for *fluxcal*

Source Name	R.A.	DEC	a1	a2	a3	a4
B0038+328	$0^h40^m55\overset{s}{.}01$	$33°10'7\overset{''}{.}3$	0.63	−0.73	−0.18	−0.09
B0428+205	$4^h31^m3\overset{s}{.}78$	$20°37'34\overset{''}{.}2$	0.59	0.07	−0.61	−0.30
B1040+123	$10^h42^m44\overset{s}{.}54$	$12°3'31\overset{''}{.}8$	0.64	−0.71	−0.03	−0.04
B1442+101	$14^h45^m16\overset{s}{.}48$	$9°58'36\overset{''}{.}0$	0.40	0.04	−0.78	−0.40
B2209+080	$22^h12^m1\overset{s}{.}58$	$8°19'16\overset{''}{.}51$	0.41	−0.83	−0.76	−0.51

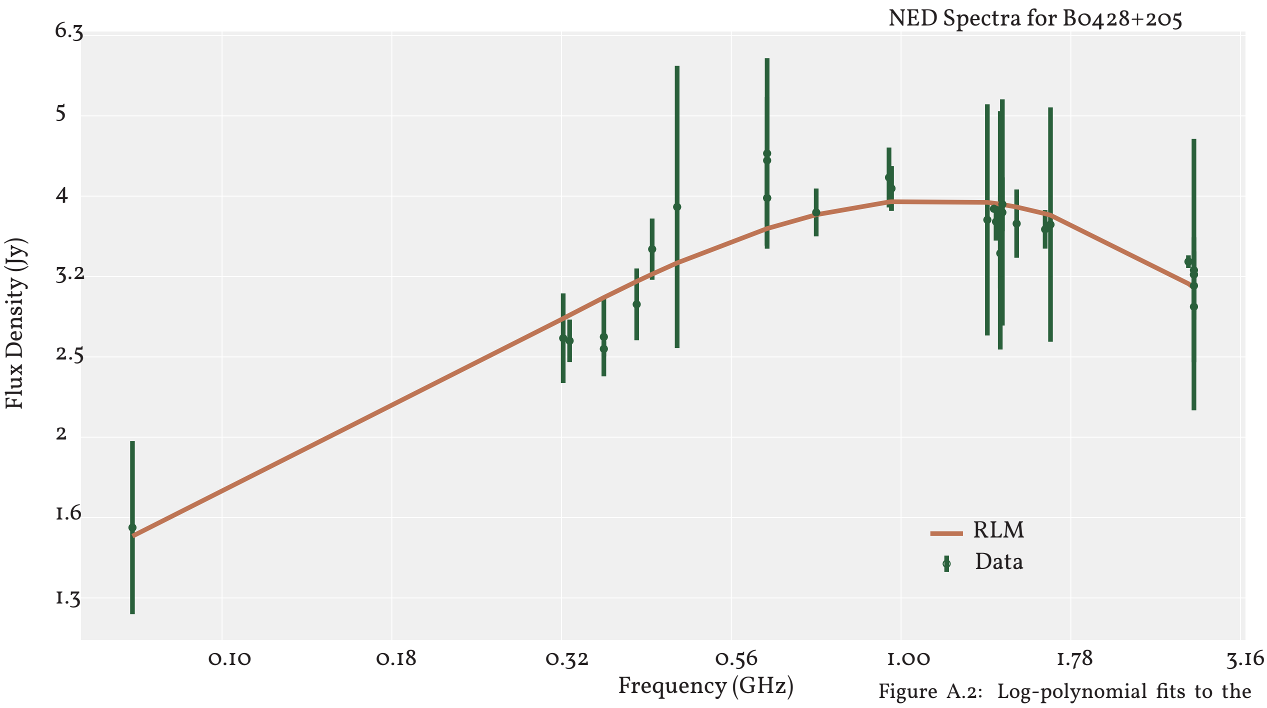

Figure A.2: Log-polynomial fits to the spectrum of B0428+205

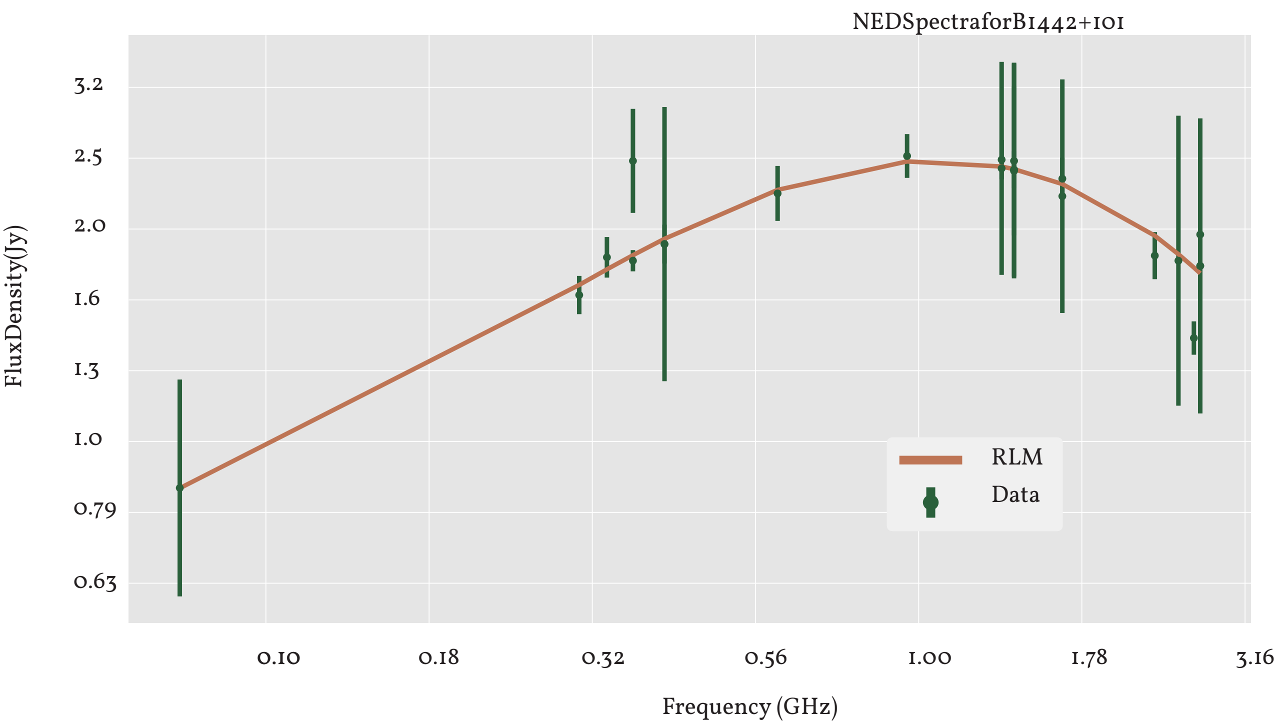

Figure A.3: Log-polynomial fits to the spectrum of B1442+101

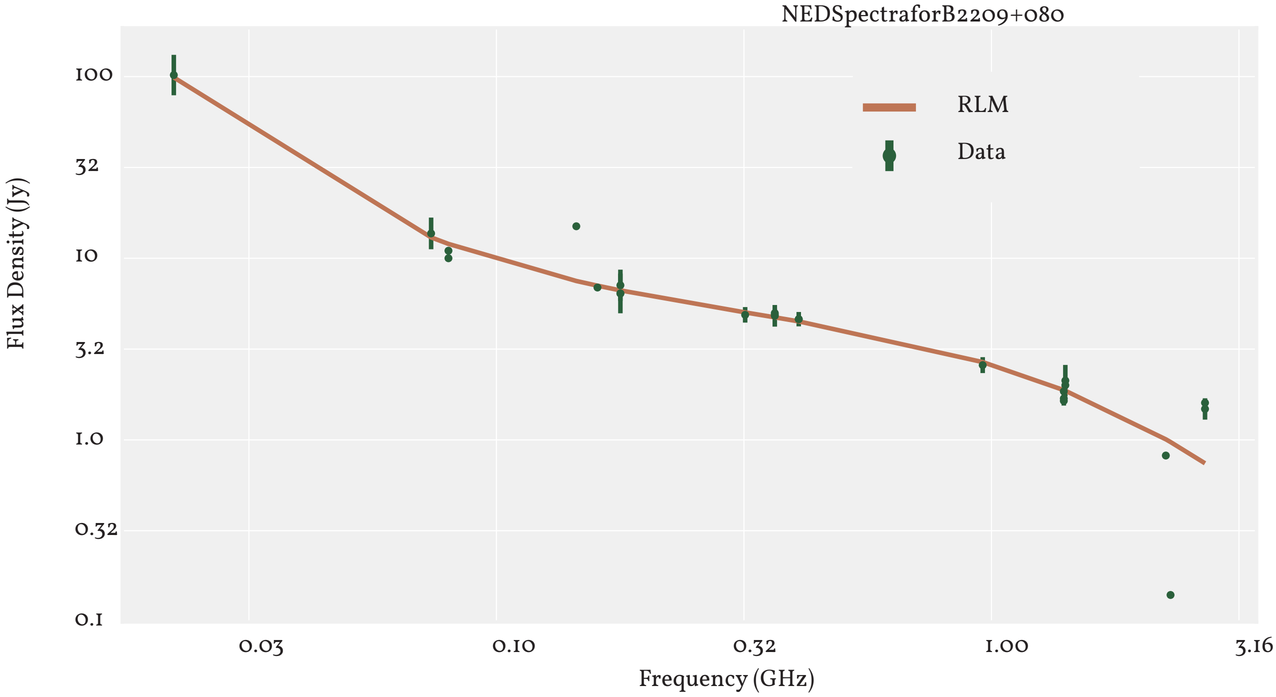

Figure A.4: Log-polynomial fits to the spectrum of B2209+080

B

Glossary of statistical tools used

This thesis utilises a number of statistical tools for which we provide a brief description below. These are by no means exhaustive and we encourage the reader to treat the original sources mentioned here and elsewhere as more definitive.

B.1 Akaike Information Criterion

Akaike's 'An Information Criterion' (Akaike, 1974, 1976) can be used for model selection from a set of fitted models if for each model, a log-likelihood value can be obtained. This is achieved by minimising the Aikake Information Criterion (AIC), i.e.,

$$argmin\left[-2(log - likelihood) + k(npar),\right] \tag{B.1}$$

where *npar* represents the number of parameters in the fitted model, and $k = 2$ for the AIC. See Burnham and Anderson (2003) for proper use and alternatives to the AIC.

B.2 Kernel Density Estimator

The kernel density estimate (see e.g., Rosenblatt, 1956) is given by:

$$\hat{f}(x; H) = n^{-1} \sum_{i=1}^{n} K_H(x - X_i) \tag{B.2}$$

where X_i are the set of n measurements, H is called the bandwidth matrix $K_H(x - X_i)$ is the normal probability distribution function with mean x. Unlike the histogram which uses discrete bins to collect the measurements, the KDE allows for a smoother estimation of the distribution and is useful in identifying small deviations. We use the KDE with a Gaussian kernel and a variable bandwidth determined from the precision of each measurement in our plots in Chapter 5.

B.3 Robust Statistics

When Gauss presented his work on the best methods to combine a number of 'independent' observations, he did so with the careful caveat that the measurements were derived from *"observations of equal accuracy"* (Gauss, 1821). However, in real data the accuracy (or precision) of each individual observation is different. This was demonstrated as early as 1886 by Newcomb (Newcomb, 1886). The problem of estimation in astronomy (or in general) then reduces to the problem of investigating *n* independent observations with a 'random error' and finding the true value from amongst these randomly distributed observations. In this problem, the distribution of the errors is an unknown. Gauss inverted this problem to investigate, given that we assume that the arithmetic mean of the independent observations is the most-likely estimate of the true value, what is the best distribution that describes the errors. In the limit of a large number of observations, this distribution turns out to be the 'normal' or Gaussian distribution. This then also allows us to apply the method of least-squares to finding the true measurement.

The central limit theorem, however only suggests approximate normality under very well-specified conditions (see empirical investigations by Bessel, 1818; Newcomb, 1886; Jeffreys, 1998, and others). This result, while implicitly ignored in many calculations, leads to the method of least-squares being very strongly influenced by any deviations from the well-specified conditions needed to ensure normality in a given problem.

Perhaps the most obvious of these errors are outliers; whose influence on the least-squares method is well-known. The most common method of handling outliers is rejection, even though in reality outliers ought to be investigated separately. Typically, this involves establishing a rule for this rejection. If this rule is basic on the statistical description of the outlier or its relation to the remaining distribution, one can test how well the rejection rule performs with increasing numbers of outliers. The point at which the rule is no longer able to reject outliers without affecting the estimation of the true value (i.e., the deviation of the measured value from the true value) is quantified by the break-down point of the rule. The break-down point is a global measure of the robustness of the rejection rules used.

