# High-Precision Observations of Relativistic Binary and Millisecond Pulsars

**Aidan Wade Hotan**

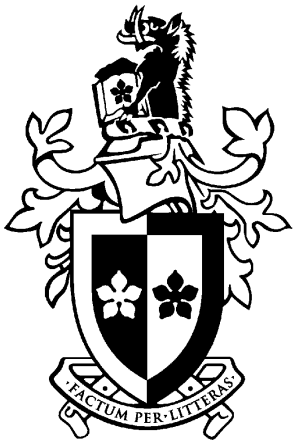


*A dissertation*
*Presented in fulfillment of the requirements*
*for the degree of Doctor of Philosophy*
*at Swinburne University of Technology*

*July 2006*

"All parts should go together without forcing. You must remember that the parts you are reassembling were disassembled by you. Therefore if you can't get them to go together again, there must be a reason. By all means, do not use a hammer."

*IBM maintenance manual, 1925*

# Abstract

The technique of pulsar timing reveals a wealth of new information when a precision of $\sim$1 $\mu$s or better is reached, but such precision is difficult to achieve. This thesis describes a series of very high precision timing observations that improve our knowledge of the targeted pulsar systems. We begin by describing a newly-developed baseband recording and coherent dedispersion system (CPSR2), along with a new object-oriented software development environment for pulsar data processing. Data obtained with this new instrument during a 3 year observing campaign at the Parkes 64 m radio telescope are analysed in a number of novel ways.

The mean profile of PSR J1022+1001 is shown to be stable on timescales of a few minutes, in contrast with previously published claims. We obtain a level of precision an order of magnitude better than any previous timing of this pulsar. In addition, we observe dramatic changes in the mean profile of the relativistic binary pulsar J1141–6545, which broadens by $\sim$50% over the time span of our observations. This is interpreted as evidence for secular evolution of the line of sight to the emission cone, caused by General relativistic geodetic precession which tilts the spin axis of the pulsar. High precision CPSR2 observations of the extraordinary double pulsar binary system are presented and we construct calibrated, mean polarimetric profiles for PSR J0737–3039A, in two frequency bands. These profiles provide a reference against which future profile evolution may be detected, given that we expect geodetic precession to alter the observed mean profile on an even shorter time scale than for PSR J1141–6545.

The bulk of this thesis involves timing a selection of millisecond pulsars whose physical characteristics should allow the highest precision to be obtained. We measure several new proper motions and parallax distances. Shapiro delay is used to constrain the inclination angles and component masses of several of the binary systems in our source list. In addition, subtle periodic variations of the orbital parameters of two nearby binary millisecond pulsars are detected and attributed to annual orbital parallax, providing additional constraints on their three-dimensional orbital geometries. Future observations of these two sources may lead to more stringent tests of post-Keplerian gravitational theories. Finally, we use the timing residuals of one very stable source (PSR J1909–3744) as a reference against which we time PSR J1713+0747 with a root-mean-square precision of 133 ns, amongst the best timing residuals ever obtained. This result is an important step in the search for long-period gravitational waves using pulsar timing arrays.

# Acknowledgments

Modern science is undoubtedly a collaborative exercise. So much has been discovered in the past 200 years that I doubt one person could understand it all. We do truly rely on the contributions of our present colleagues and those that came before us; not just the giants like Newton and Einstein, but everyone involved in the search for knowledge and understanding. Clearly, a set of acknowledgments such as this is never complete, but I will do my best to ensure that credit is given where due.

Firstly, I thank my closest colleagues; Matthew Bailes was always a prolific source of ideas and his enthusiasm for pulsars could be infectious. Willem van Straten's mastery of code is unparalleled in my experience and his gentle encouragement and solid advice were invaluable. Haydon Knight provided a constant sanity check, with his knack for asking tricky questions. Craig West's knowledge of hardware and machine-intuition made him such a valuable resource to us all that he rarely had time to get any research done. I especially thank Steve Ord, whose indomitable good nature fed discussions on topics ranging from philosophy to network architecture and most things in between. During two years of working in close proximity and sharing a house, we managed not only to stay sane, but to remain good friends. I also thank our other house mate, Daisuke Kawata, for putting up with us for that long! I also want to thank Bryan Jacoby for all his help with observing and data analysis, his visits from Caltech were always something to look forward to.

Working with Steve and the others on CPSR2 and CGSR2 has equipped me with a wealth of hardware and software development experience. Despite many sleepless nights and much trial and error, I now have an appreciation for logical deduction and problem solving that I don't think I'll ever lose. I'm also learning to accept my mistakes with a certain amount of grace (or at least a bit of fatalism), which is more important than it might seem.

During the course of my research, I became very well acquainted with the Parkes radio telescope. The combined activities of instrument maintenance and scheduled observing resulted in roughly 1/10 of my time being spent out at the observatory. I am grateful to Janette, Shirley and the rest of the lodge staff for looking after me (and the others) so well. Not to mention at least two Johns, Gina, Lewis and Brett, whose tireless support has made this research (and many other projects) possible.

In addition to my Australian Postgraduate Award, I received an ATNF post-graduate scholarship, complete with stipend, free accommodation at ATNF sites

and co-supervision. Thanks go to my co-supervisor Dick Manchester for many useful discussions and for organising a series of weekly teleconferences that help to keep pulsar research in Australia coordinated. Dick's team of researchers at the ATNF grew rapidly during my time at Swinburne, and I had the chance to share ideas with George & Maryam, Russell, Natasa, Tejinder and others during my occasional visits to Sydney.

I was fortunate enough to attend two overseas conferences during my time at Swinburne and I thank the organisers (especially Asha) for their efforts. No matter how painful, there really is no substitute for oral presentation experience and I value the opportunities I have had to describe my research to others, and to hear their feedback. In particular, I thank Andy Faulkner, Andrea Possenti and Duncan Lorimer for their thoughts.

I also value the experience I gained by marking projects for various Swinburne Astronomy Online units. In particular, I would like to thank Jennifer Russell, whose insightful and well researched reports were always interesting to read.

Closer to home, I thank my parents Edgar and Kathy for looking after me, encouraging my interest in science and not complaining too much when I pulled apart old household appliances to scrounge parts or see how they worked. One day I might even figure out how to put all the bits back together into something useful!

By some amazing stroke of luck or twist of fate, I met my fiancé Claire at Swinburne at the end of 2003. With her perspective to complement my own, I may yet understand far more than I could ever have hoped. As much as I appreciate logic and the scientific method, love and compassion must guide our minds. Claire, I pledge my heart to you in the sure knowledge that nothing in the sky could ever shine so brightly.

I would also like to thank the staff and students at the University of Tasmania, especially Karen Bradford and John Dickey. I finished the later stages of my thesis from Hobart and greatly appreciate the hospitality I received at UTas, including a desk in the visitor's room and access to a phone line. I also thank Matthew for his patience during the many long phone conversations it took to finalise this document.

Among other things, I hope that humans are always willing to listen and learn. In part, we study the the stars because they are untouched by human hands, pristine by virtue of distance and isolation. Perhaps the answers we seek can only be found where we have not yet searched.

# Declaration

This thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, this thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis. All work presented is primarily that of the author. Chapters 4, 5, 6, 7 and 8 have appeared in refereed journals. In each case, I authored the majority of the text. My supervisors, close colleagues and a number of anonymous referees helped to refine each manuscript. Readers should note that Chapter 7 has been modified since the time of its publication in MNRAS. The rotation measures and polarisation plots that appear in this thesis supersede those that appear in Hotan, Bailes & Ord (2005a).

This thesis analyses data recorded at the Parkes radio telescope during a 3 year period. The relevant observing proposals (P140, P361 and P456) are still active. Observations for P140 and P361 were made by a team of astronomers including Matthew Bailes, Steve Ord, Haydon Knight, Bryan Jacoby and myself. More recently, a group lead by Dick Manchester at the Australia Telescope National Facility (ATNF) initiated P456, which shares many sources and observing methods with P140. ATNF astronomers George Hobbs, Russell Edwards and John Sarkissian now regularly assist with P456 and P140. Tejinder Uppal also assisted with observations on a number of occasions. John Reynolds, Brett Preisig and Lewis Ball assisted with telescope operations and Gina Spratt, Dion Lewis and Craig West assisted with the administration of CPSR2.

I assisted Willem van Straten with the design and implementation of PSRCHIVE. We received additional guidance and advice from others already mentioned here.

Aidan Wade Hotan, July 2006

# Contents

# List of Figures

# List of Tables

# Chapter 1

# An Introduction to Radio Pulsar Astronomy

"Science, my lad, is made up of mistakes, but they are mistakes which it is useful to make, because they lead little by little to the truth."

*Jules Verne – Journey to the Center of the Earth*

The name "pulsar" is a contraction of "pulsating star", describing a unique class of astronomical objects that were first observed nearly 40 years ago (Hewish et al. 1968). Since their serendipitous discovery, pulsars have had a large scientific impact through studies of their nature and composition (Heiselberg & Pandharipande 2000), as astrophysical probes (Taylor & Cordes 1993) and as a means of testing fundamental physics (Stairs 2004). Qualitatively, pulsars are compact celestial objects that emit a regular series of broad-band pulses in the radio region of the electromagnetic spectrum, and sometimes at higher energies. Each pulsar has its own unique, highly regular pulse rate and characteristic mean pulse shape, or "profile" (Helfand, Manchester & Taylor 1975). Pulsar emission can in some cases be detected from distant parts of the Milky Way, including the globular clusters, and even from our satellite galaxies, the Magellanic Clouds (Crawford et al. 2001). Although there are many unanswered questions about their exact structure and evolution, the most remarkable property of the radio pulsars is the stability of their pulse rates. Radio pulsars act like ticking clocks that keep accurate time in their own part of the Galaxy. Pulsars broadcast timing signals that can be detected by any species that has mastered radio communication. By accurately measuring the arrival time of each pulse, observers can reconstruct a picture of the pulsar's environment with astonishing detail, allowing a wide range of astrophysical experiments.

## 1.1 The Discovery of Radio Pulsars

The development of Radio Detection and Ranging (RADAR) during the closing stages of World War II had a debatable influence on air supremacy at the time, but it certainly provided the technical foundation for a new window on the universe. By the mid-1960s, radio astronomy was a burgeoning field of study. Around this time, a group of Cambridge astronomers lead by Professor Antony Hewish began constructing a large "phased array" consisting of multiple dipole antennas joined to form one sensitive, electronically steerable radio telescope. This array was designed to repeatedly survey the radio sky, in an effort to use the effects of inter-planetary scintillation (mostly due to the solar wind) as a probe of the structure of distant compact sources (quasars). Observations were made at a frequency of approximately 81.5 MHz (just below the modern FM radio band) and all output was directed to a chart recorder that stored the detected signal on miles of paper tape. Until this point in time, most radio astronomers were content to study the time-averaged emission from astronomical radio sources. Radio detectors were designed to smooth the received power over many seconds to help reject terrestrial interference. The Hewish survey required shorter averaging times in order to accurately characterise the rapid brightness changes caused by scintillation, thereby opening up a new region of parameter space.

The first radio pulsar was discovered by Jocelyn Bell, at the time a graduate student with Hewish, who was assigned to analyse the chart recordings produced by the Cambridge radio telescope. In addition to the expected signatures of scintillating quasars, Bell noted an unusual trace that looked a little like Earthly interference but nevertheless kept sidereal time (which shifts by approximately four minutes per day relative to Solar time) over the course of several days. Higher time-resolution observations of this source revealed that the signal consisted of a series of regularly spaced pulses, with a separation of about 1.3 s. At the time, the origin of the new signal was a mystery. The research group even considered the possibility that it was a message from an extra-terrestrial intelligence! However, subsequent detection of three more rapidly pulsating sources scattered across the celestial sphere convinced them that these objects were a natural phenomenon. This only deepened the mystery as even the smallest, most energetic objects known to exist were unable to exhibit such rapid intrinsic intensity variations. Any source of finite size (though it may appear as a point source to a distant observer) should not intrinsically vary on a timescale shorter than the light travel time across its extent. For the majority of the stellar population, this places a lower limit of order a few seconds on the period of any variation. Smaller "dwarf" stars may vary with periods of order one second. Refractive and diffractive scintillation of a signal during its journey through the interstellar and interplanetary media could also cause intensity changes, but the spacing of the maxima should then be random. The simple regularity of these new repeating sources marked them as something entirely new. Hewish thought that the mechanism behind these objects might have something to do with physical

oscillations in the outer layers of white dwarf stars, hence the term "pulsating star".

History now reveals that the theoretical groundwork required to understand the nature of radio pulsars was in place more than 30 years prior to their discovery. The key concept was introduced by Baade & Zwicky (1934) (only two years after the discovery of the neutron!) in a paper on the nature of supernovae (the explosive demise of massive stars). This article predicted that the implosion of a massive star could leave behind a highly compact object with a degenerate collapsed core of compressed nuclear matter. Such an object might contain the mass of the Sun within a diameter of 20 km. Exotic objects like this were thought to be too small to detect with the technology of the time. Despite their likely high temperatures, the surface area of a neutron star would be too small to dissipate much thermal radiation, leaving little trace of its existence.

Theoreticians had long been aware that the magnetic flux and angular momentum of a star would have to be conserved as it evolved. Matter ejected from the outer layers of the star could only account for some of the required energy, leading to the idea that supernova-born "neutron stars" might be highly magnetised and rapidly spinning, as well as very hot. Shortly after the Hewish et al. (1968) observations, Gold (1968) and Pacini (1968) independently connected the rapid, regular pulse rates of the radio pulsars with the idea of rotating neutron stars. The resulting "lighthouse" model assumes that a pulsar is a rotating, highly magnetised neutron star that emits intense beams of radiation from its magnetic poles. These radiation beams are locked to the spinning crust and sweep across the sky, producing a pulse of radio power in any distant detector that lies in the line of sight. A number of other physical models attempted to explain the rapidity and regularity of radio pulsar signals; the possibilities included "ringing" or vibrational oscillations in the outer layers of a white dwarf star. These alternative models were refuted within two years of the discovery of the first pulsar.

As new instruments were tasked with finding ever more radio pulsars, it became clear that only very compact objects could accommodate the most rapid periods observed. The discovery of the Vela and Crab pulsars (with periods of 89 and 33 ms respectively) sealed the fate of models based on white dwarfs. Before long, observers demonstrated that the pulse periods of these two sources were systematically increasing; this fact could only be explained by models based on rotation, not vibration, as vibrations die down in amplitude over time but maintain a constant period. In addition, the spatial coincidence of the Crab pulsar with the Crab nebula (the luminous remains of a supernova recorded by observers in 1054 AD) offered clues as to the evolutionary history of the pulsar, favouring the rotating neutron star model.

### 1.1.1 Pulsar Nomenclature

The naming conventions used to identify individual pulsars are extremely simple. Unfortunately, the naming scheme has changed three times in the past and a short

history is presented here to help avoid confusion. Pulsars have always been named (somewhat unimaginatively) after their sky coordinates. Originally, the name began with a 2-letter code representing the observatory responsible for discovering the pulsar, followed by a 4-digit number representing the degrees and minutes of right ascension. For example, the first pulsar was named CP1919; "CP" standing for "Cambridge Pulsar" and "1919" representing the right ascension, which Hewish et al. (1968) reported to be 19h 19m 38 ± 3s. As the number of known pulsars increased, the observatory code was dropped in favour of the prefix "PSR" and a single letter representing the reference epoch of the coordinates. "B" is used to represent 1950 and "J" is used to represent 2000. The so-called "B" names also include two digits and a sign, representing the degrees of declination. For example, the first known pulsar became PSR B1919+21 under the new scheme. The extra position information was included to avoid confusion between pulsars with the same degrees and minutes of right ascension. More recently still, the "J" name convention includes the minutes of declination as well as the degrees. The first pulsar is now known as PSR J1921+2153. The right ascension has changed by approximately 2 minutes, due to precession of the equinoxes over a period of 50 years.

## 1.2 The Pulsar Emission Mechanism

Pulsars are thought to be rapidly rotating neutron stars with dipolar magnetic fields; energetic beams of radiation blast out from the magnetic poles of the star and are swept around the sky, provided the magnetic axis is offset from the rotation axis of the star, as shown in Fig. 1.1. Classical electrodynamics tells us that a magnetic dipole with moment $M$, rotating about an axis perpendicular to the dipole with angular velocity $\omega$, radiates an electromagnetic wave with angular frequency $\omega$ into free space, dissipating a total amount of power $P_{\rm rad}$, given by (Griffiths 1999)

$$P_{\rm rad} = \frac{1}{3} M^2 \omega^4 c^{-3} \times 10^{-7} \text{ (Watts)}. \tag{1.1}$$

The rotational kinetic energy $E_{\rm rot}$ of a spinning object is

$$E_{\rm rot} = \frac{1}{2} I \omega^2 \text{ (Joules)}. \tag{1.2}$$

Although the total power lost by the pulsar will be a combination of magnetic dipole radiation, beamed emission and perhaps some amount of particle outflow, dipole emission is a very efficient energy loss mechanism and therefore the dominant braking force. We assume that all the rotational kinetic energy lost is converted into dipole radiation

$$-\frac{dE_{\rm rot}}{dt} = P_{\rm rad}, \tag{1.3}$$

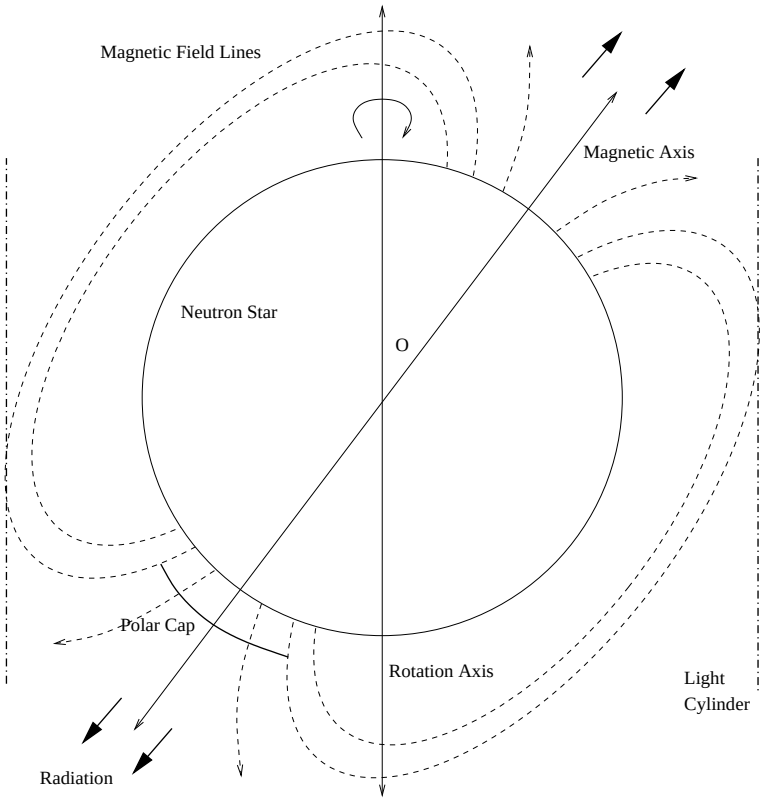


**Figure 1.1: This diagram illustrates the relative alignment of the rotation axis and magnetic field for a hypothetical pulsar. A neutron star with a dipolar magnetic field rotates about an axis that is offset from the magnetic dipole axis by an angle $\alpha$. The magnetic field is locked to the surface of the star and must co-rotate with it. At a certain distance from the surface (the light cylinder radius), the magnetic field can no longer co-rotate without exceeding the speed of light. Field lines originating too close to the pole cannot close without passing through the light cylinder and therefore remain open, creating a "polar cap". Beams of radio emission are thought to originate within or above this region.**

$$-I\omega\dot{\omega} = \frac{1}{3}M^2\omega^4 c^{-3} \times 10^{-7}. \tag{1.4}$$

Angular velocity $\omega$ is related to the rotation frequency $\nu$ and period $P$ by $\omega = 2\pi\nu = 2\pi/P$. Taking the time derivative leads to the expression

$$\frac{d\omega}{dt} = -\frac{2\pi}{P^2}\dot{P}. \tag{1.5}$$

According to Eq. 1.4,

$$4\pi^2\frac{I\dot{P}}{P^3} = \frac{16\pi^4}{3}\frac{M^2}{P^4c^3} \times 10^{-7}. \tag{1.6}$$

We then solve for $M$, to obtain the expression

$$M = \sqrt{\frac{3 \times 10^7}{4\pi^2}IP\dot{P}c^3}\ (\mathrm{m^2\,A}). \tag{1.7}$$

We also know that the field strength at a distance $l$ along the axis of a magnetic dipole of moment $M$ is given by (Griffiths 1999)

$$B_{\mathrm{axis}}(l) = \frac{\mu_0 M}{2\pi l^3}\ (\mathrm{Teslas}). \tag{1.8}$$

It is therefore possible to solve for the magnetic field,

$$B_{\mathrm{axis}}(l) = \frac{1 \times 10^{-7}}{\pi l^3}\sqrt{3 \times 10^7 IP\dot{P}c^3} \cong 9.06 \times 10^{-4}\sqrt{IP\dot{P}}\ (\mathrm{Teslas}). \tag{1.9}$$

Equation 1.9 provides a useful estimate of a pulsar's magnetic field strength, in the case of pure magnetic dipole braking. Considering that an ideal, uniform, spherical neutron star containing $1.3\,\mathrm{M}_\odot$within a diameter of 20 km has a moment of inertia approximately equal to $1{\times}10^{38}\,\mathrm{kg\,m^2}$, we see that radio pulsars are very effective flywheels. Their spin-down rates are typically of order $10^{-15}\,\mathrm{ss}^{-1}$, which makes their surface magnetic fields of order $10^8$ T, or $10^{12}$ G. Gunn & Ostriker (1969) considered in detail the theoretical case of emission from an orthogonal rotator, where the magnetic and spin axes are at $90^o$ to one another. Ostriker & Gunn (1969) applied this model to the Crab nebula and argued convincingly that dipole radiation from a spinning, magnetised neutron star was responsible for the observed luminosity.

Of course, the simple orthogonal rotator model may not be complete or correct. More generally, we can assume that the spin-down rate $\dot{\omega}$ of the pulsar is related to its current angular velocity by the expression:

$$\dot{\omega} = -k\omega^n \tag{1.10}$$

Here, $k$ is a constant and $n$ is the "braking index", which is equal to 3 in the case of pure magnetic dipole braking. Assuming that n $\neq$ 1 and that the initial angular

velocity was much larger than the current angular velocity, we can integrate Eq. 1.10 in time to obtain an expression for the characteristic age $\tau$ of a pulsar, in terms of the observable quantities $P$ and $\dot{P}$,

$$\tau = \frac{P}{(n-1)\dot{P}}. \tag{1.11}$$

Normally, we assume n = 3, reducing the expression to $\tau = P/2\dot{P}$. However, by differentiating Eq. 1.10 we can derive an expression for the braking index $n$ in terms of observable quantities,

$$n = 2 - \frac{P\ddot{P}}{\dot{P}^2}. \tag{1.12}$$

Obtaining the braking index from Eq. 1.12 requires detection of the spin period second derivative. For most pulsars, $\ddot{P}$ is far too small to measure. Very young pulsars experience the greatest braking torque and the value of $n$ has been measured for the Crab pulsar (Lyne, Pritchard & Smith 1988), PSR B1509–58 (Kaspi et al. 1994; Livingstone et al. 2005), PSR B0540–69 (Manchester & Peterson 1989; Livingstone, Kaspi & Gavriil 2005), PSR J1119-6127 (Camilo et al. 2000a) and the Vela pulsar (Lyne et al. 1996). In all cases, the measured braking index was $< 3$. Gavriil, Kaspi & Roberts (2004) used the Rossi X-ray Timing Explorer (RXTE) satellite to measure a braking index $> 3$ for the X-ray pulsar PSR J1811-1925, but this value is likely to be contaminated by timing noise. Johnston & Galloway (1999) integrate Eq. 1.10 between two arbitrary times to obtain an expression for $n$ that does not depend on $\ddot{P}$ (and does not require a phase-connected solution over a long time span) and measure 20 pulsar braking indices. The large range of values obtained suggests that most measurements are dominated by the effects of glitches or other timing instabilities. The authors call the concept of smooth spin-down with constant braking into question. In this thesis, we strive to obtain the best possible timing precision and must therefore avoid the young pulsars.

We have seen that a pulsar's magnetic field is responsible for gradually slowing the spin of the star by dissipating rotational kinetic energy into the surrounding space. However, the magnetic field is also thought to generate beamed emission originating near the polar region. It is this emission that drives the unique, regularly repeating signal that distinguishes pulsars from other astronomical sources. Unfortunately, it is also the least understood aspect of pulsar emission. The magnetic field associated with a neutron star is thought to be a remnant of the parent star's field, highly compressed and increased in strength by 10 orders of magnitude during the birth of the star. This field is locked to the body of the star by interactions with charged particles in the crust and therefore co-rotates with the solid exterior. As the field extends outwards from the surface of the star, it has to move more rapidly through space to continue co-rotating. At a certain distance (given by Eq. 1.13), the speed required for co-rotation exceeds the speed of light, preventing the magnetic field closing from pole to pole. The region around the magnetic axis

at each pole that is forced to remain open is called the "polar cap". Beamed radio emission is thought to originate within or above this region.

$$r = \frac{c}{\omega} \tag{1.13}$$

The "light cylinder" radius $r$ is the distance at which the outermost part of a closed field line would have to move at the speed of light, $c$, in order to co-rotate with the crust.

Pulsar beams are thought to be driven by synchrotron emission from relativistic electrons accelerated in their magnetic field. The exact configuration of fields and charged particles in a pulsar's magnetosphere is still a matter of debate. Goldreich & Julian (1969) and Ruderman & Sutherland (1975) presented some of the first theoretical models. Radio emission may originate near the surface of the star, or in a region much higher up, close to the light cylinder. This second scenario is known as the "outer gap" model (Cheng, Ho & Ruderman 1986). Often, the shape of the mean profile changes significantly with observing frequency, which makes it difficult to perform precision timing over a range of wavelengths. It is likely that different radio frequencies are emitted from different altitudes in the magnetosphere, so the true emission mechanism may be a combination of various existing models. The pulse profile we observe is effectively a one-dimensional cut through the emission cone. Observing at multiple frequencies adds a model-dependent second dimension, but our knowledge of the radial extent and shape of pulsar beams is severely limited.

The pulsar emission mechanism is highly broad-band; some of the most energetic (i.e. youngest) pulsars are seen to pulse in visible light and beyond. If a pulsar is accreting matter, hot spots can form on the surface or within a surrounding accretion disk, producing pulsed or continuous X-ray emission (Blumenthal & Tucker 1974). Modern research in this field is often aided by numerical simulation. The few pulsars that emit high-energy photons may provide the key to the emission mechanism puzzle (Cheng 2004; Baring 2004), but much work is still required. Within the radio region of the electromagnetic spectrum, the flux density $S$ (power received per unit collecting area per unit frequency, measured in units of Janskys, where 1 Jy = $10^{-26}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$) of a pulsar tends to decrease with increasing frequency, following a characteristic synchrotron power law of the form:

$$S \propto \nu^{\alpha} \tag{1.14}$$

Here, $\nu$ is the radio frequency at which we observe and the exponent $\alpha$ is known as the "spectral index". Based on a study of 280 pulsars, Lorimer et al. (1995b) conclude that the mean spectral index is –1.6, but extreme values range from –3.4 up to 0.2. The majority of pulsars have spectral indices close to –2.0. While pulsars are intrinsically more luminous at low frequencies, the temperature of the background sky also increases with decreasing frequency, as does dispersive smearing. For precision timing, it is best to minimise dispersion smearing by observing at high frequencies, but this must be balanced against the loss of flux density and signal-to-noise ratio

(S/N). The optimal frequency at which to observe radio pulsars can vary depending on the source, experiment, sky position and so on, but it is usually in the range 0.3 – 3.0 GHz.

## 1.3 Interstellar Propagation

Although the speed of light in a vacuum is a fundamental constant, the speed of light in other transparent media can be significantly different. In much the same way as glass or water slow down the passage of optical light, electrons in the interstellar medium (ISM) can slow down radio waves. The time delay relative to radiation with infinite frequency is given by

$$t = \frac{\mathrm{DM}}{2.41 \times 10^{-16} \nu^2} \text{ (s)}. \tag{1.15}$$

Here, DM is the "Dispersion Measure", defined to be the integral of the electron column density, and $\nu$ is the radio frequency in Hz. Low frequency signals experience more "drag" than high frequency signals, leading to refraction. Broad-band pulses are therefore spread out in time during their passage through the ISM. This effect is known as "dispersion". Figure 1.2 shows the signal from the Vela pulsar, PSR J0835–4510. The vertical axis is radio frequency, the horizontal axis is "pulse phase", which is equivalent to time, starting at zero and ending one pulse period (approximately 89 ms) later. The pulse energy is clearly distributed in time, the lower frequency components of the pulse arrive some time after the higher frequency components.

Incoherently sampling the mean pulse power over a finite bandwidth will unavoidably "smear" the profile. The amount of time by which a pulse is dispersively smeared depends on the DM, the observing frequency $\nu$ and the bandwidth $B$ and is given by

$$t_{\mathrm{smear}} = \frac{8.3 \times B \times \mathrm{DM}}{\nu^3} \text{ ($\mu$s)}. \tag{1.16}$$

Here, $\nu$ is specified in GHz, $B$ in MHz and DM in cm$^{-3}$pc. To characterise the true shape of a pulse profile it necessary to observe vanishingly narrow bandwidths, but this has a detrimental effect upon the S/N (see Eq. 3.2). In order to overcome both these problems, instrument designers have devised ways to observe many small channels that can be combined to give a large effective bandwidth. These channels can be summed together with the appropriate delays to ensure proper alignment of the pulse at each frequency.

### 1.3.1 Coherent Dedispersion

Because there are hardware limits on the number of discrete spectral channels that can be constructed in a single instrument, incoherent sampling always leaves a resid-

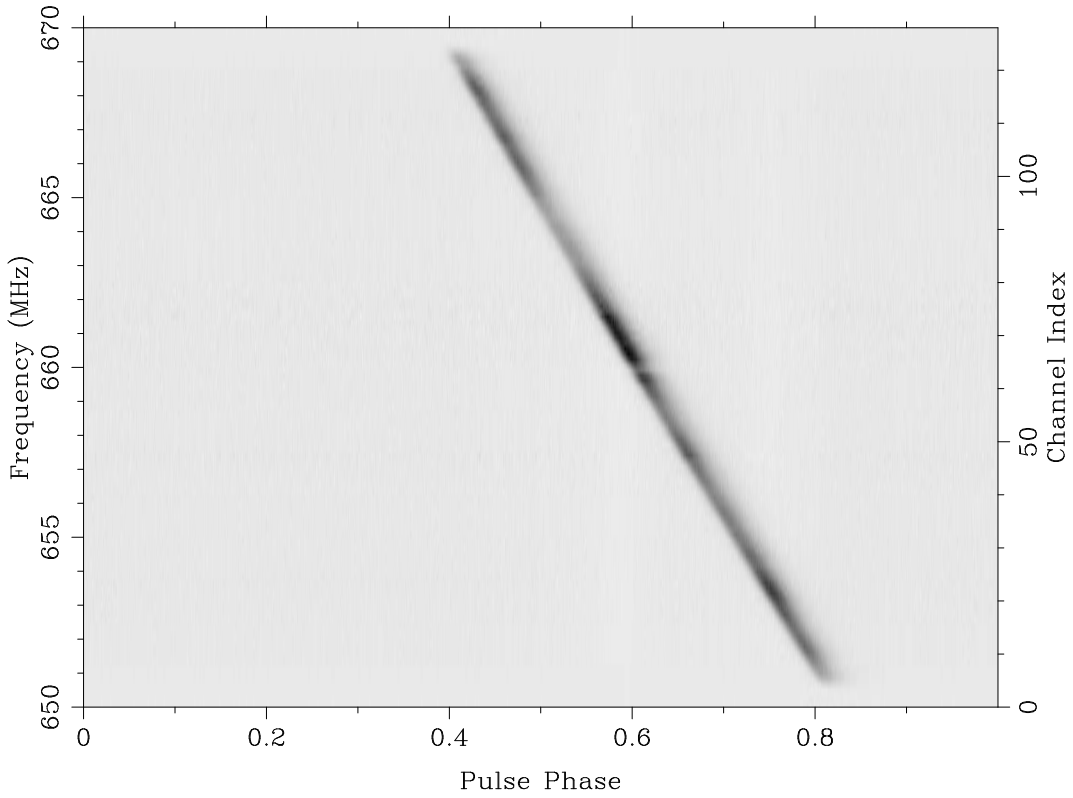


**Figure 1.2: This diagram illustrates the dispersive effect of the ISM on the mean pulse of the Vela pulsar, J0835–4510, within a 20 MHz-wide portion of the radio spectrum that has been divided into 128 narrow channels. Although the delay appears to be a linear function of frequency, Eq. 1.15 shows that the relation is actually parabolic. Deviations from linearity only become apparent with wider fractional bandwidths and shorter pulse periods.**

ual amount of dispersion smearing in the final profile, equal to the amount of smearing that occurs within the bandwidth of a single spectral channel. In theory, it is possible to remove even this residual smearing, but the operation must be performed prior to detecting the signal. Hankins & Rickett (1975) consider the ISM to be a linear filter that converts an input signal $f(t)$ into an output signal $g(t)$, and derive a suitable transfer function based on the cold plasma dispersion relation,

$$H(\nu + \nu_0) = exp\left[\frac{i2\pi D\nu^2}{\nu_0^2(\nu_0 + \nu)}\right], \nu << \nu_0. \tag{1.17}$$

Here, we have shifted by the centre frequency $\nu_0$ so that our transfer function is centred on $\nu = 0$ and the bandwidth of the observation is bounded by $\pm\nu$. The complex impulse response of the ISM, $h(t)$, is simply the Fourier transform of $H(\nu+\nu_0)$. The transfer function $H(\nu + \nu_0)$ modifies an incoming signal $f(t)$ by multiplication with the Fourier transform, $F(\nu)$. The modified signal is recovered by reverse Fourier transform, i.e. $G(\nu) = H(\nu)F(\nu)$, where $G(\nu)$ is the Fourier transform of the filtered signal, $g(t)$. In radio astronomy, we measure $g(t)$ and wish to reconstruct the form of the signal before it passed through the ISM. Because we can compute the inverse transfer function corresponding to Eq. 1.17, it is possible to reverse the filtering process. We must first compute the discrete Fourier transform of the observed signal, multiply it by the inverse transfer function in the Fourier domain and then transform back into the time domain. To do all this without introducing artifacts into the signal requires Nyquist sampling and careful accounting for the periodic nature of the Fast Fourier Transform (FFT) algorithm (Cooley & Tukey 1965). The computational load is high, and the data rate increases with observing bandwidth.

### 1.3.2 Scattering and Scintillation

Inhomogeneities in the electron content of the interstellar medium act like a phase-altering screen that can interfere with the radiation emitted by distant radio point sources, in much the same way that Earth's atmosphere interferes with the visible light from distant stars. Given the small intrinsic size of their emission regions, pulsars are highly susceptible to interstellar scintillation in the strong scattering regime. More recently, the phenomenon of intra-day variability in flat-spectrum quasars has provided a new way to study scintillation in both the strong and weak regimes (Kedziora-Chudczer et al. 1997; Dennett-Thorpe 2000; Bignall et al. 2003).

Walker (1998) summarises our understanding of radio scintillation, with a view towards extra-galactic sources and Narayan (1992) examines the physics of pulsar scintillation. Scintillation causes quasi-random changes in the intensity of the observed signal as a function of frequency and time. Astronomers describe the effects of scintillation using three parameters, the characteristic bandwidth, characteristic timescale and the modulation index of any observed variation. The characteristic bandwidth and timescale describe the structure or "dimensions" of the scintillation

pattern and the modulation index describes the amount by which the amplitude of the signal is seen to vary. It is defined as the root-mean-square (RMS) of the observed fractional flux changes.

Within the strong scattering regime, we expect two types of intensity modulation, known as refractive and diffractive scintillation. Refractive scintillation is caused by focusing and defocusing of rays due to large-scale structure within the ISM, which acts like an enormous lens. This type of scintillation is associated with wide bandwidths (of order the observing frequency used) and time scales ranging from days to years. Refractive scintillation makes it difficult to determine the intrinsic flux of radio pulsars. The mechanism of diffractive scintillation is quite different. Small-scale variations in the structure of the ISM cause diffraction of incident wavefronts, in much the same way as a slit or diffraction grating acts upon visible light. Rays diffracted from different parts of the ISM will interfere in the plane of the Solar system, mapping out an interference pattern on the ecliptic. The characteristic separation of the minima and maxima will depend on the inhomogeneity scale of the ISM and the observing frequency. As the Earth moves in its orbit, it cuts through different parts of the interference pattern, causing the source to "twinkle". Similarly, the motion of the pulsar will cause the interference pattern to change. The timescale of the resulting intensity variation depends on the velocity of Earth and the pulsar. Walker (1998) derives an expression for the characteristic bandwidth for diffractive scintillation,

$$\frac{\Delta\nu}{\nu} = \left(\frac{\nu}{\nu_0}\right)^{17/5}. \tag{1.18}$$

Here, $\nu$ is the observing frequency and $\nu_0$ is the transition frequency, which is fixed for a given line of sight. The associated timescale is given by

$$t_{\rm d} \approx 2\left(\frac{\nu}{\nu_0}\right)^{6/5} \text{ (hours)}. \tag{1.19}$$

Changes in intensity are in general only noticeable if the observing bandwidth is smaller than the scintillation bandwidth, but may still be detectable in individual spectral channels provided some frequency resolution is retained. The pulse profile is also susceptible to scatter-broadening, which is a side effect of multi-path propagation (see Fig. 1.3). High precision timing depends on precise knowledge of the observing frequency in order to accurately predict the delay caused by the ISM. Diffractive scintillation can cause power to appear off-centre within the observing band, subtly altering the effective observing frequency and introducing a systematic error.

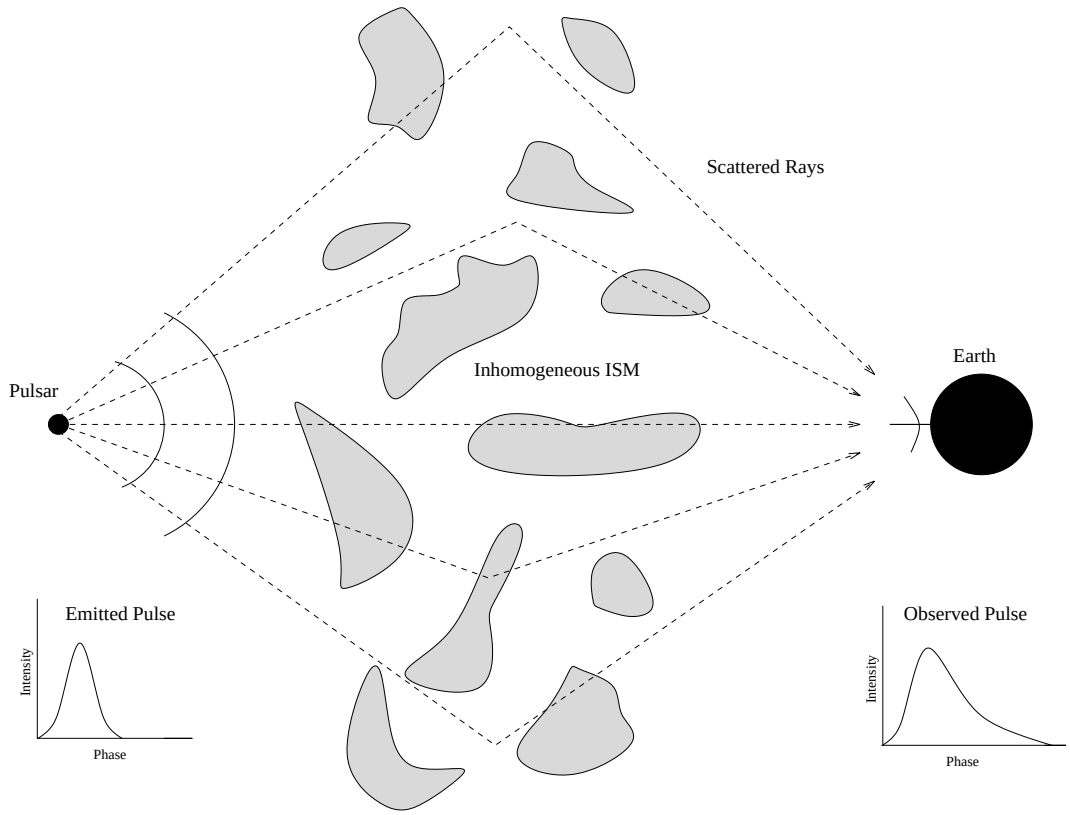


**Figure 1.3: This diagram illustrates the intrinsic broadening of a pulse due to multi-path scattering as the signal travels through an inhomogeneous interstellar medium. The first ray to arrive travels in a straight line from the pulsar to the radio telescope, subsequent rays can be scattered back from different angles, traveling a greater total distance and consequently arriving at later times. This gives the pulse an extended tail. This effect cannot be corrected for but its magnitude is a strong function of the dispersion measure and observing frequency.**

## 1.4 Radio Pulsar Timing

We now know that each individual pulse emitted by a radio pulsar varies greatly in power, structure, polarisation and even in phase. Despite this, the addition of several thousand individual pulses always produces a stable mean profile that can be used as the "fingerprint" of the source (Helfand, Manchester & Taylor 1975). This makes the technique of pulsar timing possible.

The very first observations revealed that the periods of the pulsars were stable to better than 1 part in $10^7$ (Hewish et al. 1968). Such accuracy was enough to reveal Doppler shifts due to the motion of the Earth around the Sun. Astronomers started to keep track of the arrival times of the pulses from a number of sources, in the hope that the they would reveal important information. Careful analysis of trends in these arrival times soon determined several pulsar rotation periods to better than a tenth of a nanosecond. This in turn allowed the first detection of pulsar spin-down; before long the first period derivatives of the youngest pulsars could be determined to better than a few parts in a thousand (Davies, Hunt & Smith 1969; Reichley, Downs & Morris 1970). Mathematical models describing these perturbations were constructed and used to predict when the next pulse would arrive.

Small errors in right ascension and declination can be easily detected with pulsar timing. Similarly, if the pulsar is moving across the sky (the violence of their birth leaves many with high velocities), the gradual change in position will also be detectable over time. In this way, pulsar timing can aid studies of galactic kinematics. In fact, pulsars can be used very effectively to probe the structure of the Milky Way. By measuring the relative pulse delay over several frequency channels, the dispersion measure of a pulsar can be determined. This provides an estimate of the electron column density. Alternatively, models of the galactic electron density can be used to estimate the distance to a pulsar once its dispersion measure is known (Lyne, Manchester & Taylor 1985; Taylor & Cordes 1993). In addition, if a pulsar emits any linearly polarised radiation (and most do), careful analysis of the position angle of the polarised component can yield information about the magnetic field along the line of sight, through an effect known as Faraday rotation (Smith 1968; Hamilton & Lyne 1987).

Not all of these parameters can be measured quickly or easily. After a new pulsar is discovered, 12 months of follow-up timing will usually reveal the spin period, period first derivative (and hence an estimate of the pulsar's age and magnetic field strength), sky position and dispersion measure to high accuracy. Given a few more years of observations, the proper motion of the pulsar can sometimes be determined. Very Long Baseline Interferometry (VLBI) can sometimes be used as an alternative or independent means of obtaining this information (Bartel et al. 1985; Dodson et al. 2003).

### 1.4.1 Binary Pulsars

The discovery of the first pulsar in a binary system, PSR B1913+16 (Hulse & Taylor 1975), opened up a whole new parameter space for pulsar timing experiments. This pulsar was observed to exhibit large, periodic variations in its 59 ms rotation period, interpreted as Doppler shifts induced by motion in an approximately 8 hr orbit about another object of similar mass. Pulsar timing can be used to provide an extremely detailed picture of the orbit of a binary pulsar. Shortly after a new pulsar is discovered, the orbital period, eccentricity (if significant), longitude and epoch of periastron and the projection of the semi-major axis along the line of sight can be determined. In fact, these parameters are crucial to obtaining a phase connected solution that can predict in advance when the next pulse will arrive.

## 1.5 Millisecond Pulsars

Fifteen years after the discovery of the first pulsar, Backer et al. (1982) discovered PSR B1937+21, with a rotation period of only 1.5 ms. Such rapid rotation is near the structural limit for a neutron star and the discovery came as quite a surprise. The pulsar was only identified because its emission mingled with that of a nearby HII region, mimicking a mysterious steep-spectrum source with an extended flat spectrum component. Intensive searches at a number of radio frequencies failed to identify the true nature of the steep spectrum source (even though it was suspected to be a pulsar) because the search methods applied were not sensitive to such rapid periods. Eventually, harmonics of the fundamental spin rate were identified and shortly afterwards the discovery of the fastest pulsar ever observed was confirmed.

Today it is accepted that the millisecond pulsars (MSPs) are normal pulsars that have been spun-up by a process of mass accretion from a binary companion. Bhattacharya & van den Heuvel (1991) present a comprehensive review of this theory, which is described here only in broad terms. Binary MSP progenitor systems are thought to contain at least one star that exceeds the critical mass required for a supernova. The more massive star also evolves more rapidly and provided the system stays bound after the supernova, there is a chance that a young pulsar will be left orbiting a main sequence star. The companion evolves more slowly, eventually expanding to fill its Roche lobe, at which point matter will start to transfer from the outer layers of the companion onto the pulsar. This process of accretion must conserve angular momentum, so the pulsar must spin faster as it accretes more matter. The companion may eventually become a white dwarf, or there may be a second supernova. If the system remains bound after the second supernova, we could be left with a double neutron star (DNS) or black hole neutron star binary system. Accretion also seems to reduce the pulsar's magnetic field strength by several orders of magnitude (Shibazaki et al. 1989), increasing its life expectancy up to the vicinity of the Hubble time (the age of the universe). We therefore expect these "millisecond" or "recycled" pulsars to occupy a much different evolutionary

niche to the "normal" or "slow" pulsars. Plotting the observed period against the observed period derivative for pulsars in the galactic field illustrates this fact well (Fig. 1.4). The main population of normal pulsars is concentrated in the middle of the plot. The MSPs cluster towards the bottom left, clearly isolated from the main population.

After the completion of many successful surveys during the last few decades (Large, Vaughan & Wielebinski 1968; Manchester et al. 1978; Clifton et al. 1992; Johnston et al. 1992; Manchester et al. 1996; Sayer, Nice & Taylor 1997; Lyne et al. 1998; Edwards et al. 2001), the number of known radio pulsars stands at 1513 (Manchester et al. 2005a). Nearly half of these were discovered during the highly successful Parkes multibeam survey and its high-latitude extensions (D'Amico et al. 2001; Morris et al. 2002; Kramer et al. 2003; Burgay et al. 2003; Hobbs et al. 2004a; Faulkner et al. 2004). In total, 134 of these pulsars have periods less than 20 ms. The success of modern surveys is helping to bridge the gap between the two population groups, finding more objects at various intermediate evolutionary stages.

Many of the known MSPs reside in globular clusters. Due to their compact extent, globular clusters can be surveyed with relatively few sky pointings. This allows additional integration time, which increases the chance of detecting weak sources. Deep, targeted surveys of 47 Tucanae (Manchester et al. 1991; Camilo et al. 2000b) and Terzan 5 (Ransom et al. 2005) have found close to 50 MSPs. This can complicate any analysis of the MSP population because the evolutionary histories of pulsars in globular clusters and those in the galactic field are likely to be quite different.

## 1.6 High-Precision Timing

Binary pulsars in particular provide a unique opportunity to test fundamental physics. The original binary pulsar PSR B1913+16 is mildly recycled and has a very eccentric, short-period orbit with a neutron star companion. Taylor & Weisberg (1982) determined that the orbit of the pulsar is decaying at precisely the rate predicted due to the emission of gravitational radiation (a consequence of General relativity). Chapter 2 describes the timing model parameters required to define a Keplerian orbit and several additional "post-Keplerian" parameters whose detection can be used to test theories of gravity or determine the component masses of the system.

The interior of a neutron star is governed by the equation of state for nuclear matter. Because the extreme conditions inside a neutron star cannot be replicated on Earth, our knowledge of the behaviour of matter at very high densities is limited and the exact form this equation takes is uncertain. The range of theoretically possible equations of state provide a range of acceptable neutron star masses, extending from a few tenths of a Solar mass up to about ten times this amount. Observations of neutron stars provide one of the few experimental constraints available in this field.

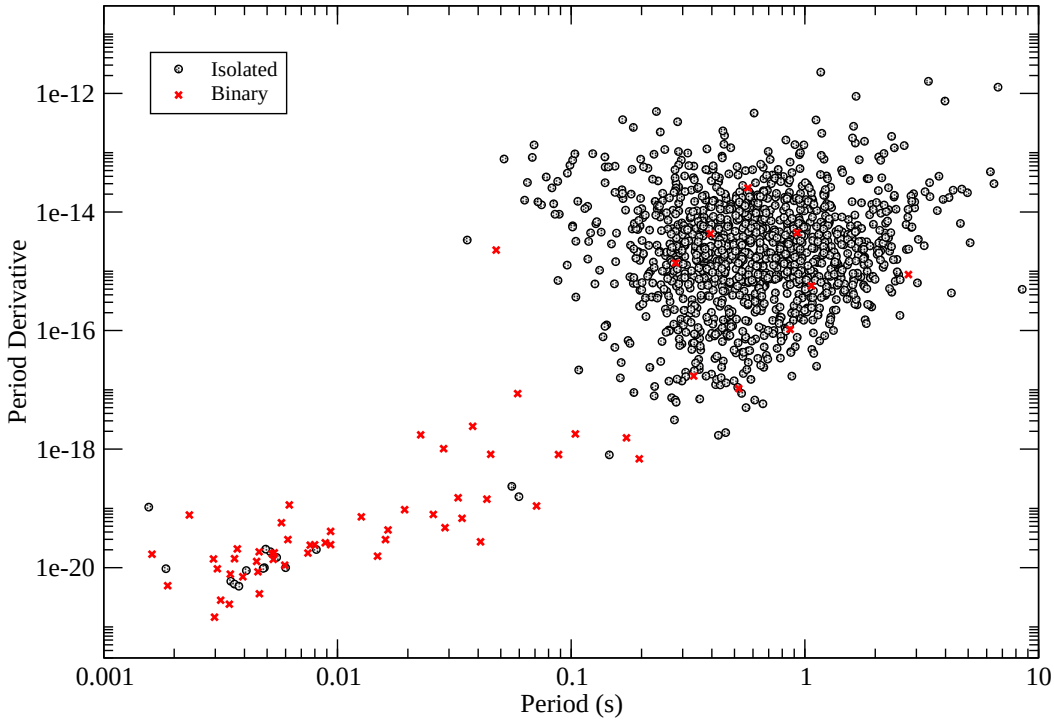


**Figure 1.4: This diagram shows period derivative plotted against period for the majority of the radio pulsar population. Pulsars known to reside in globular clusters have been excluded from the plot because their observed period derivative is likely to be corrupted by acceleration in the gravitational potential of the cluster. Solitary pulsars are shown as filled circles, pulsars with a binary companion are shown as crosses. The millisecond pulsars appear to be a separate population in the bottom left corner and are preferentially found in binary systems.**

The nature of the accretion process that recycles MSPs is also not well understood. Astronomers have few constraints on the amount of matter the pulsar must accrete to reach millisecond periods, making the precise measurement of MSP masses very important to our understanding of their evolution.

In all but the most extreme systems, relativistic influences are small and a very high level of timing precision is required to detect post-Keplerian perturbations. Because of their smaller spin-down torque and greater angular momentum, the millisecond pulsars can be timed much more precisely than the normal pulsars. When the mean arrival time uncertainty is below $\sim$1 $\mu$s, a range of interesting experiments can be performed. The better the timing precision, the greater the significance of any measured parameter. Our ability to measure the various post-Keplerian parameters also depends to a large extent on the physical characteristics of the system. For example, it is difficult to measure the rate of periastron advance if the eccentricity of the orbit is very small. The DNS binary systems offer both high timing precision and orbits most likely to exhibit post-Keplerian effects. High-precision timing of PSR B1913+16 (Taylor & Weisberg 1989), PSR B1534+12 (Stairs et al. 1998) and PSR B2127+11C (Deich & Kulkarni 1996) provided the first measurements of the masses of neutron stars and were use to test General relativity in a number of novel ways. Binary systems containing an MSP and a white dwarf tend to be made circular by tidal forces, but it is sometimes possible to detect a relativistic effect known as "Shapiro delay", that can be used to determine the component masses. The timing of PSR B1855+09 (Ryba & Taylor 1991; Kaspi, Taylor & Ryba 1994) revealed this effect. More recently, timing of PSR J1909–3744 has provided one of the most precise estimates of a recycled pulsar mass ever obtained, thanks to its favourable orbital geometry and excellent timing precision (Jacoby et al. 2005).

Lyne et al. (2004) recently discovered a double pulsar binary system, consisting of two active and observable pulsars orbiting each other. Timing of both these objects allows the orbit and both pulsar masses to be determined very accurately. The orbital period is only 2.4 hr, making post-Keplerian effects highly significant. The longitude of periastron is observed to advance at the rate of $\sim 17^o\,\mathrm{yr}^{-1}$. The orbital plane is nearly coincident with our line of sight (though the exact angle is a matter of some uncertainty) and short-duration eclipses have been observed. The emission from the 2.77 s secondary pulsar exhibits a complicated sequence of intensity variations as a function of orbital phase, providing good evidence that the magnetosphere is influenced by the emission of the primary pulsar (McLaughlin et al. 2004). The scintillation pattern of both sources can be used to investigate the properties of the ISM with unprecedented detail (Coles et al. 2005).

The MSPs do not seem to suffer from timing noise or rotational glitches to the same extent as the younger "normal" pulsars. In fact, the only example to date of a glitch-like event in a millisecond pulsar was observed very recently (Cognard & Backer 2004), and was orders of magnitude smaller than the glitches observed in Vela (Dodson, McCulloch & Lewis 2002), for example. MSPs also seem to be exempt from phenomenon like nulling and mode changing, their emission is remarkably steady.

However, Kramer et al. (1999) report that PSR J1022+1001 has a timing precision approximately 10 times worse than we might expect based on its period. These authors assert that its poor timing performance is due to shape changes in the mean profile on a timescale of several minutes. Ramachandran & Kramer (2003) interpret these changes as evidence for magnetopheric return currents. If these profile changes are real, they cast doubt on the fundamental stability of the MSPs.

Another source of profile instability is the phenomenon of precession. The rotation axis of a pulsar will precess if it is misaligned with the orbital angular momentum vector in the case of a binary system (geodetic precession) or if there is some asymmetry in the shape of the neutron star (free precession). The original binary pulsar B1913+16 exhibits secular profile changes that Weisberg, Romani & Taylor (1989) interpret as evidence for geodetic precession. Stairs, Lyne & Shemar (2000) observe correlated periodic variations in both the arrival times and pulse profile shape of the isolated pulsar PSR B1828–11, which are interpreted as evidence for free precession. Although precession has the potential to affect the timing of a pulsar, it also offers a chance to map out the complete shape of the emission beam, albeit over lengthy timescales. The geodetic precession period of the 22.7 ms pulsar in the double pulsar binary system is $\sim$70 years, but no profile shape changes have yet been detected.

In the quest for ever higher precision in order to better measure system parameters and probe fundamental physics more deeply, astronomers must find innovative ways to avoid systematic errors in pulse timing experiments. To avoid errors associated with terrestrial time standards and Solar system ephemerides, it is possible to simultaneously fit for global parameters across the timing residuals of multiple pulsars. The concept of the pulsar "timing array" in its most basic form dates back 15 years (Foster & Backer 1990). The limiting factor has always been the small number of pulsars that can be timed to the level of precision required to make a timing array useful (Jenet et al. 2005). One of the primary goals of this thesis is to better understand the fundamental limits to pulsar timing precision in an effort to increase the sample of sources that can be timed with better than 1 $\mu$s of precision.

Pulsar timing has contributed to a wealth of scientific results. For a more detailed summary, the reader can refer to two excellent books by Lyne & Smith (2004) and Manchester & Taylor (1977), and an online living review[1] that summarises recent work in the field.

## 1.7 Thesis Outline

In this thesis we describe a series of millisecond pulsar timing experiments and three studies of mean profile morphology and polarimetry, with an aim to achieve the highest possible timing precision. We begin by describing the technique of pulsar timing in some detail, including the mathematical models used to predict pulse phase as a function of time and the algorithms used to obtain arrival times.

[1] http://relativity.livingreviews.org/Articles/lrr-2001-5/index.html

Next, we describe the design and construction of the instrument used to record data. Chapter 3 begins with a description of pulsar instrumentation in general, describing the motivation for a new system, followed by a detailed account of its construction. Chapter 4 describes a new software package, designed to provide a generalised mechanism for pulsar data storage and analysis. This chapter addresses a number of problems associated with the previous generation of analytical code. Chapter 5 describes a detailed study of the shape of the mean pulse profile of PSR J1022+1001, showing that previously reported profile variations do not manifest when the pulsar is observed with CPSR2. Chapter 6 describes observations of the relativistic binary pulsar PSR J1141–6545. We demonstrate that the mean profile is broadening systematically and interpret this as evidence for geodetic precession, which should allow limited mapping of the emission cone in future years. New polarimetric profiles for PSR J0737–3039A are presented in Chapter 7, providing a reference against which future profile evolution may be detected, given that we expect this pulsar to precess on an even shorter timescale than PSR J1141–6545. The signature of post-Keplerian Shapiro delay is used to constrain the inclination angle of the J0737–3039 system. Finally, Chapter 8 describes high-precision timing observations of 15 millisecond pulsars over a period of 3 years. These sources were selected for their timing stability and several were modeled so well that the RMS timing residual was $\sim$200 ns. We measure several new parallax distances and present evidence of DM variations. Shapiro delay is used to place limits on the geometry and component masses of several binary systems in the selected list of sources. As a final test, we successfully time one very stable pulsar against the residuals of another, an essential first step towards constructing a large-scale pulsar timing array. In Chapter 9 we summarise our findings and speculate on the nature of future observing systems.

# Chapter 2

# An Introduction to Pulsar Timing

"Physics is imagination in a straight jacket."

*John Moffat*

A wealth of information is encoded in the signal from a radio pulsar. Because the spin rate is highly regular, any physical influence that is capable of varying the normal pulse rate can be detected and studied. The technique of pulsar timing is one of the most precise and widely applicable experimental methods in modern astronomy. Acceleration of either the source or the observer along their common line of sight is responsible for some of the largest timing perturbations, but more subtle relativistic effects can also be detected in rare cases. To time a pulsar, observers create a mathematical model that describes all physical perturbations acting on the spin rate of the source. When a pulse time of arrival (TOA) is measured, the predicted time is subtracted from the observed time to yield the so-called "timing residual". If the collected timing residuals display any deviation from random noise, there must be systematic errors either in the measurements or the model. Accurate measurements of pulse arrival times can be used to iteratively refine the timing model and measure more subtle effects. In this Chapter, we describe the way in which pulse arrival times are obtained and analysed, and introduce the parameters from which timing models are constructed.

## 2.1 Fundamental Concepts

Pulsar data is normally recorded along with an accurate time stamp, corresponding to the rising edge of the first phase bin at the beginning of the observation. This time is taken from the observatory clock, which must be related in some systematic way to a stable, global definition of time. Most pulsar timing is done using folded profiles that include many individual pulses. The addition of each pulse is done "in-

phase" (often at the time of the observation, by the instrument) and if the existing model is good, this addition process is accurate to well within a single rotation period. This allows the instrumental "trigger" time to be accurately modified so that it represents an epoch in the middle of the averaged pulse train, at the same phase. This epoch is treated as the "reference time" of the integration as a whole. The position of the pulse peak within the observed profile is then determined by cross-correlation with a standard template profile. This offset is added to the time stamp to yield the measured TOA.

At this point, we have what is called the "site arrival time" or SAT. The motion of the Earth in its orbit around the Sun and the perturbing influence of other bodies in the Solar system introduce Doppler shifts that are often the greatest single source of systematic deviation from a fixed pulse rate. Fortunately, these Solar system effects can be removed by converting our SAT to a "barycentric arrival time" or BAT. This is the time at which the pulse wavefront arrived at the centre of gravity of the Solar system (the barycentre), which is our best available approximation to a reference frame that is inertial with respect to the pulsar. Beyond our Solar system, dispersive propagation delays incurred as the pulses travel through the ISM introduce an offset that must be incorporated into the timing model. In addition, the apparent proper motion of the pulsar introduces an error in the estimated position that grows with time. If the pulsar orbits a binary companion, another set of corrections will be required to account for the orbital motion of the pulsar. Timing models can therefore have anywhere between approximately 5 and 15 free parameters.

In practice, pulsar astronomers are interested in the differences between the arrival times predicted by a given model and those actually observed (the timing residuals). If, for example, our knowledge of the position of a pulsar is incorrect, the timing model will predict the wrong Earth orbit travel time delay (see below), leaving a systematic and periodic "signature" in the timing residuals. Figure 2.1 shows the timing residuals of PSR J1909–3744 over a period of nearly three years. The model used to obtain the residuals was modified so that the position of the pulsar is in error by 0.1 mas in right ascension, introducing a sinusoidal signature.

The signature of an error in the model is the difference between the actual and applied correction for a given physical effect, or sometimes the presence of a completely unexpected effect that is not included in the model at all. The timing signature of an error in a well-understood parameter can be easily computed from the functional form of the model component. These signatures are what pulsar astronomers must identify and study. An unexplained signature may hold the key to new understanding. Once a set of arrival times and an initial model are available, the parameters of the timing model are adjusted (usually using a least-squares fitting procedure) to reduce the RMS residual to a minimum, thereby producing the most accurate description of the system.

The accuracy with which we can determine a given model parameter is governed by the amplitude of a particular timing signature, relative to the accuracy of our arrival time measurements. If the arrival time of a pulse can be determined observa-

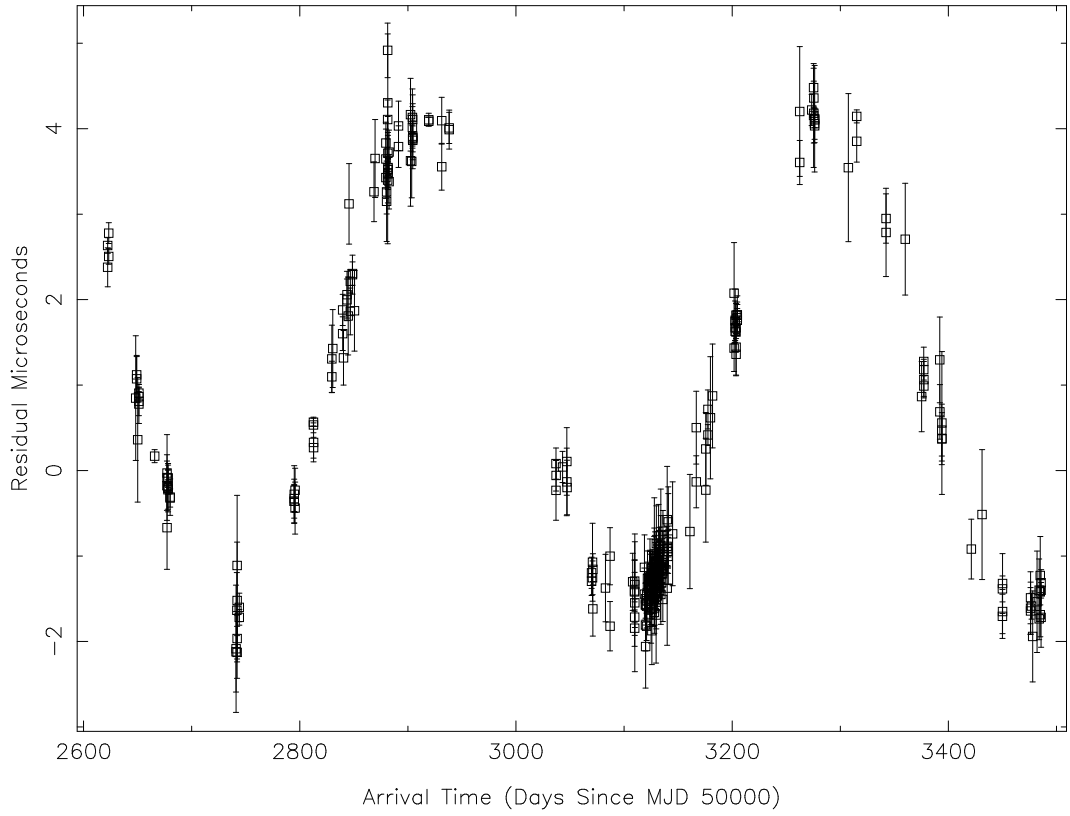


**Figure 2.1: This figure illustrates the systematic timing residual signature introduced by a 0.1 mas error in right ascension. Arrival times were obtained from PSR J1909–3744. The RMS of the fit increases by a factor of 10 when this position error is introduced.**

tionally to within 1 $\mu$s, it is possible to identify much more subtle signatures than if the experimental precision were measured in milliseconds. Unfortunately, many of the most interesting timing signatures are weak and often covariant with one other, making it essential to have a high density of accurately determined arrival times when trying to constrain the parameters of a timing model. Pulsars that can be timed precisely are therefore highly prized, especially if they reside in an interesting or unusual system. The intrinsic timing precision of a pulsar is governed primarily by its rotation period, though factors like flux density and pulse shape also contribute significantly. Because our analysis is performed with the aid of computers, pulse profiles are digitised at the telescope and recorded with a discrete number of uniformly spaced "bins" or points within a single rotation period. The time represented by each bin in the recorded profile is often longer than the RMS timing precision that can be achieved through careful analysis. In general, it is possible to determine the arrival time of a pulse to within approximately 10% of the width of an individual phase bin. We demonstrate in this thesis that some exceptional pulsars can be timed with an RMS residual better than 4% of the width of a single phase bin. If the pulsar rotates once every second, arrival times can typically be determined to within about 1 ms RMS. However, the millisecond pulsars have spin periods up to 1000 times smaller than the population of normal of "slow" pulsars and they can often be timed to $\mu$s (or better) precision.

## 2.2 Marking the Passage of Time

Throughout the course of history, humans have invented many ingenious ways of marking the passage of time. In the early days, motives may have been as simple as understanding the Earth's seasonal cycle so as to sow and harvest crops at the correct time for optimal yield. Sundials and larger monuments were constructed thousands of years ago by our ancestors to study and record the passage of the Sun through the sky, providing information about the length of the day and the timing of the Solstices and Equinoxes. The structure known as Stonehenge in Salisbury Plain, near Wiltshire in South-Western England, may have been used in just this way (Stone 1925). As craftsmanship developed, we began to build tall ships capable of crossing the oceans and the fledgling science of navigation drove the need for more precise time keeping, now using stars other than the Sun to study the rotation of the Earth. In the industrial age, we began to build mechanical clocks with complex systems of gears and developed the escapement mechanism that could convert (say) the swing of a pendulum into a quantitative measure of time. In the modern era, we count the natural oscillations of quartz crystals or even the frequency of radiation emitted during individual atomic energy level transitions. The ability to accurately record the passage of time is critical to much of modern physics; for example, we can measure the distance to the moon by accurately determining the round-trip elapsed time of laser pulses reflected from it.

Since the dawn of the industrial age, there have been several definitions of the standard unit of time. For thousands of years prior to this, time was determined by the rotation of the Earth on its axis, this being both critical to our survival and perhaps the most obvious example of a repetitive, easily observable natural cycle. The transition from day to night is however a relatively long-period cycle whose length is governed by the somewhat irregular rotation rate of the Earth. In addition, the ratio of daytime to nighttime can change by several hours over the course of a year, depending on the latitude of the observer. To measure time more precisely, we must abstract away from terrestrial rotation. The current SI unit of time is the second, defined as 9,192,631,770 cycles of the radiation corresponding to a certain hyperfine transition in the ground state of Cesium 133 (Bureau International des Poids et Mesures 1998). There are still modern timescales based on astronomical observations, but these are usually defined to suit a specific purpose and are often somehow locked to the fundamental scale of "atomic time" for long term stability.

The fundamental modern time standard is International Atomic Time (TAI), which is a weighted ensemble average of the time reported by multiple cesium atomic clocks maintained at laboratories throughout the world. The Bureau International des Poids et Mesures (BIPM)[1] compiles data from over 60 laboratories and publishes the definition of TAI in its monthly Circular T bulletin. Since the introduction of cesium fountain clocks in the last few years of the 20th century, the radiation frequency that defines TAI has been known to approximately 1 part in $10^{15}$. TAI depends only on the laws of physics and our ability to measure the frequency of a photon in a manner free of systematic errors. A single period of the radiation corresponding to the chosen cesium transition is a very short amount of time; the SI second was defined to be a certain large number (9,192,631,770) of these cycles, in order to make it almost the same as the "astronomical" second (1/86400 of a mean Solar day). TAI itself is "free floating", so any relative drift in the length of the mean Solar day since the SI second was defined can be measured independently.

In keeping with historical methods, the passage of time at a particular location on the surface of the Earth can be monitored by observing the passage of stars or distant astronomical radio sources across the sky. The modern incarnation of Solar time is known as Universal Time (UT0). This timescale is based on the rotational speed of a particular location on the surface of the Earth and is therefore contaminated by small irregularities in the Earth's rotation period. Although these errors are usually less than a tenth of a second on short timescales, the rotation rate of the Earth exhibits a secular trend of greater magnitude over many years. The International Earth Rotation and Reference Systems Service (IERS[2]) is responsible for monitoring the rotation rate of the Earth and publishes corrections in its monthly Bulletin B, accounting for any glitches. In this way, it is possible to define the UT1 timescale, which is corrected for irregularities in the Earth's spin, but still based upon the mean rotation rate. Thus TAI and UT1 provide accurate timescales that are based

[1] http://www.bipm.org/en/scientific/tai/
[2] http://www.iers.org/iers/

on two completely different physical phenomenon.

In order to unify these two schemes, the Coordinated Universal Time (UTC) scale was defined (International Telecommunication Union 2002). UTC takes advantage of the stability of TAI, and the immediate relevance of UT1 (the fact that it corresponds precisely to day and night on Earth). UTC follows TAI, but is adjusted in increments of exactly one second, once per 6 months (if required) at pre-defined times, so as to keep within 0.9 s of UT1. UTC is used throughout most of the world as the standard reference time. Variants of UTC are distributed by national standards laboratories in various countries and are kept as close as possible to the BIPM definition; measured UTC(local)-UTC(BIPM) offsets are often recorded for retrospective correction. Most institutions that require accurate time keeping can synchronise with electronically distributed references (via the Internet, satellite or radio systems), but in some cases (including pulsar astronomy), an accurate local clock is required.

Modern commercially-available clocks are based on stable electronic frequency standards. For example, many wristwatches use oscillations of a quartz crystal coupled to a feedback circuit as a means of keeping time. For any given clock, precision is limited by the stability of the mechanism used to generate periodic oscillations. Frequency stability must be measured on many different time-scales, it is quite possible for a frequency standard to remain stable within a few minutes, but drift systematically over larger time periods. It may also be the case that a mechanism with a very stable average frequency may exhibit random wanderings on shorter time scales. One must choose a mechanism whose dominant systematic errors are of least concern to the intended application. Cesium frequency standards have excellent long-term stability but lack immediate precision. Hydrogen masers exhibit the opposite behaviour, they are highly stable on timescales up to a day but must by synchronised with an external reference if long-term stability is required. Radio pulsar timing (and Very Long Baseline Interferometry) requires short-term precision in order to maintain phase connection across many pulse periods. In practice, most radio observatories operate a hydrogen maser that generates a pure reference frequency. The local oscillators used for down-conversion (see Chapter 3) are phase-locked to the maser, which is also used to drive the local clock. This in turn is kept synchronised with UTC by regular reference to time signals distributed via the Global Positioning System or the local national standards office. The GPS uses its own version of UTC that does not include any leap seconds introduced since its launch in 1980, but the offset can easily be accounted for and the correction factor is included in the GPS data stream.

### 2.2.1 Instrumental Time Assignment

In the case of our baseband recording and coherent dedispserion system (CPSR2) at the Parkes radio telescope, data acquisition is triggered by the station 1 Pulse Per Second (PPS) signal, which is provided by the Mk VI clock that is in turn

driven by a Hydrogen maser frequency standard. The difference between the Parkes timescale UTC(PKS) and UTC(GPS) is monitored regularly by the Totally Accurate Clock (TAC), consisting of a GPS receiver interfaced with a standard PC. The offset log is used to convert site arrival times (SATs) into UTC(GPS) arrival times. The GPS system itself is not synchronised with the BIPM definition of UTC, but with the United States Naval Observatory (USNO) version (which is kept within a few nanoseconds of the National Institute of Standards and Technology variant, UTC(NIST)). The GPS signal and BIPM Circular T bulletin provide all the information required to perform a 2-step conversion from UTC(GPS) to UTC(NIST) and then to UTC(BIPM)[3], from which we can convert back to TAI.

### 2.2.2 Dynamical Time

In addition to the astronomical and atomic standards described above, two convenient "dynamical" timescales are defined and used when dealing with planetary motion. The Terrestrial Dynamical Time scale (originally TDT, but now known simply as TT) defines one day as 86400 SI seconds (recall that the SI second is directly related to the tick rate of the atomic clocks used to define TAI). In addition, it must be specified that TT refers to the passage of time on an equipotential surface of Earth's gravity field that most closely matches global mean sea level. This is necessary because Einstein's General theory of relativity predicts that the passage of time is slowed in the presence of a gravitational field, by an amount that depends on the strength of the field. Experiments using satellites and other bodies within the Solar system have verified this prediction to a high level of confidence. Despite the seeming complexity of the definition, TT simply leads TAI by a constant offset of 32.184 seconds. The major source of uncertainty arises when we move our focus away from the Earth. Barycentric Dynamical Time (TDB) was defined to be a uniform timescale in the Solar system barycentre frame, equivalent to TT after removing several periodic variations. Unfortunately, the intended properties of this timescale turn out to be physically impossible. In practice, most code designed to perform the conversion to TDB actually produces a subtly different timescale known as $T_{\rm eph}$, which is the argument given to ephemerides in order to obtain the positions of the major bodies in the Solar system. This is self-consistent (pulsar timing works!) but $T_{\rm eph}$ neglects to include a secular drift due to the effects of relativistic gravity in the Solar system. This means that the passage of time under $T_{\rm eph}$ does not quite correspond to the SI second.

In order to isolate physical effects intrinsic to a pulsar system (and to simplify the mathematics of timing models), pulse arrival times should be analysed in an inertial reference frame. The motion of the Solar system barycentre (SSB) is the best local approximation to such a frame. It should be noted that the barycentre is not quite inertial with respect to other stars in the Milky Way, relative acceleration could be caused by differential Galactic rotation (usually a small effect), acceleration

[3]http://www.atnf.csiro.au/research/pulsar/psr/timing/newclkcorr.html

towards the plane of the galaxy or acceleration caused by other nearby stars. This last contribution can dominate if the pulsar is inside a globular cluster. The first step in any timing analysis is to transform the SATs recorded at an observatory to the corresponding BATs. The conversion from SAT to BAT is effectively a transformation from the observatory's local approximation of UTC to UTC(BIPM) and then to TAI (and TT by addition of a constant offset). The final conversion to TDB (or $\mathrm{T}_{eph}$) requires detailed knowledge of Solar system dynamics.

## 2.3 Shifting to the Barycentre Reference Frame

In this section, we follow the conventions used by Taylor & Weisberg (1989) to describe the conversion of a TT site arrival time $t$ to a barycentric (TDB) arrival time $t_{\mathrm{b}}$. Equation 2.1 describes the four main correction terms. The reader should note that many of the terms in Eq. 2.1 are functions of external parameters and must be computed individually for each pulsar.

$$t_b = t - \frac{D}{f^2} + \Delta_{R_\odot} + \Delta_{E_\odot} + \Delta_{S_\odot} \tag{2.1}$$

Here, $D$ is the dispersion constant (related to the dispersion measure (DM) by the expression $D = DM/2.41 \times 10^{-16}\,\mathrm{s\,Hz^2}$) and $f$ is the centre frequency of the observation (in the barycentre frame, given in Hz). This term describes the delay induced by propagation through the ISM. The other three terms are respectively the Solar system Roemer, Einstein and Shapiro delays. The $\odot$ subscript indicates that the delay results from within the Solar system; as we will see later each of these three terms has an analogue that arises due to orbital motion (if any) in the pulsar system.

The fact that the Earth orbits the Sun means that the position of the Earth with respect to the Solar system barycentre changes systematically (and cyclically) over the course of a year. This must be incorporated into the TDB transformation as it introduces a delay (or an advance) due to the modified distance the signal must travel to reach the telescope. This is known as Roemer delay ($\Delta_{R_\odot}$) and is expressed mathematically as:

$$\Delta_{R_\odot} = \frac{\vec{r} \cdot \hat{\vec{n}}}{c} \tag{2.2}$$

Here, $\vec{r}$ is the vector extending from the Solar system barycentre to the telescope at the moment the pulse arrives, $\hat{\vec{n}}$ is the unit vector extending from the SSB in the direction of the pulsar, and $c$ is the speed of light. The SAT must be modified by the time light takes to travel the distance from the SSB to the telescope, projected onto the line of sight to the pulsar. The near-circular nature of the Earth's orbit means that the required modification changes approximately sinusoidally, with a period of one year. The relative phase and magnitude of this sinusoidal function are

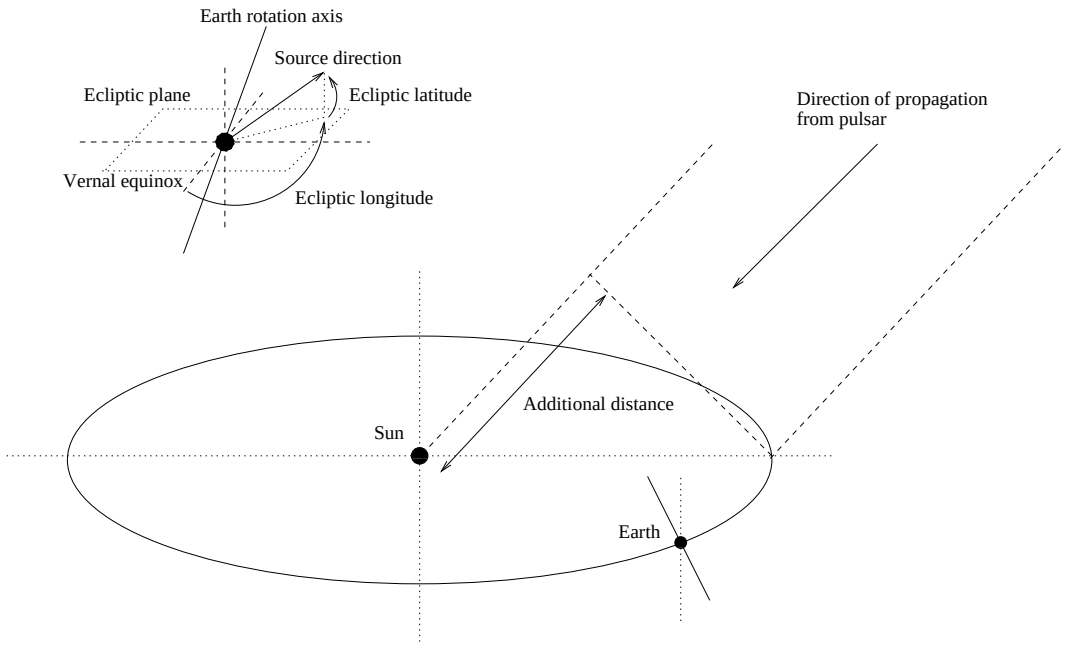


**Figure 2.2: This figure illustrates the nature of the annually-varying pulse propagation delay introduced by the Earth's orbit about the Sun. The amplitude and phase of the sinusoidal delay function depend on the ecliptic longitude (angular distance from the vernal equinox) and ecliptic latitude (angular distance from the ecliptic plane) of the source. The light travel time across 1 Astronomical Unit (AU, the mean Earth-Sun distance) is approximately 500 s. This limits the maximum amplitude of the delay, which would be experienced by a source that lies in the ecliptic plane. Higher latitude sources have a smaller delay amplitude.**

determined by the sky coordinates of the pulsar. Figure 2.2 illustrates the geometry of the Sun-Earth system and describes an ecliptic coordinate scheme that can be used to specify the position of a celestial object.

Let the ecliptic longitude be $\lambda$ and the ecliptic latitude be $\beta$. These angles define the location of a point on the celestial sphere and are (to a good approximation) fixed in time for a particular object. The projection onto the ecliptic plane of a vector with length 500 lt-s extending from the barycentre in the direction of the pulsar will approximate the amplitude of the periodic delay induced by Earth's position with respect to the Sun, if we assume that the SSB coincides with the centre of the Sun. This amplitude is determined by $\cos(\beta)$ and the phase by $\lambda$. The cumulative delay can be calculated using Eq. 2.3.

$$\tau_d = 500 \cos(\beta) \cos(\phi + \lambda) \tag{2.3}$$

Here, $\tau_d$ is the geometric Roemer delay in seconds and $\phi$ is the angular separation of the Earth's position from the vernal equinox, which is a fiducial point that lies on the line where the ecliptic plane crosses the Earth's equatorial plane. Our ability to predict pulse arrival times depends of our knowledge of the position of the pulsar,

which defines the two angles required to predict the exact amplitude and phase of the Roemer delay.

Equation 2.3 is a first-order correction to compensate for the motion of the Earth in its orbit, but there are many additional perturbative effects that should be taken into account. The orbit of the Earth is not quite circular and the Sun itself is not at the centre of mass of the Solar system. The mass of the outer planets shifts the barycentre to a point just outside the surface of the Sun. The Sun itself is very near (compared to other stars) and its gravitational field introduces additional relativistic delays (the Einstein and Shapiro delays) not predicted by Newtonian approximations. The Einstein delay term takes into account any time dilation and gravitational redshift arising from motion within the Solar system. It is computed as the integral over time of a summation over all the significantly massive bodies in the Solar system and thus depends highly on our knowledge of planetary motion. The Shapiro delay term is General relativistic in origin and arises because light takes longer to travel through a region of space containing a strong gravitational field. If our line of sight to the pulsar passes close to the Sun, the magnitude of the Shapiro delay can be of order $120\,\mu$s. For more information, see Taylor & Weisberg (1989) and references therein.

Transformation to the barycentre frame requires accurate knowledge of the position and motion of the most massive bodies in the Solar system. This information is obtained from a planetary ephemeris supplied by NASA's Jet Propulsion Laboratory (JPL). Pulsar astronomers have historically used the DE200 ephemeris but this is slowly being superseded by the more accurate DE405 version. These ephemerides can be obtained from the JPL Solar System Dynamics Group[4].

## 2.4 Fundamental Timing Model Parameters

TOAs from a solitary radio pulsar can be accurately predicted by a model that contains relatively few parameters. To compute the precise geometric Roemer delay term, the position of the pulsar must be well known. It follows that the proper motion of the pulsar (or at least the component transverse to our line of sight) must be included in the model. If proper motion is not taken into account, the assumed pulsar position at the current epoch will be in error by an amount that increases with temporal separation from the epoch at which the initial position was specified. If the pulsar is very close to the Solar system ($\sim$1 kpc), wavefronts emitted by the pulsar can retain enough curvature (on the scale of the Earth's orbit) to influence the predicted arrival times when measured over the course of a year. The electron column density of the ISM introduces an absolute delay; the dispersion measure (DM) of the pulsar must therefore be included in the timing model.

Pulsar timing models are usually expressed in terms of pulse phase, which is a measure of the location of the pulse within a single period. If the pulsar spins at a

[4]http://ssd.jpl.nasa.gov/eph_info.html

perfectly constant rate, there would be a simple scale factor separating time in the pulsar rest frame and pulse phase. Any perturbation introduced by an irregularity in the spin rate causes the phase to drift in time when compared to a stationary fiducial point. The timing model attempts to characterise complexities in the relationship between pulse phase and time, compensating for any non-random drifts. We expect the spin rate of a pulsar to slow down over time as rotational kinetic energy is converted into electromagnetic radiation. If we make no assumptions about the physics involved, we can write down a simple Taylor series expansion describing pulse phase as a function of time (Eq. 2.4).

$$\phi(t - t_0) = \nu(t - t_0) + \frac{\dot{\nu}(t - t_0)^2}{2} + \frac{\ddot{\nu}(t - t_0)^3}{6} + \cdots \tag{2.4}$$

Here, $\nu$ is the pulsar spin frequency (1/period) in Hz, $t$ is barycentric time and $t_0$ is a reference arrival time at which the pulse phase is taken to be zero. Pulse phase can be computed for an arbitrary time $t$; integer values indicate the arrival of the next pulse. Pulse phase can be thought of as how far the lighthouse beam has swung since the last pulse was observed. A phase of 0.5 means the beam is on the opposite side of its swing from our perspective. Equation 2.4 does not include any correction for the dispersive delay introduced during passage through the ISM and is therefore considered to be the *infinite radio frequency* case; this terminology can be understood by recalling Eq. 2.1, which shows that dispersive delays decrease with the square of the observed radio frequency.

Each term in the Taylor series has a different error signature in a given set of timing residuals. If the spin frequency is incorrect, the residuals will increase linearly with time as the predicted phase drifts further from the measured phase. An error in the first derivative of the spin frequency will introduce a quadratic signature and so on. Using a combination of all the parameters described above, we can construct a timing model suitable for solitary pulsars. The parameters required are summarised in Table 2.1. It is important to note that the signature associated with each of these parameters can vary greatly in amplitude (and therefore significance) depending on the nature and location of the source.

## 2.5 Incorporating Binary Motion

When a pulsar resides in a binary system, a number of additional parameters must be included in its timing model to describe the orbital motion. These additional perturbations are analogous to the corrections required to compensate for the motion of the Earth in orbit about the Sun, except that they act to disturb the rest frame of the source instead of the receiver. Often, Newtonian gravitation is sufficient to describe the motion of a binary pulsar at the level required to produce white-noise residuals. Given sufficient observational precision, it is sometimes possible to detect General relativistic effects in a pulsar binary system, analogous to the Einstein and

| Parameter | Description |
|---|---|
| $P$ | Pulsar spin period. |
| $\dot{P}$ | Spin period first derivative with respect to time. |
| DM | Dispersion measure along the line of sight. |
| RA ($\alpha$) | Right ascension. Can substitute ecliptic longitude. |
| DEC ($\delta$) | Declination. Can substitute ecliptic latitude. |
| PMRA | Proper motion in RA (or ecliptic longitude). |
| PMDEC | Proper motion in DEC (or ecliptic latitude). |
| PX | Parallax. |

**Table 2.1: Parameters that can contribute to the timing model of a solitary pulsar.**

| Parameter | Description |
|---|---|
| $x$ | Projected semi-major axis ($a\sin(i)$). |
| $e$ | Eccentricity. |
| $\omega$ | Longitude of periastron. |
| $T_0$ | Epoch of periastron. |
| $P_\mathrm{b}$ | Orbital period. |

**Table 2.2: Timing model parameters introduced by a Keplerian description of two masses that orbit a common centre of mass.**

Shapiro delays within our Solar system. The compactness of neutron stars means that the gravitational field strength experienced by binary pulsars can be far greater than that experienced by most other bodies. This is especially true in binaries where both members are neutron stars.

A simple, Keplerian description of orbital motion in a pulsar binary system introduces five additional free parameters to any pulsar timing model; these are summarised in Table 2.2.

The semi-major axis of the elliptical orbit is denoted $a$, the basic timing signature of binary motion depends on the projection of this distance onto the line of sight, $a\sin(i)$. Here, $i$ is the inclination angle of the binary system (the angle between the orbital plane and the plane of the sky). A system with an inclination angle of $90^o$ is seen edge-on. Eccentricity $e$ is a dimensionless measure of deviation from circularity. Orbits with $e \sim 0$ are very nearly circular; as eccentricity tends towards unity, the orbit becomes more elongated. The longitude of periastron $\omega$ is the angle that the point of closest approach makes with respect to the line of nodes, which is the intersection between the plane of the sky (the normal to which points along our line of sight) and the plane of the orbit. The longitude and epoch, $T_0$, of periastron define a fiducial point in the orbit.

If the binary orbit is assumed to be Keplerian, it is not possible to unambiguously extract the individual masses of the pulsar and its companion from pulse timing

| Parameter | Description |
|---|---|
| $\gamma$ | Gravitational redshift and transverse Doppler shift. |
| $r$ | Range of the Shapiro delay. |
| $s$ | Shape of the Shapiro delay. |
| $\dot{\omega}$ | Advance of periastron longitude. |
| $\dot{P}_{\rm b}$ | Orbital period time derivative. |

**Table 2.3: Theory independent post-Keplerian timing parameters, used to describe the departure of a pulsar orbit from the solution predicted by Newtonian gravitation.**

experiments alone. It is however possible to re-write Kepler's third law of orbital motion (which states that the square of the orbital period is proportional to the cube of the semi-major axis of the elliptical path) in a form that allows observational constraints to be placed on the mass of the pulsar and companion and the inclination angle of the system. Equation 2.5 is known as the "mass function":

$$f(m_{\rm p}, m_{\rm c}) = \frac{m_{\rm c}^3 \sin^3 i}{(m_{\rm p} + m_{\rm c})^2} = \frac{4\pi^2}{G} \frac{a^3 \sin^3 i}{{P_{\rm b}}^2}. \tag{2.5}$$

Here, $m_{\rm p}$ is the pulsar mass and $m_{\rm c}$ is the companion mass, all other parameters are as described above. Pulse timing allows us to measure $x = a \sin(i)$ and $P_{\rm b}$ (and thus the value of the mass function) very accurately. This allows us to determine the last of the three parameters $m_{\rm p}$, $m_{\rm c}$ and $i$ if we know any two of the others; however in most cases we cannot directly measure any of them. Fortunately, there are a number of ways to determine the missing information.

Note that the parameters presented in Table 2.2 do not fully describe the three-dimensional orientation of a binary orbit. Timing alone cannot distinguish the orientation of the line of nodes because there is no Doppler signature introduced by rotation about the line of sight. The final parameter needed to describe the three-dimensional orientation of a Keplerian orbit is the longitude of the ascending node, denoted $\Omega$. It can be measured through timing only if circumstances allow the detection of a very subtle signature known as annual-orbital parallax (Kopeikin 1995).

When the orbit is compact, with a period less than a few days, the basic Keplerian parameters presented in Table 2.2 may not properly characterise the motion. Damour & Deruelle (1986) describe an additional five post-Keplerian parameters that can be included as free parameters in a pulse timing model (Table 2.3). Any theory of gravity must provide mathematical expressions for these five parameters in terms of the mass of the pulsar and its companion. Newtonian gravitation predicts that all of these parameters are zero, but we have known for many years that it is only accurate in the weak-field limit.

Equations 2.6 to 2.10 describe the functional form of the five post-Keplerian parameters in the formalism of General relativity (Damour & Taylor 1992; Taylor

& Weisberg 1989; Damour & Deruelle 1986).

$$\gamma = e\left(\frac{\mathrm{P_b}}{2\pi}\right)^{1/3} \frac{(GM_\odot)^{2/3}}{c^2} M^{-4/3} m_2(m_1 + 2m_2) \tag{2.6}$$

$$r = \frac{GM_\odot m_2}{c^3} \tag{2.7}$$

$$s = x\left(\frac{\mathrm{P_b}}{2\pi}\right)^{-2/3} \frac{c}{(GM_\odot)^{1/3}} \frac{M^{2/3}}{m_2} \tag{2.8}$$

$$\dot{\omega} = 3\left(\frac{\mathrm{P_b}}{2\pi}\right)^{-5/3} \frac{(GM_\odot)^{2/3}}{c^2} M^{2/3}(1-e^2)^{-1} \tag{2.9}$$

$$\dot{\mathrm{P_b}} = \frac{-192\pi}{5}\left(\frac{\mathrm{P_b}}{2\pi}\right)^{-5/3}\left(1 + \frac{73}{24}e^2 + \frac{37}{96}e^4\right)(1-e^2)^{-7/2}\frac{G^{5/3}}{c^5} m_1 m_2 M^{-1/3} \tag{2.10}$$

Measurements of post-Keplerian parameters can be used to test the predictions of various theories of gravity. Accurate measurements of three or more post-Keplerian parameters "over-determines" the system and allows the self-consistency of the theory to be tested. Experiments of this kind have been performed on the double pulsar system, J0737–3039 (Lyne et al. 2004), PSR B1913+16 (Taylor & Weisberg 1982; Taylor & Weisberg 1989) and PSR B1534+12 (Stairs et al. 2002). All tests have verified the predictions of General relativity to a high level of confidence. If only two of the post-Keplerian parameters are measurable, it is possible to assume that General relativity is the correct theory of gravity and derive unambiguous values for the mass of the pulsar and its companion and the inclination angle of the system by combining two of Eqs. 2.6 to 2.10 with Eq. 2.5. In this way, astronomers can study the mass of neutron stars formed via different evolutionary paths.

Although each of the parameters that make up a pulsar timing model will have a distinct signature, some of these signatures can appear very similar to one another. This reduces the effectiveness of least-squares fitting methods because there is no longer a unique solution with minimum $\chi^2$. The ability to partially absorb the effect of an error in one parameter by modifying a different parameter is known as "covariance". This is most important in studies of relativistic binary pulsars because it can reduce the effective amplitude of an already subtle timing signature. It is also possible for a model parameter to absorb an unrelated physical effect. For example, the motion of a pulsar across the sky will introduce a small apparent period and binary period derivative. This effect cannot be distinguished from intrinsic spin-down or General relativistic orbital decay without knowledge of the pulsar's distance and space velocity (Shklovskii 1970). Bell & Bailes (1996) inverted the problem and asserted that the detection of an orbital period derivative in a clearly non-relativistic system might enable astronomers to constrain the distance to the pulsar.

Binary motion can also be described using two Laplace-Lagrange parameters, $\kappa$ & $\eta$, and the time of ascending node passage, $T_{\rm asc}$. These three parameters are related to the more familiar $e$, $\omega$ and $\mathrm{T}_0$ by Eqs. 2.11 to 2.13. If the binary system is very nearly circular, it is difficult to constrain $e$, $\omega$ and $T_0$ because they refer to a fiducial point that is not well defined (Lange et al. 2001).

$$e = \sqrt{\kappa^2 + \eta^2} \tag{2.11}$$

$$\omega = \arctan(\eta/\kappa) \tag{2.12}$$

$$T_0 = T_{\rm asc} + \frac{P_b}{2\pi}\arctan(\eta/\kappa) \tag{2.13}$$

## 2.6 Measuring Arrival Times

Extracting an arrival time from a folded profile requires two things; an accurate time stamp corresponding to a particular phase in the folded profile and a standard template profile to provide a fiducial phase against which all other observations can be compared. Pulsar instruments must perform the difficult task of folding the incoming signal synchronously with the pulse period. This requires precise knowledge of the pulse period at the time of observation. The pulse period is divided into a discrete number of time bins, which are assigned data one after the other as it is obtained from the telescope. When the last bin is reached, the recorder wraps back to the first bin, in this way building up a folded profile.

Because the period of a pulsar is susceptible to small, smooth changes that can occur on time-scales comparable to the length of an observation, accurate folding requires a good timing model to provide period predictions. However, because a timing model is designed to provide a complete description of the system, it can take significant computational power to solve the model for the pulsar period at a particular time. Unfortunately, real-time recording systems may not be able to spare this amount of processing time. For greater efficiency, instruments are usually given a much simplified model, based on a polynomial with 15 or fewer coefficients. Before observing starts, a polynomial is fitted to the timing model for the epoch immediately surrounding the intended observing time. This polynomial is only an approximation, but provided care is taken to select the size and length of the polynomial appropriately, the period estimate it provides will be accurate to the level required. The so-called "polyco", or sequence of polynomial coefficients is often stored in the resulting data archive, along with the timing model. This provides a complete record of all operations performed on the data.

The epoch associated with a pulsar archive must correspond to a particular pulse phase. This correspondence is specific to the instrument and file format used. Most instruments are triggered on an integer second boundary, which will not normally

correspond to pulse phase zero. The simplest method is to record this time in the header and let it correspond to the rising edge of the first bin. The peak of the mean profile then appears in a random place within the period. In order to ensure the peak of the profile appears in the same place every observation, it is possible to use the polynomial to find out what phase the first sample will correspond to and place it in the correct bin instead of arbitrarily assigning it to the first bin. Both methods allow for accurate timing, but the later provides phase consistency when the timing model is good. Regardless of the method used to store the start time, the TOA assigned to a fully-integrated observation is normally specified near the midpoint to minimise the detrimental effects of any phase drift.

Once the epoch of an integration is known, all that remains is to compute the offset of the observed pulse relative to a reference phase. This is done by cross-correlating the mean observed profile with a standard template profile. The standard template is a reference profile that ideally matches the detailed morphology of the mean pulse but contains no noise. The phase corresponding to the peak of the correlation function is the offset. Provided the same standard template profile is used to compute every TOA, the calculated time offsets will refer to the same fiducial phase. The TOA is simply the sum of the phase offset and the epoch of the archive.

### 2.6.1 Constructing Standard Template Profiles

In practice, it is difficult to build a standard template profile that matches the mean pulse morphology and is free of noise. The easiest way to match the shape of the mean profile is to add together all the observations of a given source, creating a high S/N mean profile from the data itself. Although a template created in this way will faithfully match the morphology of the mean pulse, it retains all the noise in the observations. This has important implications for timing experiments, which are discussed further in Appendix A. It is also possible to synthesise a synthetic, analytic standard template profile based on a mathematical function that closely matches the morphology of the mean pulse. Kramer et al. (1999) describe a technique that involves synthesising a template profile using a superposition of several Gaussian components with different widths and amplitudes. Kramer et al. (1999) even argue that it is possible to let some of the parameters (the relative amplitude for example) that define an analytic template vary when fitting, in order to compensate for profile evolution or variability.

### 2.6.2 Computing Relative Profile Shifts

Once a suitable standard template has been created, it is cross-correlated with an observed profile to determine the phase offset. Performing a simple discrete cross correlation with a number of lags equal to the number of profile bins gives a rough estimate of the phase offset, but it is only as precise as the width of a single bin. Often this is not sufficient for precision timing and it is possible to obtain sub-bin

precision by interpolating over the discrete bins. There are a number of ways to do this, Taylor (1992) describes a particularly clever algorithm that makes use of the Fourier shift theorem. This algorithm and a complementary scheme that operates in the time domain are described in more detail in Appendix A.

## 2.7 Model Fitting

Once the arrival times associated with a series of observation have been computed, they are fitted to a timing model using the standard pulsar timing package, TEMPO[5]. The original version of this program was written by R. N. Manchester and W. L. Peters in 1970 at the National Radio Astronomy Observatory (NRAO) in the United States of America. It was used to obtain the periods, period derivatives and positions of several isolated pulsars observed with the 92 m transit telescope located near Green Bank, West Virginia (Manchester & Peters 1972). When the first binary pulsar was discovered in 1974, J. H. Taylor added a number of binary models to the code so that orbital parameters could be included in the fitting process. Since then, several other people have contributed to the development and maintenance of this program, including D. J. Nice, J. M. Weisberg and N. Wex. TEMPO is now used at many observatories and institutions throughout the world. More recently, G. B. Hobbs and R. T. Edwards developed a new pulsar timing package (named TEMPO2) written in the C programming language (TEMPO is only available in FORTRAN). The new program is designed to correct for small systematic errors at the level of approximately 10-100 ns, matching the timing precision available with modern instruments[6]. For the bulk of this thesis, TEMPO2 was still in the experimental stages of development and was therefore not used.

TEMPO accepts a file containing SATs and their corresponding uncertainty, radio frequency and observatory site code. The recording site must be known in order to compensate for the light travel time across the diameter of the Earth. TEMPO also requires several clock correction files and either the DE200 or DE405 Solar system ephemeris, which are used to transform SATs into BATs. Finally, a pulsar timing model must be specified, including a list of parameters and their assumed values, along with flags that specify whether or not each parameter should be allowed to vary when fitting. The results of the fit are written to a number of files, including a binary dump of the raw residuals (resid2.tmp) and a neatly formatted text file that contains the new parameter values (tempo.lis) and their uncertainty. By default, TEMPO ignores the uncertainty in each SAT, fitting over all points with equal weight. Prepending "MODE 1" to the file containing the SATs forces TEMPO to weight each point by its uncertainty. In this mode, the $\chi^2$ of the fit can be computed and used to estimate $1\sigma$ uncertainties in each of the fitted parameters. Historically, it is rare for a pulsar timing model to fit a sequence of SATs perfectly. The reported

[5]http://www.atnf.csiro.au/research/pulsar/tempo
[6]http://www.atnf.csiro.au/research/pulsar/psrtime/tempo2/

uncertainty only represents the true $1\sigma$ error if the reduced $\chi^2$ of the fit is equal to unity. This can be forced for any data set by applying a uniform scaling to the TOA uncertainties and increasing the scaling factor until the reduced $\chi^2$ decreases to unity. Often, the uncertainty in each SAT must be multiplied by a factor of $\sim$2.

TEMPO has been in regular use for three decades and it corrects for many perturbations down to the level of a few tens of nanoseconds. Some of the largest uncertainties are due to a simplistic model of the electron content of the Solar wind and a failure to correctly account for the polar motion of the Earth. The use of TDB as the reference timescale is also contrary to IAU recommendations. One of the primary motivations for writing TEMPO2 was to correct these shortcomings and provide the user with the opportunity to carefully monitor the effect of each correction term that is applied to the TOAs. The new program will provide many other useful features including free-format input files and a number of built-in graphical user interfaces. We encourage future students of pulsar timing to experiment with the new features provided by TEMPO2.

# Chapter 3

# A New Baseband Recorder and Supercomputer

"Three years later, he left the Astronomy Department without a degree, and with nothing to show for his labors except six hundred dollars in his bank account and a staggeringly comprehensive knowledge of UNIX."

*Neal Stephenson – Cryptonomicon*

The study of millisecond pulsars is technically demanding; sophisticated instruments and techniques are required to achieve the necessary time and frequency resolution. Radio pulsars are often very weak, driving a need for large bandwidths, dual polarisations and long integration times. In order to best meet these requirements, pulsar observing systems have kept pace with rapid developments in technology over the last few decades. However, the consumer computing revolution has led to the wide availability of inexpensive, general purpose processors that can be cost-effectively applied to the computationally intensive task of real-time astronomical data processing. This chapter describes the design and implementation of the world's first wide-band on-line baseband recorder and coherent dedispersion system, the Caltech Parkes Swinburne Recorder Mk II (CPSR2).

## 3.1 Radio Telescope Sensitivity

Although beamed pulsar emission is thought to be a strong, coherent process, the radiation generated by pulsars must travel hundreds or thousands of parsecs to reach an observer on Earth. The presence of electrons in the ISM causes broadband pulses to scatter and disperse. To detect these pulses we must use a sensitive radio telescope and (for all but the brightest) correct for dispersion smearing.

The sensitivity of a radio telescope depends on a number of parameters, some of which are intrinsic to the design of the antenna (like the diameter of the reflecting surface) and others that are governed by the hardware and electronics used to receive and process the signal. Every radio system is contaminated by some amount of intrinsic noise, either due to the random motion of charged particles in the electronics, background astronomical sources, natural terrestrial interference (like thunderstorms) or interference from artificial broadcasts.

One of the most basic descriptive parameters of any observation is the signal-to-noise ratio, S/N. This is the relative amplitude of the observed signal (whatever form it may take) when compared to the background noise power in the system. It is a measure of how distinctive a particular signal appears to be. High S/N is desirable because it helps to increase the significance of a detection and allows more precise measurements of the properties of the signal. Equation 3.1 is known as the "radiometer equation", which is used to determine the S/N expected for a steady continuum source using a particular hardware configuration. It describes the functional dependence of S/N on various fundamental parameters.

$$S/N \propto \frac{G\sqrt{BtN}}{T_{\rm sys} + T_{\rm sky}} \tag{3.1}$$

Here, $G$ is the telescope gain, which is determined by the effective aperture (collecting area) of the antenna. $B$ is the bandwidth of the observation. The amount of time spent observing the source (integration time) is denoted by $t$ and the number of orthogonal polarisations received (either 1 or 2) is denoted by $N$. The total system noise is normally referred to as an equivalent blackbody temperature, the temperature a source in the field of view would have to be in order to produce the same amount of power as the background noise level. There are two components, $T_{\rm sys}$ is the noise intrinsic to the telescope hardware itself and $T_{\rm sky}$ is the noise entering the telescope from background sources in the sky that are not the object of study.

Equation 3.1 assumes that the source flux is steady in time. When observing a pulsed source, the radiometer equation must be modified slightly to take into account the pulse "duty cycle", the fraction of a pulse period in which the radio emission from the source is "switched on". Some pulsars have wide characteristic profiles, their signal is present for a large fraction of a rotation period. Other pulsars, like the recently discovered PSR J1909–3744 (Jacoby et al. 2003), have very narrow pulses that only turn on for a small fraction of the period. Whether or not this corresponds to variations in intrinsic beam width, or simply different cross-sections through similar emission cones remains an open question. Either way, we must apply a multiplicative correction to obtain the modified radiometer equation (Dewey et al. 1985) shown below.

$$S/N \propto \frac{G\sqrt{BtN}}{T_{\rm sys} + T_{\rm sky}} \cdot \sqrt{\frac{(P-w)}{w}} \tag{3.2}$$

Here, $P$ is the pulse period and $w$ is the effective pulse width (both in seconds). Note that for $w < P/2$, the pulsed nature of the source is an advantage.

Because most pulsars have very small flux densities ($\sim$1 mJy $\equiv 10^{-29}\,\mathrm{W\,m^{-2}\,Hz^{-1}}$), it is often necessary to use long integrations to obtain high S/N. Pulsar signal averaging is complicated by the fact that the fundamental time resolution of any instrument must be high in order to minimise the effective pulse width. Averaging must therefore be performed after detection and care must be taken to add in phase with the spin period of the pulsar. Stacking $n$ observed pulses increases the S/N by $\sqrt{n}$, as the signal in the off-pulse region does not add in phase. This averaging procedure is known as "folding".