If the quantity being estimated is the mean, the break-down point is zero. If this is the median, the break-down point has a value of half. That is to say, slightly less than half the measurements can be outliers without affecting the estimated mean significantly. Tukey in his seminal work in 1960 (Tukey, 1960) first presented the extreme sensitivity of some conventional statistical procedures like linear regression to seemingly minor deviations from the initial assumptions. The insight that statistical methods optimised for the conventional Gaussian model are unstable under small perturbations led to the development of the theory of stability of statistical methods, otherwise known as robust statistics. There are two main branches of this theory; the minimax approach of Huber (1981) which uses a quantitative approach and the influence functions approach of Hampel (1968), which attempts a qualitative discussion of the stability of the statistical method being applied.

In our results presented in Chapter 5; we are encountered with the problem of having a small set of data containing several scattered measurements. The influence of even a few outliers on these measurements is severe and can skew the measured spectral index very far from it's true value. Making general assumptions on the underlying shape of the distribution implies implicit assumptions which need not coincide with reality, e.g., the problem of spectral breaks in certain sources.

To avoid all of these issues, we use the robust fitting routines from the *Python* 'Statsmodels' (Wes McKinney et al., 2011) module to determine our spectral indices. The robust linear models fitting leads to more reliable estimates of the spectral index along with reasonable estimates of the uncertainties. For our fits, we use the Huber type-2 covariance matrices.

More information can be found at the Statsmodels documentation. An excellent resource for robust statistics is Huber (1964).

B.3.1 A note on upper ranges

In our spectral index measurements in Chapter 5, we have ignored upper limits for non-detections. The absence of these non-detections or the upper limits (technically, ranges) from these measurements are left out mainly because in all except a few cases, the non-detection is typically the result of RFI or technical failures. Given that the uncertainties of these upper range determination are poor at best, we felt that including those values in our fits would unfairly bias the linear regression fits while they would simply be rejected by the robust linear model fit.

Bibliography

Abbott, B. P. et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Classical Quant. Grav.*, 33, 134001 (July 2016a).

Abbott, B. P. et al. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. *Phys. Rev. D*, 93, 122003 (June 2016b).

Abbott, B. P. et al. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Phys. Rev. Lett.*, 116, 241103 (June 2016c).

Abdo, A. A. et al. Pulsed Gamma Rays from the Millisecond Pulsar J0030+0451 with the Fermi Large Area Telescope. *ApJ*, 699, 1171 (July 2009).

Abdo, A. A. et al. The Second Fermi Large Area Telescope Catalog of Gamma-Ray Pulsars. *ApJS*, 208, 17 (October 2013).

Akaike, H. A New Look at the Statistical Model Identification. *Ieee T. Automat. Contr.*, 19, 716 (1974).

Akaike, H. An information criterion (aic). *Math Sci*, 14, 5 (1976).

Alpar, M. A. et al. Giant glitches and the pinned vorticity in the Vela and other pulsars. *ApJ*, 249, L29 (1981).

Antoniadis, J. et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340, 448 (April 2013).

Antoniadis, J. et al. An Eccentric Binary Millisecond Pulsar with a Helium White Dwarf Companion in the Galactic field. *ApJ*, 830, 36 (October 2016).

Applegate, J. H. et al. Orbital period variability in the eclipsing pulsar binary PSR B1957+20: Evidence for a tidally powered star. *ApJ*, 436, 312 (1994).

Archibald, A. M. et al. A Radio Pulsar/X-ray Binary Link. *Science*, 324, 1411 (June 2009).

Arzoumanian, Z. et al. The NANOGrav Nine-year Data Set: Limits on the Isotropic Stochastic Gravitational Wave Background. *ArXiv e-prints* (August 2015).

Arzoumanian, Z. et al. The NANOGrav Nine-year Data Set: Limits on the Isotropic Stochastic Gravitational Wave Background. *ApJ*, 821, 13 (April 2016).

Baade, W. et al. Cosmic rays from super-novae. *Proc. Nat. Acad. Sci.*, 20, 259 (1934a).