## 3.2 Introduction to Pulsar Instrumentation

Advances in instrumentation are driven both by the need for higher sensitivity and a desire to eliminate systematic errors. These two demands are often mutually exclusive. Pulsar observers must be especially conscious of instrumental deficiencies that smear or distort the folded pulse profile. All pulses will be unavoidably smeared due to multi-path propagation through the ISM, but care must be taken to ensure that this resolution limit is reached.

Additional smearing can occur due to insufficient time resolution and dispersion in the largest incoherent section of bandwidth recorded (the width of a single spectrometer channel). Unfortunately, incoherent detection is limited by an inverse relationship between time and frequency resolution. This is not fundamental and is due to the fact that incoherent detection discards information about the phase of the signal. Coherent dedispersion (see below) can be used to circumvent the frequency / time resolution trade-off by eliminating dispersion smearing without the need to maintain frequency resolution (at the cost of greatly increased data rates).

The time between each digital sample is related to the number of "phase bins" into which a given pulse can be divided. A single pulse period is referred to as one unit of pulse phase, and the number of discrete points that are used to describe the pulse shape within a single period is referred to as the number of phase "bins". A detector that can rapidly sample the signal is capable of dividing a single pulse phase unit into a greater number of bins. This does, of course, reduce the flux in any given bin and can decrease the S/N of the profile if taken too far. Optimal binning is just sufficient to resolve the finest structure in the profile.

If the observed bandwidth is sampled incoherently, dispersive smearing is still present within each spectral channel. The actual amount of distortion this causes depends on the DM, width, intrinsic profile shape, pulse period and observing frequency. If the smearing time is smaller than a single phase bin, the effect is negligible. However, if the period is short and the pulse profile contains sharp features, dispersion smearing can significantly distort the observed profile. Figure 3.1 shows two mean profiles of the millisecond pulsar PSR J1600–3053, observed with two different

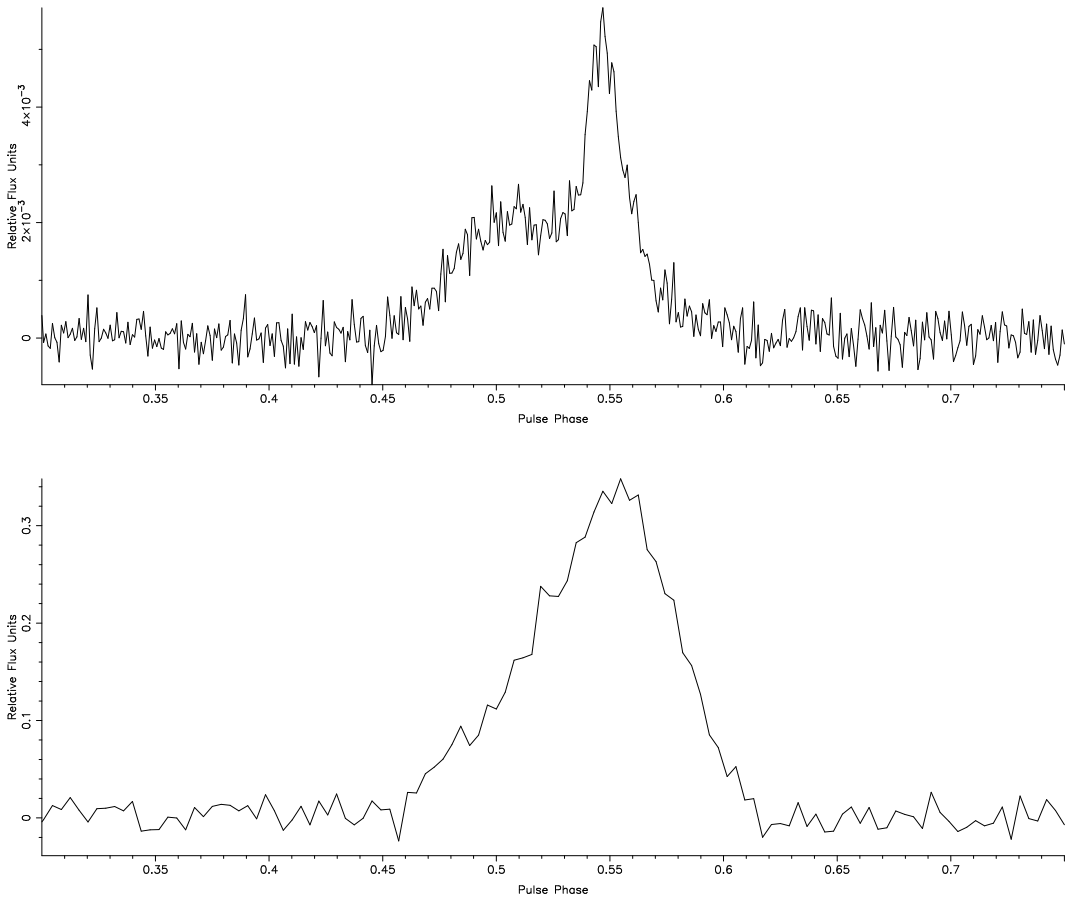


**Figure 3.1: This figure shows two integrated profiles of the millisecond pulsar PSR J1600–3053. The upper profile was observed with a coherent dedispersion system (CPSR2) using 1024 pulse phase bins. The profile in the lower panel was observed with an incoherent detector (the ATNF wideband correlator) using 256 pulse phase bins. The profile shapes are quite different.**

instruments. The lower-resolution instrument does not faithfully re-create the true pulse shape, instead convolving it with the instrumental response.

Timing a pulse means precisely determining its location within a single period. This is done using cross-correlation methods, which are more effective if the pulse profile has detailed structure. Thus it is important to use high-resolution instruments for precision timing experiments as it allowing more subtle physical effects to be studied.

## 3.3 Reception and Down-Conversion

Before pulsar signals can be recorded or analysed, they must be received and amplified by a radio telescope. This section describes the reception, amplification and frequency conversion stages required to couple an instrument to electromagnetic

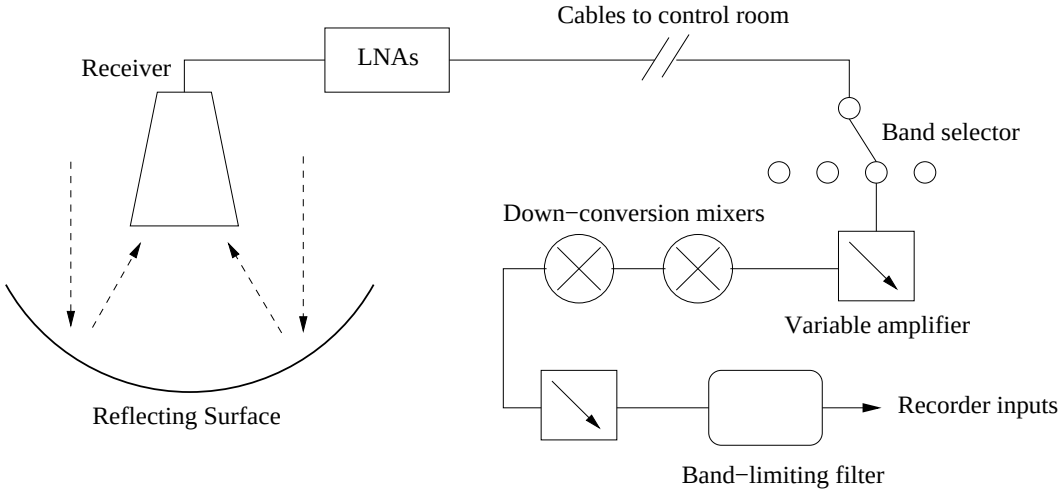


**Figure 3.2: This diagram shows the main elements in the Parkes signal path. Radio waves reflected by the 64 m diameter parabolic reflector are gathered by a feed horn and coupled to an orthogonal-mode transducer (OMT) that separates the signal into two orthogonal polarisation components. From here, the signal passes into cryogenically cooled low-noise amplifiers (known as pre-amplifiers) that compensate for transmission losses in the coaxial cable used to channel the signal away from the antenna structure and into the down-conversion system. In some cases, a first stage of down-conversion is performed close to the receiver package. The signal is further amplified and then passed through a number of mixers and filters that reduce the centre frequency of the observed passband to a lower intermediate frequency (IF). Band-limiting filters can be used to select a region of interest from the passband of the receiver. Variable amplifiers are adjusted to provide suitable input power to any of several different instruments.**

waves traveling in free space. The information relates to the Parkes radio telescope, but should be generally applicable because the basic components of any radio astronomy signal chain are similar even though antennas themselves differ from site to site.

Figure 3.2 shows a schematic diagram of the Parkes signal path. The "front-end" of any observing system is a radio receiver package that couples the incident electromagnetic radiation to a transmission line. Most radio telescopes use parabolic reflectors to concentrate the radio emissions from a small region of sky to a point at which the receiver's feed horn is mounted. The increase in signal strength that this provides (when compared to an antenna equally sensitive to radiation from all directions) is known as the telescope "gain". Unlike optical telescopes that use multi-pixel cameras, a single-dish radio telescope usually has a single receiver. At the wavelengths most commonly used for radio pulsar astronomy, the resolution of the telescope is normally limited by diffraction around the reflecting surface, not by atmospheric turbulence.

In radio astronomy, the functional dependence of gain (or diffracted intensity) on the distance and orientation of the source with respect to the central axis of

the reflector is known as the "beam shape". Because a parabolic reflector only focuses radiation incident along its central axis, we expect gain to drop sharply with distance from the central axis. In other words, a radio telescope is most sensitive to sources in its direct line of sight. The beam shape of a parabolic reflector conforms to the diffraction pattern associated with a circular aperture, provided the feed is not tapered. The intensity of the diffracted radiation is given by the square of the amplitude of the wave and corresponds to the Airy ring pattern (Eq. 3.3).

$$I(\theta) = I_0 \left[ \frac{J_1(2\pi r \sin(\theta)/\lambda)}{\pi r \sin(\theta)/\lambda} \right]^2 \tag{3.3}$$

Here, $I$ is the intensity in the focal plane, $\theta$ is the angular distance from the central axis, $I_0$ is the intensity at $\theta = 0$, $r$ is the diameter of the reflecting surface (also known as the aperture), $J_1$ is the first-order Bessel function and $\lambda$ is the wavelength of the incident radiation. In this simple case, the two-dimensional beam pattern on the sky will be circularly symmetric, allowing it to be characterised with a one-dimensional radial profile, as shown in Fig. 3.3.

The size of a radio telescope beam can also be thought of as the angular width at which the gain drops by a factor of 2 (also known as the full width at half maximum, or FWHM). For the purposes of radio astronomy, it is more useful to determine the beam width using the "Rayleigh criterion", which states that the first minimum of the diffraction pattern associated with one source should correspond with the peak of the pattern associated with a second source for both to be resolvable as individual objects. In the simple case of diffraction around a circular aperture, the Rayleigh criterion can be derived from Eq. 3.3 (by solving for the minima in $\theta$), and is given below:

$$\Theta_{\mathrm{R}} = \frac{1.22\lambda}{d} \tag{3.4}$$

Here, $\Theta_{\mathrm{R}}$ is the angular size in radians, $\lambda$ is the wavelength of interest and $d$ is the diameter of the circular reflector.

The receiver itself usually consists of an orthogonal-mode transducer (two small orthogonally-aligned antennas called "probes", which can be sensitive to either linear or circular polarisations), positioned at the end of a waveguide (the feed horn) whose opening is placed at the focus of the reflector. The receiver package often includes low-noise amplifiers and a cryogenic refrigeration system to help reduce the system temperature. The receivers used during the course of this thesis were all constructed with linear probes, consisting of small dipole antennas sensitive to transverse waves that induce electron motion along the dipole.

Before any scientific information can be extracted from the pre-amplified radio signal, it must go through a number of additional amplification and conversion stages. In radio communications terminology, this process is known as superheterodyne mixing and it offers a number of advantages. The first is that it is often difficult to amplify a signal by the full amount required in one stage. If too much gain is

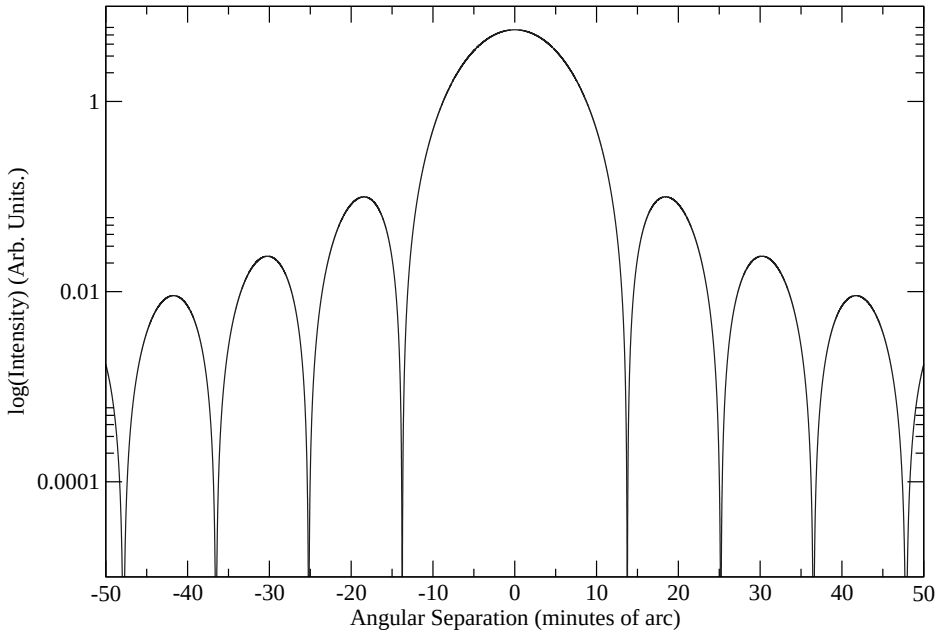


**Figure 3.3: This figure shows the Airy sensitivity pattern, calculated for a hypothetical 64 m reflector with an un-tapered feed at a wavelength of 21 cm. The first minimum in the diffraction pattern occurs at approximately $14'$ from the axis, defining the width of the primary beam. The lower-order maxima appear more significant in this logarithmic plot, the first so-called "side lobe" is actually down by a factor of 60 in intensity when compared to the primary beam.**

required, the electrical components may saturate and signal quality will be lost. Converting the signal to a different frequency makes it easier to apply several stages of amplification. In the case of modern space communications and high frequency radio astronomy, the frequencies used are often so high that amplification requires specially designed, highly advanced semi-conductor devices (like the Monolithic Microwave Integrated Circuits developed at the ATNF) at the pre-amplification stage; it would not be feasible to construct an entire signal chain from such components.

By combining an observed radio signal with the correct pure tone, it is possible to convert very different radio frequencies into the same IF. Instrument designers (or communications engineers) can then assume that any subsequent processing or demodulation will occur at the standard IF, greatly simplifying the task of designing components further down the signal chain. Information about the true sky frequency can be appended to the processed data files. Down-conversion is also necessary to reduce signal loss. The efficiency of most transmission lines decreases with increasing frequency. By down-converting a high frequency signal, the same information content can be transmitted through a given length of cable with less power attenuation.

Down-conversion of a signal is done by "mixing" it with a pure tone at the appropriate frequency. An ideal "mixer" is a device that multiplies two input frequencies and produces a combined signal containing the sum and difference frequencies of the inputs. This can be described using simple trigonometry, the multiplication of two sinusoidal functions with frequencies $A$ and $B$ produces an output signal containing the sum and difference frequencies as in Eq. 3.5.

$$\sin(At)\sin(Bt) = \frac{[\cos((A-B)t) - \cos((A+B)t)]}{2} \tag{3.5}$$

If the centre frequency of the input signal is known, it can be converted to any desired IF by mixing with an appropriate pure tone according to Eq. 3.5 and passing the output through a high- or low-pass filter to reject the unwanted component (see Fig. 3.4).

The situation is complicated by the fact that observations are usually conducted over a broad, continuous region of the radio spectrum. During down-conversion, the entire band is shifted in frequency, but the width of the original band must be maintained. If the mixer's reference tone is placed *below* the observing band, the top of the band will be converted to a higher frequency than the bottom (assuming the difference is selected), preserving order in the band. However, if the reference tone is placed *above* the band, the differences will be negative and the band will be reflected about the origin, the upper portion now appearing to have the smallest frequency. We refer to these two possibilities as upper- and lower-sideband respectively. It is important to keep track of all such conversion stages so that the final orientation of the band is known.

In radio astronomy, the term "baseband" refers to the situation where the lowest edge of the band is at zero frequency and the highest edge is at the frequency

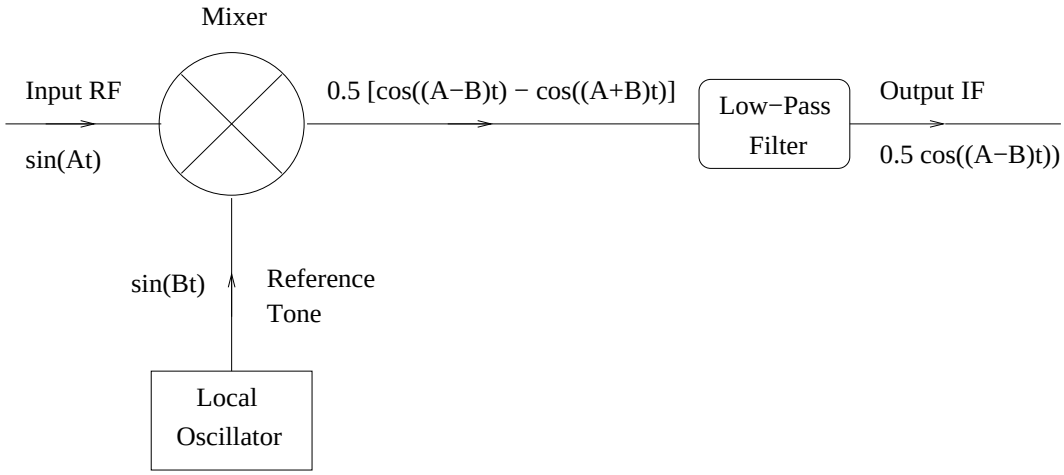


**Figure 3.4: This figure shows the basic layout of a radio frequency down-converter, using a mixer to multiply the input signal at frequency A with a specially chosen reference tone B to produce a superposed signal that contains the sum and difference frequencies, A+B and A-B. In this case, a low-pass filter is used to select the difference frequency at the output stage.**

corresponding to the bandwidth. The Nyquist sampling rate (for real-sampling of the voltage) is then twice the highest frequency component, or twice the bandwidth itself. Thus by converting a signal to baseband, it can be Nyquist sampled most efficiently. This means that the speed of digital electronics limits only the bandwidth over which we can observe. It turns out that provided the signal is band-limited (contains a well-defined portion of the radio spectrum with sharp boundaries), it is possible to Nyquist sample at a rate equal to twice the bandwidth even if down-conversion all the way to baseband is not possible. For this reason, the filters used to define observing bands need to be high quality so as to prevent aliasing any out-of-band information into the astronomy signal. It is also important that the relative phase of all reference tones is kept the same, or destructive interference may occur in the signal path. This is normally done by locking all hardware signal generators to a common reference.

It is also possible to perform "complex" sampling, which involves splitting the real-valued signal into two components via "quadrature" mixing. In this process, two copies of the signal are mixed independently, with reference signals that have the same frequency but maintain a constant phase offset of $90^o$. The two outputs are sampled simultaneously, each at half the Nyquist rate. A single time sample is stored as the complex number x + **i**y, where x and y represent the two components. For a given bandwidth, the data rate is therefore the same as Nyquist real-sampling, and the information content of both recordings is equivalent. Although two samplers are required, each can operate at half the Nyquist rate, which may be technically more feasible. Complex sampling is however more susceptible to variations in the mean power level, because the information has been distributed over two signal paths that

can vary independently. CPSR2 avoids these problems by real-sampling the signal.

## 3.4 Analogue to Digital Conversion

The Nyquist sampling theorem states that a continuous time series (such as the voltage produced by a receiver) can be digitised and later reproduced perfectly if all spectral content is preserved. This can be understood with the help of the convolution theorem, which states that multiplication in the time domain is equivalent to convolution in the frequency domain. Figure 3.5 illustrates how this theorem can be used to find the sampling rate required to preserve all the information in a given signal.

Sampling a continuous time-domain signal (Fig. 3.5a) can be thought of as multiplying it with a "comb" function that is equal to unity at every sample and zero everywhere else (Fig. 3.5c). We can then compute the Fourier transform of the digitised time series using the convolution theorem. Fourier transforming the comb sampling function produces another comb function whose "tooth" spacing is equal to the reciprocal of the sampling rate (Fig. 3.5d). Convolution of the Fourier transform of the original continuous time series with the reciprocal comb function simply replicates the pure spectrum, centering it on each tooth of the comb (Fig. 3.5f). If the comb is sufficiently broad, no two copies of the spectrum overlap and every spectral component in the original signal is preserved. We can therefore recover the original time series. However, if the repeated spectra do overlap in the Fourier domain of the sampled signal (as they do in Fig. 3.5), some of the frequency components are distorted and we cannot recover all the information. This is known as "aliasing", and it can be avoided by sampling rapidly enough to preserve the full frequency spectrum of the original signal. The sampling rate required is always twice the highest frequency component present in the band.

The above discussion assumes that the voltage at any point in time can be accurately and precisely sampled. However, digitisation involves not only a loss of continuity in time, but also a loss of precision in our knowledge of the signal amplitude. This is because we must store each voltage with a finite number of bits inside a computer. Sampling is usually performed by a circuit containing triggers that are activated if the input voltage is greater than the associated threshold. The output of a single trigger is either "on" (1) or "off" (0) and can be stored as a single bit of information. Representing an arbitrary rational number with a single bit is very imprecise, but more resolution can be retained by stacking trigger levels with ascending thresholds and storing more bits for each time sample. The loss of S/N associated with a lack of digitiser levels is surprisingly small. This is partly due to the fact that each individual sample is dominated by noise and we are interested in the statistical behaviour of the signal after a large number of samples have been combined. Even with 1-bit precision, about 80% of the S/N can be recovered using an incoherent filterbank.

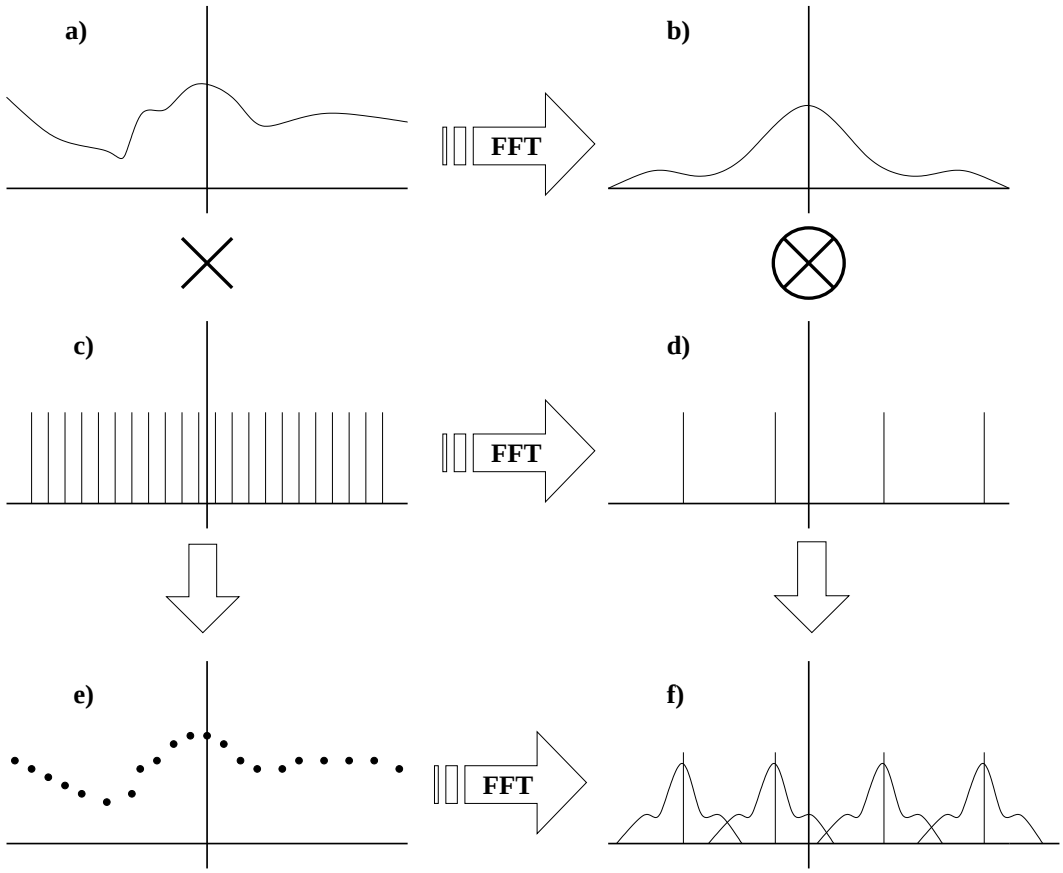


**Figure 3.5: This figure graphically represents a mathematical approach to digital sampling. The time series we wish to digitise is shown in panel a and its Fourier transform is shown in panel b. The sampling process is equivalent to multiplication with a comb function (shown in panel c), whose Fourier transform is also a comb function with reciprocal spacing (panel d). The digitised signal is shown in panel e and its Fourier transform in panel f. The left hand side panels a & c are multiplied to "sample" the signal. This is equivalent to convolution of b & d in the frequency domain. The criteria for complete information preservation is then to sample at twice the highest frequency component present in the signal.**

One of the problems encountered when sampling with a finite number of levels is the question of where to place the voltage thresholds. Even 8-bit precision will not be adequate if the levels are placed inappropriately, leading to the digital equivalent of saturation. In most circumstances, a digital sampler is designed to have a certain input range and variable attenuators are used to ensure the mean level of the incoming signal is within the required range. An informative diagnostic technique involves regularly checking the statistics of the digitiser, i.e. the number of samples recorded above or below each threshold. If the thresholds are set well and the incoming power is dominated by random noise, the number of samples at each level should conform to a Gaussian profile with more "hits" near the zero level and fewer at the extremes. The presence of interfering transmissions or a drift in the mean level will change the statistical profile.

The number of bits also defines the "dynamic range" of the system, which is a measure of the maximum variation in signal strength that can be accurately represented. The dynamic range of a 1-bit system is very low as the sampler is easily saturated. High dynamic range is important when there is strong interference competing with the astronomy signal. In order to subtract contaminated spectral components, the interference must be accurately characterised.

To correct for dispersion smearing, we must have detailed knowledge of power as a function of frequency. Instead of a single detector, we require a spectrometer. The next section describes a number of techniques commonly used to extract the spectrum of an observed signal.

## 3.5 Spectrometry

The first generation of instruments designed for pulsar astronomy were based on large arrays of analogue filters (filterbanks). The IF signal was cloned many times using a multiplexer and the streams were fed into a set of analogue filters with uniform width, each sensitive to a slightly different region of the incoming band. The output of each filter was detected in hardware and independently digitised, then written to storage media for processing. Although the idea was simple, manufacturing high-quality filters is difficult and expensive. Resolution was therefore limited by the roll-off of the filters and their response time.

It is also possible to compute a spectrum using the Wiener-Khintchine theorem, which states that "The auto-correlation function is the Fourier transform of the power spectrum". Thus, by inserting lags and correlating an observed time series with itself, it is possible to recover the Fourier transform of the power spectrum, which by reverse transform can be converted into a channelised representation of the receiver passband. Custom digital circuits can be built to perform rapid correlation, allowing the application of this method to large bandwidths. Constructing the power spectrum in this manner has the advantage that few filters interact with the observed signal, making the system less susceptible to hardware inconsistencies.

The Fourier transform can also be used to synthesise a filterbank from Nyquist sampled data. If $N$ spectral channels are desired, the real-sampled time series is split into sequential segments containing $2N$ points. An FFT is performed on each segment and the resulting complex spectral components are effectively a row of time samples in the required $N$ channels. The result of each subsequent FFT is stacked onto the last row. In this way, time resolution is traded for frequency resolution. In practice, this simple method is susceptible to spectral leakage amongst neighbouring channels, but this can be alleviated by performing a longer forwards FFT and multiple backward FFTs to construct the spectral channels.

In theory, it is possible to reduce the effects of dispersive smearing to negligibility. Hankins & Rickett (1975) described a method of coherently dedispersing across a finite band by de-convolving the signal with the inverse of a response function that described the ISM. However, this method required knowledge of the electric field and had to be performed prior to detection. One way this could be achieved was to digitally sample the analogue signal at the Nyquist rate and store the samples for off-line processing. When coherent dedispersion was first attempted, computer power and storage media capacity limited the bandwidth to approximately 250 kHz. Hankins, Stinebring & Rawley (1987) increased this bandwidth to 2 MHz by using the Reticon R5601 integrated circuit, which was designed specifically to perform the Chirp-Z transform. Fortunately, the power of consumer computing hardware increased so rapidly that by the turn of the millenium, Nyquist sampling and software-based coherent dedispersion could be performed across tens of MHz, opening the way for widespread adoption of the technique.

To allow for greater flexibility and efficiency, coherent dedispersion can be combined with the operations required to form a digital filterbank. This decreases the length of the FFTs required (lessening the computational load slightly) and results in a set of narrow channels that are internally free of dispersion smearing. The trade-off between time and frequency resolution still applies, but provided the DM of the pulsar is known accurately, dispersion smearing is eradicated regardless of the frequency resolution. The final balance can therefore be optimised to keep the time resolution just below the width of a single phase bin.

## 3.6 The Caltech, Parkes, Swinburne Recorders

In August 1998, a team of astronomers and engineers from Swinburne University of Technology and the California Institute of Technology installed a new baseband recording system at the Parkes radio telescope in New South Wales, Australia. This system (known as CPSR1) was designed to Nyquist-sample $2\times20$ MHz-wide bands, allowing coherent dedispersion and polarimetry of pulses from the bright, southern hemisphere millisecond pulsar PSR J0437–4715. Band-limited IFs were provided by the telescope down-conversion system and fed into a custom-built digitiser designed by John Yamasaki. This was connected to a Direct Memory Access (DMA) card

housed in a Sun workstation that was equipped with 4 hard disks and 4 DLT tape drives on an internal SCSI bus. The digitiser complex-sampled the signal from each polarisation with 2-bit precision. At the Nyquist rate, this generated 20 MB of data every second, which was recorded to the hard disks of the Sun machine. Unfortunately, the capacity of the hard disks was severely limited, allowing less than an hour of continuous recording. To allow extended recording, the data had to be removed from the hard disks as quickly as it arrived. Each DLT could write at a rate of 5 MB per second, using all four in parallel allowed data to be written to tape at the required rate of $20\,\mathrm{MB\,s^{-1}}$. The tapes were shipped back to the Swinburne Centre for Astrophysics and Supercomputing, then inserted into a robotic silo for processing. A cluster of high-speed servers ran a program called PSRDISP (van Straten 2003) to process the baseband samples. This program reconstructed the analytic signal, performed coherent dedispersion (and optionally formed a coherent filterbank) then square-law detection, after which the signal was synchronously folded into a binned profile using the predicted pulse period.

CPSR1 was at the forefront of coherent dedispersion technology and allowed the detection of annual-orbital parallax and Shapiro delay in PSR J0437–4715 (van Straten et al. 2001). Even so, it suffered from a number of significant limitations. The practice of recording directly to tape and forming pulse profiles well after the observations offered no real-time feedback on data quality. If the telescope was incorrectly configured, hours of data could be recorded before the mistake was noticed and rectified. In addition, the observing bandwidth of 20 MHz was an order of magnitude smaller than that typical of incoherent detectors available at Parkes at the same time; the enhanced precision associated with coherent dedispersion came at a significant cost in S/N. J0437–4715 was well suited to a precise but narrow-band system and CPSR1 was never intended to be a general purpose instrument.

Much of the expense associated with constructing a new instrument goes into development; engineers are required to design the hardware circuits, prototypes must be built and tested, refinements made and so on. Only in the last few years has it been feasible to construct a powerful processing engine out of consumer components. The number of transistors per unit area within an integrated circuit continues to double every 18 months (Moore's Law), allowing CPU power to steadily increase. Modern personal computers also have the ability to handle large volumes of input and output. They can also be programmed to perform any algorithmic task. Using PCs as processing engines has the advantage that most hardware development costs are paid by the intended consumer market, lessening the overall expense. Indeed, many modern supercomputers are constructed using relatively inexpensive PCs, linked by high-speed local area networks. Various hardware manufacturers facilitate the use of PCs for scientific research by providing expansion cards that can directly interface with data acquisition systems. The only remaining development cost comes in designing the software required to process the data. Although software development can be difficult and time consuming, well-designed code can be run on upgraded hardware with minimal modification and the cost of development must

only be paid once.

CPSR1 proved that software-based coherent dedispersion was viable, but it did not offer much scientific flexibility. PSRDISP was designed to read a pulsar ephemeris and perform the appropriate level of dedispersion and synchronous folding, making it capable of processing data from any pulsar. However, few pulsars had a short period, narrow profile and flux density large enough to benefit from coherent dedispersion with 20 MHz of bandwidth. It soon became clear that installing a local processing engine at the telescope would provide real-time feedback and do away with the cost and risk of tape transportation. It would also provide the computational power needed to extend the observable bandwidth. Benchmark tests provided an idea of how many processors were required to reduce baseband data in real time for a given bandwidth. The high-end PC systems available in 2001-2 were capable of handling significantly greater data rates than the Sun machine used to drive CPSR1, despite only four additional years of progress. 30 dual-processor servers could process 128 MB of data every second, allowing a bandwidth of 128 MHz with 2-bit real-sampling. The new instrument would run PSRDISP, and therefore inherit all the development and extensive testing already invested in the core coherent dedispersion code. A new sampler was designed to convert 4 IFs into two streams of 64 MB $s^{-1}$, near the upper limit of both computer disk and network transfer rates. The new CPSR2 system would provide an unprecedented amount of coherently dedispersed bandwidth to radio observers, offering a total sensitivity increase of 2.5 compared to CPSR1, plus all the benefits of fully on-line data reduction and real-time feedback.

## 3.7 CPSR2 Construction and Commissioning

CPSR2 combines a custom-built Fast Flexible Digitiser (FFD) board with a cluster of Intel-based rack-mountable servers. The FFD itself was the only custom digital circuit used in the construction of CPSR2, all other components were purchased commercially and configured as required. The computer cluster was obtained from Dell, in the form of 28 individual PowerEdge 2650 servers, each with dual 2.2 GHz Intel Xeon processors, 2×72 GB Ultra SCSI hard disk drives and 1 GB of Random Access Memory (RAM). Each of these servers was installed with SuSE Linux version 7.2 and equipped with the GNU Compiler Collection (GCC) and a suite of pulsar data reduction tools, including PSRDISP. In addition, two "primary nodes" were used to read data from the FFD via DMA cards, then distribute it to the cluster machines (or "secondary nodes"). These two machines were equipped with 3 GB of RAM and twice the hard disk space of the others. We used EDT–60 DMA cards, which were connected directly to the FFD by means of custom-built ribbon cables. External computer control of the FFD was facilitated by a serial connection between it and a primary node. In addition to the primary and secondary nodes, a gateway machine was installed in the Parkes radio telescope control room. This machine was connected by means of fibre-based Gb Ethernet to a Cisco switch that provided

an additional 100 Mb Ethernet layer to all the machines. This layer was used by the control and monitoring system that gathers information from all the nodes and presents it to the operator. It was considered necessary to isolate this traffic from the primary data transmission network, to maximise the bulk transfer rate.

The instrument was assembled and commissioned during an intensive, 3 week session at the Parkes radio telescope in August 2002. Astronomers and students from Caltech and Swinburne travelled to the telescope in order to assemble and test the system. First, the cluster nodes were removed from their packaging and stacked in an open-frame rack. Power and two sets of network cables were distributed throughout the rack and a VME-style crate was installed to house the FFD. The Linux operating system and all required data reduction software were installed on a single secondary node and then cloned at the device level onto each of the others, followed by manual assignment of unique Internet protocol addresses and hostnames. DMA cards were installed in the primary nodes, along with the necessary driver software. Each Dell 2650 has 5 rapid-access hard disk drive slots, facilitating easy disk duplication. A set of spare disks was cloned from a secondary node and set aside, to allow rapid recovery should any of the other internal disks fail (this has happened twice since commissioning). After the FFD had been connected to the primary nodes and the IFs, the data acquisition software was tested. These tests revealed a number of problems with the code designed to drive the FFD and DMA cards and handle data buffering and transfer. For example, the internal logic designed to re-route slow network connections suffered from a form of positive feedback that could overwhelm the network switch. These problems were slowly corrected, with the help of new cluster monitoring tools that were developed to assist the debugging process. It took several months for a stable observing system to emerge.

Most aspects of the system have remained in their initial configuration, with the important exception of the high-speed network layer. It was initially thought that Gb Ethernet would not have sufficient capacity to support sustained data transfer at a rate of $64\,\mathrm{MB\,s^{-1}}$, and each machine was installed with a Giganet card. These were linked by high-capacity CAT 7 network cable to a central switch, using a star topology. Giganet was designed for supercomputing applications and could transmit data at a sustained rate well in excess of the required $64\,\mathrm{MB\,s^{-1}}$. Unfortunately, the system seemed to have a number of stability flaws that may have been accentuated by hardware problems in the internal switching circuits. We found that at times, machines would vanish from the Giganet switch and at other times they would refuse to load card drivers at boot time or acknowledge the presence of the switch. The low-level network protocol was heavily optimised for a small number of high bandwidth connections; we found that trying to maintain a number of open connections to different devices could slow down the maximum rate of any point-to-point transfer significantly. In the end, the whole Giganet system was removed and replaced with two 24-port Gb Ethernet switches, each handling one half of the secondary nodes (but connected by a 2 Gb cross-over link). The FFD packed all recorded samples into two streams of $2\times64$ MHz 2-bit samples, usually representing

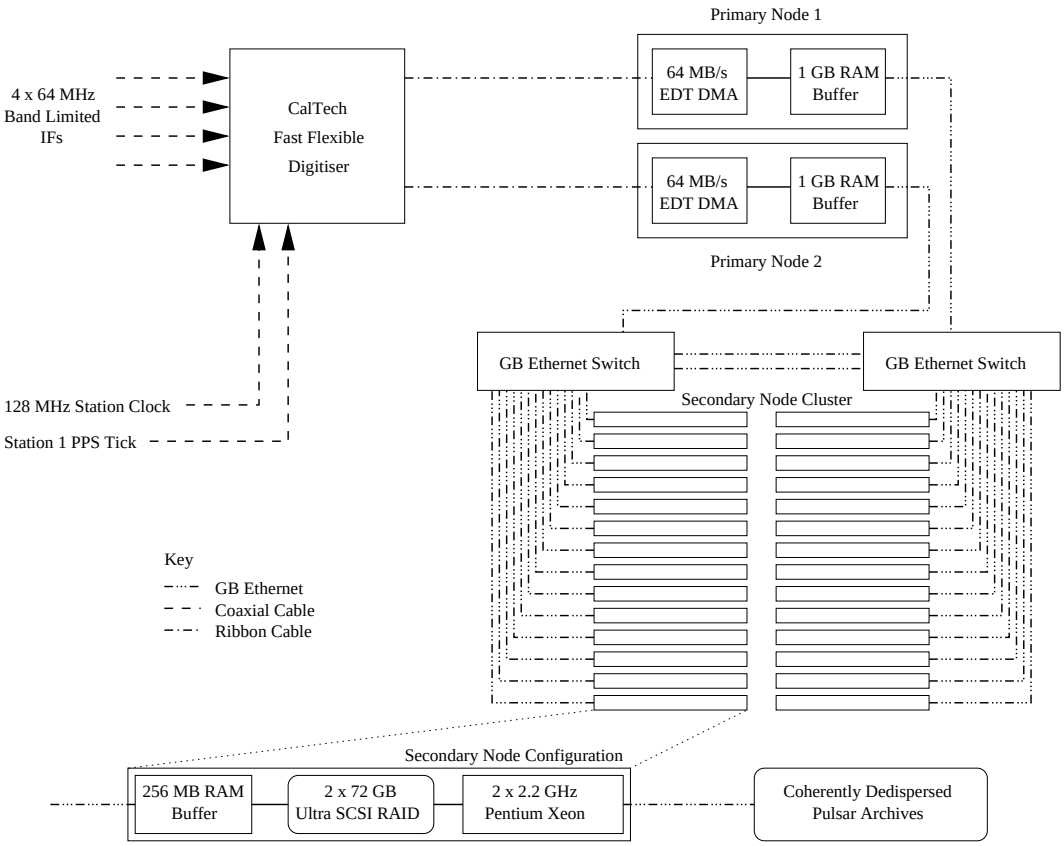


**Figure 3.6: This diagram shows the basic layout and interconnection plan of the CPSR2 baseband recorder and its processing cluster.**

orthogonal polarisations of a single, 64 MHz-wide band. Each of these streams was sent to a different switch by the associated primary node, maximising the effectiveness of our network configuration by avoiding the cross-over link. Gb Ethernet was also compatible with the TCP/IP stack within the Linux kernel, which proved to be far more reliable than the Giganet protocol. The theoretical maximum throughput is $1\times10^9$ bits s$^{-1}$, which is equivalent to 125 MB s$^{-1}$. Sustained rates in excess of the required 64 MB s$^{-1}$ were easily obtained, contrary to our initial expectation.

A schematic representation of the current system configuration is presented in Fig. 3.6. Samples recorded by the FFD were transferred from a ring buffer associated with each EDT card into a large block of shared memory (normally 1 GB) in one of the two primary nodes, depending on which of the two sides of the FFD they originated. These shared memory blocks were also configured as ring buffers, which were written to sequentially in small blocks (32 MB in most cases) of data. Ring buffers are a simple example of a first-in-first-out (FIFO) construct. A fixed number of blocks are set aside for storage, and different processes can either write to or read from these blocks. A record is kept of the state of each block and any write requests are directed to the next empty block in the sequence, wrapping back to zero after

the final block. Any read request is directed to the oldest full block, which is marked empty once the read operation is complete. As long as the reader can keep pace with the incoming data, the buffer will never reach a state where every block contains new data.

Data acquisition code read from the EDT buffers and wrote to the CPSR2 primary buffers; a second communications process read from the primary buffers and transfered data units across the network layer to a waiting secondary node, which was equipped with a smaller ring buffer and another process that wrote the data to disk. The presence of ring buffers at both ends of the network made the system robust to short glitches in flow rate. In fact, network transfer had to cease for a full 16.8 s to fill one of the primary ring buffers.

Once the raw data had been written to disk, it was processed by PSRDISP according to an ephemeris obtained from either the standard pulsar catalogue or a standard directory, pointed to by an environment variable. Each observation (the time between "start" and "stop" signals being sent to the FFD) was treated as a single entry in a job queue system, managed from a central interface. The required parameters (phase bins, spectral resolution and so on) could be set by the user on a job-by-job basis. There were default settings tailored for most of the commonly observed pulsars and a standard default for any source that had not been observed previously with the instrument. Once available, the folded pulsar archives were copied by the same data management software to a central storage location for subsequent backup and analysis, and also for inclusion in the running total constructed by the online data quality monitoring system.

The next chapter describes the configuration and operation of CPSR2. Readers interested in a detailed description of individual software components should also refer to Appendix B.

## 3.8 Operating CPSR2

The output power level of the Parkes radio telescope down-conversion chain was configured manually using a graphical user interface, logui, or automatically using the Telescope Control System (TCS). The final 64–128 MHz filters were connected by coaxial cable to the FFD inputs. Although not quite baseband, the signal was band-limited (so it contained effectively 64 MHz of spectral information) and an integer multiple of the desired range. The entire band was aliased back into the range 0–64 MHz when sampled at what was effectively half the required Nyquist rate. The information content was preserved, but higher frequencies were mapped to lower frequencies, causing a final inversion of the band. The FFD was divided into two halves, each side accepting two of the four IFs. In most cases these were configured to be orthogonal polarisations of a single 64 MHz-wide band. Each half of the FFD had an independently adjustable attenuator with a 30 dB range which was controlled by commands sent down the serial line. The 4×64 MHz inputs could

be configured to centre on a wide range of sky frequencies. We successfully observed with the central beam of the Parkes multibeam receiver and the H-OH receiver, both at 20 cm, the dual-band coaxial pulsar receiver at 10 cm and 50 cm (with the option of recording parts of both bands simultaneously), and even the 8 GHz Mars receiver.

In the normal mode of operation, the FFD real-sampled the input voltage using four levels (2-bits) of quantisation, clocked by an external 128 MHz oscillator. The FFD was also capable of 4-bit sampling, however since this doubled the data rate, two of the 4×64 MHz channels had to be discarded and it was no longer possible to send both orthogonal polarisations to the same half of the FFD. Since each half of the FFD has its own independent attenuation settings, running in this mode could introduce a systematic offset in the relative strength of the polarisations, making calibration difficult. It was also theoretically possible to halve the bandwidth in order to compensate for the increased number of bits, but this required halving the rate of the external clock signal. Unfortunately, the external clock was also used to control the serial port, meaning that any deviation from the intended 128 MHz caused the FFD control system to freeze. For these reasons, did not observe in 4-bit mode with CPSR2, although the additional bits may have been beneficial when faced with strong radio frequency interference.

The routine designed to convert a recorded bit stream into a series of floating point values (and perform any corrections for quantisation error) was called the "unpacker". The CPSR2 unpacker was also responsible for separating out the orthogonal polarisations. Assuming that the digitised signal had been recorded using a fixed set of input voltage thresholds, the unpacker needed to work backwards, deciding what floating point "voltage" to assign to a given sample. This conversion need not (and in fact *should not*) be constant (the same 2-bit number need not always represent the same voltage). Jenet & Anderson (1998) described a scheme in which the output levels were varied in order to keep the power level of the digitised signal as close as possible to the power level in the un-digitised signal. The optimal output levels were defined by estimating the power in the un-digitised signal, based on the statistical properties of 512 samples taken from the digitised signal and the assumption that the system noise was Gaussian. The greater the number of bits used to store each sample, the simpler the unpacker could be. 8-bit samples could be adequately unpacked using simple base-2 to base-10 conversion.

Of course, the input thresholds could themselves be optimised to yield the highest S/N (and least contaminated) digitised signal. The optimal relative threshold separation was constant, but the actual voltage levels were defined in terms of the power in the un-digitised signal. This could be changed using the variable attenuators built into the FFD, via commands sent through the serial control line. Based on statistical analysis of the digitised signal, the correct level of attenuation could be computed at any given time. Assuming the incoming power did not change significantly over the course of a single observation, the FFD's attenuators could be set once and then left fixed. Alternatively, if there was strong interference or a lengthy

observation was required, the attenuators could be varied more regularly in small increments, at the cost of introducing steps in the recorded power level. The process of folding averaged over these steps, but some applications (for example, giant pulse searching) depended upon the mean level being constant and were therefore unsuited to dynamic level setting.

Despite the fact that the spectral content of the signal was preserved by Nyquist sampling, 2-bit digitisation could introduce systematic errors into the measured signal. These errors were fundamental to the quantisation process and could not be avoided. They could be minimised by sampling with ever-higher numbers of bits, but this increases the data rate to an unacceptable level. Jenet & Anderson (1998) described the form of several quantisation errors and detailed ways in which they could be compensated for without increasing the data rate. PSRDISP included code designed to correct for a first-order effect that caused power loss in the vicinity of the pulse, but was in general unable to perform a second-order correction for power scattered across the band, because the operation was not compatible with coherent dedispersion.

Because pulsar timing requires precise knowledge of the time a particular pulse arrived at the telescope, it was necessary to store an accurate time stamp with each block of recorded data. To ensure the FFD started recording at exactly the right moment, the detector was "armed" by a serial command sent from a process running on the second primary node (whose clock was kept accurate to within a few tens of milliseconds via Network Time Protocol signals) and triggered by the telescope 1PPS signal, which was driven by a Hydrogen maser frequency standard whose relationship to UTC was carefully monitored. The data acquisition software responsible for reading from the EDT–60 buffer into the primary ring buffer appended 4096 bytes of header information to each file that was sent across the network. This header contained a time stamp and information about the centre frequency of the band, sky coordinates of the source, digitisation mode and so on. It was written in ASCII format for ease of interpretation, as shown in the example below.

```
CPSR2_HEADER_VERSION 0.2      # Version of this ASCII header
CPSR2_DAS_VERSION 0.1         # Version of the Data Acquisition Software
CPSR2_FFD_VERSION 5           # Version of the FFD FPGA Software

TELESCOPE    PKS                    # telescope name
PRIMARY      cpsr1                  # primary node host name

# time of the rising edge of the first time sample
UTC_START    2005-01-24-14:56:01    # yyyy-mm-dd-hh:mm:ss.fs
MJD_START    53394.6222337962963    # MJD equivalent to the start UTC

OFFSET       77277954048            # bytes offset from the start MJD/UTC

SOURCE       J0540-6919             # name of the astronomical source
RA           00:00:00.0             # Right Ascension of the source
DEC          00:00:00.0             # Declination of the source
```

```
FREQ         3000                  # centre frequency on sky in MHz
BW           64                    # bandwidth of in MHz (-ve lower sb)
TSAMP        0.00781250            # sampling interval in microseconds

NBIT         2                     # number of bits per sample
NDIM         1                     # dimension of samples (2=complex, 1=real)
NPOL         2                     # number of polarizations observed

EDGE         30
LEVEL        600
MODE         SURVEY
CALFREQ      11.123000
NMBYTES      73698
BUFFERSIZE   31457280
HALT         0
```

Information stored in the header was taken from a control process running on the cluster gateway machine. This background process was also responsible for coordinating starts and stops. Most parameters could be entered by hand using a graphical user interface, but in the normal mode of operation they were sent over the network by TCS without any user intervention.

When the FFD was triggered to start by the 1 PPS, it began filling EDT buffers. The data acquisition software on each primary node waited on a flag set by the EDT driver to indicate when new data was available. When the samples started to arrive, they were copied into the primary ring buffer as quickly as possible. A small fraction of the samples were extracted and used to monitor the digitiser statistics and adjust the FFD attenuators if necessary. An additional process thread waited on a similar flag built into the primary CPSR2 ring buffer, immediately sending any new blocks out over Ethernet to the next available secondary node.

The FFD included a check-code in its data stream by periodically over-writing the data with a string of 8 characters and a byte count from the start of the recording. This allowed the integrity and continuity of the data to be monitored. The presence and order of these codes were checked at several stages during the recording process. Any file that ended up on the hard disk of a secondary node was guaranteed to be a contiguous block and could therefore be processed independently.

Although it was theoretically possible to write continuously to a single pair of secondary nodes until their disks filled, the communications code normally sent short bursts of data to each secondary node in a continuous cycle, to ensure an even distribution of files and optimise processing efficiency. The first advantage of this approach was its robustness. If one secondary node happened to run slowly or was incapable of accepting a connection, the data flow would normally stall. However, provided the primary ring buffer was greater in size than the amount of data sent during one transmission burst, the system did not halt unless several machines in a row were incapable of receiving data at the required rate. The second advantage was that short transmissions resulted in a uniform distribution of data across the

secondary nodes, which allowed each machine to begin processing soon after the start of an observation. Short files could be rapidly processed in order to provide immediate feedback to the observer. The short archives resulting from this practice were combined into longer integrations once the observation was complete.

## 3.9 Data Processing

The CPSR2 cluster was capable of dedispersing up to a DM of about $40\,\mathrm{pc\,cm^{-3}}$ in real-time, at a wavelength of 20 cm. The primary nodes sent data to each secondary node in their target list sequentially, wrapping back to the beginning of the list after the last machine. If a secondary node had finished processing its previous file by the time it was given the next data file, processing kept pace with the observation and the average amount of free disk space remained the same. Above this DM, the cluster fell behind and each node had to continue processing the previous file for some time after the next was received. Because dedispersion is very CPU intensive, running PSRDISP on both CPUs could hinder the ability of a secondary node to accept new data from the network. For this reason, the secondary nodes were programmed to temporarily suspended all processing during data transfer. This additional loss of 16.8 processing seconds created a surprisingly sharp efficiency cutoff at the real-time transition. When the processing rate was slower than the recording rate, the cluster disks would start to fill. Each secondary node had 125 GB of scratch space within a 2-disk RAID array which stored unprocessed data, making the total cluster storage capacity approximately 3.5 TB. At a rate of $2\times1$ GB files every 16.8 s, this equated to more than 8 hours of dual band data (or 16 hours of single band data), even if no processing was attempted. Given that the average 24 hr observing session involved many different sources with a range of DMs, it was rare for the disks to fill unless data was being stored deliberately for future analysis.

As the folded, dedispersed archives were deposited in the central storage area, they were appended to an integrated profile. This profile was displayed in a graphical window, along with a summary of key observing parameters (source name, sky frequency and so on), a plot of pulse phase as a function of frequency (showing the dispersed pulse) and a history of the S/N as a function of time. This information was used to determine if the observing system was correctly configured. If, for example, the band were to be inverted, the integrated pulse may still show up in the frequency vs phase plot, but not in the dedispersed mean profile. Many pulsars with DMs less than $\sim10\,\mathrm{pc\,cm^{-3}}$ scintillate due to diffractive effects, causing their apparent flux to vary. Real-time monitoring was used to assess the strength of a source as it was being observed. This allowed dynamic scheduling and facilitated the most efficient use of limited telescope time.