Baade, W. et al. On super-novae. *Proc. Nat. Acad. Sci.*, 20, 254 (1934b).

Baars, J. W. M. et al. The absolute spectrum of Cas A; an accurate flux density scale and a set of secondary calibrators. *A&A*, 61, 99 (1977).

Babak, S. et al. European Pulsar Timing Array limits on continuous gravitational waves from individual supermassive slack hole binaries. *MNRAS*, 455, 1665 (January 2016).

Backer, D. C. et al. A millisecond pulsar. *Nature*, 300, 615 (1982).

Barr, E. D. et al. The Northern High Time Resolution Universe pulsar survey - I. Setup and initial discoveries. *MNRAS*, 435, 2234 (November 2013).

Bassa, C. G. et al. LEAP: the Large European Array for Pulsars. *MNRAS*, 456, 2196 (February 2016).

Bates, S. D. et al. The pulsar spectral index distribution. *Mon. Not. R. Astron. Soc.*, 431, 1352 (2013).

Becker, W., editor. *Neutron Stars and Pulsars*, volume 357 of *Astrophysics and Space Science Library* (2009).

Bell Burnell, S. J. Petit Four. In M. D. Papagiannis, editor, *Eighth Texas Symposium on Relativistic Astrophysics*, volume 302 of *Annals of the New York Academy of Sciences*, page 685 (1977).

Bellanger, M. et al. Digital filtering by polyphase network:application to sample-rate alteration and filter banks. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 24, 109 (Apr 1976).

Beskin, V. S. et al. The electrodynamics of a pulsar magnetosphere. *Zh. Eksp. Teor. Fiz+.+*, 85, 401 (August 1983).

Bessel, F. W. Fundamenta astronomiae (1818).

Bhattacharya, D. Evolution of neutron star magnetic fields. *J. Astrophys. Astr.*, 23, 67 (2002).

Bhattacharya, D. et al. Formation and evolution of binary and millisecond radio pulsars. *Phys. Rep.*, 203, 1 (1991).

Blandford, R. et al. Arrival-time analysis for a pulsar in a binary system. *ApJ*, 205, 580 (1976).

Boashash, B., editor. *Time Frequency Signal Analysis and Processing: A Comprehensive Reference*. Elsevier, Amsterdam (2003).

Bochenek, C. et al. The Feasibility of Using Black Widow Pulsars in Pulsar Timing Arrays for Gravitational Wave Detection. *ApJ*, 813, L4 (November 2015).

Bonnarel, F. et al. The ALADIN interactive sky atlas. A reference tool for identification of astronomical sources. *A&A Supl.S.*, 143, 33 (April 2000).

Boyles, J. et al. The Green Bank Telescope 350 MHz Drift-scan survey. I. Survey Observations and the Discovery of 13 Pulsars. *ApJ*, 763, 80 (February 2013).

Burgay, M. et al. An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system. *Nature*, 426, 531 (December 2003).

Burgay, M. et al. The parkes high-latitude pulsar survey. *MNRAS*, 368, 283 (May 2006).

Burnham, K. P. et al. *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science & Business Media (2003).

Camilo, F. et al. Discovery of two fast-rotating pulsars. *ApJ*, 412, L37 (1993).

Camilo, F. et al. A search for millisecond pulsars at galactic latitudes $-50^{\circ} ltblt - 20^{\circ}$. *ApJ*, 461, 812 (1996).

Chadwick, J. The Existence of a Neutron. *Proceedings of the Royal Society of London Series A*, 136, 692 (June 1932).

Chamel, N. et al. Physics of Neutron Star Crusts. *Living Rev. Relativ.*, 11, 10 (December 2008).

Chamel, N. et al. Neutron drip transition in accreting and nonaccreting neutron star crusts. *Phys R*, 91, 055803 (May 2015).

Champion, D. J. et al. Arecibo timing and single-pulse observations of 17 pulsars. *MNRAS*, 363, 929 (2005).

Chandrasekhar, S. The highly collapsed configurations of a stellar mass. *MNRAS*, 91, 456 (March 1931a).

Chandrasekhar, S. The Maximum Mass of Ideal White Dwarfs. *ApJ*, 74, 81 (July 1931b).

Chen, H.-L. et al. Formation of Black Widows and Redbacks—Two Distinct Populations of Eclipsing Binary Millisecond Pulsars. *ApJ*, 775, 27 (September 2013).

Cheng, A. et al. Current flow in pulsar magnetospheres. *ApJ*, 203, 209 (January 1976).

Cheng, A. F. et al. Bunching mechanism for coherent curvature radiation in pulsar magnetospheres. *ApJ*, 212, 800 (March 1977a).

Cheng, A. F. et al. Pair-production discharges above pulsar polar caps. *ApJ*, 214, 598 (June 1977b).

Cheng, A. F. et al. A theory of subpulse polarization patterns from radio pulsars. *ApJ*, 229, 348 (1979).

Cheng, A. F. et al. Particle acceleration and radio emission above pulsar polar caps. *ApJ*, 235, 576 (1980).

Cognard, I. et al. A microglitch in the millisecond pulsar psr b1821-24 in m28. *ApJ*, 612, L125 (September 2004).

Condon, J. J. et al. The NRAO VLA Sky Survey. *AJ*, 115, 1693 (1998).

Contopoulos, I. et al. The Axisymmetric Pulsar Magnetosphere. *ApJ*, 511, 351 (1999).

Conway, R. G. et al. The radio frequency spectra of discrete radio sources. *MNRAS*, 125, 261 (1963).

Cooley, J. W. et al. An algorithm for the machine calculation of complex fourier series. *Math. Comput.*, 19, 297 (1965).

Cordes, J. M. Observational limits on the limits of pulsar emission regions. *ApJ*, 222, 1006 (1978).

Cordes, J. M. et al. NE2001. II. Using Radio Propagation Data to Construct a Model for the Galactic Distribution of Free Electrons. *ArXiv Astrophysics e-prints* (Jan 2003). Http://xxx.lanl.gov/abs/astro-ph/0301598.

Crochiere, R. E. et al. On the properties of frequency transformations for variable cutoff linear phase digital filters. *Ieee T. Circuits Syst.*, 23, 684 (November 1976).

Damour, T. et al. General relativistic celestial mechanics of binary systems. II. The post-Newtonian timing formula. *Ann. Inst. H. Poincaré (Physique Théorique)*, 44, 263 (1986).

Damour, T. et al. Strong–field tests of relativistic gravity and binary pulsars. *Phys. Rev. D*, 45, 1840 (1992).

de Haan, J. M. *Filter bank design for subband adaptive filtering*. Ph.D. thesis, Department of Telecommunications and Signal Processing, University of Karlskrona and Ronneby, Sweden (2001).

Demorest, P. B. *Measuring the gravitational wave background using precision pulsar timing*. Ph.D. thesis, University of California, Berkeley (2007).

Demorest, P. B. et al. A two-solar-mass neutron star measured using Shapiro delay. *Nature*, 467, 1081 (October 2010).

Demorest, P. B. et al. Limits on the Stochastic Gravitational Wave Background from the North American Nanohertz Observatory for Gravitational Waves. *ApJ*, 762, 94 (January 2013).

Deneva, J. S. et al. Two Millisecond Pulsars Discovered by the PALFA Survey and a Shapiro Delay Measurement. *ApJ*, 757, 89 (September 2012).

Deneva, J. S. et al. Goals, Strategies and First Discoveries of AO327, the Arecibo All-sky 327 MHz Drift Pulsar Survey. *ApJ*, 775, 51 (September 2013).

Desvignes, G. et al. A new pulsar instrumentation at the Allen Telescope Array and the Nancay Radio Telescope. In M. Burgay, N. D'Amico, P. Esposito, A. Pellizzoni, and A. Possenti, editors, *American Institute of Physics Conference Series*, volume 1357 of *AIP Conf. Series*, pages 349–350 (August 2011).

Desvignes, G. et al. High-precision timing of 42 millisecond pulsars with the European Pulsar Timing Array. *MNRAS*, 458, 3341 (May 2016).

Deutsch, A. J. The electromagnetic field of an idealized star in rigid rotation in vacuo. *Annales d'Astrophysique*, 18, 1 (1955).

Dib, R. et al. 16 yr of RXTE Monitoring of Five Anomalous X-Ray Pulsars. *ApJ*, 784, 37 (March 2014).

Doroshenko, O. et al. Orbital variability of the PSR J2051-0827 binary system. *A&A*, 379, 579 (November 2001).

Duncan, R. C. et al. Formation of very strongly magnetized neutron stars: Implications for gamma-ray bursts. *ApJ*, 392, L9 (1992).

DuPlain, R. et al. Launching GUPPI: the Green Bank Ultimate Pulsar Processing Instrument. In *Advanced Software and Control for Astronomy II*, volume 7019 of *procspie*, page 70191D (August 2008).

Dyks, J. et al. The origin of the frequency-dependent behaviour of pulsar radio profiles. *MNRAS*, 446, 2505 (January 2015).

Edwards, R. T. et al. **TEMPO2**, a new pulsar timing package - II. the timing model and precision estimates. *MNRAS*, 372, 1549 (November 2006).

Espinoza, C. M. et al. A study of 315 glitches in the rotation of 102 pulsars. *MNRAS*, 414, 1679 (June 2011a).

Espinoza, C. M. et al. The Braking Index of PSR J1734-3333 and the Magnetar Population. *ApJ*, 741, L13 (November 2011b).

Faucher-Giguère, C. A. et al. Birth and evolution of isolated radio pulsars. *ApJ*, 643, 332 (May 2006).

Faulkner, A. J. et al. The Parkes Multibeam Pulsar Survey - V. Finding binary and millisecond pulsars. *MNRAS*, 355, 147 (November 2004).

Finkbeiner, D. P. A Full-Sky Hα Template for Microwave Foreground Prediction. *ApJS*, 146, 407 (June 2003).

Fliege, N. *Multiraten-Signalverarbeitung*. B. G. Teubner, Stuttgart (1993).

Folkner, W. M. et al. The Planetary and Lunar Ephemeris DE 421. *Interplanetary Network Progress Report*, 178, C1 (August 2009).

Foster, R. S. et al. A spectral study of four millisecond pulsars. *ApJ*, 378, 687 (1991).

Frail, D. A. et al. Known Pulsars Identified in the GMRT 150 MHz All-sky Survey. *ApJ*, 829, 119 (October 2016).

Freire, P. C. et al. Further results from the timing of the millisecond pulsars in 47 Tucanae. *MNRAS*, 340, 1359 (2003).

Freire, P. C. C. et al. Eight new millisecond pulsars in ngc 6440 and ngc 6441. *ApJ*, 675, 670 (March 2008).

Fruchter, A. S. et al. A millisecond pulsar in an eclipsing binary. *Nature*, 333, 237 (1988).

Fujiwara, S. et al. *A hundred verses from old Japan, being a translation of the Hyaku-nin-isshiu*. Tuttle Publishing (1985).

Gauss, C. Theoria combinationis observationum erroribus minimis obnoxiae (pars prior), commentationes societatis regiae scientiarium gottingensis recentiores, werke (1821).

Ghosh, P. et al. Disk accretion by magnetic neutron stars. *ApJ*, 223, L83 (1978).

Ghosh, P. et al. Diagnostics of disk-magnetosphere interaction in neutron star binaries. In E. P. J. van den Heuvel and S. A. Rappaport, editors, *X-ray Binaries and Recycled Pulsars*, pages 487–510. Kluwer, Dordrecht (1992).

Gold, T. Rotating neutron stars as the origin of the pulsating radio sources. *Nature*, 218, 731 (1968).

Goldreich, P. et al. Pulsar electrodynamics. *ApJ*, 157, 869 (1969).

Gonzalez, M. E. et al. High-precision Timing of Five Millisecond Pulsars: Space Velocities, Binary Evolution, and Equivalence Principles. *ApJ*, 743, 102 (December 2011).

Gottlieb, D. et al. On the gibbs phenomenon and its resolution. *Siam Rev.*, 39, 644 (1997).

Hampel, F. *Contributions to the theory of robust estimation*. Ph.D. thesis, University of California, Berkeley (1968).

Hankins, T. H. et al. Pulsar signal processing. In *Methods in Computational Physics Volume 14 — Radio Astronomy*, pages 55–129. Academic Press, New York (1975).

Henrichs, H. F. et al. Is the millisecond pulsar formed from coalescence of a close neutron–star binary? *Nature*, 303, 213 (1983).

Hessels, J. W. T. et al. A Radio Pulsar Spinning at 716 Hz. *Science*, 311, 1901 (March 2006).

Hessels, J. W. T. et al. A 350-MHz GBT Survey of 50 Faint Fermi γ-ray Sources for Radio Millisecond Pulsars. In M. Burgay, N. D'Amico, P. Esposito, A. Pellizzoni, and A. Possenti, editors, *American Institute of Physics Conference Series*, volume 1357 of *American Institute of Physics Conference Series*, pages 40–43 (August 2011).

Hewish, A. et al. Observation of a rapidly pulsating radio source. *Nature*, 217, 709 (February 1968).

Hobbs, G. B. et al. **TEMPO2**, a new pulsar-timing package - I. an overview. *MNRAS*, 369, 655 (June 2006).

Holloway, N. J. et al. Properties of gaps in pulsar magnetospheres. *MNRAS*, 194, 95 (January 1981).

Hotan, A. W. et al. PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis. *PASA*, 21, 302 (2004).

Hotan, A. W. et al. PSR J0737-3039A: baseband timing and polarimetry. *MNRAS*, 362, 1267 (October 2005).

Huber, P. J. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35, 73 (1964).

Huber, P. J. *Robust statistics*. John Wiley & Sons, New York; Chichester (1981).

Hulse, R. A. et al. Discovery of a pulsar in a binary system. *ApJ*, 195, L51 (1975).

Illarionov, A. F. et al. Why the number of Galactic X–ray stars is so small? *A&A*, 39, 185 (1975).

Jansky, K. G. Radio Waves from Outside the Solar System. *Nature*, 132, 66 (July 1933).

Jeffreys, H. *The theory of probability*. OUP Oxford (1998).

Jenet, F. A. et al. Detecting the Stochastic Gravitational Wave Background Using Pulsar Timing. *ApJ*, 625, L123 (June 2005).

Joshi, K. J. et al. Distant Companions and Planets around Millisecond Pulsars. *ApJ*, 479, 948 (1997).

Jury, E. I. *Theory and Application of the z-Transform Method*. Wiley (1964).

Kaiser, J. et al. On the use of the i0-sinh window for spectrum analysis. 28, 105 (Feb 1980).

Kaplan, D. L. et al. A Precise Proper Motion for the Crab Pulsar, and the Difficulty of Testing Spin-Kick Alignment for Young Neutron Stars. *ApJ*, 677, 1201 (April 2008).

Karastergiou, A. et al. An empirical model for the beams of radio pulsars. *MNRAS*, 380, 1678 (October 2007).

Karuppusamy, R. et al. PuMa-II: A Wide Band Pulsar Machine for the Westerbork Synthesis Radio Telescope. *PASP*, 120, 191 (February 2008).

Kondratiev, V. I. et al. A LOFAR census of millisecond pulsars. *A&A*, 585, A128 (January 2016).

Kramer, M. et al. The characteristics of millisecond pulsar emission: I. Spectra, pulse shapes and the beaming fraction. *ApJ*, 501, 270 (1998).

Kramer, M. et al. Tests of general relativity from timing the double pulsar. *Science*, 314, 97 (2006).

Kramer, M. et al. The European Pulsar Timing Array and the Large European Array for Pulsars. *Class. Quant Grav.*, 30, 224009 (November 2013).

Kuniyoshi, M. et al. Low-frequency spectral turn-overs in millisecond pulsars studied from imaging observations. *MNRAS*, 453, 828 (October 2015).

Kunkel, W. et al. Supernova 1987A in the Large Magellanic Cloud. *iaucirc*, 4316 (feb 1987).

Löhmer, O. et al. Shapiro delay in the PSR J1640+2224 binary system. *ApJ*, 621, 388 (March 2005).

Landau, L. Origin of Stellar Energy. *Nature*, 141, 333 (February 1938).

Landau, L. D. On the theory of stars. *Phys. Z. Sowjetunion*, 1, 152 (1932).

Lane, W. M. et al. VLSS redux: Software improvements applied to the Very Large Array Low-Frequency Sky Survey. *Radio Sci.*, 47, RS0K04 (January 2012).

Lange, C. et al. Precision timing measurements of PSR J1012+5307. *MNRAS*, 326, 274 (2001).

Lanza, A. F. et al. Orbital period modulation and quadrupole moment changes in magnetically active close binaries. *A&A*, 349, 887 (September 1999).