# Chapter 4

# An Open Approach to Radio Pulsar Data Storage & Analysis

"There are two major products that came out of Berkeley: LSD and UNIX. We don't believe this to be a coincidence."

*Jeremy S. Anderson*

In this chapter, a new set of software applications and libraries for use in the archival and analysis of pulsar astronomical data is introduced. Known collectively as the PSRCHIVE scheme, the code was developed in parallel with a new data storage format called PSRFITS, which is based on the Flexible Image Transport System (FITS). Both of these projects utilise a modular, object-oriented design philosophy. PSRCHIVE is an open source development environment that incorporates an extensive range of C++ object classes and pre-built command line and graphical utilities. These deal transparently and simultaneously with multiple data storage formats, thereby enhancing data portability and facilitating the adoption of the PSRFITS file format. Here, data are stored in a series of modular header-data units that provide flexibility and scope for future expansion. As it is based on FITS, various standard libraries and applications may be used for data input/output and visualisation. Both PSRCHIVE and PSRFITS are made publicly available to the academic community in the hope that this will promote their widespread use and acceptance.

## 4.1 Introduction

### 4.1.1 Collaborative Scientific Software Development

Modern, highly specialised experimental systems often require extensive original software development. This is true for all tasks from direct hardware control through

to data reduction. While individual research groups often approach such software development from an isolated perspective, the proliferation of digital hardware and wide-area networking makes global cooperative software development far more attractive, provided suitable common ground exists. Both cooperative software development and the adoption of standard packages provide a number of distinct advantages to the research community. For instance, as less effort is wasted unnecessarily duplicating the work of others, cooperative development can lead to more efficient allocation of resources. In addition, supporting the requirements of a larger user community promotes the development of basic, general purpose routines that may be used in a wider variety of situations. These influences result in more modular and extensible software.

However, it should be noted that a greater level of care and cooperation is required in collaborative software development, especially the open source approach advocated in this chapter. For example, in contrast to most commercial software, "black box" modularity is undesirable in scientific analysis, especially when the application of certain algorithms requires experienced judgment. Open, well-documented code provides researchers with an accurate understanding of third-party analytical tools. Therefore, contributing developers must be willing to put their code in the public domain, making it freely available for non-commercial use by any other academic organisation. Although this facilitates the exchange of ideas, it also raises the issue of potential loss of intellectual property, which might discourage some authors.

It is also the case that collaborative development tends to become de-centralised, especially when multiple developers have the ability to commit fundamental changes to the code. Effective communication between the core developers becomes essential to the smooth running of the project, necessitating greater attention to version control, the maintenance of stable releases, and the development of extensive and concise documentation. Also, when a wider user community is affected by modifications to the software, exhaustive methods must be employed to ensure the validity of changes and the integrity of the system as a whole. Although each of these issues tend to increase the workload of the collaborative developer, a much larger body of users will benefit from the effort.

### 4.1.2 Software Development in the Pulsar Community

The global pulsar community is ideally suited to adopt a collaborative approach to software development. It consists of a relatively small number of locally centralised groups that deal with different telescopes and instruments, leading to several parallel but incompatible software development paths. As each path tends to be built around a highly specific data storage format, cross examination of data and algorithms is problematic. In addition, because such software is generally designed for a limited purpose, it is often difficult to extend its functionality without introducing obfuscated code. This is especially true when the program develops in an experimental

fashion, as is often the case with scientific applications. In order to avoid future inflexibility, sufficient time and care must be invested during the planning stage, calling on input from both experienced software developers and pulsar astronomers.

### 4.1.3 Radio Pulsar Data

Radio pulsars are broad-band point sources of highly polarised emission that exhibit rapid pulsations with a characteristic period anywhere between one millisecond and ten seconds. They are thought to be rotating neutron stars with a strong dipolar magnetic field whose axis is not aligned with the rotation axis of the star (Gold 1968). Intense beams of emission originate at the magnetic poles, which sweep across the sky with each rotation of the star and produce the pulsed radio signal observed at Earth.

The characteristic *signature* of any radio pulsar is its integrated polarimetric pulse profile, given by the observed Stokes parameters averaged (folded) as a function of pulse longitude over several thousand individual pulses (Helfand, Manchester & Taylor 1975). Under the influence of electrons in the interstellar medium (Taylor & Cordes 1993), this pulsed signal is broadened by dispersive frequency smearing, which must be corrected in order to infer the shape of the characteristic profile at the source. This is normally done by dividing the observed bandwidth into narrow frequency channels, which are appropriately delayed relative to each other before summing the detected flux densities in each channel. However, as the dispersion measure may vary with time or may not be known with sufficient accuracy at the time of the observation, it is often necessary to store the individual pulse profiles observed in each frequency channel.

Additionally, it is possible to create a mean pulse profile only if a suitably accurate model of the pulsar's spin frequency and phase is available. The apparent pulse period is affected by a number of phenomena, including the spin-down, timing noise, and/or glitches intrinsic to the pulsar, variations in the interstellar dispersion, and Doppler effects introduced by the relative motions of the Earth and pulsar. Inaccuracies in the model that describes these effects introduce phase errors that accumulate with time and cause the integrated profile to become smeared. Therefore, it is often beneficial to store multiple, shorter integrations of the mean pulse profile instead of a single, long integration. Furthermore, when a pulsar is bright enough, a great deal of additional information about the characteristics of the pulsar emission can be obtained by recording and analysing each individual pulse. Therefore, a useful pulsar data format must be able to represent pulse profiles observed over multiple epochs of arbitrary length.

In summary, pulsar observations generally consist of a four-dimensional array of data indexed by polarization component, pulse phase, frequency, and epoch. Software support for sensible groupings in other dimensions, such as orbital phase, is also highly desirable. In addition, data from a number of telescopes can be combined to increase sensitivity and contribute to the eventual detection of new phenomena,

such as the cosmic background of stochastic gravitational radiation (Hellings & Downs 1983; Stinebring et al. 1990). Therefore, the data storage format should have a flexible structure that provides efficient access to key parameters, removed from any considerations of individual instruments or signal processing schemes.

### 4.1.4 Processing Radio Pulsar Data

Pulsars are observed for a variety of reasons, from studying the nature of their structure and emission mechanism (Dodson, McCulloch & Lewis 2002) to utilising them as highly stable clocks and astrophysical probes (Taylor & Weisberg 1982). Consequently, the same pulsar observation can be used in a number of different contexts, one focusing on the variation of polarization with frequency, another measuring general relativistic effects on pulse times of arrival, etc. Nevertheless, our experience has shown that there exist many common tasks associated with pulsar data analysis that can be standardised within a modern open source development environment.

As a demonstration of the types of operations performed on pulsar data, consider the specific example of pulse time of arrival calculation. In order to increase the S/N of each observation, data are often integrated (scrunched) by several factors in one or more of the available dimensions. Each resultant profile is then cross-correlated with a high S/N standard profile known as a template, yielding an estimate of the longitudinal offset between the two. This offset is added to the reference epoch associated with a fiducial point in the observed pulse profile, yielding an arrival time in the reference frame of the observatory, which is later converted into a barycentric arrival time using a Solar system ephemeris. This data reduction operation involves a number of typical tasks, including loading the arrays of numbers that represent the folded profiles and computing sums, products, rotations, weighted averages, and correlations of these arrays; sometimes in the Fourier domain. Most of these various operations must be performed in a manner consistent with the observational parameters, taking into account dispersive delays, observation time stamps and relative weightings of different frequency channels, for example. At each step, the software must also ensure that all parameters are updated accordingly.

### 4.1.5 Scope and Design of PSRCHIVE and PSRFITS

It should be noted that the pulsar data under consideration represents a point near the end of the typical pulsar data reduction chain. The software presented in this chapter is not intended for the direct handling of radio data such as that recorded by baseband systems, or for the purposes of performing computationally expensive offline searching, although some support for the storage of such data is provided in PSRFITS. The code is also not designed to perform any phase-coherent dispersion removal or formation of filter-bank data; these techniques are treated as separate computational tasks. Code for such data reduction is also available from the repository at the Swinburne Centre for Astrophysics and Supercomputing under

the umbrella name of BASEBAND_DSP[1], a general library for use in digital signal processing.

The PSRCHIVE and PSRFITS schemes were designed from the beginning to form an object-oriented framework into which existing algorithms and data structures could be transplanted. By introducing layers of abstraction between the various levels of responsibility, the design remains both flexible and extensible. For example, different telescopes and instruments require the storage of different types of information, including configuration parameters, observatory and instrumental status information, and other site-specific data. Because there is no way of knowing exactly what future systems might include, both PSRCHIVE and PSRFITS implement a generalised scheme for incorporating arbitrarily complex data extensions, as described in Sections 4.2.3 and 4.3.2.

In addition, a basic framework of crucial parameters common to all pulsar observations and a wide variety of fundamental data reduction algorithms, such as those described in Section 4.1.4, have been implemented. Each of these basic data structures and reduction operations may be used in the composition of more complex scientific processing algorithms. By virtue of continued development amongst the authors, the PSRCHIVE library includes an extensive array of high-level algorithms for use in the calibration, visualisation, and analysis of pulsar data; these can be used immediately on any of the supported file formats.

PSRCHIVE and PSRFITS were developed in parallel and are presented in the hope that they will promote increased data portability. The PSRFITS file format also serves as an example of how to incorporate other, pre-existing file formats into the new scheme. After two years of development, the code is now ready for formal release to the wider pulsar community. In the following sections, we describe the implementation of the new schemes and outline the specific advantages that they offer.

## 4.2 Implementation Overview

### 4.2.1 Object-Oriented Programming

The modularity and extensibility required of our new scheme suggested an object-oriented approach. Since much of the existing Swinburne analysis code had already been written in both the C and C++ programming languages, it seemed a natural step to progress in C++. The concepts of object classes and inheritance provided and enforced by the syntax of this language offer a sound foundation on which to develop. In particular, object-oriented design has aided in the realization of simultaneous support of multiple file formats. We are aware that a majority of pulsar research groups prefer to write a more procedural style of code, using FORTRAN or C. However, we feel that the benefits of an object-oriented approach to data

[1] http://astronomy.swin.edu.au/pulsar/software/libraries/dsp

processing significantly outweigh the potential learning curve involved in becoming proficient with C++.

### 4.2.2 Basic Class Structure

The required functionality of PSRCHIVE is built around a core framework of C++ object classes. The fundamental unit of all pulsar observations is the individual pulse **Profile**, a one-dimensional array of floating point numbers, indexed by pulse phase. The **Integration** is a two-dimensional vector of **Profile** instances, indexed by frequency and polarisation, as measured over a particular epoch. In turn, the **Archive** is a one-dimensional vector of **Integration** instances, indexed in one of a number of possible ways (normally time). Each of these classes implement a wide range of basic data manipulation and processing operations.

In the language of C++, we define the namespace **Pulsar**, which contains the three *base classes*: **Pulsar::Archive**, **Pulsar::Integration**, and **Pulsar::Profile**. In addition, there are other object classes in the **Pulsar** namespace that deal with specific tasks related to pulsar data analysis. For example, the **Pulsar::Calibration** class employs various mathematical models of the instrumental response to calibrate polarimetric observations (van Straten 2004).

### 4.2.3 Use of Data Abstraction

The three base classes implement a wide variety of basic algorithms, known as *methods*, that are commonly used in pulsar data analysis. However, they do not require knowledge of any specific details related to system architecture, enabling their use as templates upon which to base lower-level development. These templates define the minimum set of parameters, known as *attributes*, required to implement the data analysis methods, including observational parameters such as the name of the source, centre frequency, bandwidth, etc. At the level of the **Pulsar::Archive** and **Pulsar::Integration** base classes, nothing is known about how data are stored on permanent media or in computer memory.

The necessary task of translating between the two realms is performed by *derived classes* that *inherit* the base classes. In order to inherit a base class, it is necessary for the derived class to provide access to the required attributes and to implement the methods used to read and write the data stored on disk. Therefore, for each specific file format represented in the PSRCHIVE scheme, there corresponds a derived class that inherits **Pulsar::Archive**. The syntax for the data access and file input/output methods is defined by the base class and enforced by the C++ compiler, allowing all derived classes to be treated as equal. Therefore, high-level code can be written in the language of the base class definition without the need for considering the implementation details of the derived classes. This abstraction, which is crucial to the flexibility of the PSRCHIVE scheme, is demonstrated by the Unified Modeling Language (UML) class diagram shown in Fig. 4.1.

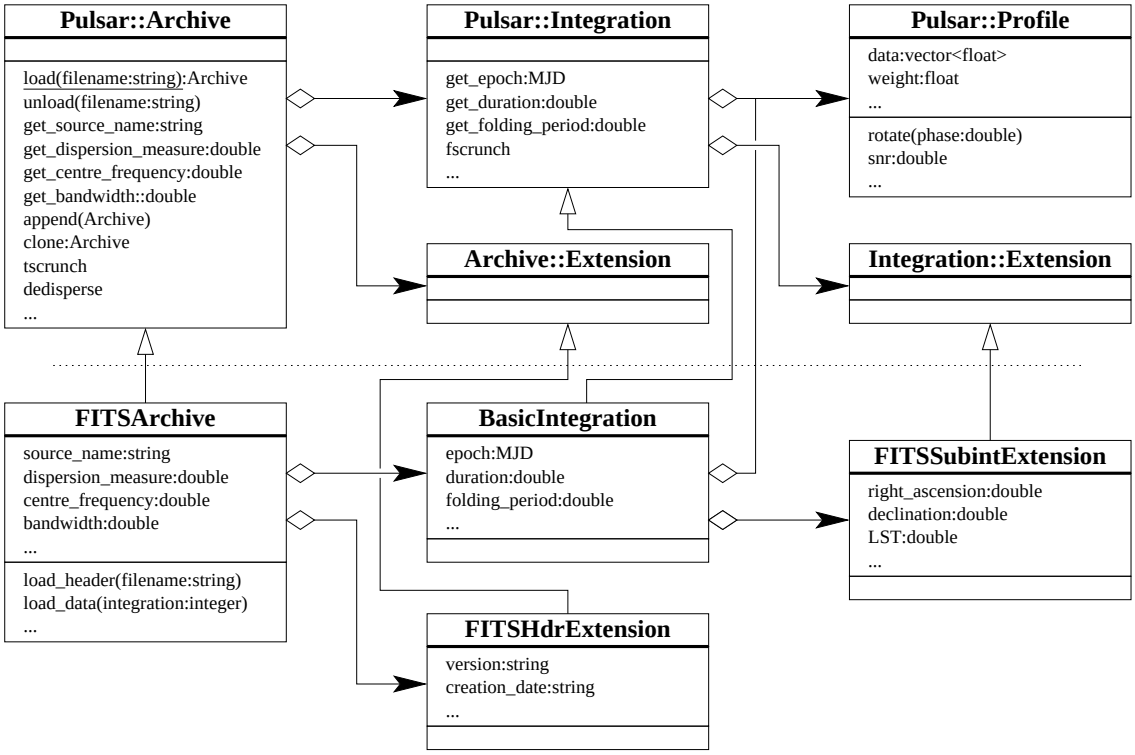


**Figure 4.1: Class diagram of a portion of the PSRCHIVE library. The abstract base classes are shown above the dotted line. Below this line, the FITSArchive class implements Pulsar::Archive attribute storage and access methods, as well as methods for loading and unloading data to and from a PSRFITS file. The combined use of composition and inheritance enables complex structures and behaviours to be constructed using modular components.**

## 4.3 PSRFITS

### 4.3.1 A Standard Format for Pulsar Data Storage

One of the motivating factors behind the development of the PSRCHIVE scheme was the alleviation of problems associated with highly specific and non-portable data storage formats. This effort has highlighted several compelling reasons for the pulsar community to move towards a more modular and standardised storage format. For instance, the existence of a standard file format would significantly decrease the amount of effort required to integrate and test new instrumentation. Historically, file formats have accreted features as they became desirable or necessary. Given the wealth of past experience available, it seems a logical step to define a new format that encompasses a wide range of features from the beginning and is written in a modular way so as to enable rapid, backwards-compatible upgrades. Indeed, one particular standard storage format has already won wide acclaim within the astronomical community; the Flexible Image Transport System (FITS) (Hanisch et al. 2001) has been in widespread use for approximately three decades and has evolved into a highly adaptable data storage scheme[2]. The format has been placed under the administration of the IAU FITS Working Group[3] and a wide array of software is available for FITS file manipulation. The NASA High Energy Astrophysics Science Archive Research Centre[4] provides useful libraries and applications for manipulation and interrogation of FITS-based files. For example, the program fv has made the process of testing and debugging the relevant PSRCHIVE software much more straight-forward.

In accordance with FITS standards, a PSRFITS file consists of a primary header-data unit (HDU) followed by a series of extension HDUs. The primary HDU contains basic information such as telescope identification and location, observation start time, etc. Extension HDUs, formatted as binary tables, contain specific information related to the observation such as the pulsar ephemeris, calibration data, and the pulsar data. Although PSRFITS is primarily designed to store folded or single-pulse profile data, it can also accommodate continuous time series data.

A useful feature of the standard FITS input/output routines is that new HDUs and header parameters may be added transparently – if they are unknown to the reading program, they are ignored. Furthermore, unused HDUs need not be written, even though they are present in the definition. This feature allows, for example, a user group to add information particular to a certain instrument without compromising use of the definition by other groups.

A novel feature of the PSRFITS definition is the inclusion of HDUs containing 'history' information. For example, the first line of the Processing History HDU contains information about the data acquisition program and the initial structure of

[2] http://archive.stsci.edu/fits/fits_standard
[3] http://www.cv.nrao.edu/fits/traffic/iaufwg/iaufwg.html
[4] http://heasarc.gsfc.nasa.gov/docs/heasarc/fits.html

the file. Subsequent lines record details of modifications to the structure or data (for example, partial or complete dedispersion or interference excision). A permanent record of the steps that have been applied during data reduction has proven to be of great value when later assessing the quality and validity of observational data.

### 4.3.2 The PSRFITS Definition

The current version of the PSRFITS definition file is available on the ATNF web pages[5]. Table 4.1 describes the basic header-data units included in the definition.

In addition to the Main and Processing History HDUs, a number of optional HDUs have been defined for general use with a variety of instrumentation. These enable the storage of important status and diagnostic information about the observation, and demonstrate the modularity and extensibility of the PSRFITS file format. The physical parameters stored in the Ephemeris History HDU are based on the pulsar timing program, TEMPO[6]. From the ephemeris parameters are derived the polynomial coefficients (polyco) used to predict the apparent pulsar period and phase at the epoch of the observation; these coefficients are stored in the Polyco History HDU. As improved physical parameters become available, the data may be reprocessed, leading to new rows in the Ephemeris and Polyco history tables. The calibration and feed cross-coupling HDUs are designed to work with the routines in the **Pulsar::Calibration** class. Owing to the intrinsic modularity of FITS, these additional HDUs are all optional; in fact, it is not even strictly necessary to include any Integration data in a PSRFITS file. For example, the polarimetric calibration modeling program creates a file containing only the feed cross-coupling, injected calibration polarisation, and flux calibration HDUs. This modularity is similar to that made available through the use of VOTable[7] XML standards and it is likely that PSRFITS could in future be incorporated into the International Virtual Observatory system with a minimum of effort.

## 4.4 Working with the PSRCHIVE Scheme

### 4.4.1 The Standard Application Set

The PSRCHIVE scheme includes an extensive set of pre-written application programs that can be used to manipulate pulsar data in various ways. These include both command line tools and graphical user interfaces built using Trolltech's Qt[8], a C++ toolkit for multi-platform GUI and application development. Table 4.2 presents a list of applications included in the package at the time of publication, with a brief description of each.

[5] http://www.atnf.csiro.au/research/pulsar/psrfits
[6] http://www.atnf.csiro.au/research/pulsar/tempo
[7] http://www.ivoa.net
[8] http://www.trolltech.com/products/qt/index.html

| HDU Title | Description |
|---|---|
| Main header | Observer, telescope and receiver information, source name and observation date and time |
| Processing history | Date, program and details of data acquisition and each subsequent processing step |
| Digitiser statistics | Digitiser mode and count statistics |
| Digitiser counts | Digitiser mode and count rate distribution |
| Original bandpass | Observed bandpass in each polarisation averaged over the observation |
| Coherent dedispersion | Parameters for coherent dedispersion of baseband data |
| Ephemeris history | Pulsar parameters used to create or modify profile data |
| Polyco history | Elements of the polyco file used to predict the apparent pulsar period |
| Flux calibration | System temperature and injected noise calibration data as a function of frequency across the bandpass |
| Injected calibration polarisation | Apparent polarisation of the injected noise calibration signal as a function of frequency |
| Feed cross-coupling | Parameters of feed cross-coupling as a function of frequency |
| Integration data | Pulse profiles or fast-sampled data as a function of time, frequency and polarisation |

**Table 4.1: PSRFITS – A summary of the current definition**

| Application | Description |
|---|---|
| pav | Archive data visualisation. Based on the PGPLOT graphics subroutine library with a simple command line interface |
| vap | Archive header examination, allowing multiple user selectable header parameters to be printed as ASCII to the terminal |
| pam | Archive manipulation, compression and processing |
| pat | Pulse profile arrival time calculation, based on cross correlation with a standard template profile |
| pas | Standard profile phase alignment, for timing with multiple standard template profiles |
| paz | Radio frequency interference mitigation tool including manual and automated channel zapping and sub-integration removal |
| pac | Archive polarimetric and flux calibration tool based on a user-selectable set of advanced algorithms |
| pcm | Polarimetric calibration modeling, creates instrumental response transformations for use with pac |
| psrgui | Interactive point-and-click data visualisation with a Qt graphical interface |
| psradd | Combination of multiple archives for formation of high S/N profiles |
| rhythm | A graphical interface for pulse arrival time fitting |

**Table 4.2: Standard applications included with PSRCHIVE**

Readers may note that the modular philosophy at the heart of PSRCHIVE extends all the way through to the user level applications. Each program tends to be small and focused on a specific task, be it data compression, timing, RFI mitigation, etc. This greatly simplifies development and maintenance compared to having one monolithic program with multiple purposes.

### 4.4.2 PSRCHIVE as a Development Environment

PSRCHIVE was designed to provide users with more than just a set of pre-made applications. The classes, libraries and examples provided are intended to simplify the task of building new processing tools. To a large extent, developers who build on the PSRCHIVE scheme do not have to directly manipulate the arrays of pulse profile amplitudes. Instead, member functions of the various classes can be called to perform basic operations like baseline removal and phase rotation. This has the dual benefit of labour saving both in the initial development phase and in the debugging phase, as both the authors and other users have already verified and tested the provided routines. In case direct access to the profile amplitudes is required, we also provide interface functions that return C style arrays. In the experience of the authors, the extra layer of abstraction provided by the PSRCHIVE scheme can cut down the time between program concept and full implementation to a matter of hours. New applications can be built with only a few lines of code. For example, to remove the system noise floor, compress all frequency channels and output the processed archive:

```
# include "Pulsar/Archive.h"

int main() {

  Pulsar::Archive* arch = 0;

  arch = Pulsar::Archive::load("filename");

  arch->remove_baseline();
  arch->fscrunch();

  arch->unload();

}
```

This simple program defines a pointer to a **Pulsar::Archive** and calls the generic **Pulsar::Archive::load** routine, which takes a filename argument, applies a number of tests to the file on disk and decides whether or not it understands the particular data format. If so, it summons the appropriate derived class to read the data from

disk. Once the data have been loaded, the **Pulsar::Archive::remove_baseline** function is called.

```
void Pulsar::Archive::remove_baseline (float phase, float width)
{
  try {

    if (phase == -1.0)
      phase = find_min_phase (width);

    for (unsigned isub=0; isub < get_nsubint(); isub++)
      get_Integration(isub) -> remove_baseline (phase, width);

  }
  catch (Error& error) {
    throw error += "Pulsar::Archive::remove_baseline";
  }
}
```

The **Pulsar::Archive::remove_baseline** function takes two arguments: the `phase` and `width` of the off-pulse baseline. Both arguments are assigned default values in the `Archive.h` header file; if `phase` is left unspecified, then the off-pulse baseline phase will be found using the **Pulsar::Archive::find_min_phase** method; if `width` is unspecified, then a default value will be used. The **Pulsar::Archive::remove_baseline** method makes multiple calls to the lower level **Pulsar::Integration::remove_baseline** routine, which performs the actual modification of amplitudes as follows:

```
void Pulsar::Integration::remove_baseline (float phase, float width)
{

  if (Pulsar::Integration::verbose)
    cerr << "Pulsar::Integration::remove_baseline entered" << endl;

  try {

    if (phase == -1.0)
      phase = find_min_phase (width);

    vector<float> phases;
    dispersive_phases (this, phases);

    for (unsigned ichan=0; ichan<get_nchan(); ichan++) {
```

```
      float chanphase = phase + phases[ichan];

      for (unsigned ipol=0; ipol<get_npol(); ipol++)
        *(profiles[ipol][ichan]) -=
          profiles[ipol][ichan] -> mean (chanphase, width);

    }

  }
  catch (Error& error) {
    throw error += "Integration::remove_baseline";
  }
}
```

This nested structure reduces the length of high-level routines, allowing actual computations to be done at the level of abstraction that best suits the task. Likewise, the **Pulsar::Integration::remove_baseline** routine calls various member functions of both the **Pulsar::Integration** and **Pulsar::Profile** classes, computing the pulse phase at which the minimum baseline level occurs in the total intensity of the entire band. Adjustments for dispersive delays in each channel are performed and the mean level at this phase is individually removed from each **Pulsar::Profile** stored in the **Pulsar::Integration**. Developers should also note the extensive use of **try/catch** blocks and a specifically designed **Error** class that carries descriptive information about any exceptions thrown back to the calling procedure.

## 4.5 Resources and Availability

### 4.5.1 Obtaining and Compiling the Code

PSRCHIVE is freely available to the worldwide academic community. It is currently held in a repository at Swinburne University of Technology in Melbourne, Australia and may be accessed via the Concurrent Versions System[9] (CVS). As it is distributed in source code form, some experience with programming and compilation is necessary; however, installation can be performed in a fairly simple step-by-step manner thanks to the GNU auto-configuration[10] system. The code is compatible with all versions of the GNU Compiler Collection[11] (GCC) between 2.95 and 3.2.2 (at the time of publication) and is routinely tested on the Linux operating system. Every effort will be made to ensure compatibility with future GCC releases. The software has also been successfully compiled on Solaris and Mac OS X.

[9] http://www.cvshome.org
[10] http://www.gnu.org/software/autoconf/
[11] http://gcc.gnu.org

The PSRCHIVE scheme makes use of several external libraries, including the Starlink Project[12] SLALIB package. It also requires at least one external FFT library and includes wrappers that provide compatibility with either FFTW[13] (available under the GPL) or Intel MKL[14] (commercially available from Intel). The PGPLOT[15] graphics subroutine library is also an integral part of the scheme.

Full documentation including instructions for download and installation of the package are available online at the following URL: http://psrchive.sourceforge.net. At present, read-only access to the CVS repository is granted upon receipt by the developers of a Secure Shell v2.0 public key that is used to allow remote access to the Swinburne gateway server. Write permission to the repository requires a computing account with the Swinburne Centre for Astrophysics and Supercomputing. Please direct all enquiries regarding access to **psrchive@astro.swin.edu.au**. In addition, users can download a pre-packaged source code distribution file that includes all the necessary files for a basic installation.

### 4.5.2 Online Documentation

All PSRCHIVE reference documentation is maintained online as part of a project hosted by SourceForge.net[16]. In addition to the online descriptions, each command line application has a **-h** option that displays a quick summary of how to use the program. The library of C++ classes is extensively documented using the Doxygen[17] system; the source code contains tagged comments from which the online manual is automatically generated. This manual is intended as a reference to programmers as it primarily describes the member functions available in each class and the syntax of their arguments.

### 4.5.3 Support Services

Although we provide no official support for the software, we are willing to assist with PSRCHIVE related problems as time permits. General queries regarding installation or operation can be addressed to **psrchive@astro.swin.edu.au**. The PSRCHIVE SourceForge project web site also provides a mechanism for reporting serious bugs and requesting new features.

[12] http://www.starlink.rl.ac.uk
[13] http://www.fftw.org
[14] http://www.intel.com/software/products/mkl/
[15] http://www.astro.caltech.edu/~tjp/pgplot/
[16] http://www.sourceforge.net
[17] http://www.doxygen.org

## 4.6 Conclusion

The task of organising astronomical data into a logical format lends itself surprisingly well to the object-oriented programming paradigm. The combination of PSRCHIVE and PSRFITS provides a powerful, ready-to-use pulsar data archival and reduction system that can be rapidly adapted to new instruments. We hope that the ready availability of an open source data reduction framework will facilitate large scale collaborative projects, such as an extended pulsar timing array (Foster & Backer 1990). Therefore, we encourage both scientists and engineers involved with pulsar data acquisition and reduction to consider taking advantage of these packages.

# Chapter 5

# PSR J1022+1001: Profile Stability and Precision Timing

"...man will occasionally stumble over the truth, but usually manages to pick himself up, walk over or around it, and carry on."

*Winston S. Churchill*

In this chapter, we present an investigation of the morphology and arrival times of integrated radio pulses from the binary millisecond pulsar PSR J1022+1001. This pulsar is renowned for its poor timing properties, which have been postulated to originate from variability in its average pulse profile. Although a sub-class of long-period pulsars are known to exhibit mode changes that give rise to very large deviations in their integrated profiles, this was the first millisecond pulsar thought to have an unstable mean profile. As part of a precision timing program at the Parkes radio telescope we observed this pulsar between January 2003 and March 2004 using a coherent dedispersion system (CPSR2). A study of morphological variability during our brightest observations suggests that the pulse profile varies by at most a few percent, similar to the uncertainty in our calibration. Unlike previous authors, we find that this pulsar times extremely well. In 5 min integrations of 64 MHz bands we obtain a weighted RMS residual of just 2.27 $\mu$s. The reduced $\chi^2$ of our best fit is 1.43, which suggests that this pulsar can be timed to high accuracy with standard cross-correlation techniques. Combining relativistic constraints with the pulsar mass function and consideration of the Chandrasekhar mass limit on the white dwarf companion, we can constrain the inclination angle of the system to lie within the range $37^o < i < 56^o$. For reasonable pulsar masses, this suggests that the white dwarf is at least 0.9 $M_\odot$. We also find evidence for secular evolution of the projected semi-major axis.

## 5.1 Introduction

Pulsar timing is a highly versatile experimental technique that is used to provide estimates of everything from the age and magnetic field strengths of pulsars, allowing population studies (Lyne, Manchester & Taylor 1985; Lorimer et al. 1993), to micro-arc-second positions and tests of relativistic gravity (van Straten et al. 2001). Central to the technique is the maxim that pulsars have average pulse profiles that are inherently stable (Helfand, Manchester & Taylor 1975), which can be cross-correlated with a template to provide accurate times of arrival (Manchester & Taylor 1977). Once corrections are made to account for the Earth's position with respect to the solar system barycentre, a model can be fit to the data, transforming pulsars from astrophysical curiosities into powerful probes of binary evolution and gravitational theories. One spectacular example is the recently-discovered double pulsar system (Lyne et al. 2004).

The discovery of the millisecond pulsars PSR B1937+21 and PSR B1855+09 (Backer et al. 1982; Segelstein et al. 1986) and their subsequent timing (Kaspi, Taylor & Ryba 1994) revealed that millisecond pulsars were very rotationally stable, rivaling the best atomic clocks. Their lower magnetic fields appeared to give rise to a fractional stability that far exceeded that of the normal pulsar population (Arzoumanian et al. 1994). This stability, combined with their small rotation periods, gives the millisecond pulsars sub-microsecond timing potential provided sufficient S/N can be achieved. This provides interesting limits on a wide range of astrophysical phenomena (Stairs 2004).

An array of accurately timed MSPs can even be used to constrain cosmological models and search for long-period (nHz) gravitational waves (Foster & Backer 1990). Such an array requires frequent observations of a selection of pulsars that are known to have sub-$\mu$s timing precision, preferably in two or more radio frequency bands. Implementation of such an array would provide an opportunity to directly detect gravitational radiation in a regime that complements the sensitivity range of ground-based detectors, to test the accuracy of solar system ephemerides and to construct an entirely extra-terrestrial timescale. The question of whether a suitable sample of millisecond pulsars positioned throughout the entire celestial sphere can be selected from the most recent pulsar catalogue (Manchester et al. 2005a) is under active assessment, through a collaboration between the Australia Telescope National Facility and Swinburne University of Technology.

The feasibility of a timing array depends in part upon the intrinsic rotational stability of the MSPs and their lack of pulse profile variation. Cognard & Backer (2004) have recently discovered a "glitch" in the millisecond pulsar PSR B1821–24, similar to those discovered in younger pulsars (Shemar & Lyne 1996). Perhaps more disturbingly, Kramer et al. (1999) report that the binary millisecond pulsar PSR J1022+1001 exhibits pulse shape variations that ruin its timing. The authors argued that this pulsar exhibits significant changes in its pulse morphology on $\sim$5 min timescales and narrow bandwidths. They interpreted these variations

as the source of the pulsar's unusually poor timing properties. By modeling the profile with Gaussian components, Kramer et al. (1999) improved the timing and argued that the trailing component of the pulse was more stable than the leading component.

The precision with which astronomers can predict pulse arrival times has been steadily improving over the past few decades with the advent of new technology and methodology. The charged interstellar medium disperses radio pulses and broadens the pulse profile across finite receiver bandwidths. Astronomers crave both the wide bandwidths that permit high S/N integrated profiles, and the resolution of sharp features that permit strong cross-correlation with a standard template in order to minimise timing errors. Pulsar astronomy has always benefited from adopting new technologies that give increased time resolution and help defeat the deleterious effect of the interstellar medium on the propagation of radio pulses.

Incoherent devices such as analogue filter-banks are prone to systematic errors due to interstellar scintillation and imperfect frequency responses. Digital spectrometers can overcome some of these inadequacies but Hankins & Rickett (1975) described a computational technique (coherent dedispersion) that removes pulse dispersion "perfectly" almost three decades ago. Unfortunately, at the time, computational hardware was not capable of processing all but the most modest of bandwidths. The exponential growth of computational power has permitted the development of new pulsar instruments that are capable of coherently dedispersing large bandwidths in near real time. Such an instrument is CPSR2, a baseband recorder that permits near-real time coherent dedispersion of $2\times64$ MHz bands using a cluster of high-end workstations. These new instruments represent the best opportunity to study small variations of integrated pulse profiles because they deliver an immunity against systematic errors induced by narrow-band scintillation and dispersion measure smearing. They also provide full polarimetry and multi-bit digitisation, radio frequency interference rejection and the opportunity to apply accurate statistical corrections that help eliminate systematic errors induced by discretisation (Jenet & Anderson 1998).

Motivated by the arrival of CPSR2, we have commenced a timing campaign of the best millisecond pulsars in an effort to make progress towards the aims of the timing array. PSR J1022+1001 was included because of its high flux density, and the paucity of better candidates at the hour angle it is visible from the Parkes radio telescope. PSR J1022+1001 is a recycled or millisecond pulsar with a pulse period of approximately 16 ms. It orbits once every 7.8 days in a binary system with a companion whose minimum mass can be derived from the pulsar mass function, Eq. 5.1, if we assume that the orbit is face-on and the pulsar has a mass of $1.35\,\mathrm{M}_\odot$.

$$f(m_\mathrm{p}, m_\mathrm{WD}) = \frac{m_\mathrm{WD}^3 \sin^3 i}{(m_\mathrm{p} + m_\mathrm{WD})^2} = \frac{4\pi^2}{G}\frac{a^3 \sin^3 i}{P_\mathrm{b}^2}. \tag{5.1}$$

Here $m_\mathrm{p}$ and $m_\mathrm{WD}$ refer to the mass of the pulsar and white dwarf respectively, $i$ is the inclination angle of the orbit, $P_\mathrm{b}$ is the orbital period and $G$ is Newton's

gravitational constant. For PSR J1022+1001, this corresponds to a mass of at least $0.7\,M_\odot$.

Measurements of the dispersion measure along the line of site to the pulsar, combined with the Taylor & Cordes (1993) galactic electron density model, place PSR J1022+1001 at a distance of roughly 600 pc from the sun, making it a relatively nearby source. Accurate timing of this pulsar could therefore lead to a greater knowledge of its orientation, companion masses, distance (via parallax) and proper motion. Unfortunately, this pulsar is very near the ecliptic plane and hence pulsar timing is only good at accurately constraining its position in one dimension. It is therefore a good target for very long baseline interferometry, as this could accurately determine the position and proper motion in ecliptic latitude.

In this chapter we present an analysis of the first 15 months of CPSR2 observations of PSR J1022+1001, demonstrating that this pulsar can be timed to high accuracy using standard techniques. In section 5.2 we describe the observing system and methodology, followed by a description of our data reduction scheme, including polarimetric calibration. Section 5.3 describes the results of a search within our data set for variations of the type reported by Kramer et al. (1999). Pulse arrival times calculated from this data set are analysed in section 5.4, followed by a summary of the newly derived pulsar spin and binary system parameters, including those that may become significant in the future. In section 5.5 we discuss the implications of our work for this pulsar and precision timing programs in general.

## 5.2 Observations

### 5.2.1 Instrumentation

The second Caltech Parkes Swinburne Recorder (CPSR2) is a baseband recording and online coherent dedispersion system that accepts 4×64 MHz IFs. Using the in-built supercomputer, data can be processed either in real time if the pulsar has a DM less than approximately $40\,\mathrm{pc\,cm^{-3}}$ at a wavelength of 20 cm, or recorded to disk and processed offline. This allows a large amount of flexibility in observing methodology. CPSR2 was commissioned in August 2002 and has been recording data on a regular basis since about November 2002. Observing sessions were primarily conducted using a standard configuration consisting of two independent 64 MHz bands, each with two orthogonal linear polarisations. The down-conversion chain was configured to make both bands contiguous, at centre frequencies of 1341 and 1405 MHz. In addition, a small number of observations were made at the widely separated frequencies of 3000 and 685 MHz. After down-conversion and filtering to create band-limited signals, each IF was fed into the CPSR2 Fast Flexible Digitiser board which performed real, 2-bit Nyquist sampling. The digitisation process generates a total of 128 MB of data every second. These samples were fed via DMA cards to two high-speed computers that divided the data into discrete blocks which were distributed via Gb Ethernet to a cluster of client machines for reduction us-

ing a software program called PSRDISP (van Straten 2003). PSRDISP performed software-based coherent dedispersion (Hankins & Rickett 1975), assuming prior knowledge of the DM along the line of sight to the pulsar and the validity of the tenuous plasma dispersion law.

### 5.2.2 Summary of Observations

Observations were conducted at the Parkes radio telescope over a period of around 400 days between January 2003 and February 2004. The central beam of the Parkes multibeam receiver (Staveley-Smith et al. 1996) was originally used at the front end, providing a system temperature of approximately 21 K at 20 cm. The multibeam receiver was removed for maintenance in October 2003 and all 20 cm data since this date were recorded using the refurbished, wide-band H-OH receiver whose observable band encompasses that of the multibeam receiver. The H-OH system is slightly less sensitive than the multibeam, with a system temperature of 25 K at 20 cm wavelengths. In addition, a new coaxial dual-band 10/50 cm receiver system was installed in place of the multibeam. During the 2004 observing sessions, data were taken with this system in order to expand our frequency coverage for the purposes of dispersion measure monitoring. Preliminary system temperature measurements of the coaxial system yield values of 30 K and 40 K for the 10 cm and 50 cm feeds respectively.

Scheduled observing sessions were typically a few days in duration, occurring once or twice a month. Individual tracks of PSR J1022+1001 varied from approximately 15 min to a few hours in duration. The data were somewhat biased towards episodes of favourable scintillation as this allowed the most efficient use of limited telescope time. In addition to the on-source tracks, most observations were immediately preceded or followed by a short (2 min) observation, taken $1^o$ South of the pulsar position, during which the receiver noise source was driven with a square wave at a frequency of $\sim$11 Hz. These calibration tracks were used to characterise the polarimetric response of the signal chain so that corrections to the observed Stokes parameters could be made during data reduction. Observations of the radio galaxy 3C218 (Hydra A) were also taken at approximately monthly intervals (though often only at 20 cm) and used to calibrate the absolute flux scale of the observing system.

### 5.2.3 Data Reduction

Data from each dual-polarisation band were reduced to a coherent filterbank with 128 spectral channels, each 0.5 MHz wide, corresponding to an effective sampling time of 2 $\mu$s. 1024 phase bins were stored across the pulse profile, yielding a time resolution for PSR J1022+1001 of about 15 $\mu$s per bin. Each individual data block recorded by CPSR2 represented 16.8 s of data (1 GB of 2-bit samples). These blocks were later combined to form integrated profiles. For this purpose, we used the application set provided with the PSRCHIVE (Hotan, van Straten & Manchester

2004) scheme. Due to the IERS[1] practice of retrospectively publishing corrections to the rotation rate of the Earth, the data presented in this chapter are those for which corrections to the Solar system barycentre could be accurately made at the time of preparation. Observations of PSR J1022+1001 are ongoing.

### 5.2.4 Processing & Calibration

The set of 16.8 s PSR J1022+1001 integrations were summed to a total integration time of 5 min. These data were calibrated to account for the polarimetric response of the observing system, then all frequency channels and Stokes parameters were combined to form a single total intensity profile corresponding to each 5 min integration. All of the receiver systems used during observations of PSR J1022+1001 were equipped with orthogonal linear feed probes, so the recorded polarisation products were corrected for relative gain and phase. Off-source calibrator observations were used to compute the relative gain and phase terms for each receiver system at the epoch of the observation, using a simple case of the scheme described by van Straten (2004). We did not attempt to correct for more subtle errors arising from cross-contamination between the two probes, primarily because sufficiently accurate models of all three receiver systems were not yet available. In the case of the central beam of the multibeam receiver, existing models indicate that imperfections in the orthogonality of the two feed probes may be as large as several degrees (van Straten 2004). These imperfections break the fundamental assumption of orthogonality that is made when applying a simple complex gain correction and induce errors in the calculated total intensity of order 1–2% (Ord et al. 2004). Errors of this magnitude are only present in the profile at phases where the fractional polarisation is close to unity.

## 5.3 Pulse Profile Stability

PSR J1022+1001 has an interesting profile morphology. At 20 cm the profile consists of two sharp peaks (the principal components), separated by about 0.05 phase units but joined by a bridge of emission (Fig. 5.1). This characteristic double-peaked shape was critical to the analysis performed by Kramer et al. (1999), who made an effort to characterise variations in the relative amplitude of the two components by normalising against one or the other and computing ratios of their height. The trailing pulse component was observed to possess almost 100% linear polarisation (Ord et al. 2004; Stairs, Thorsett & Camilo 1999; Xilouris et al. 1998), while the leading component corresponded to a local minimum in the linear polarisation fraction which is near zero (Fig. 5.1). The profile also changes rapidly as a function of frequency, only the leading component is visible at a wavelength of 10 cm (Fig. 5.2) and the amplitude ratio at 50 cm (Fig. 5.3) is nearly half that measured at

[1] http://www.iers.org

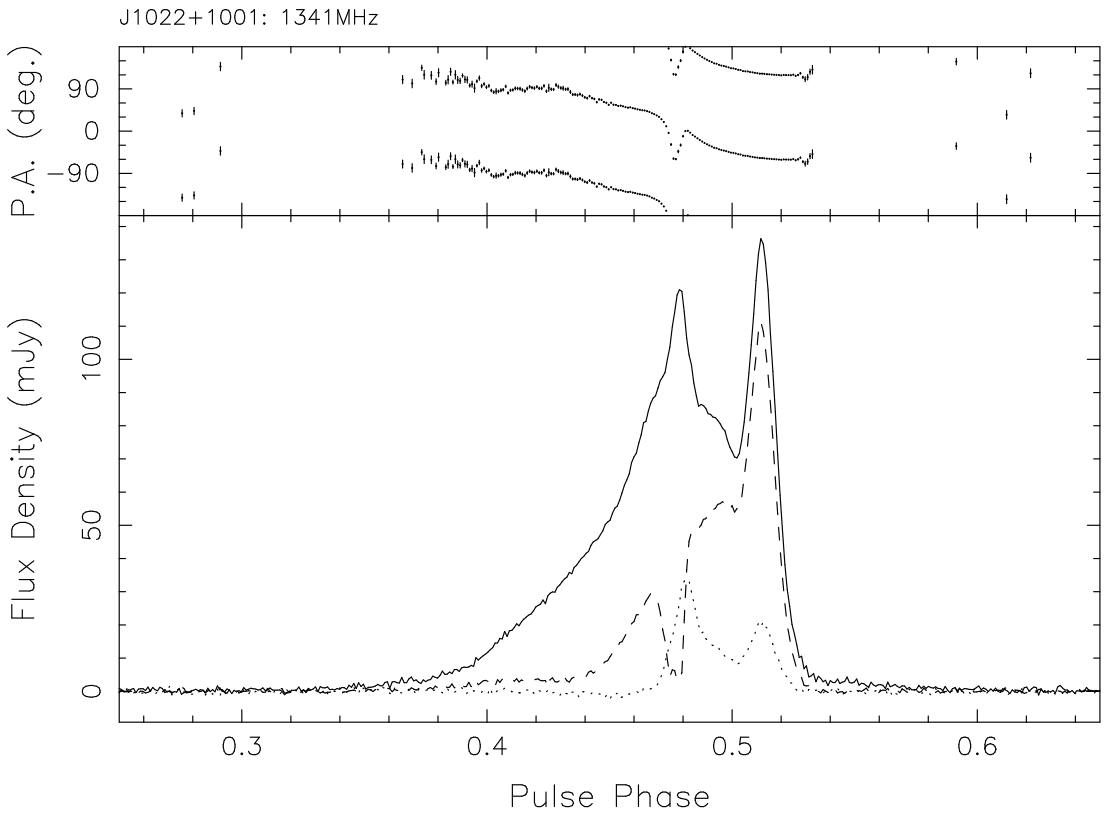


**Figure 5.1: PSR J1022+1001 average profile at a wavelength of 20 cm, formed from data taken with the central beam of the Parkes multibeam receiver. The profile in the lower panel has been calibrated for relative instrumental gain and phase. It has a centre frequency of 1341 MHz and covers a bandwidth of 64 MHz. The solid curve represents total intensity (Stokes I), the dashed curve represents linearly polarised emission and the dotted curve represents circularly polarised emission. Note the high degree of linear polarisation in the trailing component. The relative (not absolute) position angle of the linearly polarised radiation is shown in the top panel. This profile has been calibrated against the flux of Hydra A (3C218).**

20 cm (Fig. 5.1). This spectral index variation within the profile and the highly asymmetric nature of the polarised flux in the principal components, combined with the fact that PSR J1022+1001 is prone to strong episodes of scintillation, means that arrival time analysis over broad frequency bands is easily subject to the introduction of systematic errors. It also complicates any search for variability by exacerbating the effect of instrumental imperfections, especially those associated with polarimetry.

### 5.3.1 Profile Normalisation

In order to directly compare multiple integrated profiles from the same pulsar, brightness changes due to interstellar scintillation must be taken into account. This involves scaling or normalising each profile by finding a characteristic value associated with the source flux during the observation and adjusting the measured ampli-

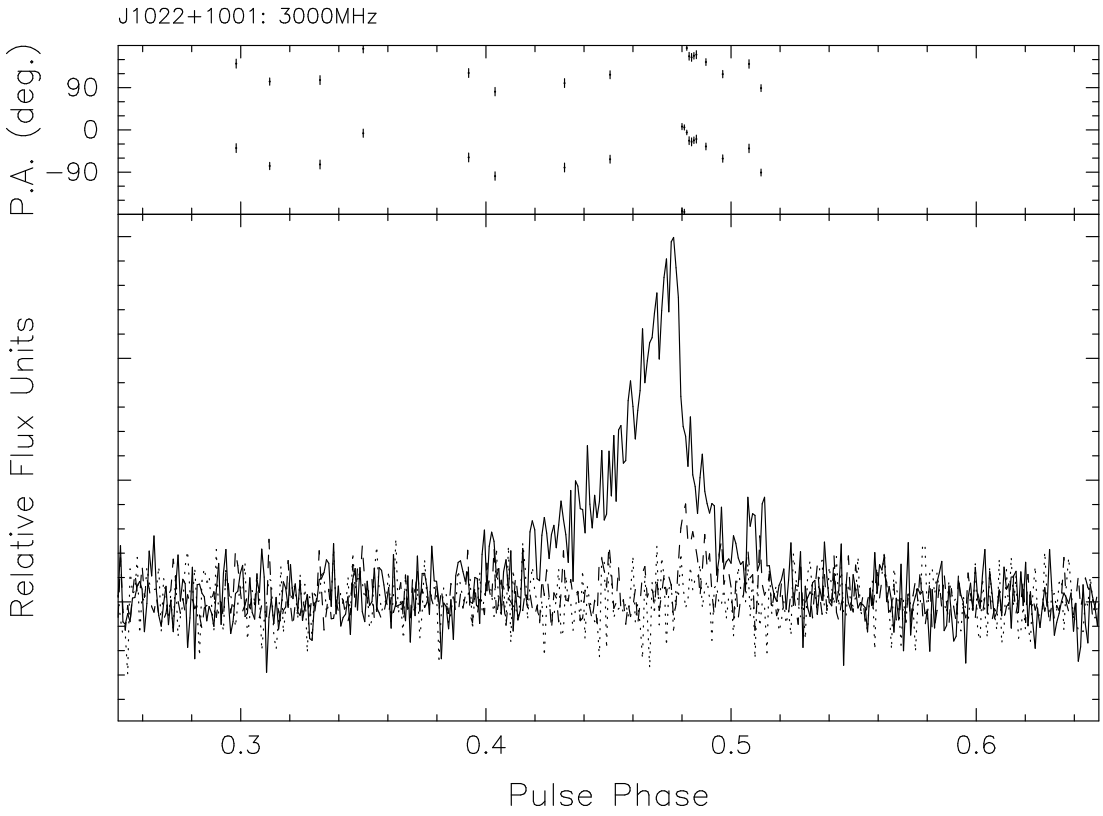


**Figure 5.2: PSR J1022+1001 average profile at a wavelength of 10 cm, formed from data taken with the new coaxial 10/50 cm receiver system at Parkes. Again, the profile has been calibrated for relative gain and phase. It has a centre frequency of 3000 MHz and a bandwidth of 64 MHz. Note the small degree of polarisation and the complete absence of the trailing pulse component. Flux calibrator observations were not made at the time these data were recorded, thus the amplitude scale is arbitrary.**

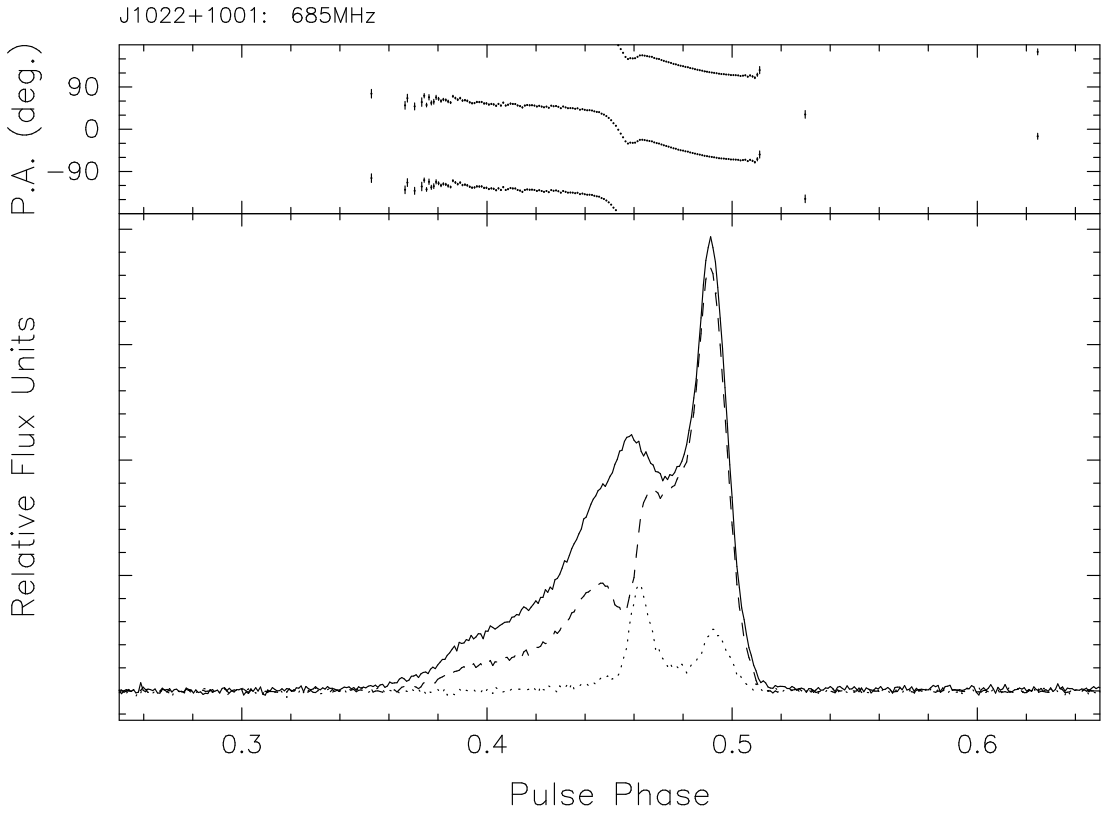


**Figure 5.3: PSR J1022+1001 average profile at a wavelength of 50 cm, formed from data taken with the new coaxial 10/50 cm receiver system at Parkes. Again, the profile has been calibrated for relative gain and phase. It has a centre frequency of 685 MHz and a bandwidth of 64 MHz. Note that the trailing component is almost 100% linearly polarised and significantly stronger than the leading component. Flux calibrator observations were not made at the time these data were recorded, thus the amplitude scale is arbitrary.**

tudes to ensure this value remains fixed across all observations. It should be noted that applying a simple relative scaling to all observed profiles masks any intrinsic variability in the total flux of the pulsar; but that is not the subject of this chapter. Using the PSRCHIVE (Hotan, van Straten & Manchester 2004) development environment, software was constructed to perform tests similar to those described by Kramer et al. (1999). In the process of analysis and testing, we developed a normalisation scheme based on the concept of difference profiles (Helfand, Manchester & Taylor 1975).