Lanza, A. F. et al. The residual orbital eccentricity in close binaries. An application to millisecond binary pulsars. *A&A*, 376, 165 (September 2001).

Lattimer, J. H. et al. The physics of neutron stars. *Science*, 304, 536 (2004).

Lazaridis, K. et al. Evidence for gravitational quadrupole moment variations in the companion of PSR J2051-0827. *MNRAS*, 414, 3134 (July 2011).

Lazarus, P. et al. Arecibo Pulsar Survey Using ALFA. IV. Mock Spectrometer Data Analysis, Survey Sensitivity, and the Discovery of 40 Pulsars. *ApJ*, 812, 81 (October 2015).

Lazarus, P. et al. Prospects for high-precision pulsar timing with the new Effelsberg PSRIX backend. *MNRAS*, 458, 868 (May 2016).

Lentati, L. et al. European Pulsar Timing Array limits on an isotropic stochastic gravitational-wave background. *MNRAS*, 453, 2576 (November 2015).

Levin, L. et al. The high time resolution universe pulsar survey –viii. the galactic millisecond pulsar population. *MNRAS*, 434, 1387 (2013).

Liu, K. et al. Measuring pulse times of arrival from broad-band pulsar observations. *MNRAS*, 443, 3752 (October 2014).

Livingstone, M. A. et al. New phase-coherent measurements of pulsar braking indices. *Ap&SS*, 308, 317 (April 2007).

Livingstone, M. A. et al. Post-outburst Observations of the Magnetically Active Pulsar J1846-0258. A New Braking Index, Increased Timing Noise, and Radiative Recovery. *ApJ*, 730, 66 (April 2011).

Lommen, A. N. et al. New Pulsars from an Arecibo Drift Scan Search. *ApJ*, 545, 1007 (December 2000).

Lorimer, D. R. et al. Multifrequency flux density measurements of 280 pulsars. *MNRAS*, 273, 411 (1995).

Lorimer, D. R. et al. Discovery of 10 pulsars in an Arecibo drift-scan survey. *MNRAS*, 359, 1524 (June 2005).

Lorimer, D. R. et al. *Handbook of Pulsar Astronomy*. Cambridge University Press (2005).

Lorimer, D. R. et al. PSR J1453+1902 and the radio luminosities of solitary versus binary millisecond pulsars. *MNRAS*, 379, 282 (July 2007).

Lovelace, R. V. E. et al. Intensity Variations of the Pulsar CP 1919. *Nature*, 220, 875 (November 1968).

Lynch, R. S. et al. The Timing of Nine Globular Cluster Pulsars. *ApJ*, 745, 109 (February 2012).

Lynch, R. S. et al. The Green Bank Telescope 350 MHz Drift-scan Survey II: Data Analysis and the Timing of 10 New Pulsars, Including a Relativistic Binary. *ApJ*, 763, 81 (February 2013).

Lyne, A. et al. *Pulsar astronomy*. 48. Cambridge University Press (2012).

Lyne, A. G. Glitches and timing noise. In Z. Arzoumanian, F. van der Hooft, and E. P. J. van den Heuvel, editors, *Pulsar Timing, General Relativity, and the Internal Structure of Neutron Stars*, pages 141–150. North Holland, Amsterdam (1999).

Lyne, A. G. et al. 23 years of Crab pulsar rotational history. *MNRAS*, 265, 1003 (1993).

Lyne, A. G. et al. Very low braking index for the Vela pulsar. *Nature*, 381, 497 (1996).

Mallat, S. *A Wavelet Tour of Signal Processing: The Sparse Way*. Elsevier Science (2008).

Malofeev, V. M. et al. Mean spectra for 39 pulsars, and the interpretation of their characteristic feaures. *Sov. Astron.*, 24, 54 (1980).

Manchester, R. N. et al. The Australia Telescope National Facility Pulsar Catalogue. *AJ*, 129, 1993 (April 2005).

Maron, O. et al. Observations of millisecond pulsars at 8.35 GHz. *A&A*, 413, L19 (January 2004).

Martinez, J. G. et al. Pulsar J0453+1559: A Double Neutron Star System with a Large Mass Asymmetry. *ApJ*, 812, 143 (October 2015).

McKee, J. W. et al. A glitch in the millisecond pulsar J0613-0200. *MNRAS*, 461, 2809 (September 2016).

McLaughlin, M. A. et al. Transient radio bursts from rotating neutron stars. *Nature*, 439, 817 (2006).

Meltzer, D. W. et al. Normal Modes of Radial Pulsation of Stars at the End Point of Thermonuclear Evolution. *ApJ*, 145, 514 (August 1966).

Mestel, L. et al. The axisymmetric pulsar magnetosphere. I. *MNRAS*, 188, 385 (August 1979).

Mestel, L. et al. The axisymmetric pulsar magnetosphere. II. *MNRAS*, 188, 799 (September 1979).

Michel, F. C. Rotating Magnetosphere: a Simple Relativistic Model. *ApJ*, 180, 207 (February 1973a).

Michel, F. C. Rotating Magnetospheres: an Exact 3-D Solution. *ApJ*, 180, L133 (March 1973b).

Michel, F. C. A pulsar emission model: observational tests. *ApJ*, 322, 822 (1987).

Miller, M. C. et al. Implications of the PSR 1257+12 Planetary System for Isolated Millisecond Pulsars. *ApJ*, 550, 863 (April 2001).

Mintzer, F. Filters for distortion-free two-band multirate filter banks. 33, 626 (June 1985).

Morris, D. et al. Pulsar detection at 87GHz. *A&A*, 322, L17 (June 1997).

Morrison, I. et al. Ska csp ska1-mid array non-imaging processor pulsar timing sub-element signal processing matlab model. Technical report, SKA (August 2015).

Mostow, J. S. *Pictures of the Heart: The Hyakunin Isshu in Word and Image*. University of Hawaii Press (1996).

Navarro, J. et al. The arecibo 430 mhz intermediate galactic latitude survey: Discovery of nine radio pulsars. *ApJ*, 594, 943 (September 2003).

Newcomb, S. A generalized theory of the combination of observations so as to obtain the best result. *Am. J. Math.*, 8, 343 (1886).

Newcomb, S. Prof. Newcomb's new planetary tables. *The Observatory*, 18, 86 (February 1895).

Ng, C. et al. The High Time Resolution Universe pulsar survey - X. Discovery of four millisecond pulsars and updated timing solutions of a further 12. *MNRAS*, 439, 1865 (April 2014).

Nice, D. J. et al. Two newly discovered millisecond pulsars. *ApJ*, 402, L49 (1993).

Nice, D. J. et al. Binary eclipsing millisecond pulsars: A decade of timing. In M. Kramer, N. Wex, and R. Wielebinski, editors, *ASP Conf. Ser. 202: IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond*, pages 67–72 (2000).

Nyquist, H. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47, 617 (1928a).

Nyquist, H. Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32, 110 (1928b).

Ogura, J. et al. Some Properties of an Axisymmetric Pulsar Magnetosphere Constructed by Numerical Calculation. *Prog. Theor. Phys.*, 109, 619 (April 2003).

Okamoto, I. Force-free pulsar magnetosphere. II - The steady, axisymmetric theory for a normal plasma. *MNRAS*, 170, 81 (January 1975).

Olausen, S. A. et al. The McGill Magnetar Catalog. *ApJS*, 212, 6 (May 2014).

Oppenheim, A. et al. *Signals and Systems*. Always Learning. Pearson (2013).

Oppenheim, A. V. et al. Discrete representation of signals. *P. Ieee*, 60, 681 (June 1972).

Oppenheim, A. V. et al. Effects of finite register length in digital filtering and the fast fourier transform. *P. Ieee*, 60, 957 (1972).

Oppenheimer, J. R. et al. On massive neutron cores. *Phys. Rev.*, 55, 374 (1939).

Ostriker, J. P. et al. On the nature of pulsars. I. Theory. *ApJ*, 157, 1395 (1969).

Özel, F. et al. On the Mass Distribution and Birth Masses of Neutron Stars. *ApJ*, 757, 55 (September 2012).

Pacini, F. Energy emission from a neutron star. *Nature*, 216, 567 (1967).

Parks, T. W. et al. Chebyshev approximation for nonrecursive digital filters with linear phase. *IEEE Trans. Circuit Th.*, 19, 189 (1972).

Peacock, J. A. Two-dimensional goodness-of-fit testing in astronomy. *MNRAS*, 202, 615 (February 1983).

Pennucci, T. T. et al. Elementary Wideband Timing of Radio Pulsars. *ApJ*, 790, 93 (August 2014).

Pétri, J. Theory of pulsar magnetosphere and wind. *J. Plasma Phys.*, 82, 635820502 (October 2016).

Phillips, J. A. Radio emission altitudes in the pulsar magnetosphere. *ApJ*, 385, 282 (1992).

Press, W. H. et al. *Numerical Recipes: The Art of Scientific Computing, 2^{nd} edition*. Cambridge University Press, Cambridge (1992).

Pringle, J. E. et al. Accretion disc models for compact x-ray sources. *A&A*, 21, 1 (1972).

Rabiner, L. R. et al. *Digital processing of speech signals*. Prentice Hall (1978).

Rankin, J. M. Toward an Empirical Theory of Pulsar Emission. XI. Understanding the Orientations of Pulsar

Radiation and Supernova “Kicks”. *ApJ*, 804, 112 (May 2015).

Ransom, S. M. et al. Three Millisecond Pulsars in Fermi LAT Unassociated Bright Sources. *ApJ*, 727, L16+ (January 2011).

Ransom, S. M. et al. A millisecond pulsar in a stellar triple system. *Nature*, 505, 520 (January 2014).

Ray, P. S. et al. A survey for millisecond pulsars. *ApJ*, 470, 1103 (1996).

Reardon, D. J. et al. Timing analysis for 20 millisecond pulsars in the Parkes Pulsar Timing Array. *MNRAS*, 455, 1751 (January 2016).

Reid, M. J. et al. Trigonometric Parallaxes of High Mass Star Forming Regions: The Structure and Kinematics of the Milky Way. *ApJ*, 783, 130 (March 2014).

Remez, E. Y. Sur le calcul effectif des polynômes d’approximation de tschebyscheff. *C. P. Paris*, pages 337–340 (1934).

Rickett, B. J. Interstellar scattering and scintillation of radio waves. *Ann. Rev. Astr. Ap.*, 15, 479 (1977).

Rickett, B. J. et al. Slow scintillation in the interstellar medium. *A&A*, 134, 390 (1984).

Roberts, M. S. E. Surrounded by spiders! New black widows and redbacks in the Galactic field. In J. van Leeuwen, editor, *IAU Symposium*, volume 291 of *IAU Symposium*, pages 127–132 (March 2013).

Robinson, B. J. et al. Measurements of the pulsed radio source cp 1919 between 85 and 2700 mhz. *Nature*, 218, 1143 (June 1968).

Rosenblatt, M. Remarks on some nonparametric estimates of a density function. *Ann. Math. Statist.*, 27, 832 (09 1956).

Roy, J. et al. Observations of four glitches in the young pulsar J1833-1034 and study of its glitch activity. *MNRAS*, 424, 2213 (August 2012).

Ruderman, M. A. et al. Theory of pulsars: Polar gaps, sparks, and coherent microwave radiation. *ApJ*, 196, 51 (1975).

Ruderman, M. A. et al. Fate of very low-mass secondaries in accreting binaries and the 1.5-ms pulsar. *Nature*, 304, 425 (1983).

Ryle, M. et al. Optical Identification of the First Neutron Star? *Nature*, 217, 907 (March 1968).

Scharlemann, E. T. et al. Aligned Rotating Magnetospheres. General Analysis. *ApJ*, 182, 951 (June 1973).

Scheuer, P. A. G. Amplitude variations of pulsed radio sources. *Nature*, 218, 920 (1968).

Schmalz, R. et al. On the self-consistent description of axisymmetric pulsar magnetospheres. *MNRAS*, 189, 709 (December 1979).

Schmalz, R. et al. On the self-consistent description of axisymmetric pulsar magnetospheres. II - A method of solution. *MNRAS*, 192, 409 (August 1980).

Sesana, A. Prospects for Multiband Gravitational-Wave Astronomy after GW150914. *Phys. Rev. Lett.*, 116, 231102 (June 2016).

Shaifullah, G. et al. 21 year timing of the black-widow pulsar J2051-0827. *MNRAS*, 462, 1029 (October 2016).

Shannon, C. E. Communication in the presence of noise. *Proceedings of the IRE*, 37, 10 (1949).

Shapiro, I. I. Fourth test of general relativity. *Phys. Rev. Lett.*, 13, 789 (1964).

Shapiro, S. L. et al. *Black Holes, White Dwarfs and Neutron Stars. The Physics of Compact Objects*. Wiley–Interscience, New York (1983).

Shapley, H. et al. The Scale of the Universe. *Bulletin of the National Research Council, Vol. 2, Part 3, No. 11, p. 171-217*, 2, 171 (May 1921).

Shemar, S. L. et al. Observations of pulsar glitches. *MNRAS*, 282, 677 (1996).

Siemens, X. et al. The stochastic background: scaling laws and time to detection for pulsar timing arrays. *Class.*

Quant Grav., 30, 224015 (November 2013).

Smith, M. et al. Exact reconstruction techniques for tree-structured subband coders. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 34, 434 (Jun 1986).

Spitkovsky, A. Electrodynamics of pulsar magnetospheres. In F. Camilo and B. M. Gaensler, editors, *Young Neutron Stars and Their Environments*, volume 218 of *IAU Symposium*, pages 357–364 (2004).

Spitkovsky, A. Pulsar Magnetosphere: The Incredible Machine. In C. Bassa, Z. Wang, A. Cumming, and V. M. Kaspi, editors, *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, volume 983 of *American Institute of Physics Conference Series*, pages 20–28 (February 2008).

Spitkovsky, A. et al. Simulations of Pulsar Wind Formation. In P. O. Slane and B. M. Gaensler, editors, *Neutron Stars in Supernova Remnants*, page 81. Astronomical Society of the Pacific, San Francisco (2002).

Splaver, E. M. *Long-term Timing of Millisecond Pulsars*. Ph.D. thesis, Princeton University, Princeton, N. J., U.S.A. (2004).

Stairs, I. H. et al. A baseband recorder for radio pulsar observations. *MNRAS*, 314, 459 (2000). (astro-ph/9912272).

Stappers, B. W. et al. Detection of an Irradiated Pulsar Companion. *ApJ*, 473, L119 (December 1996).

Stappers, B. W. et al. The Orbital Evolution and Proper Motion of PSR J2051-0827. *ApJ*, 499, L183 (1998).

Stappers, B. W. et al. Intrinsic and reprocessed optical emission from the companion to PSR J2051-0827. *ApJ*, 548, L183 (2001).

Stappers, B. W. et al. A State Change in the Missing Link Binary Pulsar System PSR J1023+0038. *ApJ*, 790, 39 (July 2014).

Stoica, P. et al. *Spectral Analysis of Signals*. Prentice Hall, Upper Saddle River, New Jersey (2005).

Stovall, K. et al. The Green Bank Northern Celestial Cap Pulsar Survey. I. Survey Description, Data Analysis, and Initial Results. *ApJ*, 791, 67 (August 2014).

Strang, G. et al. *Wavelets and filter banks*. SIAM (1996).

Strohmer, T. et al. Implementations of shannon's sampling theorem, a time-frequency approach. *Sampling Theory in Signal and Image Processing*, 4, 1 (2005).

Strohmer, T. et al. Fast reconstruction methods for bandlimited functions from periodic nonuniform sampling. *Siam J. Numer. Anal.*, 44, 1073 (2006).

Sturrock, P. A. A model of pulsars. *ApJ*, 164, 529 (1971).

Swiggum, J. K. et al. PSR J1930-1852: a Pulsar in the Widest Known Orbit around Another Neutron Star. *ApJ*, 805, 156 (June 2015).

Takata, J. et al. X-Ray and Gamma-Ray Emissions from Rotation Powered Millisecond Pulsars. *ApJ*, 745, 100 (January 2012).

Tauris, T. M. Five and a Half Roads to Form a Millisecond Pulsar. In L. Schmidtobreick, M. R. Schreiber, and C. Tappert, editors, *Evolution of Compact Binaries*, volume 447 of *Astronomical Society of the Pacific Conference Series*, page 285 (September 2011).

Tauris, T. M. et al. Runaway velocities of stellar components originating from disrupted binaries via asymmetric supernova explosions. *A&A*, 330, 1047 (1998).

Tauris, T. M. et al. *Compact stellar X-ray sources.*, chapter Formation and evolution of compact stellar X-ray sources, pages 623–665. Number 39 in Cambridge Astrophysics Series. Cambridge University Press, Cambridge, UK (April 2006). ISBN 978-0-521-82659-4, ISBN 0-521-82659-4.