Kramer et al. (1999) described a simple method of normalisation in which the characteristic quantity associated with each integration is taken to be the amplitude of one of the two component peaks. All the amplitudes in the profile are scaled by the constant factor required to give the chosen phase location the value of unity. In the case of PSR J1022+1001, there is a choice as to which peak will be used as the reference. An alternative characteristic quantity is the total flux under the profile, or within a particular phase range. This quantity is simply the sum of all the profile amplitudes within the region of interest. Normalising by flux has the advantage that it uses no morphological information (other than the duty cycle of the pulsed emission in our case) and may therefore be better suited to detecting subtle profile variations.

In this chapter we computed fluxes after subtracting a baseline level from each profile. A running mean with a phase width of 0.7 units was used to estimate the baseline flux, which was then subtracted from each bin. In order to further reduce our sensitivity to radio frequency interference and other factors that can distort the uniformity of an observed baseline, we normalised only by the flux in the on-pulse region of the profile. This region was defined by an edge detection algorithm that measured when the total flux crossed a threshold level. This edge detection was performed on the standard template profiles at each wavelength and the phase windows so defined were held fixed for the remainder of the analysis. Unfortunately, normalising profiles by flux tends to artificially amplify noisy observations because a noise-dominated profile has a mean (and therefore a total flux after baseline removal) approaching zero. In the data reduction stage it is therefore advantageous to perform a cut on the basis of S/N. This also helps reduce contamination by corrupted profiles. A S/N threshold of 100 was deemed the best compromise between retaining a large fraction of the observed profiles and rejecting noise. This cut was applied to the set of 5 min CPSR2 integrations before any analysis was commenced.

Kramer et al. (1999)'s technique of normalising to the trailing feature has a number of limitations:

- It cannot place any limits on the variability of the feature that is used to normalise the profiles.

- Polarimetric calibration errors are maximised if the chosen component is highly polarised.

- Finite signal to noise profiles will vary due to noise even if there is no intrinsic variability in the pulsar emission, this is true of all components in the pulse.

We prefer our (more robust) flux-based normalisation scheme which normalises by the total flux in the on-pulse region. Flux normalisation allows us to properly identify which (if any) component is varying as it makes no prior assumptions. It also reduces the effect of poor calibration by averaging over the entire profile. There will still be a background of variability due to noise, but it will be spread uniformly.

### 5.3.2 The Difference Profiles

In order to detect subtle changes in pulse shape, it is possible to subtract the amplitudes of a high S/N standard template profile from an appropriately scaled and aligned copy of a given integration. This procedure should yield Gaussian noise if the observed profile precisely matches the standard; any morphological deviations will protrude above the resulting baseline. Before differences can be computed however, the observed profiles must be normalised to a standard template. We chose to normalise all profiles according to the amount of flux in the on-pulse region, as described above. In addition, cross-correlation methods were used to determine relative shifts, after which the observed profiles were rotated to align with the standard template. Because profile variability studies are sensitive to errors in polarimetric calibration, the difference profile analysis was performed on both the calibrated and uncalibrated profiles. Only the 20 cm and 50 cm wavelengths had a sufficient number of high S/N profiles to make a difference profile study feasible.

Figs. 5.4 – 5.7 show difference profiles computed from the CPSR2 observations at a wavelength of 20 cm, while Figs. 5.8 & 5.9 show the 50 cm results. The number of calibrated profiles is less than or equal to the corresponding number of uncalibrated profiles, as not all observations have corresponding calibration tracks.

As all figures show, the difference profiles are almost consistent with noise, with no alarming peaks near the leading or trailing components. This can be seen from both the difference profiles and their standard deviations. The only evidence for variability comes from Fig. 5.5 where $\sim$2% variability above the noise floor can be seen near the trailing component. However, these variations are not present in Fig. 5.4, which represents the same data set in the absence of polarimetric calibration. It would seem that the calibration procedure itself is capable of inducing profile instability, which is perhaps not surprising considering the simplicity of the model used to correct for receiver imperfections. In addition, CPSR2 has dynamic level setting that ensures the mean counts are equal in both polarisations. Provided the effective system temperature in the two polarisations is similar, calibration is almost unnecessary to form an accurate total intensity profile. It therefore seems a bizarre coincidence that these variations could be removed by simply neglecting to calibrate the data, and that they correspond to the maximum polarisation fraction in the profile. The simplest interpretation is that across 64 MHz bandwidths, this pulsar's

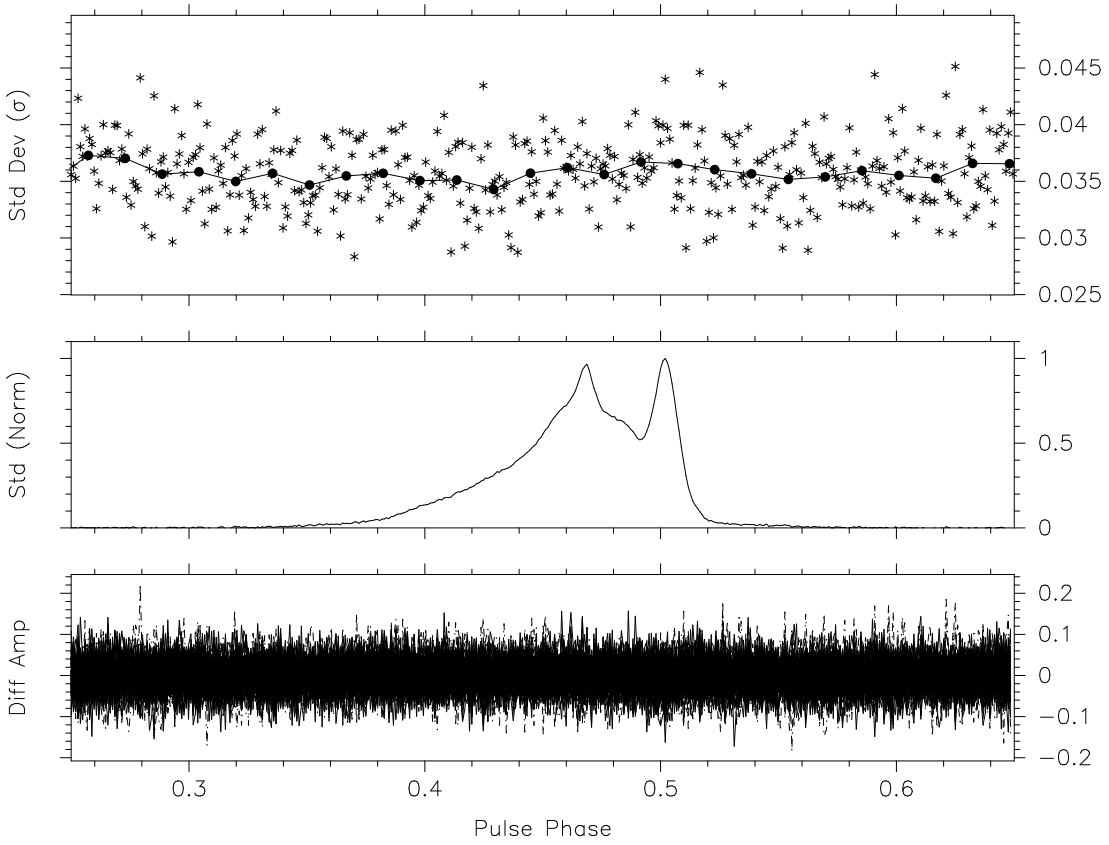


**Figure 5.4: The bottom panel shows superposed morphological difference profiles of 97, 5 min total intensity (uncalibrated) integrations at a sky frequency of 1405 MHz. The middle panel shows the mean profile. The top panel shows the standard deviation of each phase bin (stars) in the difference profiles displayed in the bottom panel and a series (filled circles) representing the mean value of the standard deviation in 64 windows, connected by lines, to aid perception of trends in the data.**

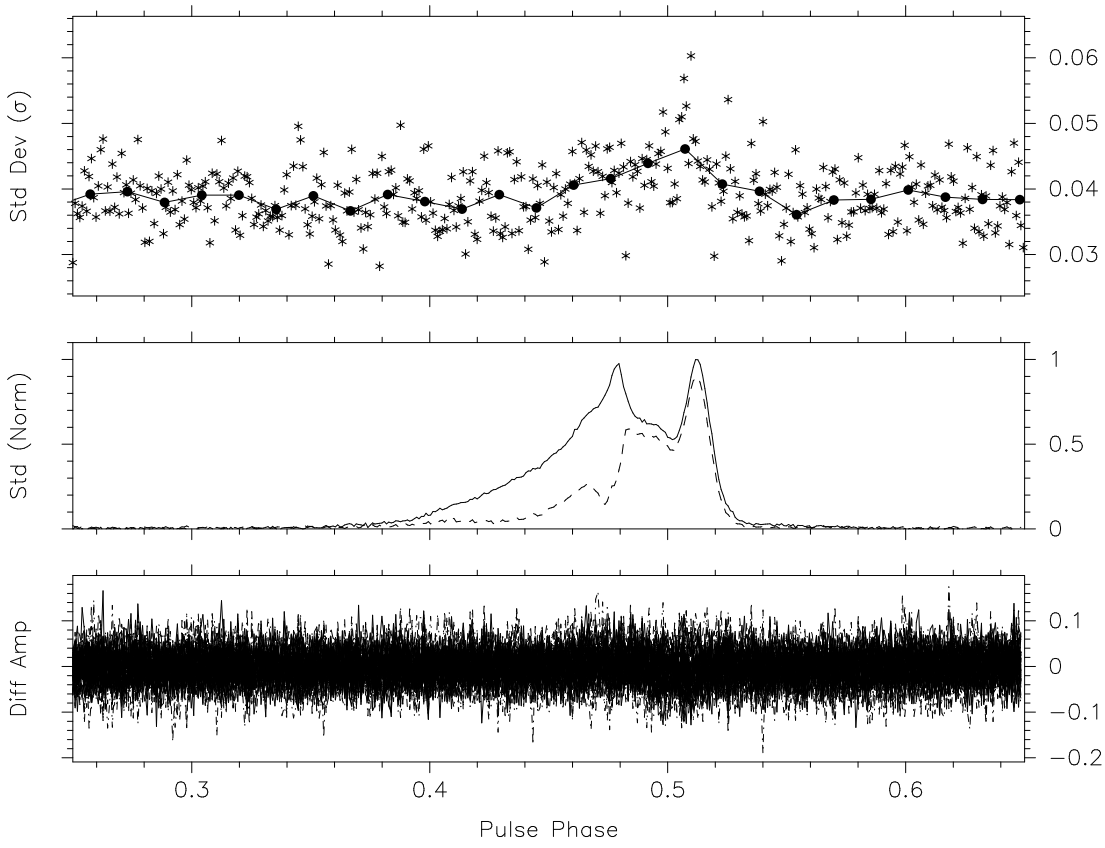


**Figure 5.5: As in Fig. 5.4, superposed morphological difference profiles of 50 total intensity (calibrated) integrations at a sky frequency of 1405 MHz. The dashed line under the solid curve in the middle panel represents the total polarised emission in the mean profile.**

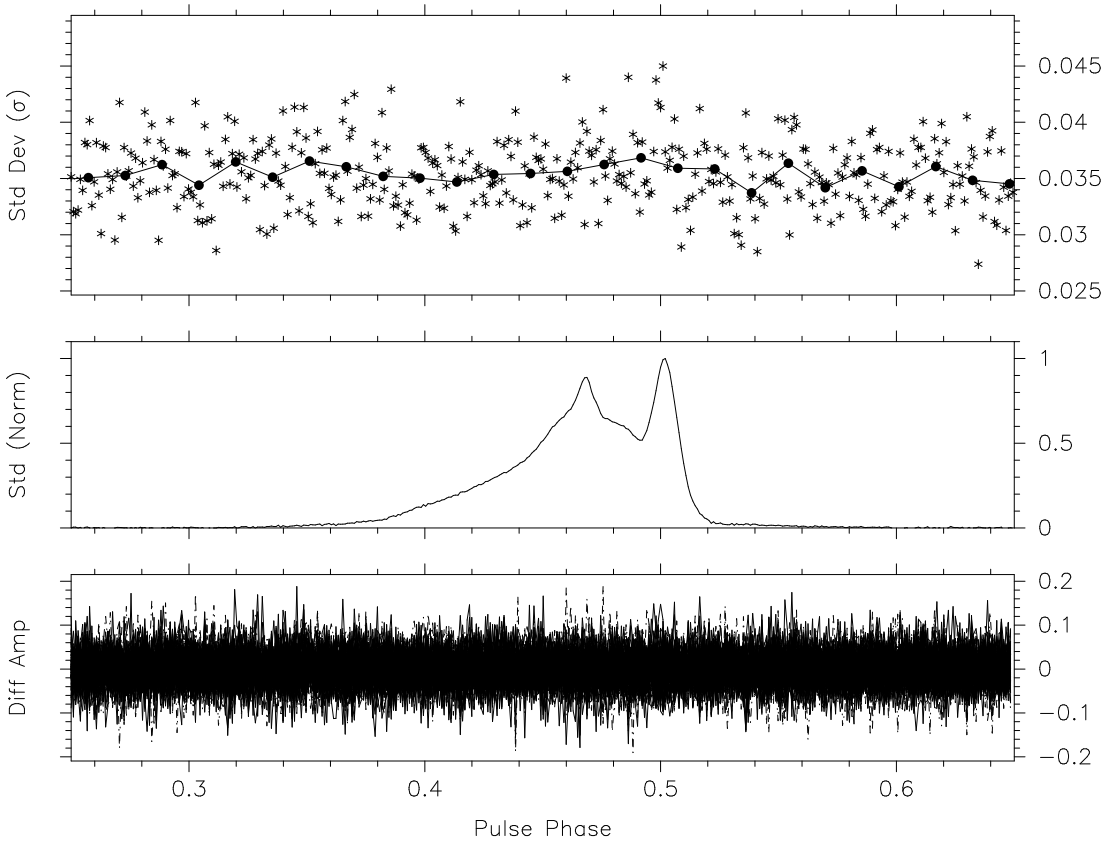


**Figure 5.6: As in Fig. 5.4, superposed morphological difference profiles of 102 total intensity (uncalibrated) integrations at a sky frequency of 1341 MHz**

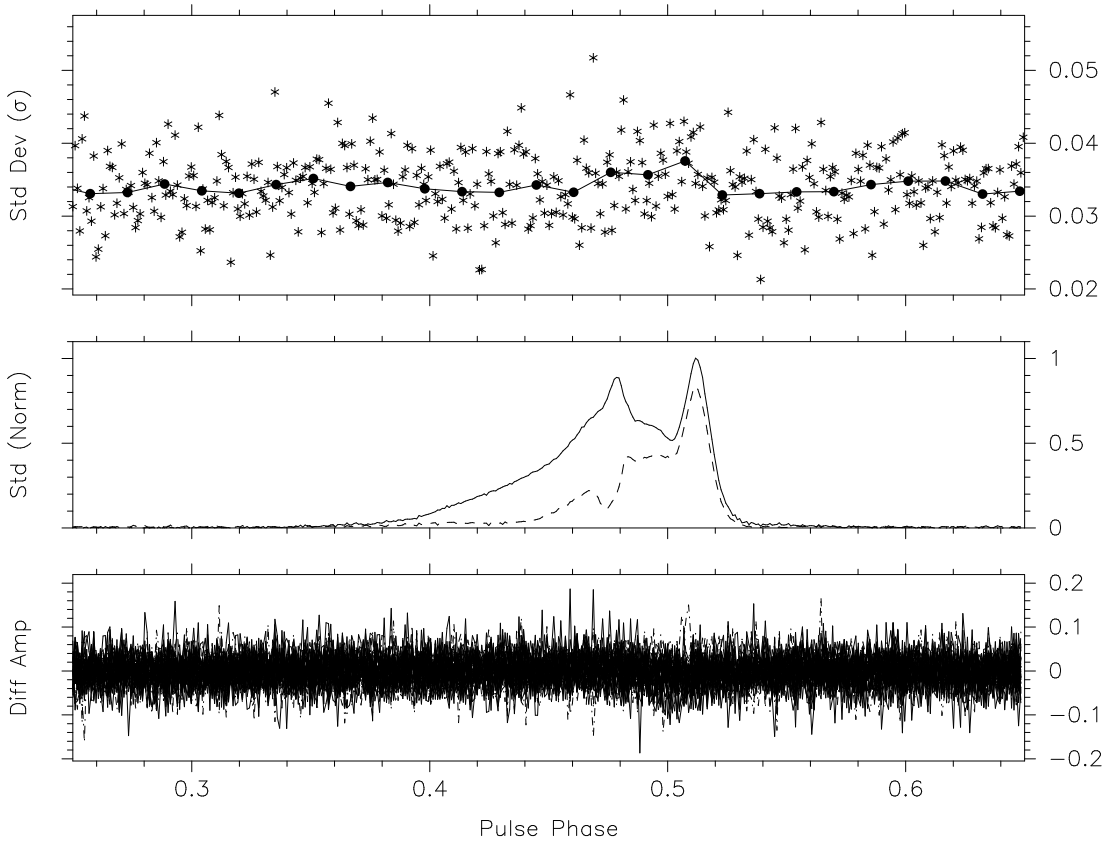


**Figure 5.7: As in Fig. 5.5, superposed morphological difference profiles of 37 total intensity (calibrated) integrations at a sky frequency of 1341 MHz**

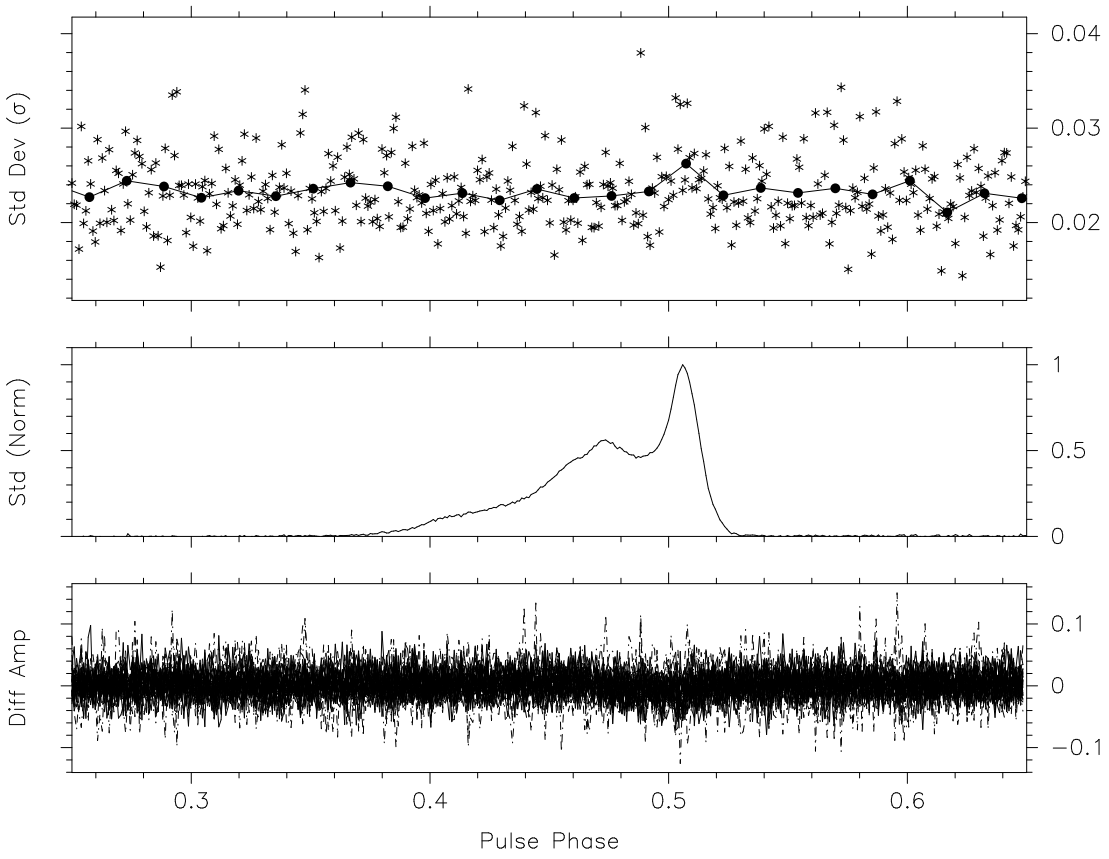


**Figure 5.8: As in Fig. 5.4, superposed morphological difference profiles of 31 total intensity (uncalibrated) integrations at a sky frequency of 685 MHz**

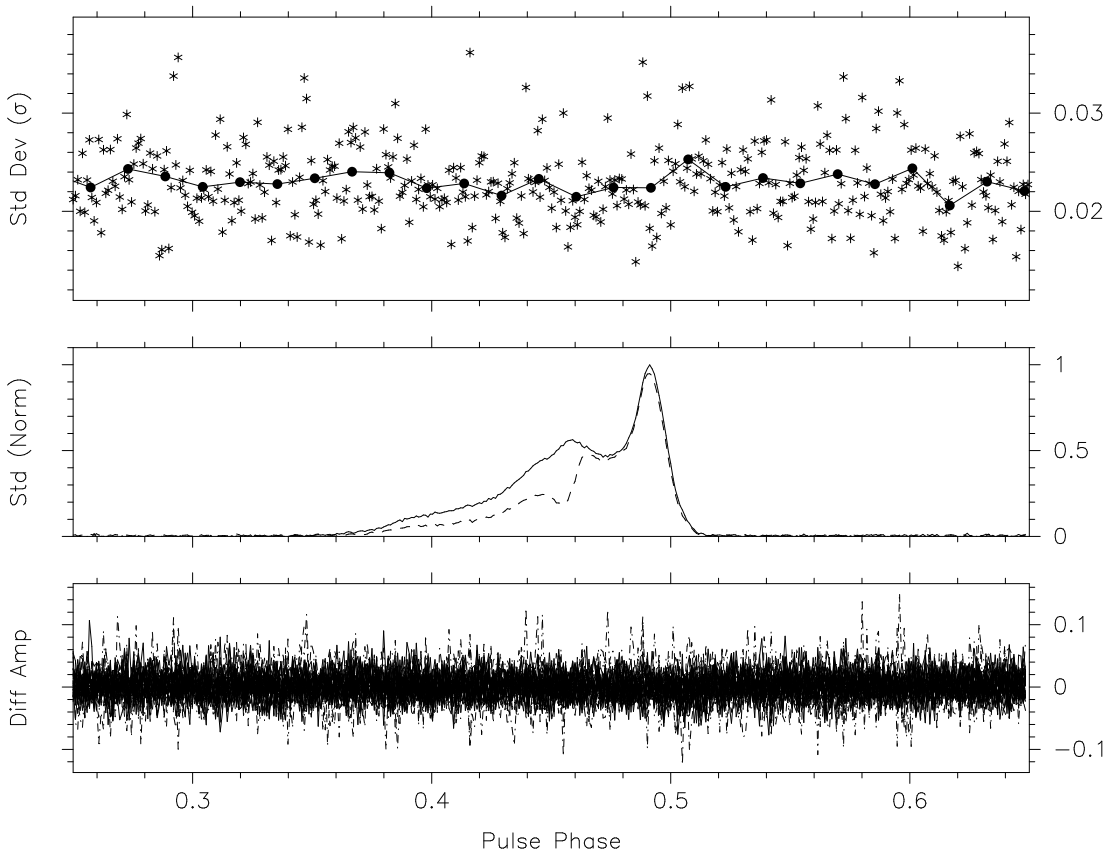


**Figure 5.9: As in Fig. 5.5, superposed morphological difference profiles of 31 total intensity (calibrated) integrations at a sky frequency of 685 MHz**

| Frequency (MHz) | Ratio | Std Dev |
|---|---|---|
| 1405.0 | 0.97 | 0.07 |
| 1341.0 | 0.91 | 0.07 |
| 685.0 | 0.58 | 0.02 |

**Table 5.1: Ratio of leading to trailing component for the calibrated data set.**

| Frequency (MHz) | Ratio | Std Dev |
|---|---|---|
| 1405.0 | 0.97 | 0.06 |
| 1341.0 | 0.90 | 0.06 |
| 685.0 | 0.58 | 0.02 |

**Table 5.2: Ratio of leading to trailing component for the uncalibrated data set.**

mean profile is stable and that we are simply seeing the limitations of our calibration procedure.

### 5.3.3 Peak Ratio Evolution

Kramer et al. (1999) infer the presence of smooth variations in the relative amplitudes of the principal components on timescales of a few minutes to an hour or more by calculating peak amplitude ratios and demonstrating that they evolve at a level significantly above the uncertainty in the measurement. They also note that instances of such smooth variation may not be common and that the timescales involved can change from one data set to another.

As a secondary check, we computed similar amplitude ratios to that of Kramer et al. (1999) for our high S/N profiles. A S/N threshold of 30 was deemed sufficient for this analysis because it is less susceptible to baseline corruption than the difference profile test. The RMS of the off-pulse region in each profile was used as an indication of the error in peak amplitude. The ratio of leading to trailing component amplitude was computed for all profiles and plotted against observation time (see Fig. 5.10 for an example). In addition, mean ratios and their associated standard deviation were computed over the 15 month time span of our data set. These values are presented in Tables 5.1 & 5.2.

The evolution of the profile's components with frequency is quantified in Tables 5.1 & 5.2. It is interesting to note that calibration of the data increases the scatter in the measured ratios at higher frequencies, this further supports the notion that the variability observed in Fig. 5.5 and by Kramer et al. (1999) is likely due to errors in the simple model used for polarimetric calibration. This does not seem to be the case at lower radio frequencies, as the calibration procedure has no detectable influence on the computed component ratios at 685 MHz. The 50 cm system may simply be better suited to the receiver model used for calibration, however the

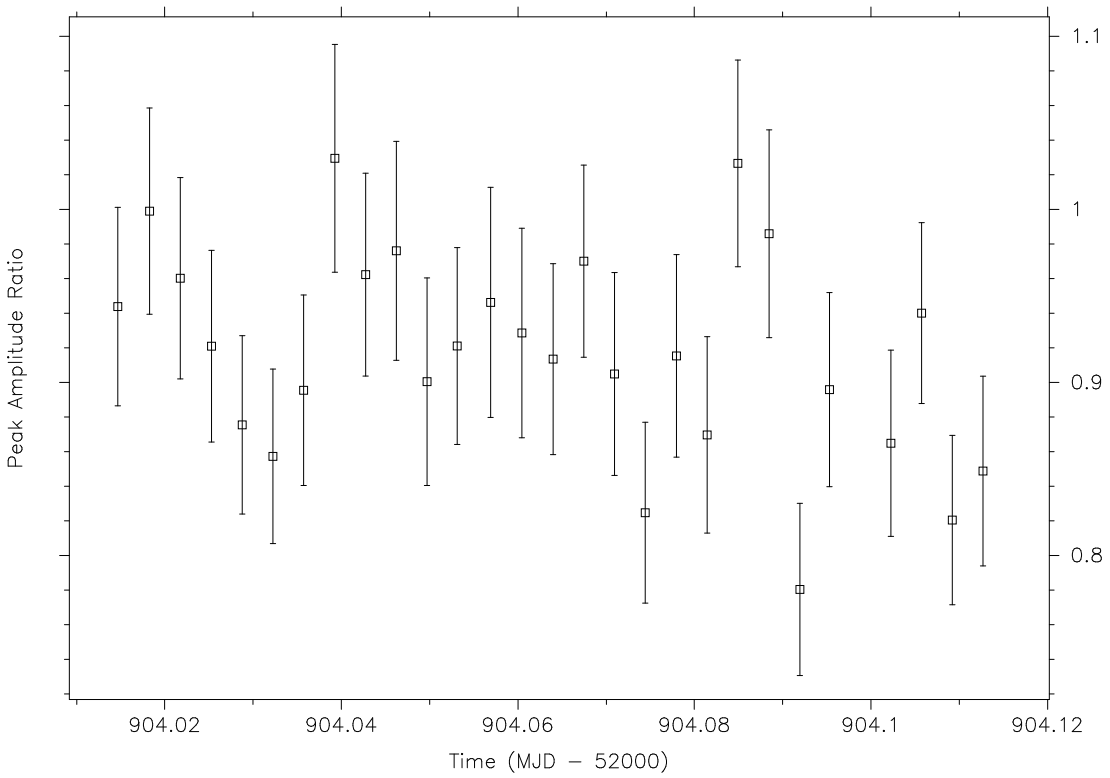


**Figure 5.10: Ratios of leading to trailing peak amplitudes during a single 2.5 hr observation taken on September 22, 2003, at a frequency of 1341 MHz. The uncertainties were computed based on the assumption that the amplitude of each peak could vary by the RMS of the noise in the baseline of the profile.**

intrinsic scatter at this frequency is also much smaller so the effect of calibration may be imperceptible. To investigate the evolution of these component ratios in time, the data set was examined by eye in the hope of finding clear indications of non-random evolution between 5 min integrations. Fig. 5.10 shows the observation that exhibited the most variation. Calibration has little effect on these points. The variation about the mean is small and perhaps insignificant.

### 5.3.4 Summary

The CPSR2 data set suggests that across a bandwidth of 64 MHz, PSR J1022+1001 does not vary significantly. Any instabilities, if present, must also be transitory in nature and therefore extremely difficult to characterise. It is possible that variability with a random or quasi-periodic structure, or small ($< 64$ MHz) characteristic bandwidth might still be present, but if so it would appear to have little effect on the profile when all frequency channels are combined. Profile instabilities should also be reflected in the timing of the pulsar, which we investigate in the next section.

## 5.4 Arrival Time Analysis

Kramer et al. (1999) reported that arrival times recorded at the Arecibo, Jodrell Bank and Effelsberg radio telescopes yielded a best fit RMS residual of 15-20 $\mu$s when applied to a model of PSR J1022+1001 and the binary system in which it resides. To achieve this RMS residual, a complicated standard template profile consisting of five Gaussian components with floating amplitudes was used, to compensate for supposed intrinsic profile variability. In contrast, the CPSR2 data set yields arrival times (from each 5 min integration, across multiple radio frequencies) that fit our timing model with an RMS residual of 2.27 $\mu$s. These TOAs were obtained using a frequency-domain cross-correlation process. Simple, static standard template profiles were constructed from the sum total of multiple integrations and the baseline noise was flattened to zero, to decrease spurious self-correlation of the timing profiles with the noise in the standard, which can lead to an RMS that is artificially low (see Appendix A). Separate standards were used for each frequency band and were aligned to a common fiducial point. Our 5 min (uncalibrated) integrations were fit to a model of PSR J1022+1001 using the standard pulsar timing package TEMPO[2]. Fig. 5.11 shows the corresponding timing residuals. It is interesting to note that although the front-end receiver system was changed mid-way through the data set, there does not seem to be any large systematic offset between the two receivers and no jumps were used to fit across the boundary. Changes in cable length or amplifier response between the two systems must have some impact on the assignment of arrival times, however it would appear that the offset is too small to measure in this data set.

A binary model of the type described by Blandford & Teukolsky (1976) was used to model the spin-down characteristics of PSR J1022+1001 and the perturbations introduced by its companion. A list of fitted parameters and their corresponding values (with error estimates) is presented below (Table 5.3). The errors in each TOA were uniformly scaled by a factor of 1.195 before fitting, to ensure the reduced $\chi^2$ was equal to unity. This was necessary to produce better error estimates for the fitted parameters. Sky position is shown in both Ecliptic and Equatorial (J2000) coordinates. PSR J1022+1001 lies in the ecliptic plane, making it difficult to determine the ecliptic latitude ($\beta$) accurately through timing measurements. This accounts for the relatively high error in our measurement of $\beta$ in Table 5.3. In addition, the CPSR2 data has a time baseline of only 15 months, whereas the data presented by Kramer et al. (1999) extends over more than 4 years. As a consequence, we do not fit for a proper motion in our model, choosing instead to hold this parameter fixed at the value of $-17 \pm 2\,\mathrm{mas\,yr^{-1}}$ quoted by Kramer et al. (1999). With proper motion thus constrained, we obtain a significant estimate for the parallax of the system. Given the short temporal baseline of our data set, this parallax estimate should be considered preliminary and it is included because it appears significant.

[2]http://www.atnf.csiro.au/research/pulsar/tempo/

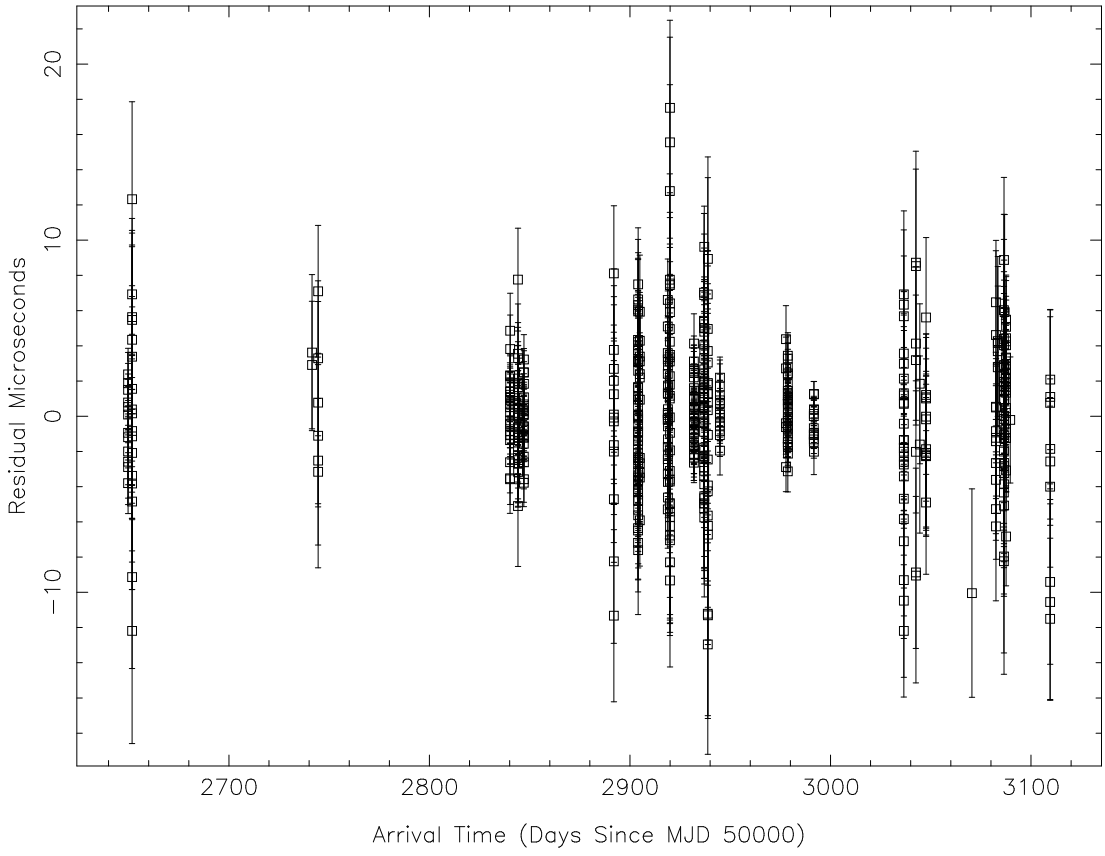


**Figure 5.11: Timing residuals from all 5 min integrations with S/N in excess of 30 and timing error less than 7 $\mu$s. The RMS residual is 2.27 $\mu$s and the reduced $\chi^2$ of the weighted fit is 1.43. The vertical scale is an order of magnitude better than arrival times presented by Kramer** et al. **(1999) who saw residuals as large as 400 $\mu$s.**

| Parameter | Value |
|---|---|
| Ecliptic Lon. ($\lambda$) (deg) | 153.86589029 (4) |
| Ecliptic Lat. ($\beta$) (deg) | -0.06391 (6) |
| Proper Motion in $\lambda$ (mas yr$^{-1}$) | -17 (2) * |
| Parallax (mas) | 3.3 (8) |
| Period (ms) | 16.4529296931296 (5) |
| Period Derivative ($10^{-20}$) | 4.33 (1) |
| Period Epoch (MJD) | 52900 |
| Dispersion Measure (cm$^{-3}$pc) | 10.25180 (7) |
| Projected Semi-Major Axis (lt-s) | 16.7654148 (2) |
| Eccentricity | 0.00009725 (3) |
| Time of Periastron Passage (MJD) | 52900.4619 (3) |
| Orbital Period (days) | 7.805130160 (2) |
| Angle of Periastron (deg) | 97.73 (1) |
| Right Ascension ($\alpha$) | 10:22:58.015 (5) |
| Declination ($\delta$) | +10:01:53.2 (2) |
| Number of TOAS | 555 |
| Total $\chi^2$ | 545.74 |
| RMS Timing Residual ($\mu$s) | 2.27 |
| MJD of first TOA | 52649 |
| MJD of last TOA | 53109 |
| Total Time Span (days) | 460 |

**Table 5.3: PSR J1022+1001 Blandford & Teukolsky (1976) timing model parameters derived from 15 months of CPSR2 observations. The error in the last significant digit is given in parentheses after the value. (* This value was given by Kramer** et al. **(1999))**

Figure 5.11 shows that PSR J1022+1001 has benefited from more regular observations in recent months. There are only a few small groups of points in the earliest part of the data set. This is another reason for not including proper motion in the model, as the small number of points observed during the beginning of 2003 would unreasonably dominate the fit. Parallax measurements are sensitive to day of year coverage more than total observing time, so it is still possible that the value of 3.3±0.8 mas is meaningful. It is interesting to note that our parallax measurement implies a distance of only $300\ ^{+100}_{-60}$ pc, as opposed to the value of 600 pc derived from the Taylor & Cordes (1993) electron density model. Many pulsars within 1 kpc of the Sun are inconsistent with the Taylor & Cordes (1993) model at this level. If the system is indeed only 300 pc from the sun, other secular changes may soon be detectable in the orbital parameters.

The accuracy with which we have been able to time this pulsar affords the chance to derive new limits on several physical parameters. Of particular interest is the possible presence of orbitally modulated Shapiro delay as the distance each pulse must travel into the companion star's gravity well changes throughout the 7.8 day orbit. While fitting for the Shapiro delay parameters $m_{\rm WD}$ and $i$ is not directly possible due to the small amplitude of their timing signature, a statistical analysis of the effect these parameters have on the $\chi^2$ of the fit can still reveal important information. The parameters derived from our Blandford & Teukolsky (1976) model were imported into a model (Damour & Deruelle 1985; Damour & Deruelle 1986) that includes the two Shapiro delay parameters $r$ and $s$, in a theory-independent formulation. The range $r$ and shape $s$ of the Shapiro delay signature are equivalent to the companion mass and $\sin(i)$ if we assume that General relativity is the correct theory of gravity. In this chapter we plot $\cos(i)$ instead of $\sin(i)$ because we expect it to represent a uniform probability distribution over the space of all possible inclination angles. The new model was used to create a $\chi^2$ map (Fig. 5.12). This map allowed us to place upper and lower bounds on the inclination angle, although it did not tightly constrain the companion mass. Based on $2\sigma$ Shapiro delay contours, $\cos(i) > 0.56$. Assuming that the companion is a white dwarf (a valid assumption given the small eccentricity of the system), the Chandrasekhar mass limit for the companion also constrains $\cos(i) < 0.8$, leading to an acceptable range of inclination angles between $37^o$ and $56^o$. If we assume a pulsar mass of 1.35 $M_\odot$, the Newtonian mass function adds an additional constraint and we can estimate that the mass of the companion is $> 0.9\,M_\odot$. We therefore predict that the initial binary was near the limit required to produce two neutron stars, and that the white dwarf is composed of heavier elements (maybe even O-Ne-Mg) than many of the millisecond pulsar companions (which are most often He white dwarfs).

The model parameters presented in Kramer et al. (1999) include a value for the projected semi-major axis of the orbit, $x = a\sin(i) = 16.765409(2)$ s. Given that the reference epoch of our data set is some 7 years ahead of the corresponding Kramer et al. (1999) epoch, it is possible to compare our value of $x$ in the hope of detecting

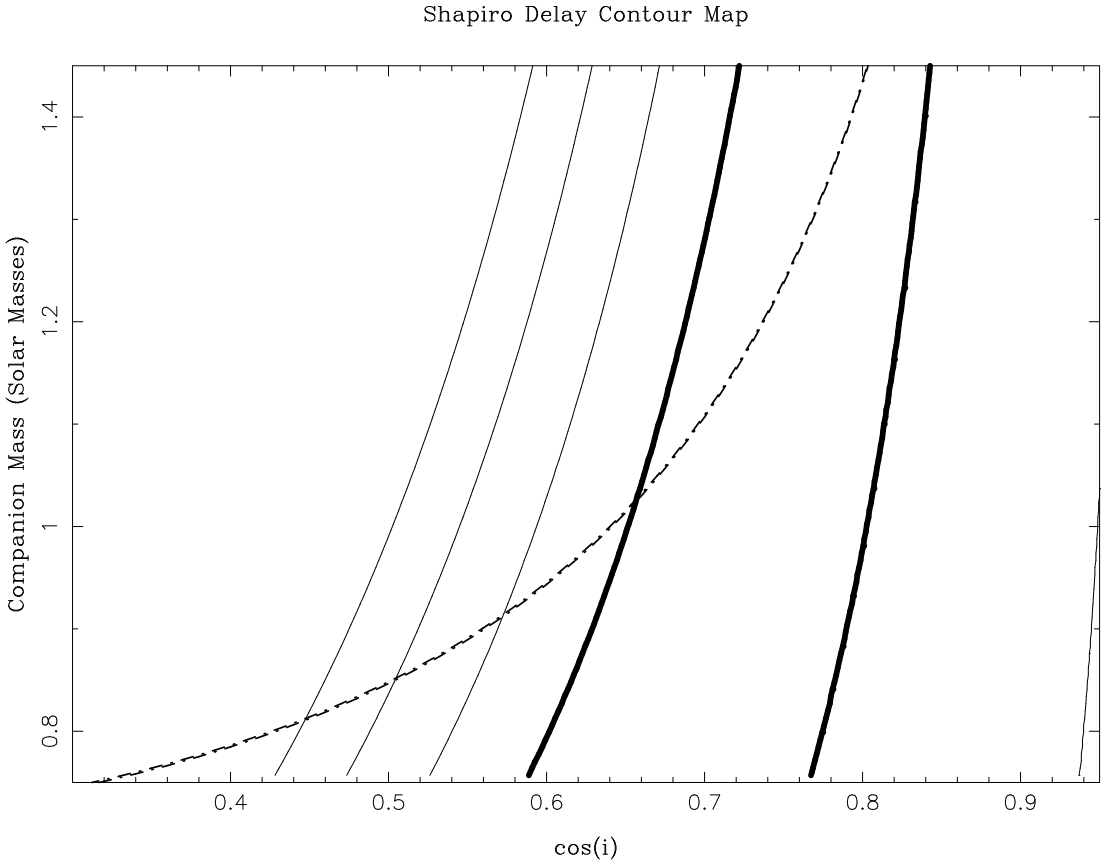


**Figure 5.12: Shapiro delay contour map, produced by constructing a $\chi^2$ surface over the range and shape parameters. The thick lines are $1\sigma$ contours, followed by 2,3 and $4\sigma$ contours on the left hand side. These contours do not represent true two-dimensional confidence regions but have been drawn so that their projection on the vertical and horizontal axes corresponds to the true one-dimensional confidence interval for either of the two parameters. The dashed line represents the constraint due to the mass function of the system, assuming a $1.35\,M_\odot$ neutron star.**

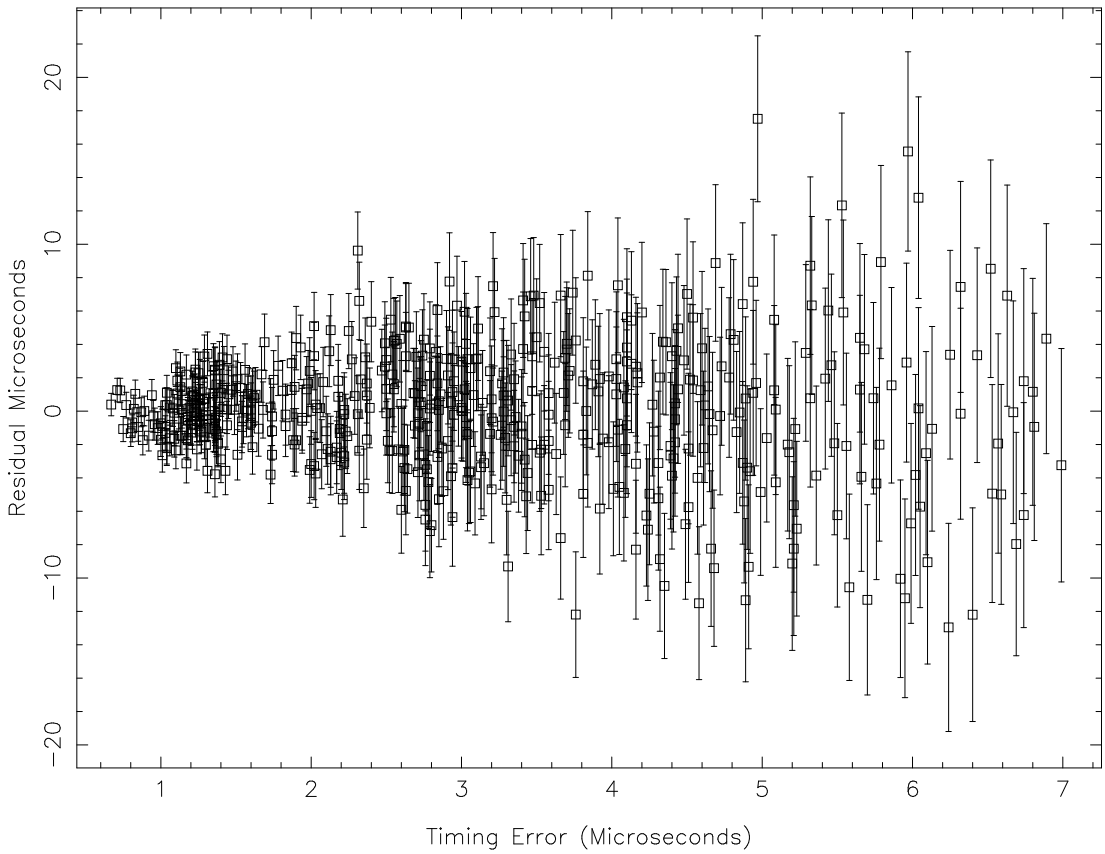


**Figure 5.13: Model residual plotted as a function of error in arrival time.**

a significant change as the proper motion of the system alters our line of sight to the plane of the orbit. We find $\dot{x} = 0.8 \pm 0.3\,\mu\mathrm{s\,yr}^{-1}$. Sandhu et al. (1997) show that

$$\dot{x} = x \cot(i)[-\mu_\alpha \sin(\Omega) + \mu_\delta \cos(\Omega)], \tag{5.2}$$

where $\mu_\alpha$ and $\mu_\delta$ represent the components of proper motion in RA and Dec and $\Omega$ is the longitude of the ascending node. Only the component of $\vec{\mu}$ in ecliptic longitude was measured by Kramer et al. (1999), but given our constraints on the inclination angle $i$, the only two unknowns in the equation are now $\mu_\delta$ and $\Omega$. If VLBI measurements could provide a value for $\mu_\delta$ (and perhaps confirm our detection of parallax), we would be able to place limits on the angle $\Omega$, further constraining the three-dimensional orientation of the orbit on the sky.

The reduced $\chi^2$ of 1.43 indicates that the arrival times are relatively free from unmodelled systematic effects. Increasing the error estimates by 20% gives a $\chi^2$ of unity, which is low by precision pulsar timing standards. Another way to view this low level of systematics is to plot the residual against the arrival time measurement uncertainty for each TOA (Fig. 5.13). With perhaps one exception, there are almost no points many standard deviations from zero.

Kramer et al. (1999) argued that the leading component of PSR J1022+1001 varies across a characteristic bandwidth of approximately 8 MHz. In order to test this, a separate timing analysis was performed on the 20 cm wavelength data before complete summation of the frequency channels, reducing the bandwidth in each integrated profile to 8 MHz. If the pulsar is varying on these bandwidths and 5 min timescales, this should lead to an enormous increase in the RMS residual that grossly exceeds the expected factor of $\sqrt{8}$. We find that the RMS residual of these 8 MHz bandwidth arrival times is 6.65 $\mu$s, which is very close to $\sqrt{8}\times 2.27 = 6.42\,\mu$s, implying that the increased scatter is consistent with the loss of S/N associated with the reduction in bandwidth. This analysis is strong evidence that there are no profile variations with a characteristic bandwidth of approximately 8 MHz and timescales of a few minutes.

## 5.5 Conclusions

PSR J1022+1001 has an interesting pulse profile morphology in terms of characteristic shape, polarimetric structure and evolution with radio frequency. These factors conspire to make it a difficult source to calibrate and analyse. Extensive studies of the morphological differences between 5 min integrations observed with CPSR2 at the Parkes 64 m radio telescope are consistent with a stable pulse profile. This bodes well for the future of precision timing of both this source and millisecond pulsars in general. Although the RMS residual presented is probably not small enough to make an immediate contribution to any pulsar timing array, lengthier integrations and continued monitoring may yet push the timing of this pulsar below the 1 $\mu$s mark.

We have used our improved timing to place some interesting limits on the geometry of this source which demonstrate that the inclination angle lies within the range $37^o < i < 56^o$. If our parallax is confirmed, this system will make a good target for VLBI observations which could both improve the distance to the pulsar and determine the as yet unknown component of proper motion in the direction of ecliptic latitude. This would provide additional limits on the three-dimensional orientation of the orbit via consideration of the change in projected semi-major axis. In the future, limits on $\dot{x}$ will improve and on a 10-year baseline, other relativistic observables may become measurable. We anticipate that with another 20 months of timing, we should be able to obtain a more meaningful parallax and independent proper motion for this source in right ascension. Our error in $\omega$ is only $0.01^o$, which suggests that the expected rate of advance of periastron for this source $\dot{\omega} \sim 0.01^o\,\mathrm{yr}^{-1}$ may be measurable on a 10-year timescale.

# Chapter 6

# Geodetic Precession in PSR J1141–6545

"Truth is stranger than fiction; fiction has to make sense."

*Leo Rosten*

We present observations that show dramatic evolution of the mean pulse profile of the relativistic binary pulsar PSR J1141–6545 over a period of 5 years. This is consistent with precession of the pulsar spin axis due to relativistic spin-orbit coupling. Observations made between 1999 and 2004 with a number of instruments at the Parkes radio telescope demonstrate a steady, secular evolution of the mean total intensity profile, which increases in width by more than 50% during the 5 year period. Analysis of the changing position angle of the linearly polarised component of the mean profile suggests that our line of sight is shifting closer to the core of the emission cone. We find that the slope of the position angle swing across the centre of the pulse steepens with time and use a simplified version of the rotating vector model to constrain the magnitude and direction of the change in our line of sight angle relative to the pulsar magnetic axis. The fact that we appear to be moving deeper into the emission cone is consistent with the non-detection of this pulsar in previous surveys.

## 6.1 Introduction

Soon after the discovery of the binary pulsar PSR B1913+16 (Hulse & Taylor 1974) it was pointed out that if the pulsar's spin axis and orbital angular momentum vector were misaligned, they should precess around their common sum on a timescale of around 300 years due to General relativistic spin-orbit coupling (Damour & Ruffini 1974; Hari Dass & Radhakrishnan 1975; Barker & O'Connell 1975b; Esposito &

Harrison 1975). This phenomenon is often referred to as "geodetic precession". Because radio pulsars are thought to possess lighthouse-like beams of emission that only beam to a small fraction of the sky, such precession should lead to observable pulse shape changes once a significant fraction of the precession period has passed.

Secular changes in the mean pulse profile of PSR B1913+16 were first reported by Weisberg, Romani & Taylor (1989). The changes became more pronounced in later years, allowing limited modeling of the emission cone (Kramer 1998; Weisberg & Taylor 2002). The variations seen in PSR B1913+16 suggest a small misalignment angle between the spin and orbital angular momentum vectors, consistent with a natal kick being imparted to the most recently formed neutron star (Bailes 1988). Large misalignment angles raise the possibility that we might be able to, for the first time, map out the entire emission cone of a pulsar until it disappears completely from view (Kramer 1998).

For many years, PSR B1913+16 was the only binary pulsar with the right combination of relativistic parameters and an evolutionary history that would make detection of geodetic precession possible. However, shortly after geodetic precession was first hinted at, new relativistic pulsars were uncovered in pulsar surveys that would become suitable for measurements in the future. Anderson et al. (1990) discovered a near clone of the original binary pulsar in the globular cluster M15C, which although weak, might one day exhibit the phenomenon. The much closer PSR B1534+12 (Wolszczan 1990) is a relativistic pulsar in a 10 hour binary orbit that has recently been shown to exhibit pulse shape changes. These changes have been combined with polarimetric models to make the first reliable estimate of the rate at which the spin axis of the pulsar is tilting away from us (Stairs, Thorsett & Arzoumanian 2004). Thus, measurements of geodetic precession are allowing new tests of General relativity and constraining evolutionary models (Wex, Kalogera & Kramer 2000; Konacki, Wolszczan & Stairs 2003). These new tests complement the earlier pioneering work of authors such as Taylor & Weisberg (1982), that examined other aspects of the theory.

In principle, it should be possible to use the relativistic pulsars to map both shape and intensity changes across a pulsar emission cone. However, this is complicated by the fact that many of the pulsars have random variations in their total intensity because of refractive and diffractive interstellar scintillation. Pulsars with small DM often have their fluxes change by factors of several on time scales that vary between minutes and days (Stinebring et al. 2000). However, large dispersion measure pulsars ($DM > 100\,\mathrm{pc\,cm^{-3}}$) have relatively stable fluxes when observed at high frequencies ($\nu > 1.4\,\mathrm{GHz}$) if long integration times are used and the observing system samples a large ($> 100\,\mathrm{MHz}$) bandwidth. The discovery of a relativistic pulsar at a large DM with a rapid geodetic precession timescale might then offer the first hope of determining how rapidly pulsar emission varies in both intensity and shape across the emission cone. This discovery would have important ramifications for pulsar statistics, where one of the great uncertainties is the pulsar "beaming fraction", often guessed at by observing the duty cycles of radio pulsars. Various

arguments have been made in favour of both meridional compression (Biggs 1990) and the elongation of pulsar beams (Narayan & Vivekanand 1983), making the actual measurement very important.

In the future it seems as though we will have a much larger number of relativistic systems with which to both test General relativity and map pulsar emission cones. Recent surveys have been finding relativistic pulsars at a welcome rate (Faulkner et al. 2004) and the total pulsar population now exceeds 1500 objects (Manchester et al. 2005a). In particular, the Parkes multibeam surveys have discovered a large number of pulsars in relativistic orbits. The spectacular "double pulsar" is a 2.4 hour binary with two active pulsars, and a geodetic precession period of just 80 years (Lyne et al. 2004). Faulkner et al. (2005) report the recent discovery of the 7.7 hour binary pulsar PSR J1756–2251 with an eccentricity of 0.18 that almost certainly consists of two neutron stars.

The first relativistic binary pulsar discovered by the multibeam surveys was however the 4.8 hour binary pulsar PSR J1141–6545 (Kaspi et al. 2000a). This pulsar orbits what is most likely a white dwarf companion, but the system is unique in that it still possesses a significant orbital eccentricity ($e = 0.17$). An eccentricity of this magnitude suggests that the system was put in its final configuration in an explosive event that may have given the pulsar a significant kick. Recent timing of the system is consistent with a $1.3\,M_\odot$ pulsar orbiting a $1.0\,M_\odot$ white dwarf companion (Bailes et al. 2003). It is likely that the system originated as a binary containing two main sequence stars, both below the critical mass required for a supernova. The initially more massive star begins to transfer matter onto its companion as it evolves, causing it to exceed the critical mass. If the system remains bound after the resulting supernova, a young neutron star is left orbiting a white dwarf companion. In the case of symmetric supernovae, the eccentricity of the orbit is induced by the sudden mass loss, and allows us to determine the pre-supernova mass uniquely (Radhakrishnan & Shukre 1985). This ejected mass can be related to the expected runaway velocity of the system, which in the case of PSR J1141–6545, is less than $50\,\mathrm{km\,s^{-1}}$.

Significant progress has been made in understanding the geometry and location of PSR J1141–6545 through a range of timing and spectroscopic studies. Ord, Bailes & van Straten (2002a) demonstrated via HI absorption analysis that the pulsar is at least 3.7 kpc distant. In addition, PSR J1141–6545 was the first pulsar to exhibit dramatic changes in its scintillation timescale as a function of orbital phase, which enabled an independent estimate of both its orientation and velocity. Ord, Bailes & van Straten (2002b) used the orbital modulation of the scintillation timescale to calibrate the usually unknown scale factor that relates the scintillation timescale, scattering screen distance and intrinsic motion to pulsar velocity. Their subsequently determined space motion was greater than that expected from a symmetric supernova (Ord, Bailes & van Straten 2002b).

Several years ago we commenced an observing campaign of PSR J1141–6545 to study its HI, scintillation and timing properties over long baselines. It was clear that such observations could also be used to search for relativistic effects such as orbital

decay and precession. The DM of the pulsar is 116 pc cm$^{-3}$, and its scintillation properties are well understood. Ord, Bailes & van Straten (2002b) showed that the scintillation bandwidth is much smaller than the 256 MHz observing band used by the analogue filterbank at the Parkes radio telescope, resulting in fairly stable fluxes when integrated over the 4.8 hour orbit.

In this chapter we demonstrate that PSR J1141–6545 is undergoing rapid secular evolution of both its total intensity and polarimetric profiles in a manner consistent with geodetic precession. In section 6.2 we describe the many instrumental systems used to observe this pulsar since its discovery in 1999, along with the associated data reduction methods. Section 6.3 describes the parameters of the binary system in greater detail and includes a calculation of the expected precession rate. It also introduces two polarimetric profiles that are considerably different from each other, providing the first evidence of profile evolution. Section 6.4 describes in detail the observed secular changes in the total intensity profile. Polarimetric evolution is considered in section 6.5, which shows that there has been a convergence of the linear and circular components of the pulsar profile in the last twelve months and that the slope of the position angle swing is steepening, presumably as we approach the emission pole of the pulsar. Finally, in section 6.6 we discuss the implications of our observations for pulsar surveys and pulsar emission models.

## 6.2 Observations

Observations were made at the Parkes radio telescope between July 1999 and May 2004, at centre frequencies ranging from 1318.25 MHz to 1413.50 MHz. Two different receiver packages were used to record data during this period; the central beam of the Parkes multibeam receiver and the wide band H-OH receiver. The multibeam receiver (Staveley-Smith et al. 1996) has a system temperature of approximately 21 K at 20 cm, which was about 5 K cooler than the H-OH receiver before it was upgraded near the end of 2003. Flux calibration of both systems using the radio galaxy 3C218 (Hydra A) suggests that the post-upgrade difference is only 1–2 K. Our data were recorded with three different instruments, each designed for high time-resolution observations across the widest possible bandwidths. Due to the rapid development of digital electronics within the past decade, each new observing system differed significantly from its predecessor.

The Caltech Fast Pulsar Timing Machine (FPTM), described by Navarro (1994), was a hardware-based auto-correlation spectrometer that performed incoherent dedispersion of dual orthogonal polarisations across two bands, each 128 MHz wide. Although the sampling rate was high enough to observe millisecond pulsars (MSPs) with only a few microseconds of smearing (at low DM), this instrument suffered from a number of artifacts induced by radio frequency interference and some deterioration in the correlator boards themselves. In some pulsars this led to oscillations in the passband that contaminated the pulsar profile. Nevertheless, many observations

with this instrument were not affected by these problems and it successfully timed many MSPs to high accuracy (Toscano et al. 1999b). The FPTM 2-bit sampled the raw data and formed auto-correlation functions that were binned at the apparent spin period of the pulsar. We were able to apply routine 2-bit corrections to enable accurate polarimetry. Being an incoherent detector with a finite number of lags, the FPTM could divide the passband into 4×128×1 MHz channels, leaving a residual dispersion smear ($t_{\rm smear}$) given by Eq. 1.16. For the configuration used to observe PSR J1141–6545, this corresponds to 350 $\mu$s of smearing in the detected pulse profile. Given that the FPTM uses 1024 phase bins across a single pulse period and that PSR J1141–6545 rotates once every 394 ms, each phase bin represents 384 $\mu$s of time. The detrimental effects of dispersion smearing are therefore confined to within a single phase bin.

However, if the pulsar spin period is two orders of magnitude shorter (as is typical of the millisecond pulsar population), dispersion smearing can significantly reduce the resolution of an incoherent detector. Motivated by a desire to overcome this problem for MSPs, CPSR1 was commissioned in August, 1998. This system implemented a technique called coherent dedispersion (Hankins & Rickett 1975), which requires Nyquist sampling of the observed band, followed by deconvolution with a response function characteristic of the interstellar medium (ISM). While this approach effectively eliminates dispersion smearing in the detected profiles, it is highly computationally intensive both in terms of the initial data rate and subsequent reduction. CPSR1 streamed digital samples to a striped set of four DLT drives (analogous to the method implemented for the S2 VLBI recorder) whose tapes were shipped to the Swinburne Centre for Astrophysics and Supercomputing for processing. Even with four striped tape drives, CPSR1 was limited to a bandwidth of 20 MHz. Rapid advances in consumer digital electronics soon made it feasible to upgrade the capabilities of the system and in August, 2002, CPSR2 was installed at Parkes. CPSR2 performs coherent dedispersion in near real-time, using a cluster of 30 server-class computers located at the telescope. It is capable of recording 2×64 MHz dual-polarisation bands simultaneously, providing a total bandwidth comparable to that of the previous generation of incoherent detectors, like the FPTM. The coherent dedispersion method employed by both CPSR machines allowed essentially arbitrary spectral resolution and reduced the dispersion smearing in each channel to a minute fraction of PSR J1141–6545's period, giving an effective sampling time of a few microseconds.

Individual observations of PSR J1141–6545 ranged in duration from a few minutes to several hours. In recent years, our strategy has been to maximise orbital phase coverage by observing in concentrated sessions during which the pulsar is tracked continuously for two whole orbits (∼9.6 hours). To calibrate the data we point the telescope one degree south of the pulsar and drive the in-built receiver noise source with a square wave at a frequency of ∼11 Hz at least once per orbit, to characterise the polarimetric response of the system. In addition, at least once per month we observe the flux calibration source 3C218 (Hydra A).

Our highest density of observations were taken with CPSR2 during 2003 and 2004 (MJD 52845 - 53134), during which time we have a full record of polarimetry, flux and profile morphology. The CPSR1 recorder was designed primarily to observe the bright southern millisecond pulsar PSR J0437–4715 (van Straten et al. 2001), however in January, 2001 it took data on PSR J1141–6545 for a total of 30 hours beginning on MJD 51922. The resulting calibrated mean profile provides important, early epoch information. We have selected three profiles at the widely spaced epochs of MJD 51381, 51781 and 52087, representing high quality FPTM data to further supplement our temporal coverage. It is fortuitous that our earliest FPTM pointing (MJD 51381) dates all the way back to 1999, extending our time baseline by almost two full years. For this reason, we include a 1999 profile despite the fact that the observation was only 12 min in duration. Fortunately, PSR J1141–6545 is a bright pulsar with an average flux density of approximately 4 mJy, so the S/N ratio of this 12 min observations is 125, quite sufficient for our analysis.

All data were processed using the tools included with the PSRCHIVE (Hotan, van Straten & Manchester 2004) software package, with the addition of several extra routines specific to pulse variability analysis.

## 6.3 PSR J1141–6545

PSR J1141–6545 was discovered in the first Parkes multibeam survey (Kaspi et al. 2000b). It resides in an unusual relativistic binary system, orbiting what is most likely a heavy white dwarf companion once every 4.8 hours. The pulsar does not appear to be recycled and is assumed to be the most recently evolved member of the system. Positioned close to the galactic plane, it is one of the few pulsars whose distance can be estimated by analysis of neutral hydrogen absorption features in its frequency spectrum. Ord, Bailes & van Straten (2002a) obtain a lower limit of 3.7 kpc using this method. In addition, the signal from this pulsar exhibits diffractive scintillation over small bandwidths ($\sim$1 MHz) and timescales of a few minutes (Ord, Bailes & van Straten 2002b), which can be used to place timing-independent constraints on the binary parameters. Ord, Bailes & van Straten (2002b) report a significant detection of orbital modulation in the observed scintillation velocity (due to the motion of the pulsar in its orbit) and infer both a relative velocity of $\sim$115 km s$^{-1}$ and an orbital inclination angle $i = 76\pm2.5^\circ$. This velocity is large compared to the value of $< 50$ km s$^{-1}$ expected to result from a symmetric supernova. Uncertainties in our knowledge of the true distance to the pulsar, and hence the relative velocity of the Earth's standard of rest, combined with uncertainties introduced by the Earth's orbital motion and bulk flows or anisotropies in the ISM, mean that we cannot convincingly state that the pulsar has an anomalous velocity due to an asymmetric supernova. However, if the profile evolution reported here is due to spin and orbital angular momentum misalignment and geodetic precession, we would expect the pulsar to have received a kick at birth, which would increase

its runaway velocity over that of a symmetric explosion.

The rotation period of PSR J1141–6545 is 394 ms, so the precision obtainable through pulse timing experiments is somewhat limited when compared to results (van Straten et al. 2001) obtained by timing millisecond pulsars, whose spin periods are of order 100 times shorter. Despite this, several post-Keplerian parameters are measurable. Bailes et al. (2003) describe a timing solution that includes significant detections of periastron advance ($\dot{\omega}$ $\sim$5.3° yr$^{-1}$), combined transverse Doppler and gravitational redshift ($\gamma$) and a marginal detection of orbital period derivative ($\dot{P}_{\rm b}$). Despite the lack of any Shapiro delay measurement, we can still derive a good estimate of the component masses. The post-Keplerian parameters $\dot{\omega}$ and $\gamma$ are related to the pulsar mass ($m_{\rm p}$) and companion mass ($m_{\rm c}$) by Eq. 6.1 and Eq. 6.2 respectively. In addition, pulse timing accurately determines the quantity, derived from Newtonian gravitation, known as the mass function (Eq. 6.3). These equations allow determination of the inclination angle and component masses.

$$\dot{\omega} = 3\left(\frac{2\pi}{P_{\rm b}}\right)^{5/3}\left(\frac{G(m_{\rm p}+m_{\rm c})}{c^3}\right)^{2/3}(1-e^2)^{-1}. \tag{6.1}$$

$$\gamma = e\left(\frac{P_{\rm b}}{2\pi}\right)^{1/3}\frac{G^{2/3}}{c^2}m_{\rm c}(m_{\rm p}+2m_{\rm c})(m_{\rm p}+m_{\rm c})^{-4/3}. \tag{6.2}$$

$$f(m_{\rm p},m_{\rm c},i) = \frac{m_{\rm c}^3\sin^3 i}{(m_{\rm p}+m_{\rm c})^2} = \frac{4\pi^2}{G}\frac{a^3\sin^3 i}{P_{\rm b}^2}. \tag{6.3}$$

Here, $G$ is Newton's gravitational constant, $P_{\rm b}$ is the pulsar orbital period and $a\sin(i)$ is the projected semi-major axis. According to Bailes et al. (2003), $m_{\rm p}$ = 1.30$\pm$0.02 M$_\odot$ and $m_{\rm c}$ = 0.986$\pm$0.02 M$_\odot$. The timing-derived inclination angle limit ($i > 75^o$) compares well with the value obtained from scintillation experiments (Ord, Bailes & van Straten 2002b).

Assuming that General relativity is the correct theory of gravity, Barker & O'Connell (1975a) present an expression (Eq. 6.4) for the expected, time-averaged precession rate of the pulsar spin axis

$$\Omega_{\rm p} = \frac{1}{2}\left(\frac{G}{c^3}\right)^{2/3}\left(\frac{P_{\rm b}}{2\pi}\right)^{-5/3}\frac{m_{\rm c}(4m_{\rm p}+3m_{\rm c})}{(1-e^2)(m_{\rm p}+m_{\rm c})^{4/3}}. \tag{6.4}$$

Here, $c$ is the speed of light and $e$ is the eccentricity of the system. For PSR J1141–6545, this evaluates to an intrinsic precession rate of 1.36$^o$ yr$^{-1}$, which implies a precession period of 265 years.

Bailes (1988) showed graphically that the maximum observable rate of precession may be significantly less than the intrinsic value. The geometry of the system and our viewing angle have a significant impact on our ability to detect geodetic precession. The observable quantity is the rate at which the angle $\delta$ between the observer's line of sight and the pulsar spin axis changes, as this will manifest itself as a changing cut through the emission cone. Bailes (1988) and Cordes, Wasserman

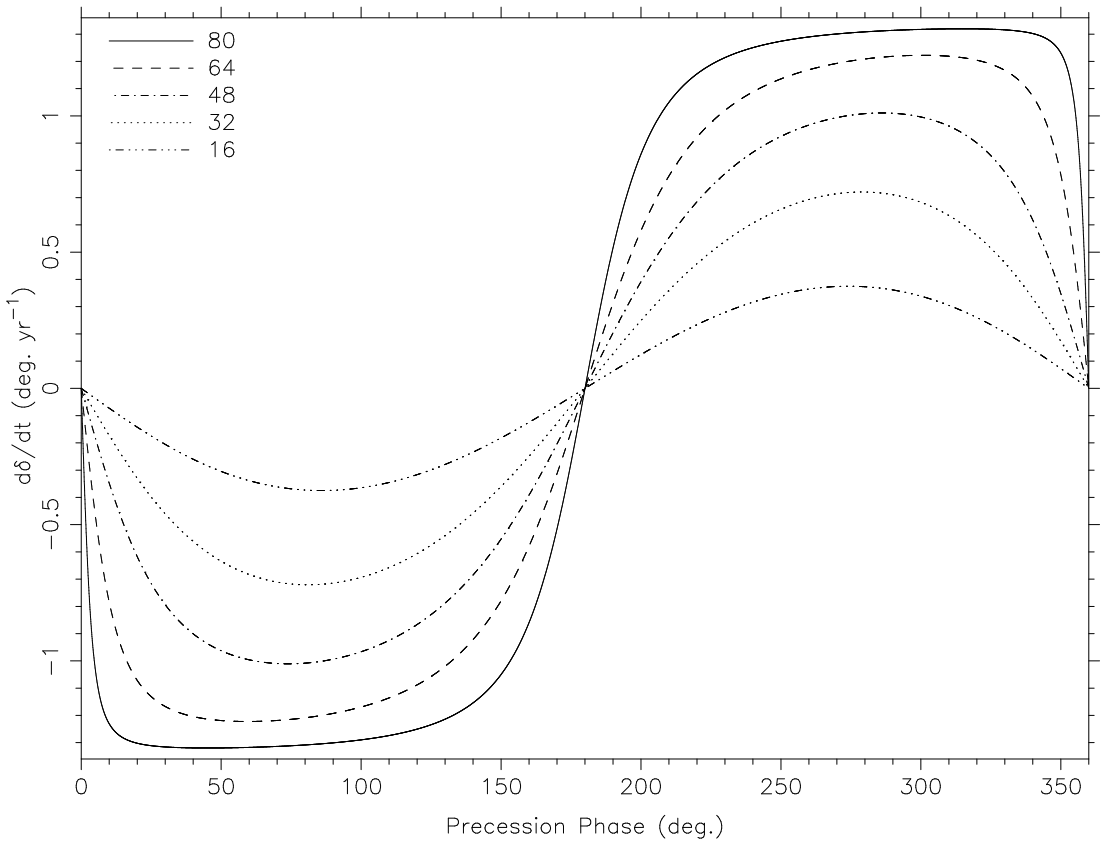


**Figure 6.1: Rate of change of the angle between our line of sight and the spin axis of the pulsar ($d\delta/dt$) as a function of precessional phase, for various misalignment angles (shown inset top left, in degrees) for PSR J1141–6545. We have assumed an inclination angle of 76$^o$, consistent with scintillation measurements. The amplitude of the observable precession signature is highly dependent on both the misalignment angle and our current position in the precession cycle, but any derived value of $d\delta/dt$ gives a minimum misalignment angle for the system.**

& Blaskiewicz (1990) present expressions for the rate of change of $\delta$, of the form reproduced in Eq. 6.5. The most important parameters in the expression are the misalignment angle between the spin axis and the orbital angular momentum vector, and the precessional phase at the current epoch, neither of which are known.

$$d\delta/dt = \Omega_p \vec{n} \cdot (\vec{s} \times \vec{j})(1 - [\vec{n} \cdot \vec{s}]^2)^{-1/2}. \tag{6.5}$$

Here, $\vec{n}$ is a unit vector along the line of sight to the observer, $\vec{s}$ is a unit vector along the pulsar spin axis and $\vec{j}$ is a unit vector in the direction of the orbital angular momentum. To evaluate this equation, we must know the misalignment angle, the current precessional phase and the orbital inclination angle $i$. The first two parameters are unknown for the PSR J1141–6545 system, but we can assume the value of $i$ derived from scintillation studies and plot one precessional period of $d\delta/dt$ for various misalignment angles (see Fig. 6.1).

Unfortunately, it is difficult to translate an observed difference in mean pulse profile morphology into a quantitative measure of angular shift. This requires knowledge of the intrinsic beam shape, which of course we do not possess, making the result model dependent. It is possible that polarimetric studies may offer a key alternative method, however it is still necessary to assume some model of pulsar polarisation as a function of impact angle. In this chapter we restrict ourselves, where possible, to a quantitative description of the observed evolution of the mean pulse profile.

### 6.3.1 Selected Observations

Figure 6.2 shows a coherently dedispersed, polarimetrically calibrated mean profile, observed with the multibeam receiver and CPSR1 at a frequency of 1413 MHz in January, 2001 (MJD 51922). The pulse profile is morphologically quite simple, consisting of a single component flanked on the left by a shoulder of emission. Note the small "bump" high on the leading edge of the profile, which is also present in the FPTM and analogue filterbank data. The peak fractional polarisation is of order 20% in both linear and circular. The position angle swing does not seem to fit the predictions of the rotating vector model and is similar to that seen by Kaspi et al. (2000a), although lacking the orthogonal mode change that is present in the leading shoulder of the Kaspi et al. profile.

Figure 6.3 shows our most recent fully calibrated mean profile, observed in May, 2004 (MJD 53134) with the H-OH receiver and CPSR2, at a centre frequency of 1341 MHz. There are a number of striking differences when compared to Fig. 6.2, most notably an overall broadening of the profile, which now has an extended trailing component; loss of the leading "bump" and general steepening of the position angle swing, which now has an identifiable slope.

Pulsar profiles can be corrupted by systematic errors associated with instrumentation. Typical effects include baseline artifacts due to improper binning and radio frequency interference, insufficient quantisation capabilities in the samplers (leading to the removal of power near the pulse as the samplers attempt to maintain an optimum mean level), and smearing due to insufficient time and frequency resolution. Fortunately, in the case of PSR J1141–6545 we have a variety of instruments, each of which has sufficient time resolution to over-sample the profile, and that give self-consistent results at similar epochs. We would be concerned if observed changes only coincided with equipment upgrades, but this has not been the case.

The mean pulse profile has changed significantly in the space of 3 years. This is confirmed in the next two sections where we present an analysis of data taken during (and before) the epochs presented in Figs. 6.2 & 6.3. We observe a smooth secular change in the characteristics of the mean pulse.

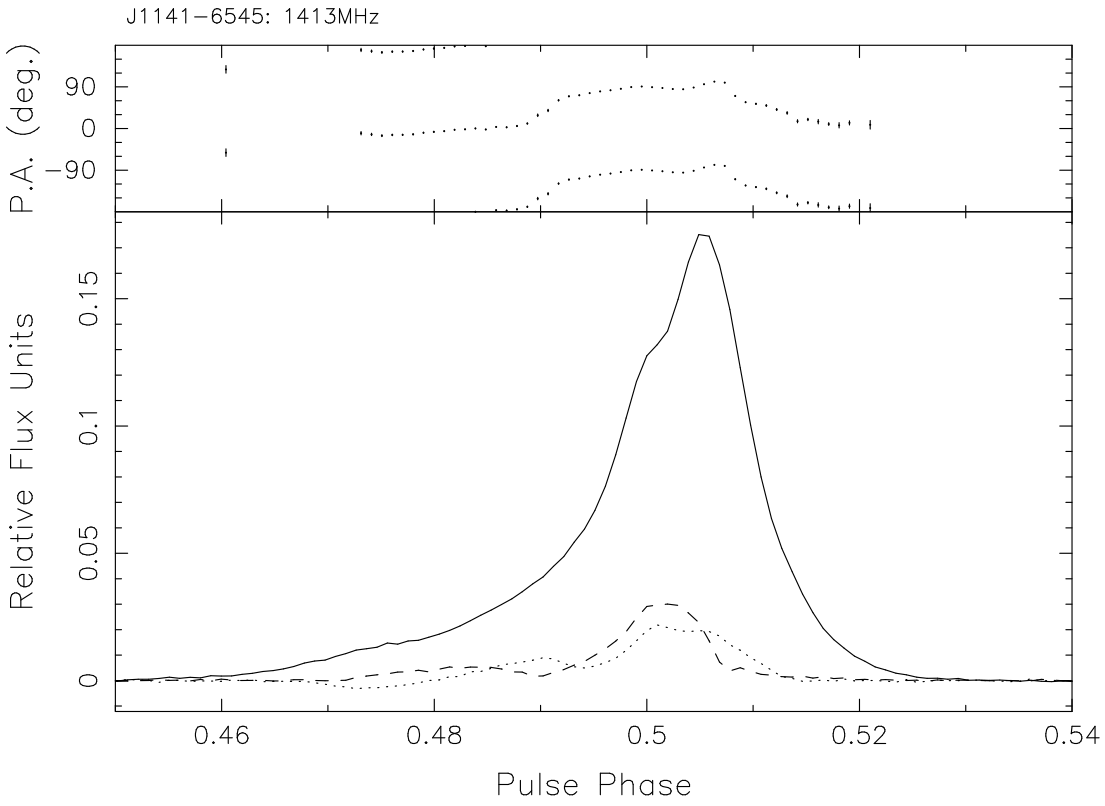


**Figure 6.2: PSR J1141–6545 mean profile, obtained from 28 hours of data taken in January, 2001 (MJD 51922), at a centre frequency of 1413 MHz. The solid line represents total intensity, the dashed line total linear and the dotted line total circularly polarised emission. There are 1024 phase bins across the profile, which has been polarimetrically calibrated using a simple model of relative gain and phase in the orthogonal linear receiver probes. Note the slight "bump" on the leading edge of the profile and the absence of any steep position angle evolution across the phase range shown.**

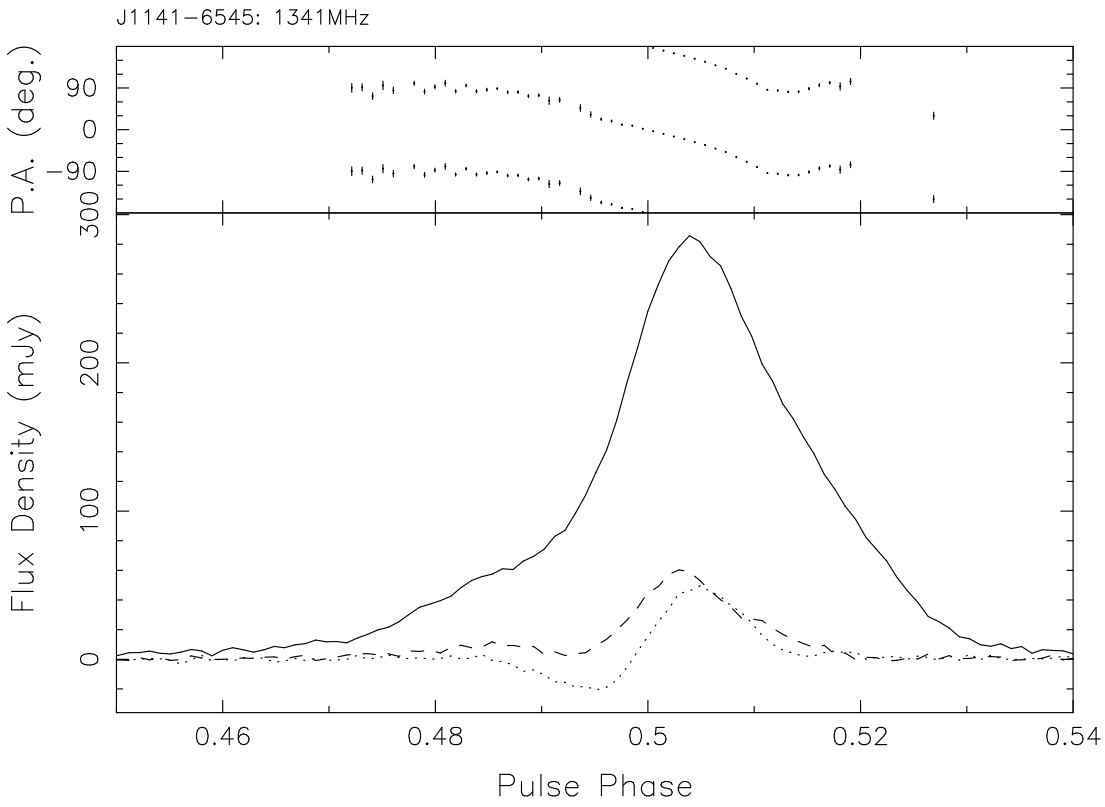


**Figure 6.3: May, 2004 (MJD 53134) PSR J1141–6545 profile (as in Fig. 6.2) obtained from 2 hours of data taken with CPSR2 at a centre frequency of 1341 MHz. Note the smooth leading edge and extended trailing component, as well as the more pronounced position angle sweep.**

| Instrument | Date (MJD) | Freq. (MHz) | S/N |
|---|---|---|---|
| FPTM | 51381 | 1318.25 | 125 |
| FPTM | 51781 | 1413.50 | 791 |
| CPSR1 | 51922 | 1413.00 | 778 |
| FPTM | 52087 | 1413.50 | 743 |
| CPSR2 | 52614 | 1405.00 | 1051 |
| CPSR2 | 52845 | 1341.00 | 427 |
| CPSR2 | 52902 | 1341.00 | 822 |
| CPSR2 | 52902 | 1341.00 | 1354 |
| CPSR2 | 52920 | 1341.00 | 390 |
| CPSR2 | 53109 | 1341.00 | 346 |
| CPSR2 | 53130 | 1341.00 | 728 |
| CPSR2 | 53133 | 1341.00 | 537 |
| CPSR2 | 53134 | 1341.00 | 553 |

**Table 6.1: List of parameters associated with the 13 observations used to characterise the secular evolution of the PSR J1141–6545 mean pulse profile.**

## 6.4 Evolution of the Total Intensity Profile

To examine the evolution of PSR J1141–6545's mean pulse profile in greater detail, we sum all the data within each observing session to produce a set of 13 well spaced, high S/N profiles. These mean profiles typically span 1–2 days, with total integrated times of a few hours. Table 6.1 summarises the most important parameters of each profile including observing system, start date, observing frequency and S/N.

We demonstrate that even though our points are not evenly spaced in time, the data describe a clear trend in profile evolution (Figs. 6.4 & 6.5). The changes are so great that visual inspection of the profiles can reveal much qualitative information including an overall broadening, extension of the trailing shoulder and smoothing of the leading edge (Fig. 6.4). In addition, we compute a quantitative measure of the width of each mean profile using an algorithm that defines thresholds in pulse phase based on where the flux under the pulse exceeds 10% of the peak value for the first time, and drops below this value for the last time. We choose the 10% threshold because it includes the majority of the on-pulse region, thus incorporating the important leading and trailing components of the pulse, which are seen to evolve significantly. The secular trend remains if the level is set to 50%, however the $\chi^2$ of the linear fit worsens marginally as we would expect from this narrower region of the profile.

A simple linear least-squares fit to the width data shows that the rate of profile broadening is well approximated by a straight line with slope $1.3 \pm 0.06\,\mathrm{ms\,yr^{-1}}$. Unlike PSR B1913+16, PSR J1141–6545 has only a single pulse profile component. It is therefore difficult to constrain the angular extent of the beam, or the angle

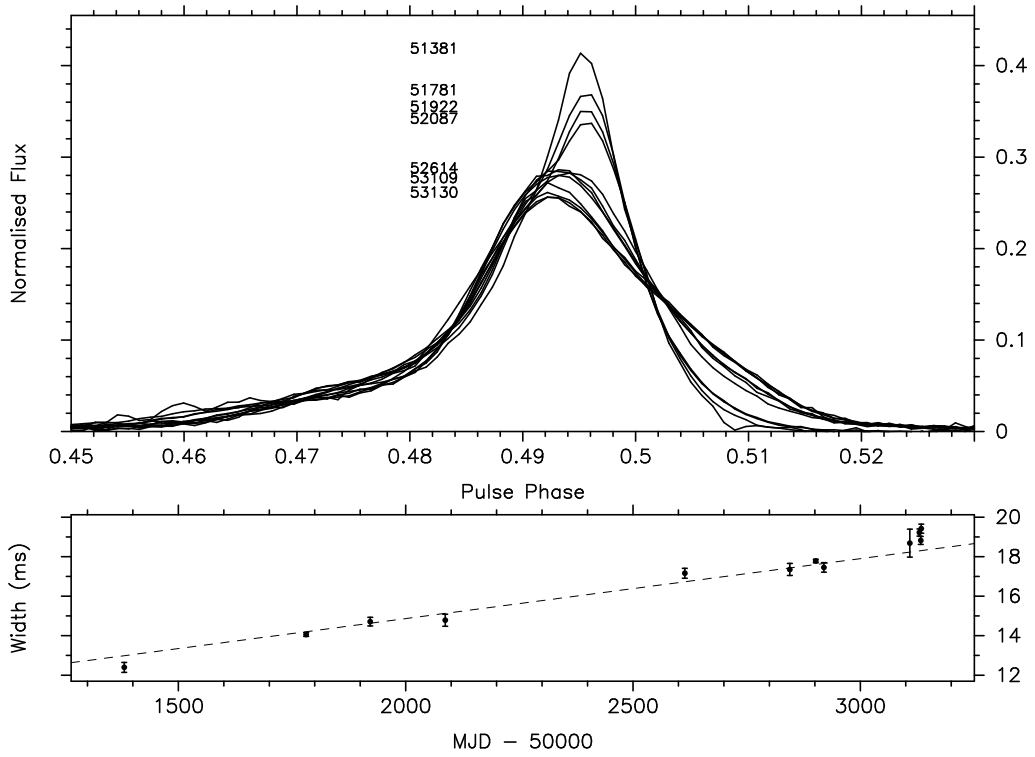


**Figure 6.4: The upper panel shows superposed mean total intensity profiles from observations taken at various epochs during the past five years. Each profile was observed at a wavelength near 20 cm. The MJD of each observation is marked on the plot, in line with the peak of the corresponding profile. For example, at pulse phase 0.495, the MJD of each observation increases monotonically with decreasing amplitude. The flux under each profile in the plotted region has been normalised to the flux under the earliest profile to allow direct comparison, regardless of the amplitude scaling schemes used by individual instruments. Each mean profile was aligned using an ephemeris obtained from a global timing solution across the entire data set, effectively maximising the cross-correlations between each profile. The lower panel shows the evolution of 10% width (see text) as a function of time, with one point for each profile in the upper panel. The profile clearly broadens over the span of our observations. Uncertainties are derived from consideration of the RMS noise level in each profile and represent $1\sigma$ uncertainties. The line of best fit, obtained using a linear least-squares method, is shown (dashed).**

between the magnetic axis and the line of sight, using total intensity information alone. A similar total intensity profile can be produced by intersecting the centre of a narrow beam or the edge of a wider beam, which introduces a degeneracy in the interpretation. Stairs, Thorsett & Arzoumanian (2004) describe a method for determining the full geometry of the system through the detection of a secondary precession effect caused by orbitally modulated aberration. Unfortunately we have been unable to detect this in PSR J1141–6545. Given a sufficient time baseline, the profile variations should eventually deviate from the currently observed linear trend. This will also provide an opportunity to constrain the three-dimensional geometry of the pulsar system.

In order to characterise the rate of change in more detail, we perform a difference profile analysis on our data set. This involves using (arbitrarily) the first profile in the series as a standard template whose amplitudes are subtracted from the remaining profiles after their flux and alignment have been normalised to the standard. Measurement of the remaining flux in the difference profile gives an indication of how much the profiles vary across a particular epoch (Fig. 6.5). This method is similar to the technique of principal component analysis performed on PSR B1534+12 by Stairs, Thorsett & Arzoumanian (2004), where only the mean profile and one orthogonal component are taken into account. We use this numerical measure of profile difference to examine quantitatively the rate at which evolution is occurring.

Figure 6.5 shows that the fractional difference trend is well approximated by a straight line with slope $2.4 \pm 0.07 \times 10^{-4}$ fractional difference units per day. The mean pulse profile of PSR J1141–6545 is therefore changing at a steady rate of approximately 9% per year. The profile evolution seen in Figs. 6.4 & 6.5 is unlikely to be instrumental in origin because it occurs smoothly over the entire time span, instead of jumping discontinuously at the points when new hardware was introduced. In addition, the instrumental upgrades always decreased systematic smearing of the observed profile, whereas we observe the profile width increasing with time.

Given that our observations span a frequency range of almost 100 MHz, it is possible that intrinsic evolution of the profile shape with frequency might contaminate the result. Such contamination is however unlikely to be responsible for the observed secular trend in pulse width, because as Table 6.1 shows, the changes in observing frequency have not been linear in time. To demonstrate that PSR J1141–6545 does not exhibit significant profile evolution over the range of observing frequencies in our data, we compare two CPSR2 profiles observed on the same day (MJD 53204) in July, 2004. These two profiles were observed at 1405 MHz and 1341 MHz respectively and we analyse them in a similar fashion to Fig. 6.5, using the 1405 MHz profile as a standard template and constructing a difference profile. Figure 6.6 shows that the profile does not evolve significantly across a bandwidth of 64 MHz, therefore we can be confident that frequency evolution does not contaminate our result.

A number of different authors have reported observing the mean profile of a pulsar changing in some way. Some slower pulsars have been observed to emit in two or more "modes" of pulse shape, that each remain stable for a time before switching

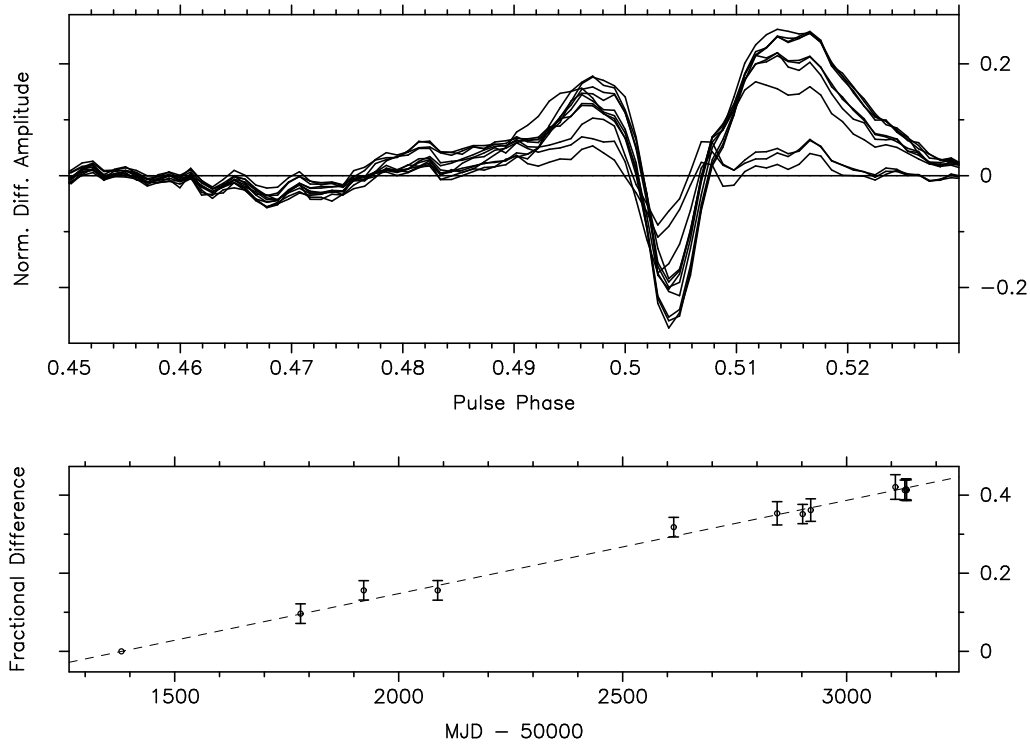


**Figure 6.5: The upper panel shows superposed difference profiles, constructed from the data shown in Fig. 6.4, using the earliest profile as a standard template (shown as a horizontal line through zero). The vertical scale is in units of the standard profile flux, indicating that some individual components of the profile have changed by up to 20%. There is a monotonic increase in difference amplitude with time, at pulse phase 0.515 for example, the amplitude increases with the MJD of the observation. The lower panel shows the fractional difference between each difference profile and the standard template, as a function of time. Each point represents a single difference profile. The fractional difference is found by summing the absolute values of the amplitude in each pulse phase bin in the difference profile and dividing by the flux in the standard template. Errors are based on measurements of the off-pulse RMS. While this technique is sensitive to changes in S/N as well as morphology, we are confident that our choice of high S/N observations and a narrow phase window allows the morphological information to dominate. The line of best fit is shown (dashed), obtained using a linear least-squares method. Note that the earliest profile has a fractional difference (and error) of zero, providing a reference point.**

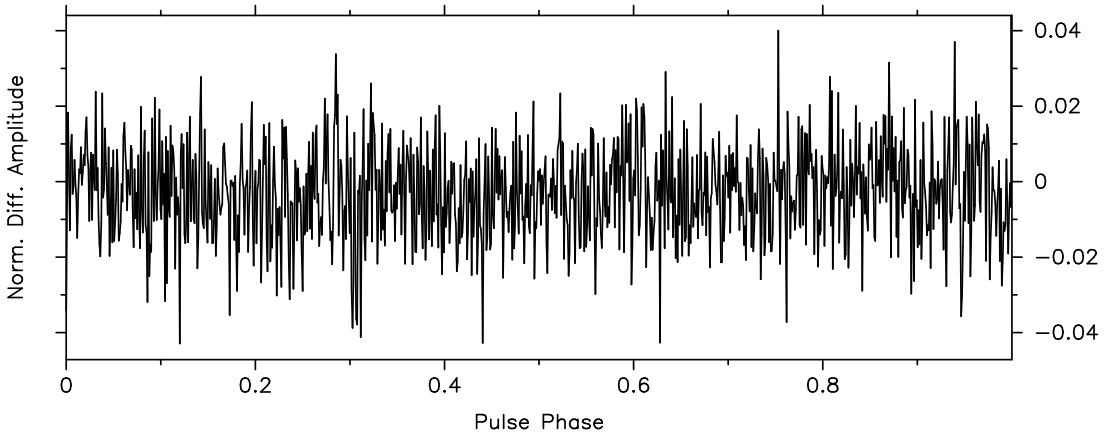


**Figure 6.6: Difference profile constructed from two CPSR2 observations taken on MJD 53204 at two frequencies separated by 64 MHz. There is no systematic morphological difference above the level of the noise that might be expected if the pulsar profile was rapidly evolving as a function of radio frequency.**

rapidly to another mode (Lyne 1971). In more recent years, various authors have reported observing random profile variations in some of the millisecond pulsars, most notably PSR B1821–24 (Backer & Sallmen 1997) and PSR J1022+1001 (Kramer et al. 1999). None of the reported variations, or the mode changing phenomenon, match the steady secular change we have observed in PSR J1141–6545, which is itself a slow pulsar and therefore may not suffer from the erratic variations seen in a small number of the millisecond pulsars. It is still unclear whether or not random variations in MSP profiles are intrinsic to the pulsar as a recent analysis of PSR J1022+1001 observations made by Hotan, Bailes & Ord (2004) finds no instabilities that might induce strange timing behaviour.

Can we hope to distinguish between geodetic precession and, say, free precession of the pulsar? Pulse shape variations are attributed to free precession in observations of PSR B1828–11 (Stairs, Lyne & Shemar 2000) and PSR B1642–03 (Shabanova, Lyne & Urama 2001). This evidence manifests in the form of profile shape changes correlated to variations in the pulse arrival times. These changes are cyclic, but with no clear *a priori* timescale. The geodetic precession timescale is already well determined to be near 265 years. From Fig. 6.1 we would hope to see the emission cone tilt at a rate $< 1.36^o\,\mathrm{yr}^{-1}$ and continue any secular trend for decades unless we are at one of the two crossing points encountered every geodetic precession period. The overwhelming majority of slow pulsars exhibit no evidence for free precession whatsoever. On the other hand, we expect any binary pulsar that has received a misaligning kick to precess to some degree. We therefore assert that geodetic precession is responsible for the observed secular variation in the mean profile of PSR J1141–6545 and we discuss this interpretation in section 6.6. Long-term monitoring of the precession will ultimately determine a timescale and hence the true mechanism.

We now present an analysis of the polarimetry of this pulsar, providing further

evidence that our line of sight to the emission cone is changing steadily with time.

## 6.5 Evolution of Polarised Emission

The standard interpretation of pulsar polarimetry is the rotating vector model (RVM), put forward by Radhakrishnan & Cooke (1969). This assumes a dipolar magnetic field whose central axis is offset from the neutron star rotation axis. As the emission cone sweeps past our line of sight, the changing orientation between the observer and the magnetic field is expected to produce a characteristic "S" shaped curve in the measured position angle of any linearly polarised components. In addition, the rate at which the linear polarisation vector rotates as the beam crosses the observer is dependent on whether or not the line of sight cuts close to the centre of the beam. Polarimetric observations may therefore offer a sensitive indicator of both the rate at which the beam is precessing past the observer and where in the emission cone we are at the present time.

First we consider the morphology of the polarised component of the mean pulse. The polarimetric changes observed between Figs. 6.2 & 6.3 are extreme. To convince ourselves that rapid polarimetric evolution is taking place, we focus on the most recent data with a high time density of observations and the same instrument (CPSR2). The polarimetric capabilities of CPSR2 and the PSRCHIVE software package (Hotan, van Straten & Manchester 2004) have recently been verified by comparison with published results with good agreement (Ord et al. 2004). Figure 6.7 shows profiles of the polarised emission of PSR J1141–6545, recorded using CPSR2 and ordered consecutively in time.

Figure 6.7 shows that the polarimetric profile is certainly changing, even over a period of less than 10 months. The linear and circular emission appears to converge during this time. Significant morphological differences in polarised emission appear between two widely separated observations that also correspond to a change in receiver, when the multibeam system was replaced by the refurbished H-OH. Therefore it is still possible that the differences in polarised emission could be instrumental in origin, although comparison with Fig. 6.2 suggests that position angle evolution is also occurring. Observations over a longer time span are required to make more conclusive (and perhaps quantitative) statements. Regular polarimetric observations of this pulsar will be a high priority in future years.

Next we attempt a more quantitative analysis of the position angle of the linearly polarised component of the pulsar beam. In the formalism of the RVM, measured position angle (PA, $\psi$), is presented as a function of the angle between the spin and magnetic axes ($\alpha$), the minimum angle between the magnetic axis and the line of sight ($\beta$) and pulse phase ($\phi$)

$$\tan(\psi(\phi) - \psi_0) = \frac{\sin\alpha \sin(\phi - \phi_0)}{\cos\alpha \sin\delta - \sin\alpha \cos\delta \cos(\phi - \phi_0)}. \quad (6.6)$$

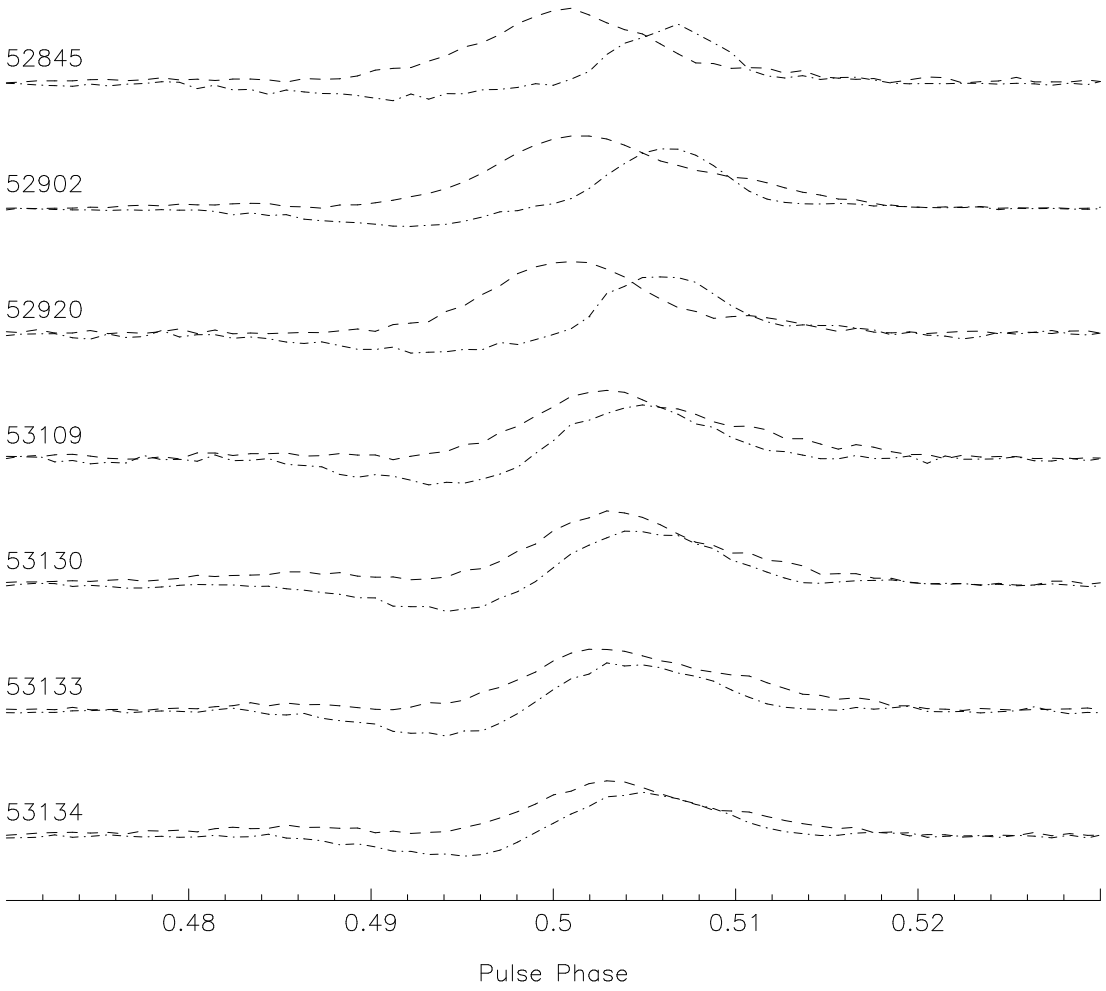


**Figure 6.7: Consecutive profiles of the polarised emission from PSR J1141–6545, observed with CPSR2 over a period of 289 days beginning in July, 2003 (MJD 52845). The dashed line represents total linear polarisation and the dashed-doted line represents total circular. The vertical axis (unlabeled for simplicity) is in mJy, the vertical range is kept constant across all sub-panels. All profiles have been polarimetrically calibrated using a simple model of relative gain and phase for two orthogonal linear receiver probes and flux calibrated against 3C218 (Hydra A). The peaks of the linear and circular components can be seen to move closer together as time progresses.**

Here $\delta = \alpha + \beta$ is the angle between the spin axis and the line of sight, as in section 6.3; $\phi_0$ is the pulse phase of steepest PA swing and $\psi_0$ is a constant position angle offset.

Unfortunately, application of this method to pulsars with narrow duty cycles and shallow or complicated PA swings does not well constrain $\alpha$ or $\beta$ independently. In common with the analysis presented by Kaspi et al. (2000a) we find it impossible to fit the RVM model to the early observations of PSR J1141–6545 with any degree of confidence. More recent observations are better suited to a partial RVM analysis, in that the recent PA behaviour includes a sweep in the central region of the pulse profile. The narrow duty cycle still restricts the applicability of the RVM and it is therefore impossible to constrain $\alpha$. Under the assumption that the steepest PA evolution is still providing an indication of the orientation of the magnetic field lines with respect to both the spin axis and the line of sight we have applied a simplification of the RVM model with the sole purpose of determining the general evolution of $\beta$. Within the RVM formalism the rate of change of PA as a function of pulse phase has a maximum value,

$$\left(\frac{d\psi}{d\phi}\right)_{\max} = \frac{\sin\alpha}{\sin\beta}. \tag{6.7}$$

We have evaluated the gradient of the steepest PA swing in those observations for which a straight line fit can be obtained from the same central region of the pulse profile. This procedure could only be applied to epochs between which the linear behaviour of the PA swing was considerably wider in pulse phase than any possible translation, allowing the profiles to be aligned by a suitable ephemeris. The applicable observations are those obtained with CPSR2 since mid-2003. Earlier observations cannot be subject to this analysis as they display PA behaviour which is too complicated. This fact alone indicates that the detected emission represents a different cut through the emission region than was evident in earlier observations.