Tauris, T. M. et al. Formation of millisecond pulsars with CO white dwarf companions - II. Accretion, spin-up, true ages and comparison to MSPs with He white dwarf companions. *MNRAS*, 425, 1601 (September 2012).

Taylor, J. H. Pulsar Timing and Relativistic Gravity. *Philosophical Transactions of the Royal Society of London Series*

A, 341, 117 (October 1992).

Taylor, J. H. et al. A new test of general relativity: Gravitational radiation and the binary pulsar PSR 1913+16. *ApJ*, 253, 908 (1982).

Taylor, J. H. et al. Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16. *ApJ*, 345, 434 (1989).

Taylor, J. H. et al. Pulsar distances and the Galactic distribution of free electrons. *ApJ*, 411, 674 (1993).

Taylor, S. R. et al. Limits on Anisotropy in the Nanohertz Stochastic Gravitational Wave Background. *Phys. Rev. Lett.*, 115, 041101 (July 2015).

Tielens, A. G. G. M. Origin and Evolution of the Interstellar Medium. *Astrophysics Space*, 10, 271 (2009).

Timmes, F. X. et al. The neutron star and black hole initial mass function. *ApJ*, 457, 834 (1996).

Tolman, R. C. *Relativity, Thermodynamics, and Cosmology*. Oxford: Clarendon Press (1934).

Toscano, M. et al. Spectra of southern pulsars. *The Astrophysical Journal*, 506, 863 (1998).

Tukey, J. W. A survey of sampling from contaminated distributions. In I. O. G. G. H. G. M. B. Mann, editor, *Contributions to Probability and Statistics*, pages 448–485. Stanford, Calif. (1960).

Vaidyanathan, P. Theory of optimal orthonormal subband coders. *Signal Processing, IEEE Transactions on*, 46, 1528 (Jun 1998).

Vaidyanathan, P. P. *Multirate Systems and Filter Banks*. Prentice-Hall, Upper Saddle River, New Jersey (1993).

van Straten, W. Radio astronomical polarimetry and point-source calibration. *ApJ*, 152, 129 (2004).

van Straten, W. Radio Astronomical Polarimetry and High-Precision Pulsar Timing. *ApJ*, 642, 1004 (May 2006).

van Straten, W. Radio astronomical polarimetry and high-precision pulsar timing. *ApJ*, 642, 1004 (2006).

van Straten, W. et al. Pulsar Data Analysis with PSRCHIVE. *Astronomical Research and Technology*, 9, 237 (July 2012).

Vary, P. On the design of digital filter banks based on a modified principle of polyphase. *Archiv für Elektronik und Übertragungstechnik (AEÜ), Electronics and Communications*, 33, 293 (1979).

Verbiest, J. P. W. et al. Precision timing of PSR J0437–4715: An accurate pulsar distance, a high pulsar mass, and a limit on the variation of newton's gravitational constant. *ApJ*, 679, 675 (May 2008).

Verbiest, J. P. W. et al. Timing stability of millisecond pulsars and prospects for gravitational-wave detection. *MNRAS*, 400, 951 (December 2009).

Verbiest, J. P. W. et al. Lutz-Kelker bias in pulsar parallax measurements. *MNRAS*, 405, 564 (June 2010).

Verbiest, J. P. W. et al. The International Pulsar Timing Array: First data release. *MNRAS*, 458, 1267 (May 2016).

von Fraunhofer, J. Neue modifikation des lichtes. *Denkschriften der Königlichen Akademie der Wissenschaften zu München für das Jahre 1821 und 1822*, 8, 1 (1822).

Weinstein, C. J. *Quantization Effects in Digital Filters*. Ph.D. thesis, Department of Electrical Engineering, Massachusetts Institute of Technology (July 1969a).

Weinstein, C. J. Quantization effects in digital filters. Technical Report TR-468, Lincoln Laboratory Report (November 1969b).

Welch, B. L. The generalization of 'student's' problem when several different population varlances are involved. *Biometrika*, 34, 28 (1947).

Weltevrede, P. et al. The glitch-induced identity changes of PSR J1119-6127. *MNRAS*, 411, 1917 (March 2011).

Wes McKinney et al. Time series analysis in python with statsmodels (2011).

Wilson, T. L. et al. *Tools of Radio Astronomy*. Springer-Verlag, Berlin, Heidelberg (2013). ISBN: 978-3-642-39949-7.

Wolszczan, A. (1995). Unpublished.

Wu, J. et al. Discovery of γ-ray pulsation and x-ray emission from the black widow pulsar PSR J2051–0827. *ApJ*, 748, 141 (2012).

Yao, J. M. et al. A New Electron-density Model for Estimation of Pulsar and FRB Distances. *ApJ*, 835, 29 (January 2017).

Acronyms

ADC	analog-to-digital convertor
AFB	Analogue Filter-bank system
AIC	Aikake Information Criterion
ATNF	Australia Telescope National Facility
BAT	barycentred arrival time
BCRS	barycentric celestial reference system
BH	black hole
BIPM	Bureau International des Poids et Mesures
BIPM15	Latest realisation of the TAI to TT conversion, computed by the BIPM
BON	Berkeley-Orleans-Nançay
BT	Blandford-Teukolsky pulsar timing model, see Blandford and Teukolsky (1976)
BTX	extended Blandford-Teukolsky pulsar timing model
BW	bandwidth
BWP	black-widow pulsar
CE	common envelope
CO/ONeMg-WD	Carbon-Oxygen/Oxgen-Neon-Magnesium white dwarf
CO-WD	Carbon-Oxygen white dwarf
CT	continuous-time
DEC	declination
DFB	Digital Filter-bank system
DM	dispersion measure
DNS	double neutron star
DSP	digital signal processing
EBPP	Effelsberg-Berkeley Pulsar Processor
ELL1	Laplace-Lagrange parameter based pulsar timing model for nearly cicruclar binaries, see Lange et al. (2001)
EM	electromagnetic
EPTA	European pulsar timing array
FB	filterbank
FDM	Fourier Domain with Markov-chain Monte-Carlo
Fermi-LAT	*Fermi* Large Area Telescope
FFT	fast Fourier Transform
FIR	finite impulse response filter
FPGA	field programmable gate array
GCRS	Geocentric Celestial Reference System
GOF	goodness-of-fit
GQ	gravitational quadrupole
GUPPI	Green Bank Ultimate Pulsar Processing Instrument

GW	gravitational wave
He-WD	Helium white dwarf
HMXB	high mass X-ray binary
HTRU	High Time Resolution Universe
IAU	International Astronomical Union
ICS	inverse Compton scattering
IIR	infinite impulse response filter
IISM	ionised ISM
IMXB	intermediate mass X-ray binary
IPS	inter-planetary scintillation
IR	infra-red radiation
ISM	interstellar medium
KDE	kernel density estimator
LIGO	Laser Interferometer Gravitational-Wave Observatory
LMC	Large Magellanic Cloud
LMS	least-mean-squares
LMXB	low mass X-ray binary
LNA	low noise amplifier
LOFAR	LOw Frequency ARray
LTV	linear time variant
MHD	magneto-hydrodynamic
MJD	modified Julian date
MSP	millisecond pulsar
NED	NASA/IPAC Extragalactic Database
NPR	near-perfect reconstruction
NRAO	National Radio Astronomy Observatory
NS	neutron star
NUPPI	Nancay Ultimate Pulsar Processing Instrument
NVSS	NRAO VLA Sky Survey
OFD	orbital-frequency derivative
P.A.	position angle
PFB	polyphase filterbank
PR	perfect reconstruction
PSR	pulsar
PSRIX	PSRIX
PTA	pulsar timing array
PUPPI	Puerto Rico Ultimate Pulsar Processing Instrument
QMF	quadrature modulated filterbank
R.A.	right ascension
RBP	red-back pulsar
RFI	radio frequency interference
RFM	radius-to-frequency mapping
RLO	Roche-lobe overflow
ROACH	Reconfigurable Open Architecture Computing Hardware
RP	recycled pulsar
S/N	signal-to-noise ratio
SDSS	Sloan Digital Sky Survey
SKA	the Square Kilometre Array
SMBH	super massive black hole

SN	supernova
SNR	supernova remnant
SOC	spin-orbit coupling
SRT	Sardinia Radio Telescope
SSB	Solar system barycentre
T2EFAC	Tempo2 error scaling factor
TAI	Temps Atomique International
TCG	Geocentric Coordinate Time
TGSS	the GMRT Southern Sky
ToA	time of arrival
TT	Terrestrial Time
UL	ultra-light
UTC	Universal Coordinated Time
VLA	Jansky Very Large Array
WD	white dwarf
WSRT	Westerbork Synthesis Radio Telescope
XMM	X-ray Multi-Mirror Mission
ZAMS	zero-age main-sequence

List of Symbols used

Δ_{A_B}	aberration due to proper motion
Δ_{IS}	delay from SSB to pulsar
Δ_{R_B}	Römer delay due to companion of pulsar
$\Delta_{R_\odot}$	Römer delay due to the Sun
Δ_{S_B}	Shapiro delay due to companion of pulsar
$\Delta_{S_\odot}$	shapiro delay due to the sun
$\Delta_\odot$	delay due to sun
Δ_{SW}	delay due to the solar wind
Δ_{VP}	vacuum propagation delay
Δ_{Atm}	delay due to propgation through the Earths atmosphere
Δ_{BB}	delay from Earth to SSB
$\Delta_{E_{SSB,BB}}$	delay due to relativistic motion of SSB and pulsar (or binary barycentre)
Δ_{E_B}	Einstein delay due to pulsars companion
$\Delta_{E_\odot}$	Einstein delay due to the Sun
Δ_{FDD}	delay due frequency dependent effects
Δ_{ISD}	delay due to dispersion by the interstellar medium
e	elementary charge
$\gamma_{Einstein}$	Einstein parameter
f_p	plasma frequency
G_r	receiver gain
G_T	telescope gain
H	on-source flux-density
L	off-source flux-density
m_e	mass of the electron
μ_δ	proper motion in declination
μ_α	proper motion in right ascension
n_e	Electron number density
Ω	Longitude of ascending node
e	Orbital eccentricity
P_b	Orbital period
$\dot{P}$	First derivative of spin period
δ_θ	orbital shape correction due to Shapiro delay
r	range parameter of Shapiro delay
s	shape parameter of Shapiro delay
S	spectral density function
P	Spin period
T_0	Epoch of periastron
T_{asc}	Epoch of ascending node passage

t_a^{earth}	time of arrival at earth
t_e^{psr}	time of emission at pulsar
T_{sky}	Temperature of 1 $k\Omega$ resistor when supplied power equivalent to flux-density of sky excluding the source
T_{source}	temperature of 1 $k\Omega$ resistor when supplied power equivalent to flux-density of source
T_{diode}	temperature of 1 $k\Omega$ resistor when supplied power equivalent to flux-density of the noise diode
T_{std}	temperature of 1 $k\Omega$ resistor when supplied power equivalent to flux density of standard candle
T_{sys}	Temperature of observing system ($T_{rec} + T_{sky}$)
v_g	group velocity
x	Projected semi-major axis

Acknowledgements

I am deeply grateful to my supervisor, Dr. Joris P. W. Verbiest. Joris, thank you very much! From the day I arrived in Germany, I have constantly received help and kindness from you. As I struggled through days of darkness, you have made it known that there was a comforting voice I could call upon and, provided me with practical advice without prejudice. More than anything else, you have been exemplary to me, in your hard work and efficiency, your persistence for careful, vetted science and in matters of personal life. I am also deeply grateful to Prof. Michael Kramer, my supervisor at the Max Planck Institute for Radioastronomy, Bonn. Thank you Michael, for your kind and generous support and being approachable for the most trivial of questions.

I also owe thanks to Dr. Ramesh Karuppsamy (MPIfR-Bonn) and Prof. Dominik Schwarz (Universität Bielefeld), whose support with the administrative issues meant I had few worries with the University or at the MPIfR. Thank you, to Stefan Osłowski for your frank advice and willingness to teach me some of the mysteries of PSRCHIVE and to Jörn Künsemöller, for tolerating my often inane questions on computing, German banking and the dicey topic of how to return used bottles to the supermarket. To my fellow students at Bonn, Alessandro Ridolfi, Eleni Graikou, Nataliya Porayko, Nicolas Caballero and Patrick Lazarus; thank you for your encouragement and support, even though our interactions remained limited. To my colleagues in the EPTA, thank you reminding me cutting edge science isn't always cut-throat science, for treating my naive questions with patience and for your constant encouragement.

I am extremely grateful to the past and present operators and other staff at the respective radio telescopes, without whose cheerful and unflagging support much of the observations presented here would have been impossible. Thank you for your cheerful help with my quaint requests; Alexander Hochguertel, Axel Jessner, Klaus Schlich, Marcus Keseberg, Majid Esmailpour Motlagh, Norbert Tacken, Peter Vogt, Ralf Kisky, Thomas Wedel and Wilfried Schmitz at Effelsberg and Arun Venkataraman, Daniel Padilla, Edwin Muniz, Efrain Matos and Jonathan Irizarry at Arecibo.

Two people without whose constant support my life and work would have come to a standstill during the three years I have lived in Germany were Frau Susi Reder (Universität Bielefeld) and Frau Kira Kühn (MPIfR-Bonn). Kira, for having patience with my incorrectly filled forms and getting my travel claims processed at the MPIfR, thank you! To Susi, who has been a paragon of efficiency and whose careful organisation helped me in multiple ways, from finding a house in Bielefeld, to getting thesis preparation advice and to whom I have turned to for help in the smallest of matters, not enough words of gratitude can express my sentiments. Thank you so much and more.

Aakanksha, thank you for being the best friend you are, for reading my mind across oceans and knowing when I was floundering and needed that gentle chiding or reminder. Abir, your cheerful laughter and your impish pestering has helped me tide over days of not knowing where terra-firma lay. To both of you, thank you for everything, not least of which, for being the voices at the end of the telephone as my world has gone to pieces. Asmita, though you are younger, you stand taller and I continue to hope you grow far beyond my imagination. Caterina, you have been a friend who has dragged me out from behind the shuttered windows of my house to art exhibitions, jaunts across Bielefeld and hours of conversation. Thank you, for you have been kind and giving even though you barely knew me. Niruj, thank you for helping me abandon my archaic notions and my narrow view of the world. Thanks also for those three invites to dinner. I can only hope to be as good a friend to all of you as you are to me, though in fairness I have always leaned upon you far more than you have looked to me for support.

To Mayumi, thank you for your patience with my queries on Japanese literature and culture. To Shintaro and Kaori, thank you for being the little humans who reminded me why I mustn't lose hope because of the grown-ups.

To my siblings; Mahmuda, Mahjabeen, Tamanna and Ghazi, and my father, thank you for supporting me as I struggled through my ways and for always being there for me. These few words do justice to neither my gratitude nor all that I owe you.

To my mother and to Ishita, I offer no words of gratitude for your love is too great and I am too little.

Translations of quoted poetry

The translations offered below are by no means definitive, neither do I claim mastery over the languages they were originally written in. I would hope the keen reader forgives my occasional mistake.

Japanese Poems

Chapter 1
Even if I leave
and go to Inaba Mountain,
whose peak lies covered in
pines, should I hear you pine for me,
I will return swiftly to you.

Chapter 2
Spring has passed and summer is here,
And I can just cry out to
Amanokagu's peak,
Where heavenly angels
Spread their white robes to dry.

Chapter 3
Must you so avoid others' eyes
that not even at night,
along the road of dreams,
will you draw near like the waves
to the shore of Suminoe Bay?

Chapter 4
When I look at the moon
I am overcome by the sadness
of thousands of things
even though it is not Fall
for me alone.

Chapter 5
O waves crashing upon the rocks,
fanned by the violent wind
it is I alone
who breaks, whenever
I think of her

Chapter 6
For the one who doesn't come
I wait at the Bay of Matsuo
in the patient evening
as they boil the seaweed for salt,
I, too, burn with longing

Better translations of the Japanese poems can also be found in Fujiwara et al. (1985) and Mostow (1996)

Other languages

Chapter 4
Aye, ay, ay,ay!
Take this waltz that dies in my arms.
—Federico García Lorca; Little Viennese Waltz

Chapter 5
The greatest artist has no concept
that marble doesn't conscribe
in its expanse, so that only the hand
guided by intellect could reveal
—Michelangelo; Sonnet, circa. 1538

Chapter 6
Friend, what is thought?
Friend, what is it to care?
Is it only that which is anxious,
is it only in tears?
—Rabindranath Tagore; 1881, Gitobitan (anthology)