The data have been grouped into two epochs (2003 and 2004) separated by approximately 0.7 years. Average position angle profiles were formed from the constituent observations at each epoch. These average PA profiles include observations from both 64 MHz observing bands and within each epoch are separated by several days. Although there were very slight variations between the PA profiles within each epoch, there was no systematic trend. A linear least squares minimisation of a straight line fit to the PA profiles demonstrates a significant deviation in the gradient obtained at the two epochs. Epoch one (2003) displays a gradient of $-15.1\pm0.3^o$ of PA swing per degree of phase. Epoch two (2004) displays a gradient of $-17.1\pm0.3^o$ of PA swing per degree of phase. A straight line fit to a difference PA profile, formed by the subtraction of the epoch two profile from the epoch one profile, is presented in Fig. 6.8. The gradient of this difference fit is $2.3\pm0.4^o$, which is consistent with the simple subtraction of the best fit for epoch two from that of epoch one.

Equation 6.7 was then evaluated for all $\alpha$ and a rate of change of $\beta$ between the two epochs was determined. The calculated value of $d\beta/dt$ is a strong function

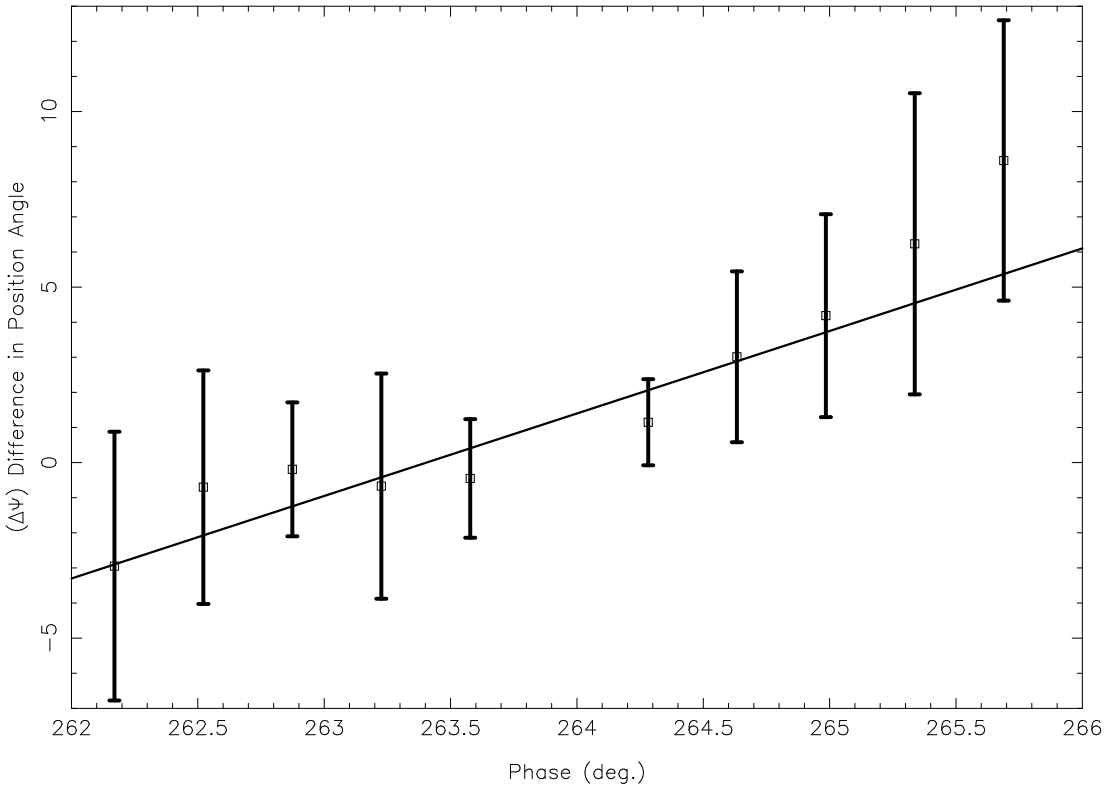


**Figure 6.8: Difference PA as a function of phase. The average PA from the mid-2004 epoch was subtracted from that of the mid-2003 epoch. A straight line fit has been applied to the residual. The steepening of the PA swing as a function of time is an indication that the pulsar emission beam may be precessing into, and not away from our line of sight.**

of $\alpha$ but suggests that $\beta$ is increasing. The peak rate of change is $0.8^o\,\mathrm{yr}^{-1}$ (68% confidence), but this interpretation requires the pulsar spin axis to be $90^o$ from the magnetic axis, which is unlikely if we assume the orientation is random. It is therefore probable that $d\beta/dt$ is less than $0.8^o\,\mathrm{yr}^{-1}$. The gradient of PA also indicates that $\beta$ is currently negative for all possible $\alpha$. We therefore infer that the beam is precessing into the line of sight at a rate less than $0.8^o\,\mathrm{yr}^{-1}$. It should be noted that $d\beta/dt \equiv d\delta/dt$ as the two angles are related by a constant offset. Thus we can use this result, coupled with Fig. 6.1, to make some general statements about the unknowns in Eq. 6.5. It is clear that either the misalignment angle is smaller than approximately $30^o$, or we are currently at a special precessional phase where $d\delta/dt$ is changing rapidly and has not attained its maximum value in the span of our observations. Conversely, unless $\alpha$ is near $0^o$ or $180^o$, we expect that the misalignment angle is greater than approximately $15^o$.

The mean flux of the pulsar in these observations is $4.0\pm0.5\,\mathrm{mJy}$. There does not appear to be a significant trend in mean flux within the past year, but this may change as the time baseline of flux calibrated data grows.

## 6.6 Discussion and Conclusion

Our observations indicate that the line of sight to the emission cone of PSR J1141–6545 is precessing deeper into the core, producing a steeper position angle swing. Depending on the chosen beam model, this could also explain why the pulse profile is seen to increase in width. Our result has implications both for the study of pulsar emission beams and for the detection rates of relativistic pulsars in large-scale surveys. PSR J1141–6545 is the slowest pulsar (by an order of magnitude) for which geodetic precession is observable, providing a unique means of examining the emission cone of a normal (un-recycled) pulsar. Further polarimetric observations, extended time baselines or the detection of orbitally modulated aberration may allow determination of our present location in the emission cone and the geometry of the beam as a whole. Given the expected precession period of order 265 years and the fact that we seem to still be moving towards the central axis of the beam, it is possible that this pulsar may only have precessed into view within the past few decades. As Kaspi et al. (2000b) noted, precession might explain the non-detection of this pulsar in early surveys with flux limits well below the required threshold (Johnston et al. 1992; Manchester et al. 1996). If this is the case, it could be argued that surveys of the sky for relativistic pulsars should continue on a regular basis. By their very nature, the most interesting objects are likely to be visible for the least amount of time.

# Chapter 7

# PSR J0737–3039A: Baseband Timing & Polarimetry

"In physics, you don't have to go around making trouble for yourself – nature does it for you."

*Frank Wilczek*

In this chapter, we describe May 2004 observations of the 22.7 ms "A" pulsar in the double pulsar binary system, J0737–3039. Data were obtained with a coherent dedispersion system at 20 cm and 50 cm wavelength bands, during an intensive 15 day observing session at the Parkes radio telescope. High S/N polarimetric profiles of the "A" pulsar are presented; these profiles provide templates against which to search for evolution of the pulse profile in coming years. We measure flux densities for the "A" pulsar of 1.2±0.1 mJy at 1373 MHz and 4.2±0.5 mJy at 685 MHz. Faraday rotation is also detected in both bands; the rotation measure of the pulsar is $120\pm4\,\mathrm{rad\,m^{-2}}$ at 1373 MHz and $118.4\pm0.3\,\mathrm{rad\,m^{-2}}$ at 685 MHz, implying a mean longitudinal component of the magnetic field of approximately $3\,\mu$G along the line of sight to the pulsar. Pulse arrival times from individual 2 min integrations are analysed and we detect a clear signature of Shapiro delay in the timing residuals. Assuming a mass of $1.25\,\mathrm{M}_\odot$ for the companion "B" pulsar, we derive a limit of $88.5^{+0.8}_{-1.1}$ degrees on the inclination angle of the system, almost consistent with the most recent limit of $0.29\pm0.14^o$ away from $90^o$, obtained from scintillation studies. The need to obtain many arrival times near inferior conjunction forces us to use low S/N profiles for our analysis. This can lead to unrealistic error estimates if standard timing techniques are used. In an appendix, we describe how to obtain more realistic error estimates and how to avoid some catastrophic consequences of building a standard template profile from low S/N data.

## 7.1 Introduction

The double pulsar binary system J0737–3039 is the only known example of a double neutron star binary in which both members are active radio pulsars. The discovery of the 22.7 ms "A" pulsar was first reported by Burgay et al. (2003). Pulse timing revealed that it orbited a companion star once every 2.4 hours, making the system an ideal laboratory for testing relativistic gravity, similar to the original binary pulsar system, PSR B1913+16 (Hulse & Taylor 1975) and the more recently discovered PSR B1534+12 (Wolszczan 1990). Within two months, Lyne et al. (2004) reported the detection of 2.8 s pulsations from the companion (PSR J0737–3039B). The presence of two pulsars in the same system enables timing measurements to over-determine many of the parameters required to describe the relativistic orbit, allowing unique tests of General relativity. The component masses of the system are already known to extremely high precision; according to Lyne et al. (2004), $m_A = 1.337 \pm 0.005\, M_\odot$ and $m_B = 1.250 \pm 0.005\, M_\odot$.

The diffractive interstellar scintillation timescale of the "A" pulsar is modulated by the large ($\sim$300 km s$^{-1}$) orbital velocity, allowing estimation (Ransom et al. 2004) of the space velocity and inclination angle of the system, assuming a thin scattering screen and isotropic interstellar medium. A recent analysis (Coles et al. 2005) of the two-dimensional correlation between the scintillation patterns of both the "A" and "B" pulsars indicates that the ISM is not isotropic on the scale of the projected semi-major axis of the orbit ($\sim$1.4 lt-s). Coles et al. (2005) correct for this and infer an inclination angle limit of $0.29 \pm 0.14^o$ away from edge-on ($90^o$), but caution that refractive effects in the magnetospheres of both pulsars or even gravitational lensing (Lai & Rafikov 2005) may introduce systematic errors in measurements of the inclination angle. Comparison of the values derived using various methods may in future help to isolate such effects.

Relativistic influences within the J0737–3039 system are so great that the longitude of periastron advances by $\sim 17^o\,\mathrm{yr}^{-1}$. If there is any misalignment between the spin and orbital angular momentum axes, we should expect the spin axes of both pulsars to precess on timescales of $\sim$70 years (Manchester et al. 2005b). This "geodetic precession" (Hari Dass & Radhakrishnan 1975; Barker & O'Connell 1975b; Esposito & Harrison 1975) has been detected in PSR B1913+16 (Weisberg, Romani & Taylor 1989; Kramer 1998; Weisberg & Taylor 2002), PSR B1534+12 (Stairs, Thorsett & Arzoumanian 2004) and PSR J1141–6545 (Hotan, Bailes & Ord 2005b). Manchester et al. (2005b) found no significant evidence of shape changes in the mean profile of "A" over a one-year period that might indicate a changing line of sight relative to the emission cone. It is possible that we are simply at a special precessional phase where changes to our line of sight are minimised, or that there is no angular momentum misalignment.

In this chapter we briefly describe the baseband recording system and coherent dedispersion process (Section 7.2) used to obtain data at the Parkes radio telescope. We also present high S/N polarimetric profiles for the "A" pulsar at centre frequen-

cies of 685 MHz and 1373 MHz (Section 7.3), both of which exhibit slightly different characteristics to the 820 MHz profile presented by Demorest et al. (2004). These profiles define the morphology of the mean pulse at the current epoch, providing a valuable reference that can be used to detect future pulse shape changes. Our data also reveal a value of rotation measure along the line of sight, similar in magnitude but opposite in sign to that obtained by Demorest et al. (2004). Section 7.4 describes our analysis of 2 min average arrival times obtained from observations of the "A" pulsar. We fit for several binary parameters, the spin period of the pulsar and the dispersion measure along the line of sight over a time span of 15 days, during which the "A" pulsar was observed for several hours per day. The astrometric position given by Lyne et al. (2004) is assumed in our model of the binary system. The signature of Shapiro delay is clearly present in our data and we use this to place limits on the inclination angle of the system. Section 7.5 summarises our results. Appendix A describes a new algorithm for measuring the relative phase shift between two pulse profiles and compares it with another commonly-used method. Simulated timing experiments were used to determine the most appropriate algorithm to apply to our observations.

## 7.2 Observations

PSR J0737–3039A was observed at the Parkes radio telescope as part of a study of several relativistic binary pulsars. The aim of this study was to obtain precise measurements of the binary parameters for a small number of pulsar systems, by observing them daily over a period of three weeks in April-May, 2004. Obtaining a high density of observations over several orbits is essential when characterising the signature of perturbations that only appear in small orbital phase ranges (like Shapiro delay in an edge-on system). During the course of this experiment, the "A" pulsar was tracked for several hours every day over a continuous period of 15 days, between MJD 53120 and MJD 53135. Data were recorded using CPSR2, which records 4×64 MHz bands, normally representing two dual-polarisation bands at different sky frequencies. Observations were made across two bands centred on 1341 MHz and 1405 MHz, using the wide-band H-OH receiver, and across a single band centred on 685 MHz, using the low-frequency feed from the coaxial dual-band 10/50 cm pulsar receiver. Both the 10/50 and H-OH systems use orthogonal linear probes to record full polarimetric information.

CPSR2 real-samples each IF at the Nyquist rate with two bits of precision, writing raw samples to the disks of a processing cluster. This cluster of servers performs coherent dedispersion (Hankins & Rickett 1975), writing out archives folded using a single ephemeris. In the standard mode of operation, we synthesise a filterbank with 128 coherently dedispersed spectral channels, corresponding to an effective sampling time of 2 $\mu$s. Individual integrations correspond to 1 GB of baseband data or approximately 16.8 s of observing time.

Individual tracks lasted for approximately one hour and were often performed sequentially for a period of several hours each day. Between these on-source tracks, short (2 min) calibration pointings were taken, $1^o$ South of the pulsar position. During these calibration observations, the receiver noise source was driven with a square wave at a frequency of $\sim$11 Hz, injecting a linearly polarised signal at a known angle to both receiver probes. This characterises the polarimetric response of the signal chain so that corrections to the observed Stokes parameters can later be made. Observations of the radio galaxy 3C218 (Hydra A) were also taken several times each week to calibrate the absolute flux scale of the observing system. All calibration and data reduction was performed using the methods provided within PSRCHIVE (Hotan, van Straten & Manchester 2004). We use a simple case of the polarimetric model described by van Straten (2004), performing a correction for relative gain and phase between the two orthogonal receiver probes.

## 7.3 Polarimetric Profiles

All 16.8 s archives produced by PSRDISP were averaged within each 64 MHz band to form 2 min integrations. These integrations were then calibrated and corrected for relative gain and phase and any changes in parallactic angle induced by the altitude-azimuth antenna mount. Because the 1341 and 1405 MHz bands were recorded simultaneously, each using two of the four CPSR2 IFs, it was possible to find pairs of archives corresponding to the same time period within a given 20 cm observation. These pairs were combined (after calibration) to form new archives spanning 128 MHz of bandwidth. The addition was performed without any loss of spectral resolution. These frequency-added archives were then summed in time to form a grand average archive representing the mean polarimetric state during our two week session. The 685 MHz integrations were also summed in time to form a single grand average archive with 64 MHz of bandwidth. The grand average archives were subjected to a fitting procedure designed to determine the Faraday rotation measure in the direction of the pulsar. The fitting algorithm applies various trial rotation measure (RM) corrections and measures the total polarised flux density (the quadrature sum of Stokes Q, U and V) after aligning and combining all spectral channels. Because Faraday rotation can only decrease the amount of observed polarised flux density, it is reasoned that the trial RM leading to a maximum level of polarised emission must represent our best estimate of the true physical value. A Gaussian curve is fitted to the flux density values obtained using a discrete sample of trial rotation measures and the centroid taken to be the measured RM. The uncertainty is defined to be the distance by which the centroid must shift to increase the $\chi^2$ of the Gaussian fit by unity. The error at each trial RM is defined to be the standard deviation of the baseline noise in the polarised emission profile, multiplied by the square root of the number of bins used to compute the flux density.

At 20 cm, we find a rotation measure of $120\pm4\,\mathrm{rad\,m^{-2}}$. At 50 cm the value

is 118.4±0.3 rad m$^{-2}$. These values are consistent with each other to within $1\sigma$ but the 50 cm observations are more precise, simply because depolarisation due to Faraday rotation is a strong function of observing frequency. Demorest et al. (2004) obtain values for rotation measure for both the "A" (–112.3±1.5 rad m$^{-2}$) and "B" (–118±12 rad m$^{-2}$) pulsars independently, at a centre frequency of 820 MHz. It should be noted that Hotan, Bailes & Ord (2005a) published negative rotation measure values but these were subsequently discovered to be in error due to a mistake in the calibration process. We are confident that the true rotation measure is positive. Even ignoring the sign difference, the rotation measure quoted by Demorest et al. (2004) for the "A" pulsar is not consistent with our own measurement, although the difference is within $3\sigma$. The discrepancy may be due to the different methods used to fit for a rotation measure value. The method described in this chapter makes no assumptions about the linearity of the position angle sweep and may therefore be more robust. The appropriate correction was applied to the grand average archives and all frequency channels summed to produce Figs. 7.1 & 7.2[1]

The overall polarimetric structure of Figs. 7.1 & 7.2 is consistent with the 820 MHz profile presented by Demorest et al. (2004), except for the sign errors noted previously. The mean profile clearly evolves with frequency, the relative strength of the two components is nearly the same at 50 cm, while the trailing component dominates at 20 cm. Manchester et al. (2005b) present mean profiles that show similar frequency evolution. The 3030 MHz profile presented by Manchester et al. (2005b) reveals the presence of an additional, narrow outer-edge component that begins to dominate the profile at higher frequencies. Both the main profile components contain a highly linearly polarised region with approximately half the duty cycle of the host component. In all cases, approximate mirror-symmetry about pulse phase 0.5 is preserved. Figures 7.1 & 7.2 show that the structure of the profile is quite complicated on small timescales. The position angle of the profile at both frequencies is consistent with data presented by Demorest et al. (2004) at 820 MHz. The regions of strongest linear polarisation have an almost flat position angle profile, while the inner edges exhibit more complicated behaviour that, at least in the leading component, may involve competing orthogonal modes.

Polarimetric structure is a sensitive indicator of emission region geometry and secular profile evolution (Stairs, Thorsett & Arzoumanian 2004). Since the geodetic precession period of "A" is only ∼70 years, Figs. 7.1 & 7.2 will be useful as a future reference to study any changes in the polarimetric profile.

We obtain mean flux densities for the "A" pulsar of 1.2±0.1 mJy at 1373 MHz and 4.2±0.5 mJy at 685 MHz. Assuming a simple power law relation between flux and frequency, we measure a spectral index of –1.8±0.3. Our 20 cm flux is consistent with the value published by Lyne et al. (2004), but the spectral index we measure

[1]The plots shown in this thesis differ from the plots that were published by Hotan, Bailes & Ord (2005a). An error in the receiver orientation parameters assumed during the calibration process meant that the sign of Stokes V and the direction of the position angle swing were incorrectly reported in the previous publication.

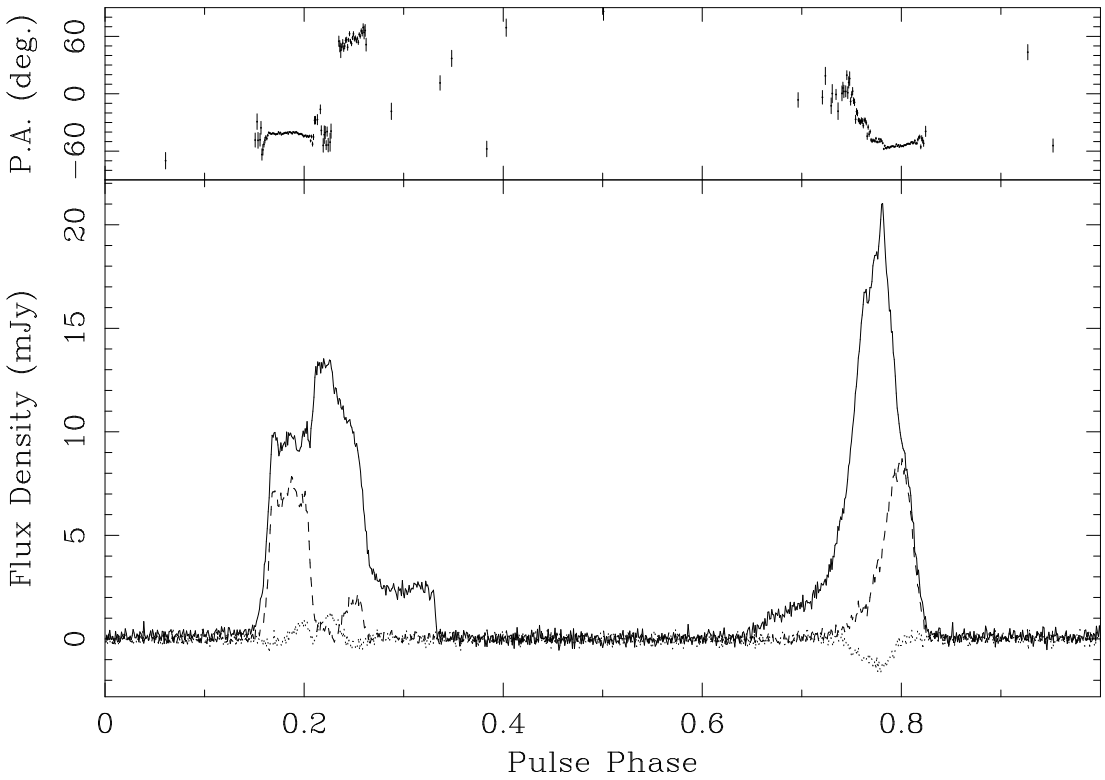


**Figure 7.1: PSR J0737–3039A average profile at a centre frequency of 1373 MHz, formed from 46 hours of data taken in two dual-polarisation bands using the H-OH receiver and combined in software for a total bandwidth of 128 MHz. The profile in the lower panel has been calibrated for relative instrumental gain and phase. The solid curve represents total intensity (Stokes I), the dashed curve represents linearly polarised emission and the dotted curve represents circularly polarised emission. There are two main pulse components, both with linearly polarised sub-components near the outer edge of the profile. The relative position angle of the linearly polarised radiation is shown in the top panel. Absolute flux calibration was performed using 3C218 (Hydra A) as a reference.**

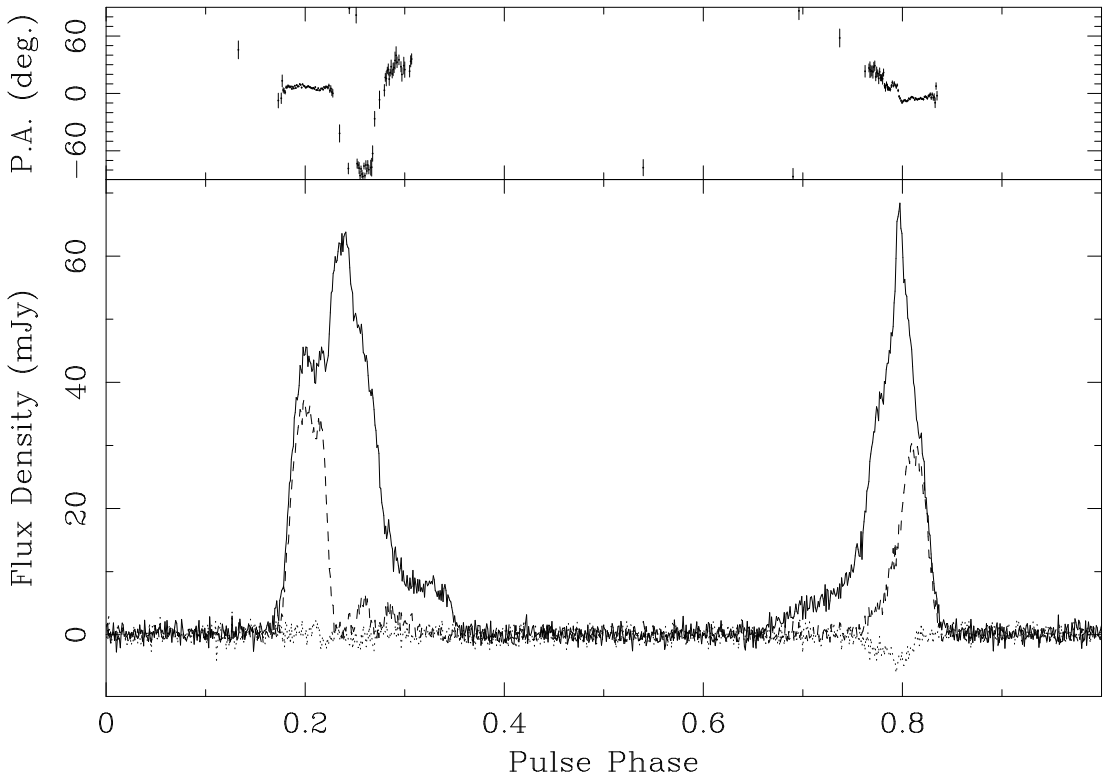


**Figure 7.2: PSR J0737–3039A average profile, as in Fig. 7.1, formed from 14 hours of data taken at a centre frequency of 685 MHz. Data were taken with the low-frequency feed of the coaxial 10/50 cm pulsar receiver using a single 64 MHz wide, dual-polarisation band.**

appears to be at odds with Fig. 1 of Demorest et al. (2004), where the mean flux density at 820 MHz is shown to be approximately 5 times higher than we would predict.

## 7.4 Arrival Time Analysis

The S/N obtainable from the "A" pulsar in 2 min integrations with 128 MHz of bandwidth and a 64 m diameter telescope is of order 10, quite low for precision timing purposes. We chose to apply the time-domain algorithm described in Appendix A to measure the relative shift of each profile with respect to the standard template, to avoid severely underestimating the arrival time uncertainties.

All 2 min, calibrated archives were summed in frequency and polarisation to produce a single mean profile. Arrival times were obtained from each profile with S/N greater than 7 and fit to a model (Damour & Deruelle 1985; Damour & Deruelle 1986) of the binary system (henceforth known as the "DD" model) that includes theory-independent post-Keplerian corrections. Total intensity versions of Figs. 7.1 & 7.2 were used as standard template profiles, after mutual alignment. The standard pulsar timing package TEMPO was used to perform the model fit. Because our time baseline only covers 15 days, we do not fit for any parameters that induce long-term signatures, instead adopting the values published by Lyne et al. (2004). Table 7.1 summarises the parameters of the timing model and the properties of our data set. All errors quoted represent the $1\sigma$ uncertainty reported by TEMPO.

To obtain a reduced $\chi^2$ of unity, a scaling factor of 1.72 was applied to all errors. All of our fitted parameters are consistent to within $1\sigma$ of the corresponding values given by Lyne et al. (2004) and are in some cases of higher precision. In order to search for the signature of Shapiro delay in our data, we constructed a $\chi^2$ map over two timing model parameters (the companion mass $m_B$ and cosine of the inclination angle $i$) that define the range and shape of the Shapiro delay. This map is shown in Fig. 7.3.

Although we have relatively few data points, there is a broad but clearly defined constraint on the companion mass and inclination angle of the system. If we assume the companion mass value reported by Lyne et al. (2004) it is possible to take a one-dimensional slice through the map and constrain the inclination angle alone (Fig. 7.4).

The $1\sigma$ limit shown in Fig. 7.4 constrains the inclination angle to be $88.5^{+0.8}_{-1.1}$ degrees. The $2\sigma$ limit allows any inclination angle up to and including $90^o$. This result is consistent with the (less precise) limit of $87\pm3^o$ (also based on Shapiro delay) published by Lyne et al. (2004). Ransom et al. (2004) use a combination of timing and scintillation velocity methods to derive a limit of $88.7\pm0.9^o$, which is also consistent with our result. Coles et al. (2005) use scintillation velocity methods alone to derive a stronger limit of $0.29\pm0.14^o$ away from $90^o$, which is still within our $2\sigma$ region. These results are in good agreement, especially considering the very

| Parameter | Value |
|---|---|
| **Right Ascension** ($\alpha$) | 07:37:51.247 (2) |
| **Declination** ($\delta$) | –30:39:40.74 (3) |
| **Period Derivative** ($10^{-18}$) | 1.74 (5) |
| **Periastron Advance** ($\dot{\omega}$) (deg yr$^{-1}$) | 16.90 (1) |
| **Gravitational Red-shift** ($\gamma$) (ms) | 0.38 (5) |
| **Mass of B** ($M_\odot$) | 1.250 (5) |
| Period (ms) | 22.6993785566 (1) |
| Period Epoch (MJD) | 52870.0000 |
| Dispersion Measure (cm$^{-3}$pc) | 48.9151 (5) |
| Projected Semi-Major Axis (lt-s) | 1.415044 (2) |
| Eccentricity | 0.087784 (3) |
| Time of Periastron Passage (MJD) | 53127.6859965 (5) |
| Orbital Period (days) | 0.1022515621 (5) |
| Angle of Periastron (deg) | 85.726 (2) |
| Inclination Angle (deg) | 88.5 (+8 –11) |
| MJD of first TOA | 53120 |
| MJD of last TOA | 53135 |
| Total Time Span (days) | 15 |
| Number of TOAS | 529 |
| RMS Timing Residual ($\mu$s) | 30.933 |

**Table 7.1: PSR J0737–3039A "DD" timing model parameters derived from 15 days of CPSR2 observations. The error in the last significant digit is given in parentheses after the value. All parameters in bold were adopted from Lyne** et al. **(2004).**

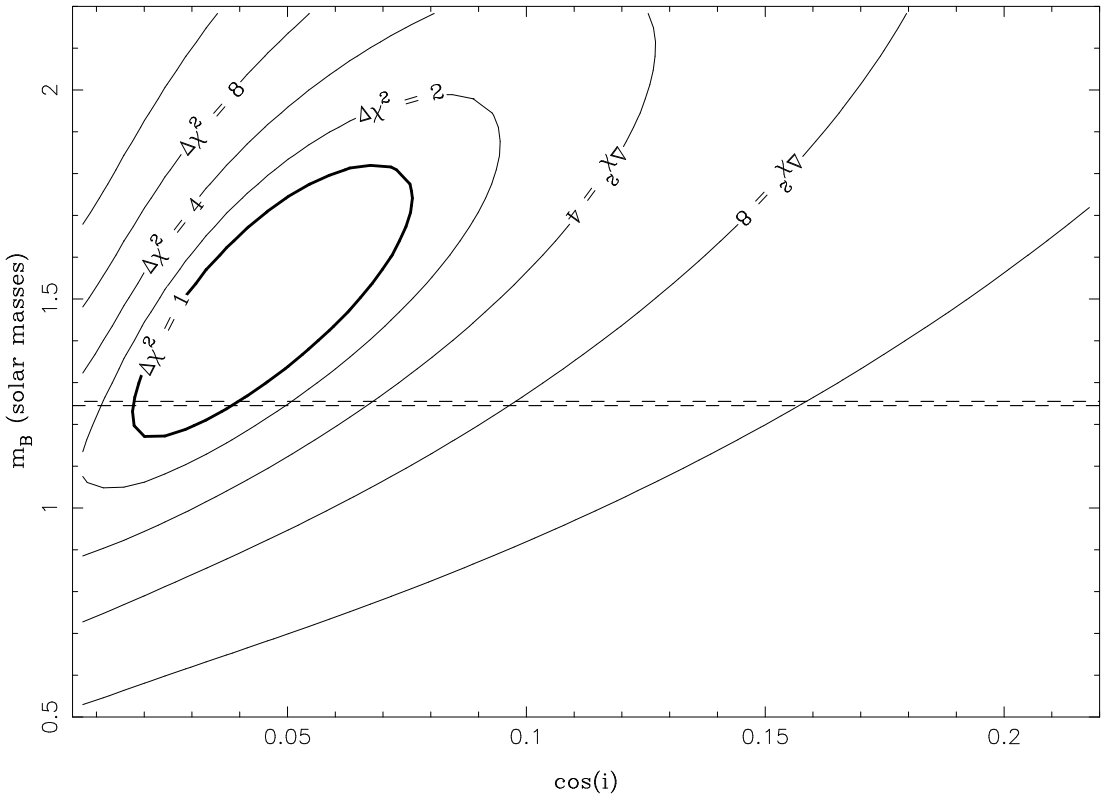


**Figure 7.3: PSR J0737–3039A Shapiro delay contour map. We show the $\chi^2$ reported by TEMPO when holding the values of $m_B$ and $\cos(i)$ fixed at each point in the map and fitting for the Keplerian binary parameters, spin period and dispersion measure. Curves of constant $\Delta\chi^2$ (with respect to the minimum) are shown. The two dashed lines represent the companion mass limits published by Lyne et al. (2004).**

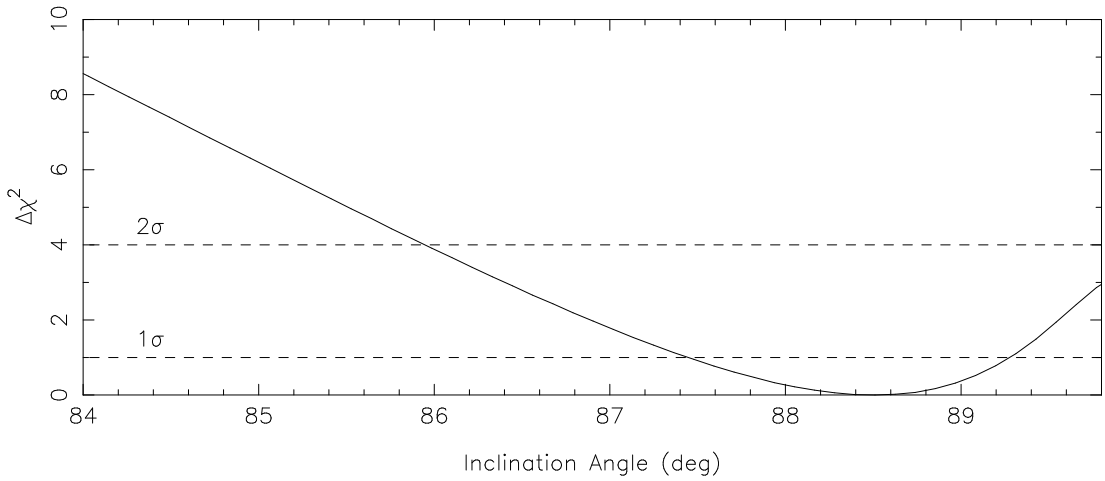


**Figure 7.4: One-dimensional slice through the Shapiro delay contour map shown in Fig. 7.3, assuming the mass of "B" is 1.250±0.005,$M_\odot$ as given by Lyne** et al. **(2004). Note that the error in the mass of B is so small that it moves the curve by approximately the width of the line. The minimum $\chi^2$ occurs at an inclination angle of 88.5$^o$. Lines of 1$\sigma$ (68%) and 2$\sigma$ (95%) confidence are drawn horizontally. The curve presented in this figure has a slightly different confidence interval compared to that obtained from a simple cut through Fig. 7.3 because we have re-computed the extent of the confidence interval based on the new minimum $\chi^2$ value associated with our assumed companion mass. Note that the "DD" model includes a degeneracy in the inclination angle, Shapiro delay measurements alone cannot determine if the projection of the orbital angular momentum vector along our line of sight is pointing towards or away from the observer. We give values for $i$ assuming that the projection points towards Earth and the inclination angle is therefore less than 90$^o$.**

different methods used.

Lai & Rafikov (2005) use the inclination angle estimate of Coles et al. (2005) to predict an additional $\sim$4 $\mu$s contribution to the Shapiro delay as a result of gravitational lensing. More interestingly, Lai & Rafikov (2005) also predict the presence of a new timing perturbation due to the apparent displacement of the pulsar during lensing episodes, which last for $\sim$4 s. Detection of either of these effects will require much smaller RMS timing residuals and higher S/N observations, to enable studies of the polarimetric profile and pulse arrival times over small orbital phase ranges. The fact that the "A" pulsar is eclipsed during the 4 s lensing event makes detection of any effect even more challenging.

## 7.5 Conclusions

Our calibrated mean polarimetric profiles of PSR J0737–3039A may be used as a reference for future studies of pulse profile evolution. The rotation measure in the direction of the system is found to be 120$\pm$4 rad m$^{-2}$ at 1373 MHz and 118.4$\pm$0.3 rad m$^{-2}$ at 685 MHz. The flux density of this pulsar is measured to be 1.2$\pm$0.1 mJy at 1373 MHz and 4.2$\pm$0.5 mJy at 685 MHz, giving a spectral index of –1.8$\pm$0.3 (assuming a power law relation). The detection of Shapiro delay in our pulse timing residuals is used to place a limit of $88.5^{+0.8}_{-1.1}$ degrees on the inclination angle of the system, near that obtained by other authors from both timing and scintillation experiments.

# Chapter 8

# High-Precision Baseband Timing of 15 Millisecond Pulsars

"I believe that there is a subtle magnetism in Nature, which, if we unconsciously yield to it, will direct us aright."

*Henry David Thoreau*

In this chapter, we describe extremely precise timing experiments performed on 5 solitary and 10 binary millisecond pulsars during the past 3 years, with the CPSR2 coherent dedispersion system at the Parkes 64 m radio telescope. 12 of our sources have RMS timing residuals below 1.5 $\mu$s and 4 are below 200 ns. The quality of our data allows us to measure 8 parallaxes and 9 proper motions, from which we conclude that models of galactic electron density still have limited predictive power for individual objects. We derive a mean transverse velocity of $87\ ^{+31}_{-14}$ km s$^{-1}$ for these pulsars, in good agreement with previous authors. We demonstrate that unless multi-frequency observations are made, typical variations in DM could introduce an additional drift in arrival times of $\sim$1 $\mu$s per year at 20 cm wavelengths. Our high timing precision means that Shapiro delay can be used to constrain the inclination angles and component masses of all but two of the selected binary systems. The signature of annual orbital parallax is detected in the timing of PSR J0437–4715 and PSR J1713+0747, providing additional geometric constraints. The timing of PSR J1909–3744 is used to demonstrate that the DE405 ephemeris is a better model of the Solar system than the earlier DE200. In addition, we show that pulsar astrometric parameters measured using DE200 and DE405 often differ significantly. In order to use pulsars to search for a cosmological gravitational wave background, it is desirable to time them against each other to eliminate Earth-based time standards. We demonstrate that PSR J1909–3744 can be used as a reference against which we obtain a very small RMS residual of 133 ns for PSR J1713+0747. Although the gain

of the Parkes antenna is small compared to other telescopes involved in precision timing, we obtain some of the lowest RMS residuals ever measured, highlighting the importance of good instrumentation such as CPSR2 and good analysis software.

## 8.1 Introduction

The pulse periods of the radio pulsars are so regular that measuring ensemble average pulse arrival times and fitting them to a mathematical model allows the pulsar's position and several other parameters to be determined very accurately (Reichley, Downs & Morris 1970). Pulsars can therefore be used as remote clocks that allow Earth-based observers to probe the physics of distant and in some cases extraordinary binary systems. The discovery of the first pulsar in a binary system, PSR B1913+16, opened the way to new timing-based studies of relativistic 2-body gravitation (Hulse & Taylor 1975). Fortuitously, PSR B1913+16 resides in a compact, eccentric "double neutron star" (DNS) binary system. Taylor & Weisberg (1982) showed that the pulsar's orbit was decaying at precisely the rate predicted due to energy loss by gravitational radiation, confirming an important prediction of General relativity. The discovery of this system spurred early development of advanced timing models (Blandford & Teukolsky 1976; Damour & Deruelle 1985; Damour & Deruelle 1986).

Several years later, the first millisecond pulsar PSR B1937+21 was discovered (Backer et al. 1982). The millisecond pulsars represent a distinct class of objects, with much smaller period derivatives and correspondingly smaller inferred magnetic fields and larger characteristic ages. The majority of MSPs reside in binary systems, supporting the notion that they are members of the normal pulsar population that have accreted matter from a companion (Bhattacharya & van den Heuvel 1991). In general, they are very stable rotators and can be timed with precision levels far higher than their more slowly rotating counterparts.

Since the discovery of PSR B1937+21, several successful surveys have found more than 100 new MSPs. The Princeton-Arecibo declination-strip survey (Camilo et al. 1996) and the Green Bank fast pulsar survey (Sayer, Nice & Taylor 1997) both discovered 2 (although PSR J1022+1001 was found almost simultaneously in both surveys), and the highly successful Parkes 70 cm survey (Manchester et al. 1996; Lyne et al. 1998; D'Amico et al. 1998) discovered 17. The Parkes multibeam survey (D'Amico et al. 2001) was not optimised for rapidly spinning sources but recent reprocessing (Faulkner et al. 2004) of the data using advanced search algorithms added another 15 MSPs to the catalogue. The Swinburne intermediate-latitude survey discovered a further 8 (Edwards & Bailes 2001a; Edwards & Bailes 2001b) of these objects. Roughly half the MSP population resides in the two globular clusters 47 Tucanae (Manchester et al. 1991; Camilo et al. 2000b) and Terzan 5 (Ransom et al. 2005). Unfortunately, globular cluster pulsars tend to have small flux densities and are unsuitable for precision timing as the motion of the pulsar in

the gravitational potential of the cluster is difficult to model (Blandford, Romani & Applegate 1987).

The increased timing precision that one can obtain from an MSP allows the detection of subtle timing perturbations, making them ideal for the study of post-Keplerian gravitation. One such perturbation is known as Shapiro delay and was first observed in the binary millisecond pulsar PSR B1855+09. The orbital plane of this system is nearly edge-on to our line of sight, leading to a gravitationally-induced propagation delay that can be detected whenever the low-mass companion comes close to our line of sight (Ryba & Taylor 1991). Kaspi, Taylor & Ryba (1994) measure a median time of arrival (TOA) uncertainty below $1\,\mu$s for this pulsar, allowing the signature of Shapiro delay to be studied despite the relatively long 12 day orbital period.

In general, the detection of post-Keplerian perturbations can provide additional constraints on the equations that describe orbital motion under a given theory of gravity. If several additional parameters can be measured within the same system, the component masses can be determined and the self-consistency of gravitational theories can be tested (Taylor & Weisberg 1989; Stairs et al. 1998; Bailes et al. 2003; Stairs 2004).

The more precisely we measure arrival times, the better our knowledge of the pulsar system. Sandhu et al. (1997) obtained an RMS timing residual of $0.5\,\mu$s for PSR J0437–4715. This allowed the proper motion and parallax to be measured with a high degree of precision. They also detected a secular change in the projected semi-major axis of the orbit, due to the proper motion of the system gradually altering our line of sight to the orbital plane. van Straten et al. (2001) used a coherent dedispersion system to make very precise arrival time ($0.13\,\mu$s RMS) measurements and were able to detect a more subtle kinematic effect known as annual orbital parallax. This allowed the three-dimensional orbital geometry of the system to be determined. The authors used General relativity to predict the signature of Shapiro delay (based on the orbital geometry) and subsequently detected it in the timing residuals, providing further unique tests of this theory of gravity. In this chapter, we confirm these outstanding results for the first time with a completely independent data set.

More recently, Jacoby et al. (2003) report the discovery of PSR J1909–3744, a binary MSP with a very narrow pulse profile that promises to be a good target for precision timing. Its highly inclined orbit has allowed the most precise determination of the mass ($1.438 \pm 0.024\,\mathrm{M}_\odot$) of a millisecond pulsar (Jacoby et al. 2005). In this chapter we extend the timing baseline and perform new experiments on these data.

The plan of this chapter is as follows: Section 8.2 summarises the characteristics of our 15 selected sources. Section 8.3 describes the observing system used to record pulse profiles at the Parkes radio telescope and outlines the methods used to obtain arrival times. Section 8.4 describes the model fitting process. In Section 8.5 we present standard template profiles and updated ephemerides for each of the pulsars observed, along with a brief description of our findings. Section 8.6 discusses

pulsar distance estimates and proper motion measurements, Section 8.7 discusses Shapiro delay and annual-orbital parallax and Section 8.8 discusses various factors that influence the timing stability of a pulsar, including variations in the observed DM. We also use our best pulsar as a time reference to obtain 133 ns RMS residuals for PSR J1713+0747. This is an essential first step towards using pulsars to search for gravitational radiation in a manner independent of terrestrial time standards. Finally, Section 8.9 summarises the chapter as a whole.

## 8.2 Targeted Millisecond Pulsars

We have concentrated our efforts on 15 MSPs, selected due to a combination of their interesting physical characteristics and precision timing potential. Ord et al. (2006) describe the timing of a 16$^{\rm th}$ object, PSR J1600–3053, in more detail. For simplicity of presentation, we divide our source list into two groups on the basis of whether or not the MSP has a binary companion. Table 8.1 shows the discovery reference, RMS timing precision and the mean flux density at 20 cm wavelengths for each pulsar, along with the spin and orbital (where applicable) periods.

Isolated MSPs are thought to be recycled pulsars that have lost their companions, perhaps via tidal disruption or ablation caused by the pulsar itself (Fruchter, Stinebring & Taylor 1988). In this chapter, we present observations of 5 isolated MSPs. On the other hand, binary MSPs typically reside in the system responsible for their recycling. The companion star is likely to have evolved significantly since the spin-up epoch, becoming either a white dwarf or another neutron star. In this chapter we also present observations of 10 binary MSPs.

## 8.3 Observations and Analysis

Three years ago, the second Caltech Parkes Swinburne Recorder (CPSR2) was commissioned at the Parkes radio telescope. During the period spanning August 2002 – April 2005, regular observing sessions ($\sim$1 per month) were conducted, each a few days in duration. These produced approximately 1.5 TB of folded data, which were stored on the Swinburne supercomputer to allow rapid, automated re-processing. Data were archived as coherently dedispersed time-series, folded at the apparent pulse period of each source and averaged with 16.8 s time resolution. We occasionally conducted intensive observing sessions of up to three weeks in duration, when several selected binary MSPs were observed daily. During these 3 years, we obtained many orbits of each pulsar, giving us good binary phase and day-of-year coverage which allows accurate measurement of astrometric and binary parameters.

CPSR2 is a general-purpose baseband recorder that was constructed primarily from consumer components. It samples the received signal after down-conversion and performs all subsequent processing in software. The only specially-designed hardware component is the Fast Flexible Digitiser (FFD) board that links 4$\times$64 MHz

| PSR | Discovery | Timing | RMS (old) $\mu$s | RMS (new) $\mu$s | Flux mJy | Period ms | $P_b$ days |
|---|---|---|---|---|---|---|---|
| J0711–6830 | Bailes et al. (1997) | Toscano et al. (1999b) | 1.3 | 4.7 | 3.7 | 5.5 | – |
| J1024–0719 | Bailes et al. (1997) | Toscano et al. (1999b) | 2.9 | 1.1 | 1.5 | 5.2 | – |
| J1744–1134 | Bailes et al. (1997) | Toscano et al. (1999a) | 0.3 | 0.3 | 5.3 | 4.1 | – |
| B1937+21 | Backer et al. (1982) | Kaspi, Taylor & Ryba (1994) | 0.2* | 0.1 | 12.5 | 1.6 | – |
| J2124–3358 | Bailes et al. (1997) | Toscano et al. (1999b) | 3.3 | 2.7 | 4.5 | 4.9 | – |
| J0437–4715 | Johnston et al. (1993) | van Straten et al. (2001) | 0.1 | 0.1 | 137 | 5.8 | 5.7 |
| J0613–0200 | Lorimer et al. (1995a) | Toscano et al. (1999b) | 2.2 | 0.6 | 1.3 | 3.1 | 1.2 |
| J1022+1001 | Camilo et al. (1996) | Hotan, Bailes & Ord (2004) | 0.7 | 1.5 | 10.9 | 16.5 | 7.8 |
| J1045–4509 | Bailes et al. (1994) | Toscano et al. (1999b) | 10.2 | 5.9 | 2.5 | 7.5 | 4.1 |
| J1603–7202 | Lorimer et al. (1996) | Toscano et al. (1999b) | 1.0 | 1.6 | 3.2 | 14.8 | 6.3 |
| J1713+0747 | Foster, Wolszczan & Camilo (1993) | Splaver et al. (2005) | 0.2 | 0.1 | 10.4 | 4.6 | 67.8 |
| B1855+09 | Segelstein et al. (1986) | Kaspi, Taylor & Ryba (1994) | 0.8* | 1.0 | 3.1 | 5.4 | 12.3 |
| J1909–3744 | Jacoby et al. (2003) | Jacoby et al. (2003) | 0.1 | 0.1 | 2.1 | 2.9 | 1.5 |
| J2129–5721 | Lorimer et al. (1996) | Toscano et al. (1999b) | 0.8 | 1.2 | 1.7 | 3.7 | 6.6 |
| J2145–0750 | Bailes et al. (1994) | Loehmer et al. (2004) | 0.5 | 1.3 | 9.4 | 16.1 | 6.8 |

**Table 8.1: Summary of selected MSPs. The best previously published RMS residuals are listed beside the corresponding journal reference in columns 3 and 4. Column 5 shows the RMS residual we obtain from the data presented in this chapter. The RMS residuals shown in columns 4 & 5 have been normalised to a standard integration time of 60 min, scaling simply by the square root of the integration time where necessary. Readers should note that in some cases, systematic errors reduce the effectiveness of time averaging and the theoretical 60 min RMS cannot be reached in practice. The 20 cm flux densities in column 6 were obtained from Ord** et al. **(2004). * Kaspi, Taylor & Ryba (1994) do not quote RMS residuals, choosing instead to give typical values for the uncertainty in an individual day-long average TOA, which could be significantly smaller than the RMS residual over the entire data set. These values are included in the table for comparative purposes and have not been scaled to the standard 60 min integration time.**

bands to a cluster of high-speed computers. In the standard mode of operation, these four bands are configured as orthogonal polarisations from two different centre frequencies. We used three different cryogenically-cooled receiver packages to observe 64 MHz-wide bands centred on 685, 1341, 1405 and 3000 MHz. Wavelengths near 20 cm were recorded with the wide-band H-OH receiver and the central beam of the Parkes multibeam receiver (Staveley-Smith et al. 1996). Bands centred on 685 MHz and 3000 MHz were observed (often simultaneously) with the coaxial dual-band pulsar receiver, installed near the end of 2003.

The CPSR2 FFD performs 2-bit Nyquist sampling on each band and transfers the samples to two high-speed servers (known as "primary nodes") via EDT–PCD60 Direct Memory Access (DMA) cards. Two bands are sent to each primary node, allowing orthogonal polarisations to be packed together. Data acquisition software monitors the sampler's statistics and adjusts attenuators on the FFD to ensure optimal digitisation. Once in RAM, the samples are distributed via gigabit Ethernet to a cluster of 28 dual Intel Xeon processing (or "secondary") nodes, located in shielded racks at the telescope. In the standard mode of operation, each secondary node receives 1 GB ($\sim$16.8 s) of raw data which it must process before the next segment arrives. The cluster is divided in two, each half processing a single dual-polarisation band. The secondary nodes perform real-time coherent dedispersion (Hankins & Rickett 1975) using a program called PSRDISP (van Straten 2003). This program averages the data synchronously with the predicted pulse period and constructs a coherent filterbank with (typically) 128 channels. The four Stokes parameters are constructed from the orthogonal polarisations. PSRDISP corrects for 2-bit quantisation errors (Jenet & Anderson 1998) and discards any data with bad digitisation statistics. The resulting "folded" archives are binned using (typically) 1024 phase bins.

The PSRCHIVE (Hotan, van Straten & Manchester 2004) software package was used to sum the archives to produce 5 min integrations. 12.5% of the passband was given zero weight at both edges to reduce the effects of filter roll-off and aliasing. Unless otherwise specified, all remaining frequency channels were dedispersed and summed to increase S/N and arrival times were computed from the uncalibrated total intensity (Stokes I) mean profile. If necessary, further time averaging was performed to increase the S/N.

The existing CPSR2 cluster can operate in real-time at a wavelength of 20 cm, provided the DM of the source is less than approximately 40 $\mathrm{cm}^{-3}$pc. Monitoring software displays an integrated profile for the whole observation that updates as each short integration is processed. This allows the observer to immediately assess data quality and fix hardware configuration problems with minimal loss of observing time.

Standard template profiles were constructed for each wavelength band by summing together all the observed profiles. Templates from distinct wavelength bands were aligned to a common fiducial phase by finding the maximum of their cross correlation function. The two 64 MHz-wide bands centred on 1341 and 1405 MHz

were combined and all frequency channels integrated to yield a single mean profile with a centre frequency of 1373 MHz. This profile was used to independently time the 1341 and 1405 MHz data. We found that very few points needed to be discarded from the timing residuals due to "corruption". The Parkes 20 cm band is relatively interference free and the statistical tests performed by CPSR2 and PSRDISP ensure a high level of data quality. In most cases, a simple selection on the basis of S/N is sufficient to reach the RMS residuals presented in this chapter, without the arbitrary deletion of "bad" profiles required by many of our former instruments.

## 8.4 Timing Models

The algorithm described by Taylor (1992) was used to measure each individual arrival time. In cases of low S/N we found it necessary to use a procedure that involves zeroing the last 7/8 of the discrete Fourier transform of the template profile before fitting for the relative shift in the frequency domain. The time-domain algorithm described by Hotan, Bailes & Ord (2005a) was used for PSR J1045–4509 because it was better suited to the large duty cycle and low S/N of this pulsar. We applied standard pulsar timing techniques to the measured TOAs. Model fitting was performed using the TEMPO[1] software package.

### 8.4.1 Galactic Electron Density Models

Broad-band pulses are subject to dispersive smearing by electrons in the interstellar medium. Lyne, Manchester & Taylor (1985) describe a model of galactic electron density that was used to estimate pulsar distances based on the observed DM. This model consisted of two nearly isotropic disk components and a contribution due to the Gum Nebula. It predicted that electron densities in the plane of the galaxy should in most cases be greater than $0.04\,\mathrm{cm}^{-3}$. Taylor & Cordes (1993) published a new multi-component model of the galactic electron density that included spiral arms and had a characteristic disk density approximately half that used by Lyne, Manchester & Taylor (1985). More recently, Cordes & Lazio used various independent distance estimates to produce an updated model known as NE2001 (Cordes & Lazio 2002; Cordes & Lazio 2003).

The original Taylor & Cordes (1993) electron density model is still widely used to estimate pulsar distances. Due to a limited number of free parameters and the patchy nature of the ISM, it has large uncertainties along some lines of sight. In this chapter we use the more recent NE2001 electron density model to estimate the distance to each of our sources. Although this model has not been published in a refereed journal, a comprehensive description and FORTRAN implementation can be found at the following URL: http://rsd-www.nrl.navy.mil/7213/lazio/ne_model/

[1] http://www.atnf.csiro.au/research/pulsar/tempo

### 8.4.2 Binary Motion

Keplerian binary motion adds 5 basic parameters to the timing model. These are the orbital period $P_{\rm b}$, the eccentricity of the orbital ellipse $e$ and the length of its semi-major axis projected along the line of sight $x$, the longitude of periastron $\omega$ and an epoch $T_0$ when the pulsar is at periastron. Damour & Deruelle (1985) and Damour & Deruelle (1986) describe a model that includes several post-Keplerian corrections, defined in a theory-independent manner that allows observers to test the predictions of any post-Newtonian gravitational theory. We use this model (henceforth known as "DD") and a similar model known as "ELL1" (Lange et al. 2001), that parameterises the orbit using two Laplace-Lagrange parameters $\kappa$ & $\eta$, and the time of ascending node $T_{\rm asc}$. ELL1 is used when the product of the eccentricity and semi-major axis of the system is very small (and the periastron not well defined).

## 8.5 Results

In this section, we summarise the timing of our selected millisecond pulsars, beginning with the isolated sources. In all tables, the error in the last significant figure is given in parentheses. These errors represent twice the formal $1\sigma$ uncertainty returned by TEMPO. The minimum profile width recorded in each table is defined as the time between half-power points (also known as the 50% width) if there is a single, central component in the mean pulse profile. When the profile consists of multiple distinct components, a visual estimate of the width is made, using the narrowest significant feature in any frequency band. The quoted $\dot{P}$ or $\dot{P}_{\rm b}$ distance limits (see Section 8.6) represent upper bounds, computed using the value of the relevant parameter plus twice the formal $1\sigma$ uncertainty returned by TEMPO. $\dot{P}_{\rm G}$ represents the component of the observed period derivative that is due to relative acceleration in the galactic potential (Paczyński 1990). $\dot{P}_{\rm S}$ is the component of the observed period derivative that we estimate is due to the Shklovskii effect (Shklovskii 1970), which depends on the transverse velocity of the source. $\dot{P}_{\rm I}$ is the estimated intrinsic period derivative of the pulsar, taking $\dot{P}_{\rm G}$ and $\dot{P}_{\rm S}$ into account (see Section 8.6). The S/N threshold is the minimum S/N of any profile included in the timing model fit.

### 8.5.1 Isolated Millisecond Pulsars

Fig. 8.1 shows standard template profiles for PSR J0711–6830, PSR J1024–0719, PSR J1744–1134, PSR B1937+21 and PSR J2124–3358. Table 8.2 shows the corresponding timing model parameters and several derived quantities.

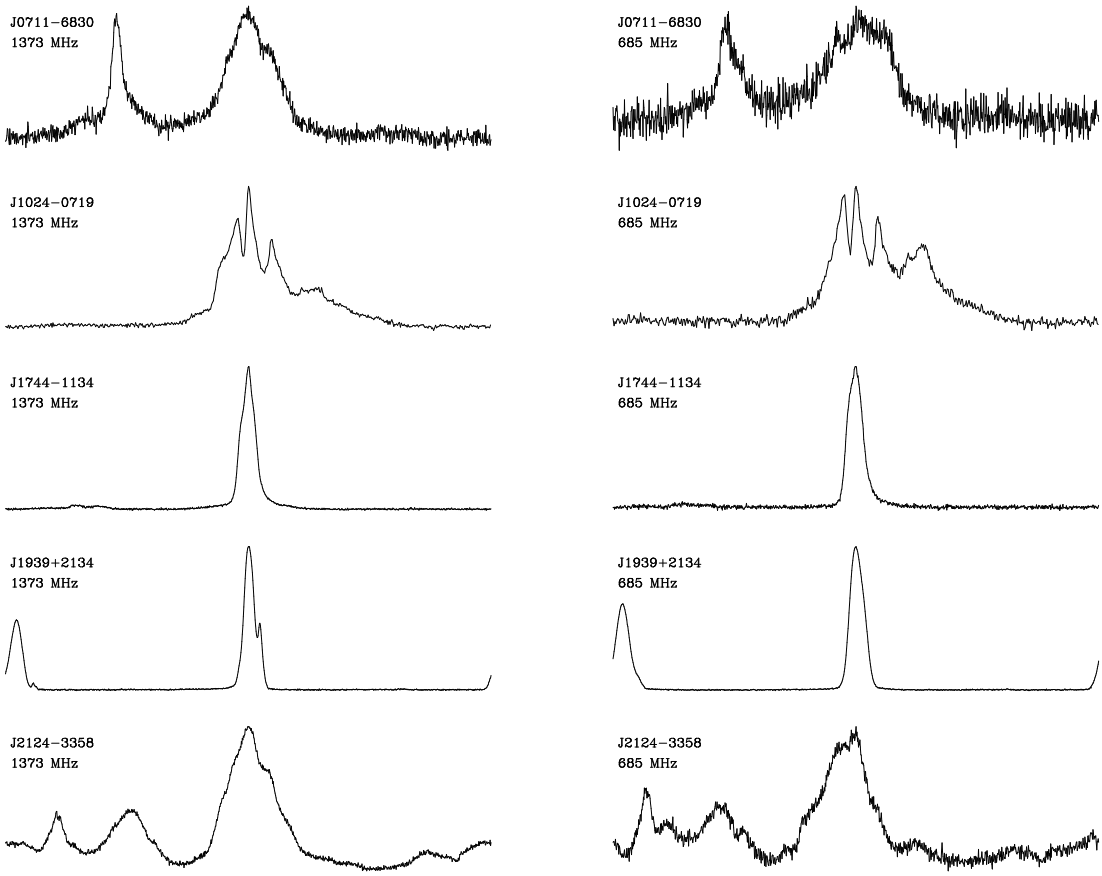


**Figure 8.1: Standard template Stokes I profiles for PSR J0711–6830 (first row), PSR J1024–0719 (second row), PSR J1744–1134 (third row), PSR B1937+21 (fourth row) and PSR J2124–3358 (fifth row), at centre frequencies of 1373 (left) and 685 (right) MHz. The vertical axis represents arbitrarily scaled amplitude and the horizontal axis represents one unit of pulse phase per panel.**

| **Data Reduction Parameters** | **J0711–6830** | **J1024–0719** | **J1744–1134** | **B1937+21** | **J2124–3358** |
|---|---|---|---|---|---|
| MJD of First TOA | 52619 | 52622 | 52622 | 52623 | 52649 |
| MJD of Last TOA | 53484 | 53484 | 53483 | 53484 | 53485 |
| Total Time Span (days) | 865 | 862 | 861 | 861 | 836 |
| Total Number of TOAS | 44 | 92 | 611 | 231 | 140 |
| Integration Time (min) | 60 | 60 | 5 | 5 | 60 |
| S/N Minimum | 10 | 20 | 30 | 30 | 20 |
| RMS Residual ($\mu$s) | 4.7 | 1.1 | 0.9 | 0.2 | 2.7 |
| Reduced $\chi^2$ | 5.9 | 1.1 | 1.9 | 2.9 | 1.4 |
| **Measured Parameters** | | | | | |
| Right Ascension ($\alpha$) | 07:11:54.2007 (6) | 10:24:38.68845 (2) | 17:44:29.401230 (6) | 19:39:38.561308 (1) | 21:24:43.85350 (4) |
| Declination ($\delta$) | –68:30:47.473 (2) | –07:19:19.1702 (6) | –11:34:54.6474 (6) | 21:34:59.12943 (4) | –33:58:44.668 (1) |
| Proper Motion in $\alpha$ (mas yr$^{-1}$) | –11 (4) | –34.9 (4) | 19.6 (4) | –0.41 (4) | –14.4 (8) |
| Proper Motion in $\delta$ (mas yr$^{-1}$) | 19 (2) | –47 (1) | –7 (1) | –1.5 (1) | –50 (2) |
| Period (ms) | 5.490968419123 (1) | 5.16220455707167 (8) | 4.07454587724659 (8) | 1.55780651922290 (4) | 4.9311148642820 (2) |
| Period Derivative ($\times 10^{-20}$) | 1.498 (8) | 1.8531 (6) | 0.8925 (6) | 10.51160 (5) | 2.059 (2) |
| Period 2$^{\rm nd}$ Derivative ($\times 10^{-31}$) | – | – | – | –3.4 (5) | – |
| Parallax (mas) | – | 1.9 (8) | 2.1 (4) | – | 4 (2) |
| Period Epoch (MJD) | 53000.0 | 53000.0 | 53184.0 | 53050.0 | 53174.0 |
| Dispersion Measure (cm$^{-3}$pc) | 18.4066 (6) | 6.48520 (8) | 3.13908 (4) | 71.0226 (2) | 4.5956 (2) |
| DM 1$^{\rm st}$ Derivative | – | – | – | –0.0012 (4) | – |
| Minimum Component Width ($\mu$s) | 110 | 50 | 140 | 30 | 730 |
| Minimum Duty Cycle (%) | 2.0 | 1.0 | 3.4 | 1.9 | 14.8 |
| Mean Profile S/N @ 1341 MHz | 19.7 | 43.0 | 39.4 | 48.3 | 55.9 |
| **Derived Parameters** | | | | | |
| Composite PM (mas yr$^{-1}$) | 22 (3) | 59 (1) | 20.9 (6) | 1.5 (1) | 53 (2) |
| Distance (Dispersion) (pc) | 860 | 390 | 410 | 3560 | 270 |
| Distance (Parallax) (pc) | – | $520^{+390}_{-150}$ | $470^{+120}_{-70}$ | – | $250^{+250}_{-80}$ |
| Distance ($\dot{P}$ Limit) (pc) | 2320 | 430 | 2060 | 12000 | 610 |
| $\dot{P}_{\rm G}$ ($\times 10^{-20}$) | –0.07[2] | –0.06 (1) | 0.010 (2) | –0.04[2] | –0.048 (5) |
| $\dot{P}_{\rm S}$ ($\times 10^{-20}$) | 0.56[2] | 2.2 +(18) –(6) | 0.20 +(5) –(3) | 0.003[2] | 0.8 +(8) –(2) |
| $\dot{P}_{\rm I}$ ($\times 10^{-20}$) | 1.01[2] | ¡ 0.3 | 0.69 +(2) –(6) | 10.55[2] | 1.3 +(2) –(8) |
| Characteristic age (Gyr) | 8.6[2] | ¿ 27 | 9.3 +(9) –(2) | 0.23[2] | 6 +(9) –(1) |
| Surface magnetic field ($\times 10^8$ G) | 6.7[2] | ¡ 3.6 | 4.80 +(10) –(1) | 11.62 | 7.2 +(6) –(27) |

**Table 8.2: Timing model parameters for the isolated millisecond pulsars. The uncertainty in the last significant figure represents twice the formal $1\sigma$ uncertainty returned by TEMPO and is given in parentheses after the value. Values marked with ([1]) were adopted from Toscano et al. (1999b). The error in any value marked with ([2]) is dominated by the ill-determined uncertainty in the pulsar distance as derived from the NE2001 galactic electron density model and these values should be taken as a guide only.**

### PSR J0711–6830

This pulsar suffers from a low density of successful observations and we were forced to include points with S/N as low as 10 to fit for position and spin-down parameters. Day-of-year coverage is patchy and we therefore do not attempt to fit for proper motion or parallax, choosing instead to adopt the proper motion published by Toscano et al. (1999b). The RMS residual we obtain is twice as large as the Toscano et al. (1999b) value, indicating that for this pulsar, additional bandwidth is more beneficial than coherent dedispersion. Unless the S/N can be increased, this pulsar is not a suitable timing array candidate.

### PSR J1024–0719

We obtain 92 individual, 1 hr integrated profiles with S/N in excess of 20 for PSR J1024–0719, largely due to favourable scintillation. The timing model requires both proper motion and parallax (which we detect for the first time), yielding an RMS residual of approximately 1 $\mu$s. Our proper motion is very different to that published by Toscano et al. (1999b). This pulsar's complicated, multi-component mean profile contains sharp $\sim$50 $\mu$s wide features that allow TOAs to be precisely determined.

### PSR J1744–1134

CPSR2 observations of PSR J1744–1134 have produced 611 individual 5 min integrations with S/N in excess of 30 and the RMS timing residual is less than 1 $\mu$s. The pulse period is short and the mean flux density exceeds 5 mJy (Ord et al. 2004), allowing precise timing even with short integrations. Toscano et al. (1999a) obtained an RMS residual of 470 ns using 24 min integrations. Scaling this by the square root of the integration time leads us to expect $\sim$1 $\mu$s RMS with 5 min integrations. Despite having 1/4 the bandwidth of Toscano et al. (1999a) in each independent band, CPSR2 TOAs have an RMS timing residual of 900 ns. Our density of observations is low in the first $\sim$400 days and we detect apparent long-period timing noise with an amplitude of $\sim$2 $\mu$s, which may jeopardise this pulsar's contribution to a timing array.

### PSR B1937+21

This pulsar is plagued by timing noise and dispersion measure variations. Kaspi, Taylor & Ryba (1994) observe systematic drifts in the timing residuals with amplitudes of a few $\mu$s. Using CPSR2, we achieve an RMS timing residual of 211 ns with only 5 min integrations. Integrating to 1 hr reduces the RMS residual to 141 ns, which is significantly worse than we would expect if no systematic errors were present in the residuals. This is still one of the smallest RMS residuals ever measured, but we appear to be limited by DM variations and intrinsic rotational instabilities. If the

second frequency derivative is left out of the model, the residuals exhibit a clear systematic drift over a period of three years, with a maximum amplitude of $\sim$1 $\mu$s. The larger drifts observed by Kaspi, Taylor & Ryba (1994) may be a result of their longer timing baseline, or simply the intrinsic randomness of the mechanism responsible for this timing noise. We must also fit for a single dispersion measure derivative, which provides somewhat poor compensation for the quasi-random DM variations detected by Kaspi, Taylor & Ryba (1994) and Ilyasov et al. (2005). Cordes et al. (1990) argue that variable interstellar scattering may also contribute to long-term timing noise in this pulsar.

#### PSR J2124–3358

This is one of the few pulsars for which emission is visible over nearly the entire rotation period. Although this has important implications for the width of the pulsar's beam, the lack of any sharp features in the mean profile is detrimental to timing precision. The mean flux density at 1400 MHz is only $\sim$2.5 mJy and this energy is distributed over nearly the entire pulse period. Lengthy integrations are therefore necessary to achieve high S/N. We obtain 140 individual 1 hr integrations with S/N in excess of 20. Fitting these to a timing model yields an RMS residual of 2.7 $\mu$s, which is sufficient to measure proper motion and parallax, given the proximity of this source.

#### Comments

We have achieved small timing residuals for PSR J1024–0719 and PSR J2124–3358, allowing us to measure parallax and more precise proper motions than ever before (see Section 8.6). Toscano et al. (1999b) benefit from additional bandwidth when timing PSR J0711–6830 and PSR J1744–1134, but the exceptional stability of CPSR2 defeats the expected factor of 2 decrease in S/N, at least in the case of PSR J1744–1134. PSR B1937+21 has very high timing precision, but only on short timescales. More importantly, the reduced $\chi^2$ we obtain from 3 of these timing model fits is close to unity, demonstrating our relative freedom from systematic errors.

### 8.5.2 Binary Millisecond Pulsars

We divide the observed binary MSPs into two groups of 5 on the basis of timing precision and spin period. Timing model parameters for the most precise timing sources are shown in Table 8.3. The remaining sources are displayed in Table 8.4. Figs. 8.2, 8.3 and 8.4 show standard template profiles for the binary MSPs, at a variety of wavelengths.

| **Data Reduction Parameters** | **J0437–4715** | **J0613–0200** | **J1713+0747** | **B1855+09** | **J1909–3744** |
|---|---|---|---|---|---|
| MJD of First TOA | 52618 | 52618 | 52622 | 52651 | 52619 |
| MJD of Last TOA | 53485 | 53485 | 53485 | 53483 | 53485 |
| Total Time Span (days) | 867 | 867 | 863 | 832 | 866 |
| Total Number of TOAS | 6517 | 1258 | 635 | 152 | 2859 |
| Integration Time (min) | 5 | 5 | 5 | 5 | 5 |
| S/N Minimum | 50 | 10 | 30 | 10 | 20 |
| RMS Residual ($\mu$s) | 0.4 | 2.1 | 0.2 | 3.3 | 0.3 |
| Reduced $\chi^2$ | 2.4 | 1.5 | 1.1 | 1.6 | 1.2 |
| **Measured Parameters** | | | | | |
| Right Ascension ($\alpha$) | 04:37:15.8478747 (8) | 06:13:43.97509 (2) | 17:13:49.531524 (2) | 18:57:36.39161 (4) | 19:09:47.4378773 (8) |
| Declination ($\delta$) | –47:15:08.822580 (8) | –02:00:47.1712 (4) | +07:47:37.51449 (6) | +09:43:17.2322 (8) | –37:44:14.32397 (6) |
| Proper Motion in $\alpha$ (mas yr$^{-1}$) | 121.31 (1) | 2.1 (2) | 4.97 (6) | – | –9.49 (2) |
| Proper Motion in $\delta$ (mas yr$^{-1}$) | –71.53 (1) | –10.5 (6) | –3.7 (1) | – | –36.06 (6) |
| Period (ms) | 5.757451840111090 (4) | 3.06184403897281 (6) | 4.57013652589336 (1) | 5.3621004627635 (4) | 2.947108021715178 (2) |
| Period Derivative ($\times 10^{-20}$) | 5.72933 (4) | 0.9591 (4) | 0.8533 (1) | 1.771 (2) | 1.40241 (2) |
| Parallax (mas) | 6.3 (2) | 2.1 (6) | 1.1 (1) | – | 0.88 (4) |
| Period Epoch (MJD) | 53019.0 | 53012.0 | 53100.0 | 53186.0 | 53055.0 |
| Dispersion Measure (cm$^{-3}$pc) | 2.6445 (1) | 38.77919 (6) | 15.9916 (2) | 13.2952 (6) | 10.3940 (1) |
| Minimum Component Width ($\mu$s) | 85 | 46 | 91 | 160 | 42 |
| Minimum Duty Cycle (%) | 1.5 | 1.5 | 2.0 | 3.0 | 1.4 |
| Mean Profile S/N @ 1341 MHz | 262.0 | 18.6 | 144.9 | 33.7 | 28.2 |
| Orbital Period (days) | 5.7410423638 (1) | 1.1985125564 (2) | 67.825129880 (6) | 12.32717115 (4) | 1.53344945052 (2) |
| $\dot{P}_{\rm b}$ ($\times 10^{-12}$) | – | – | – | – | 0.6 (1) |
| $x \equiv a\sin(i)$ (lt-s) | 3.36670403 (2) | 1.0914443 (2) | 32.34242156 (6) | 9.230788 (1) | 1.89799111 (2) |
| $\dot{x}$ ($\times 10^{-15}$) | 79.7 (8) | – | –5 (4) | – | – |
| Eccentricity ($\times 10^{-5}$) | 1.918 (1) | 0.55 (4) | 7.4959 (6) | 2.09 (2) | 0.013 (1) |
| Epoch of Periastron (MJD) | 52985.830642234 (400000) | 53019.27 (2) | 53082.782202417 (400000) | 53200.36 (4) | 53114.70 (4) |
| Angle of Periastron (deg) | 1.294455 (20000) | 47 (3) | 176.200576 (2000) | 275 (1) | 177 (8) |
| $\kappa$ ($\times 10^{-8}$) | – | 401 (40) | – | –2083 (20) | 0.8 (20) |
| $\eta$ ($\times 10^{-8}$) | – | 370 (30) | – | 193 (40) | –13 (2) |
| Time of Asc. Node (MJD) | – | 53019.11370768 (4) | – | 53190.9310259 (4) | 53113.950587418 (2) |
| Inclination Angle (deg) | 42.6 +(1) –(3) | – | 77 (2) | $60 < i < 88$ | 86.6 (2) |
| Companion Mass ($M_\odot$) | 0.21 (2) | – | 0.25 (3) | 0.21 (7) | 0.207 (2) |
| $\Omega$ (deg) | 237 (4) | – | 81 (9) | – | – |
| **Derived Parameters** | | | | | |
| Composite PM (mas yr$^{-1}$) | 140.83 (1) | 10.7 (6) | 6.21 (8) | – | 37.26 (6) |
| Distance (Dispersion) (pc) | 140 | 1705 | 890 | 1170 | 460 |
| Distance (Parallax) (pc) | 159 (5) | $480^{+190}_{-110}$ | $910^{+90}_{-70}$ | – | 1140 (50) |
| Dist. ($\dot{P}$) Limit (pc) | 210 | 11160 | 19800 | – | 1410 |
| Dist. ($\dot{P}_{\rm b}$) Limit (pc) | – | – | – | – | 1570 |
| $\dot{P}_{\rm G}$ ($\times 10^{-20}$) | –0.035 (5) | 0.003 (2) | –0.029 (1) | – | –0.001 (4) |
| $\dot{P}_{\rm S}$ ($\times 10^{-20}$) | 4.4 +(1) –(2) | 0.04 +(2) –(1) | 0.039 +(4) –(2) | – | 1.13 (5) |
| $\dot{P}_{\rm I}$ ($\times 10^{-20}$) | 1.4 (2) | 0.92 +(1) –(3) | 0.843 +(3) –(5) | – | 0.27 (5) |
| Characteristic age (Gyr) | 6.5 (10) | 5.3 (1) | 8.60 (4) | – | 17 (3) |
| Surface magnetic field ($\times 10^8$ G) | 8.1 (6) | 4.81 +(2) –(8) | 5.62 (1) | – | 2.6 (2) |
| Pulsar mass ($M_\odot$) | 1.3 (2) | – | 1.1 (2) | $0.4 < m_{\rm p} < 1.7$ | 1.47 +(3) –(2) |

**Table 8.3: DD timing model parameters for the binary millisecond pulsars PSR J0437–4715, PSR J0613–0200, PSR J1713+0747 and PSR B1855+09. The ELL1 model was used to time PSR J1909–3744. For PSR J0437–4715 and PSR J1713+0747, the DD model was extended to include annual orbital parallax. Where given, companion masses were estimated from Shapiro delay and pulsar masses were derived from the mass function for the system, using our knowledge of the inclination angle and companion mass.**

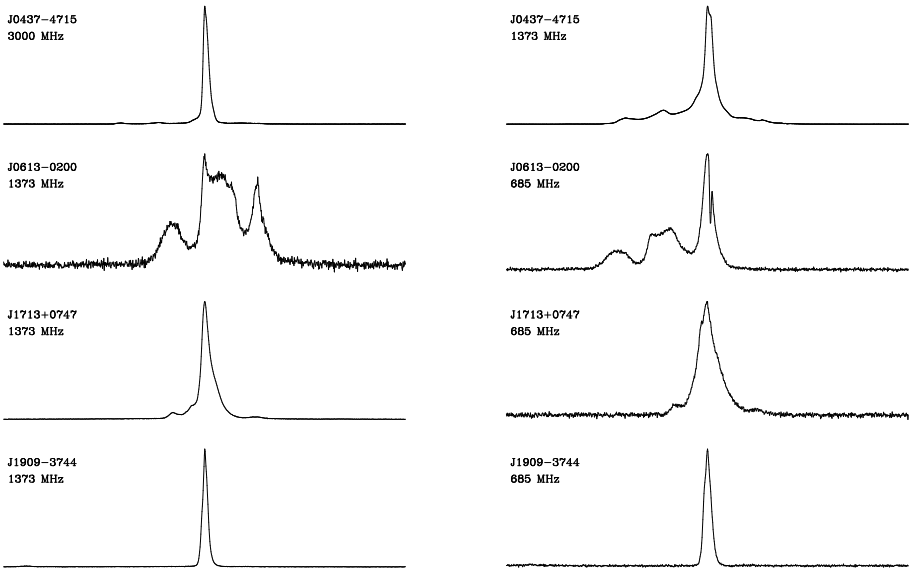


**Figure 8.2: The first row shows two standard template profiles for PSR J0437–4715, at centre frequencies of 3000 MHz (left) and 1373 MHz (right). The 1373 MHz profile represent the polarimetric invariant interval. The second row shows Stokes I templates for PSR J0613–0200, the third PSR J1713+0747 and the fourth PSR J1909–3744, all at centre frequencies of 1373 (left) and 685 (right) MHz.**

### PSR J0437–4715

The best previous timing solution for PSR J0437–4715 (van Straten et al. 2001) had an RMS residual of 130 ns at a wavelength of approximately 20 cm. However, this RMS residual was obtained from 1 hr integrations and several different instruments linked with 4 arbitrary phase jumps, which could have helped to absorb timing noise. The authors detected the signature of annual orbital parallax, with an expected peak-to-peak amplitude close to our current RMS residual. In addition, van Straten et al. (2001) time the polarimetric "invariant" profile in the hope that it will reduce the impact of observing at different parallactic angles. The invariant interval is defined in terms of the Stokes parameters $I$, $Q$, $U$ and $V$ to be:

$$S_{\rm inv} = I^2 - Q^2 - U^2 - V^2 \tag{8.1}$$

Britton (2000) demonstrates that this quantity is invariant to most of the transformations an observing system can inadvertently apply to a polarised signal. $S_{\rm inv}$ is of little use when the pulsar is highly polarised (as it removes most of the flux) but PSR J0437–4715 has a large component of un-polarised emission. We time the invariant profile of PSR J0437–4715 in this chapter. All other pulsars were timed using Stokes I.

We now have a continuous, unbroken series of CPSR2 observations spanning nearly 3 years. Only the receiver package has changed during this time, yet we

encounter a number of problems. At 50 cm, there is significant radio frequency interference both within and just outside the observing band. The profile of PSR J0437–4715 broadens at lower frequencies, reducing timing precision. In this chapter, we therefore consider only 10 cm and 20 cm data, which is relatively interference free.

Observations at 20 cm suffer from a systematic error that is most likely induced by the large intrinsic flux density of the pulsar itself, 2-bit quantisation errors, or imperfect filters. van Straten (2003) notes that a significant fraction of the pulses emitted by PSR J0437–4715 cannot be adequately characterised using 2-bit sampling and CPSR2 is therefore likely to suffer from quantisation errors. Qualitatively, we note that within a 64 MHz-wide band near a wavelength of 20 cm, the time delay between successive frequency channels is not representative of the DM of the pulsar. There is an additional slope that can extend over several $\mu$s and is perhaps due to aliasing or scattered power (Jenet & Anderson 1998). To combat this problem, we discard all but the central 5 MHz of the band. PSR J0437–4715 is so bright at 20 cm that the associated S/N penalty is not significant. Our data quality varies with the receiver used; with the multibeam system we note that the timing residual for any observation depends on the parallactic angle. This dependence is not completely removed by timing the invariant interval, though it is significantly reduced. The residual dependence may arise from imperfect estimation of the off-pulse signal, quantisation errors or small non-linear polarisation transformations introduced by the receiver system itself. The new arrival time fitting method described by van Straten (2006) may overcome these problems.

During the past ∼400 days, regular observations of PSR J0437–4715 were also made at a centre frequency of 3000 MHz. This frequency band offers relative freedom from the systematic errors encountered at 20 cm (due to a smaller intrinsic flux density), a sharper pulse profile, a cleaner observing band and better polarimetric purity. Using 1 hr integrations at a centre frequency of 3000 MHz, we obtain an RMS residual close to 200 ns. Future observations will be made primarily at 10 cm, but in this chapter we rely on 20 cm observations to extend our timing history and thus constrain the astrometric parameters of the pulsar.

### PSR J0613–0200

This is one of the few pulsars that can be timed to higher precision at lower frequencies (despite interference), where the mean profile evolves a narrow spike. The RMS residual quoted in Table 8.3 is comparatively large due to the fact that we must include 20 cm TOAs to extend our timing baseline in order to measure proper motion and parallax (which we detect for the first time). We also maintain 5 min time resolution to increase the number of TOAs in the fit. After integrating to 1 hr and removing all 20 cm data, fitting only for position, $P$ and $\dot{P}$ and all binary parameters, we obtain an RMS residual of 620 ns.

### PSR J1713+0747

This is one of the most precisely timed pulsars ever observed. With 5 min integrations and 64 MHz-wide bands centred on 1341 and 1405 MHz, we obtain a very small RMS residual of 250 ns. Using 1 hr integrations this drops to 125 ns, one of the smallest RMS residuals ever measured. Even this value is larger than we might expect based on the scaling of the integration times, indicating that systematic errors may be present at a low level. We do not include any 50 cm data because the profile broadens at lower frequencies, reducing the precision of arrival time measurements. Our data reveals a wealth of information about the binary system. We use this very precise timing to detect Shapiro delay and annual orbital parallax (see Section 8.7).

### PSR B1855+09

We include PSR B1855+09 in this group even though its RMS residual is above 3 $\mu$s. Given that observations of this source commenced only recently, we must use 5 min integrations to maintain a large number of timing points. Lengthier integrations and more intensive observing should reduce the RMS residual significantly, given that Kaspi, Taylor & Ryba (1994) have already demonstrated sub-$\mu$s timing for this pulsar. Despite the small number of observations, we have reasonable orbital phase coverage and use Shapiro delay to constrain the orbital inclination and companion mass, though we cannot yet confirm previous detections of parallax or proper motion.

### PSR J1909–3744

PSR J1909–3744 is one of the best timing sources ever discovered. This pulsar's orbit is extremely circular and we therefore adopt the ELL1 timing model described by Lange et al. (2001). Again, we discard all 50 cm data to improve the RMS residual, which is 300 ns with 5 min integrations and 150 ns with 1 hr integrations. Our line of sight is almost edge-on to the orbital plane, allowing the most precise measurement of Shapiro delay ever made. The observed period and binary period derivatives are dominated by kinematic effects, providing tight limits on the distance to the pulsar. Because our time span is slightly longer than that of Jacoby et al. (2005), our values for proper motion and orbital period derivative are more precise.

### PSR J1022+1001

This pulsar was once thought to have an unstable mean profile that corrupted its timing at a level of 20 $\mu$s (Kramer et al. 1999). The timing residual presented in Table 8.4 (1.5 $\mu$s RMS) is consistent with the work of Hotan, Bailes & Ord (2004); we find that this pulsar times very well given its relatively long, 16 ms period. We obtain stronger limits on the proper motion, parallax and Shapiro delay in this chapter and demonstrate its continued timing stability.

| Data Reduction Parameters | J1022+1001 | J1045–4509 | J1603–7202 | J2129–5721 | J2145–0750 |
|---|---|---|---|---|---|
| MJD of First TOA | 52649 | 52619 | 52622 | 52618 | 52618 |
| MJD of Last TOA | 53485 | 53468 | 53483 | 53477 | 53485 |
| Total Time Span (days) | 836 | 849 | 861 | 859 | 867 |
| Total Number of TOAS | 171 | 120 | 128 | 158 | 105 |
| Integration Time (min) | 60 | 60 | 60 | 5 | 60 |
| S/N Minimum | 50 | 30 | 10 | 10 | 45 |
| RMS Residual ($\mu$s) | 1.5 | 5.9 | 1.6 | 4.1 | 1.3 |
| Reduced $\chi^2$ | 3.0 | 1.9 | 1.2 | 2.7 | 4.1 |
| **Measured Parameters** | | | | | |
| Ecliptic Longitude ($\lambda$) (deg) | 153.86589022 (6) | – | – | – | 326.0246384 (1) |
| Ecliptic Latitude ($\beta$) (deg) | –0.06400 (6) | – | – | – | 5.313066 (2) |
| Proper Motion in $\lambda$ (mas yr$^{-1}$) | –15.6 (4) | – | – | – | –11.9 (8) |
| Proper Motion in $\beta$ (mas yr$^{-1}$) | – | – | – | – | –7 (8) |
| Right Ascension ($\alpha$) | 10:22:57.998 (4) | 10:45:50.18941 (8) | 16:03:35.6798 (1) | 21:29:22.7618 (2) | 21:45:50.4640 (2) |
| Declination ($\delta$) | 10:01:52.5 (1) | –45:09:54.143 (1) | –72:02:32.6984 (8) | –57:21:14.169 (6) | –07:50:18.439 (6) |
| Proper Motion in $\alpha$ (mas yr$^{-1}$) | –16.8 (4) | –7 (1) | –2.8 (6) | 7[2] | –9 (2) |
| Proper Motion in $\delta$ (mas yr$^{-1}$) | – | 8 (2) | –6 (1) | –4[2] | –10 (6) |
| Period (ms) | 16.4529296938470 (2) | 7.4742241103250 (6) | 14.841952018961 (1) | 3.726348423287 (1) | 16.0524236655171 (6) |
| Period Derivative ($\times 10^{-20}$) | 4.334 (4) | 1.755 (6) | 1.574 (8) | 2.09 (1) | 2.981 (8) |
| Parallax (mas) | 2.5 (8) | – | – | – | – |
| Period Epoch (MJD) | 53100.0 | 53050.0 | 53024.0 | 52944.0 | 53070.0 |
| Dispersion Measure (cm$^{-3}$pc) | 10.2521 (1) | 58.1662 (4) | 38.0471 (1) | 31.86 (1) | 9.0015 (2) |
| Minimum Component Width ($\mu$s) | 165 | 373 | 223 | 298 | 241 |
| Minimum Duty Cycle (%) | 1.0 | 5.0 | 1.5 | 8.0 | 1.5 |
| Mean Profile S/N @ 1341 MHz | 199.0 | 55.6 | 103.8 | 19.5 | 160.5 |
| Orbital Period (days) | 7.805130162 (1) | 4.083529191 (6) | 6.308629577 (4) | 6.6254931 (1) | 6.838902508 (2) |
| $x \equiv a\sin(i)$ (lt-s) | 16.7654154 (4) | 3.015131 (2) | 6.8806611 (4) | 3.500570 (6) | 10.1641068 (8) |
| $\dot{x}$ ($\times 10^{-14}$) | 4 (2) | – | – | – | – |
| Eccentricity ($\times 10^{-5}$) | 9.728 (6) | 2.4 (1) | 0.93 (1) | 1.9 (6) | 1.93 (1) |
| Epoch of Periastron (MJD) | 53095.5905 (6) | 53048.98 (4) | 53312.29 (2) | 53195.6 (1) | 53056.114 (8) |
| Angle of Periastron (deg) | 97.75 (2) | 243 (2) | 169 (1) | 184 (6) | 200.9 (4) |
| $\kappa$ ($\times 10^{-6}$) | 96.40 (6) | –21.51 (100) | 1.8 (2) | –1.42 (400) | –6.9 (1) |
| $\eta$ ($\times 10^{-6}$) | –13.14 (6) | –10.78 (100) | –9.1 (1) | –19.36 (2000) | –18.1 (1) |
| Time of Asc. Node (MJD) | 53093.47118169 (4) | 53046.2232147 (4) | 53309.3306255 (1) | 53192.22261 (2) | 53052.29703625 (8) |
| Companion Mass ($M_\odot$) | $0.8 < m_c < 1.3$[1] | – | 0.14 (1)[1] | – | $> 0.5$[1] |
| Inclination Angle (deg) | $41 < i < 53$[1] | – | $59 < i < 68$[1] | – | $< 45$[1] |
| **Derived Parameters** | | | | | |
| Composite PM (mas yr$^{-1}$) | – | 11 (1) | 7 (1) | 8[2] | 14 (6) |
| Distance (Dispersion) (pc) | 450 | 1960 | 1170 | 1360 | 570 |
| Distance (Parallax) (pc) | $400^{+190}_{-100}$ | – | – | – | – |
| Dist. ($\dot{P}$) Limit (pc) | – | 8700 | 8900 | 35500 | 4070 |
| $\dot{P}_G$ ($\times 10^{-20}$) | – | –0.13 | –0.08 | –0.07 | –0.23 |
| $\dot{P}_S$ ($\times 10^{-20}$) | – | 0.4 | 0.2 | 0.08 | 0.4 |
| $\dot{P}_I$ ($\times 10^{-20}$) | – | 1.49 | 1.45 | 2.09 | 2.81 |
| Characteristic age (Gyr) | – | 7.9 | 16.2 | 2.8 | 9.0 |
| Surface magnetic field ($\times 10^8$ G) | – | 9.6 | 13.3 | 8.0 | 19.2 |

**Table 8.4: DD timing model parameters for the binary millisecond pulsars PSR J1022+1001, PSR J1045–4509, PSR J2129–5721 and PSR J2145–0750. The ELL1 model was used to time PSR J1603–7202. Where coordinates are given in two reference frames, the best fit for each frame was obtained independently. Values marked with ([1]) assume a pulsar mass of 1.3 $M_\odot$. Values marked with ([2]) were adopted from Toscano et al. (1999b). The errors in $\dot{P}_G$ and all derived parameters following it in the table are dominated by the ill-determined uncertainty in the pulsar distance as derived from the NE2001 galactic electron density model and uncertainties are therefore not reported for these values, which should be taken as a guide only. For PSR J1022+1001, we have only measured the proper motion in one dimension and therefore cannot estimate $\dot{P}_G$, $\dot{P}_S$ or $\dot{P}_I$**

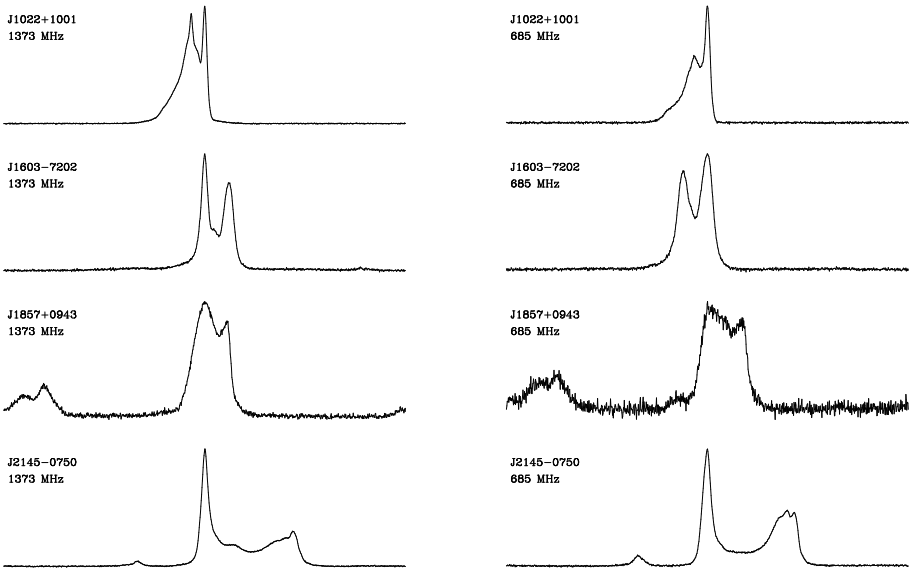


**Figure 8.3: Standard template Stokes I profiles for PSR J1022+1001 (first row), PSR J1603-7202 (second row), PSR B1855+09 (third row) and PSR J2145–0750 (fourth row) at centre frequencies of 1373 (left) and 685 (right) MHz.**

### PSR J1603–7202

PSR J1603–7202 has a narrow double-peaked profile that allows precise arrival time measurements. We obtain an RMS residual of the same order as PSR J1022+1001. With a DM of approximately $38\,\mathrm{cm}^{-3}\mathrm{pc}$, this pulsar is one of the more distant in our sample and we detect only a small proper motion and no parallax. This pulsar's companion has a mass of order $0.5\,\mathrm{M}_\odot$ but we do not detect any signature of Shapiro delay. We infer that the inclination angle of the system must be less than $\sim 45^o$, if the pulsar mass is assumed to be $1.3\,\mathrm{M}_\odot$.

### PSR J2145–0750

PSR J2145–0750 is another relatively long-period MSP that times very well ($1.3\,\mu$s RMS) due to a favourably narrow mean pulse feature. Although Loehmer et al. (2004) measure a parallax that is consistent with the DM distance of 500 pc, we do not detect any parallax signature in the timing residuals. Our $2\sigma$ limit on parallax (see Section 8.6) constrains the pulsar to be further than 1.1 kpc away.

### PSR J1045–4509

PSR J1045–4509 appears to exhibit timing noise with an amplitude of approximately 10 $\mu$s, which may explain the large RMS residual. It has a low-mass companion and we cannot place any interesting limits on the system from Shapiro delay.

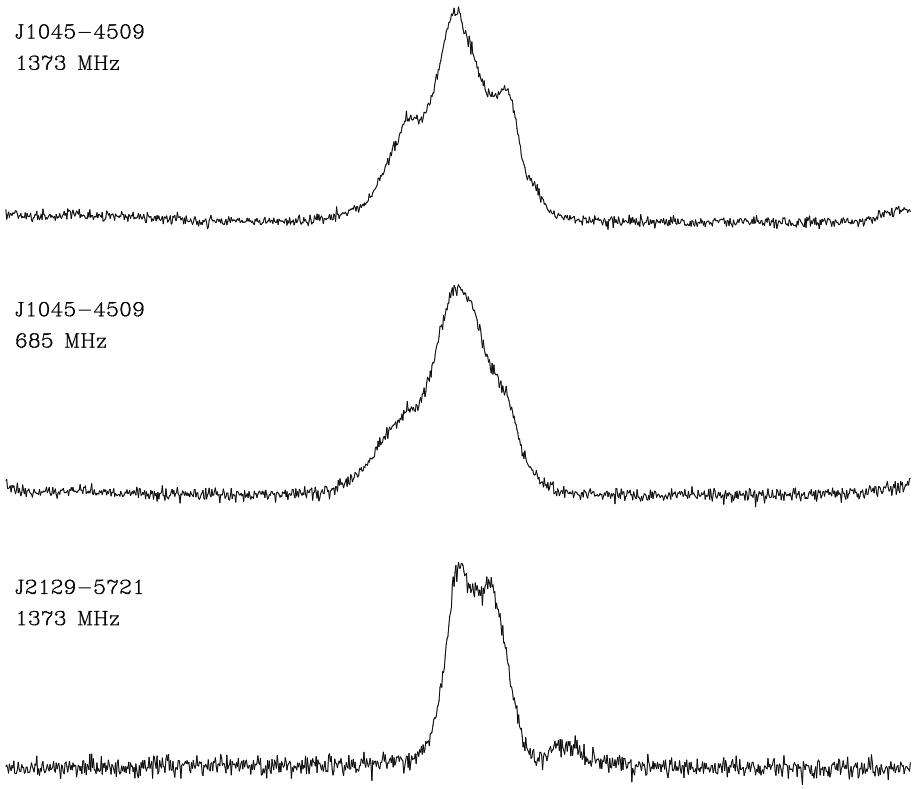


**Figure 8.4: Standard template Stokes I profiles for PSR J1045–4509 at a centre frequency of 1373 MHz (first row) and 685 MHz (second row), and for PSR J2129–5721 at a centre frequency of 1373 MHz (third row). The frequency evolution of PSR J1045–4509's mean profile may contribute to its relatively poor timing.**

### PSR J2129–5721

Observations of PSR J2129–5721 are limited and we are forced to maintain 5 min time resolution, which greatly reduces the S/N. This contributes to the comparatively large RMS residual; lengthy observations may yet lead to greater timing precision.

### Comments

Most of the MSPs in Table 8.4 have periods of order 3 times larger than those in Table 8.3 and their RMS residuals are higher by approximately the same order, demonstrating a proportionality between spin period and timing precision. In the next section, we test our newly-derived timing models for self-consistency by interpreting the fitted parameters in a physical context.

## 8.6 Proper Motions, Distances and Initial Spin Periods

Isolated millisecond pulsars can be modeled using relatively few free parameters. Aside from their intrinsic period $P$, spin-down rate $\dot{P}$ and DM, only position, proper motion and parallax must be accounted for. These last two often have small timing signatures and are only detectable if the pulsar is nearby. Detection of both allows the distance and transverse velocity of the pulsar to be uniquely determined, aiding studies of Galactic kinematics. The observed spin-down rate is also connected to the motion of the pulsar. Constant motion along the line of sight introduces a Doppler shift that changes the observed period by a fixed amount, but this is absorbed as a constant offset. If the velocity along the line of sight changes, the offset will shift, causing an apparent change in period. Thus, any acceleration along the line of sight will contribute to the observed spin-down rate. In most cases, the acceleration experienced by a pulsar (excluding binary motion, which is modeled separately) will be a combination of differential Galactic rotation (usually small), acceleration towards the plane of the Galaxy and perhaps acceleration due to nearby stars. If the pulsar happens to reside in a globular cluster, the gravitational potential of the cluster will dominate; sometimes to the point where the observed period derivative is negative.

Shklovskii (1970) pointed out that constant transverse motion can induce an apparent acceleration *along* the line of sight, contributing to the observed $\dot{P}$. The magnitude of this effect is

$$\frac{\dot{P}_{\rm s}}{P} = \frac{|\vec{v}|^2}{dc} = \frac{|\vec{\mu}|^2 d}{c}. \tag{8.2}$$

Here, $\dot{P}_{\rm s}$ is the Shklovskii period derivative (which adds to the intrinsic $\dot{P}$), $\vec{v}$ is the transverse velocity of the system, $d$ is the distance to the pulsar, $c$ is the speed of light, $P$ is the period and $\vec{\mu} = \vec{v}/d$ is the composite proper motion.

Knowledge of the composite proper motion and observed period derivative allows an upper limit to be placed on the distance to the pulsar by assuming that all the observed $\dot{P}$ is kinematic in origin. This provides a consistency check for distances determined using electron density models or parallax. It should be noted that Eq. 8.2 applies to any periodic signal whose point of origin is in motion and an equivalent expression can be written down for the observed binary period $P_{\rm b}$ (if any) and its derivative. Although the signature of $\dot{P}_{\rm b}$ is much more difficult to detect, it can lead to a better distance estimate (Bell & Bailes 1996).

Alternatively, if the distance and proper motion of the pulsar can be precisely determined by observation, we can calculate the kinematic contribution to the observed period derivative. Using models of the Galactic gravitational potential (Paczyński 1990), we can also estimate the component of the observed period derivative due to relative acceleration along the line of sight. By subtracting both these contributions

from the observed period derivative, we can determine the intrinsic period derivative, which can be used to estimate the characteristic spin-down age of the pulsar, assuming a constant rate of change and an initial period close to zero.

If the pulsar resides in a binary system, relative motion between the source and the observer slowly changes the line of sight to the orbital plane, introducing either a secular or periodic (depending on the cause of the motion) time dependence into several of the binary parameters. These effects are subtle and rarely observable, but offer the chance to constrain the parameters of the system in a manner completely independent of post-Keplerian gravitational theory.

Kopeikin (1996) derived expressions for the secular kinematic contribution to $x_{\rm obs} = x + \delta x$ and a similar effect upon the longitude of periastron $\omega_{\rm obs} = \omega + \delta\omega$.

$$\delta x = \frac{x}{\tan(i)}(-\mu_\alpha \sin(\Omega) + \mu_\delta \cos(\Omega))(t - t_0), \tag{8.3}$$

$$\delta\omega = \frac{1}{\sin(i)}(\mu_\alpha \cos(\Omega) + \mu_\delta \sin(\Omega))(t - t_0). \tag{8.4}$$

Here, $i$ is the inclination angle of the system, $\mu_\alpha$ is the component of proper motion in right ascension, $\mu_\delta$ is the component of proper motion in declination and $\Omega$ is the longitude of the ascending node, which defines the orientation of the orbit with respect to rotations about the line of sight. Taking the time derivative of Eq. 8.3 gives

$$\frac{\dot{x}}{x} = \cot(i)(-\mu_\alpha \sin(\Omega) + \mu_\delta \cos(\Omega)), \tag{8.5}$$

which can be re-written as

$$\frac{\dot{x}}{x} = |\vec{\mu}| \cot(i) \sin(\theta_\mu - \Omega). \tag{8.6}$$

Here, $|\vec{\mu}|$ is the magnitude of the proper motion vector and $\theta_\mu$ is the corresponding position angle on the sky. Often, $i$ and $\Omega$ are unknown, making it difficult to estimate the observable signature of $\dot{x}$. Knowledge of $i$ allows an upper limit to be placed on the absolute value of $\dot{x}$ by assuming $|\sin(\theta_\mu - \Omega)| = 1$. Alternatively, rearranging the equation allows detection of $\dot{x}$ to place a lower limit on the magnitude of the composite proper motion and an upper limit on the inclination angle. If the detection of Shapiro delay can be used to measure $\sin(i)$, detection of $\dot{x}$ can constrain $\Omega$.

In this section, we analyse the self-consistency of our astrometric measurements and compare them with other published results.

We measure the parallax of PSR J0613–0200 for the first time; the value of $2.1 \pm 0.6$ mas corresponds to a distance of only $\sim$500 pc, 3 times closer than the distance predicted by NE2001. Additional high-precision 50 cm observations should reduce the parallax uncertainty. The proper motion we measure in declination ($-10.5 \pm 0.6$ mas yr$^{-1}$) is not consistent with the value of $-7 \pm 1$ mas yr$^{-1}$ measured by Toscano et al. (1999b).

PSR J1022+1001 lies very near the ecliptic plane; timing experiments cannot strongly constrain the position or proper motion in a direction normal to the ecliptic, but the uncertainty is minimised by fitting for position in ecliptic coordinates. Our value for proper motion in ecliptic longitude agrees with the value of –17 ± 2 mas yr$^{-1}$ published by Kramer et al. (1999), but the proper motion in ecliptic latitude is not constrained. We have an estimate of the inclination angle from Shapiro delay measurements ($41^o < i < 53^o$) and also a measurement of $\dot{x}$. We can therefore use Eq. 8.6 to calculate a lower limit on the composite proper motion. We require $|\vec{\mu}| > 7$ mas yr$^{-1}$, which means that our model is consistent with $\mu_\beta$ being near zero. Because we only have a lower limit on the proper motion we cannot constrain the maximum $|\dot{x}|$. However, if we assume that $|\vec{\mu}| = 7$ mas yr$^{-1}$, Eq. 8.6 can be used to place an upper limit of $2\times10^{-14}$ lt-s s$^{-1}$ on $|\dot{x}|$. This is consistent with the value of $4 \pm 2\times10^{-14}$ lt-s s$^{-1}$ that we measure. It is therefore very likely that our measured $\dot{x}$ is purely kinematic and the composite proper motion is only slightly larger than 7 mas yr$^{-1}$. Sources in the ecliptic plane experience the largest signature of parallax; we measure a value of 2.5 ± 0.8 mas, compared to the Hotan, Bailes & Ord (2004) value of 3.3 ± 0.8 mas. The previous value was obtained by fixing the proper motion to the value published by Kramer et al. (1999), but the more recent value was obtained from a longer timing baseline which allowed us to fit for proper motion and parallax simultaneously. Future observations will further refine this value. The parallax distance of $400^{+190}_{-100}$ pc is consistent with the NE2001 distance of 450 pc.

The composite proper motion of PSR J1024–0719 published by Toscano et al. (1999b) (81 ± 4 mas yr$^{-1}$) exceeds our measurement of 59 ± 1 mas yr$^{-1}$. Because we have more than three times the number of TOAs and 1/4 the RMS timing residual, we feel confident that the new values for $\mu_\alpha$ and $\mu_\delta$ presented in Table 8.2 are correct. The NE2001 distance for this pulsar is 390 pc. The Shklovskii distance upper limit is very interesting because it implies that the intrinsic spin-down rate would be negative if the pulsar were more than 430 pc away, meaning that nearly all of the observed spin-down rate is kinematic in origin. We measure a parallax of 1.9 ± 0.8 mas that is consistent with the Shklovskii limit at the $1\sigma$ level and constrains the distance to the system. Using this distance measurement, we calculate an upper limit on the intrinsic spin-down rate of $0.3 \times 10^{-20}$ s s$^{-1}$, corresponding to a characteristic age greater than 27 Gyr. Riess et al. (2005) report that recent Hubble space telescope observations are consistent with a Hubble constant of 73 ± 5 km s$^{-1}$ Mpc$^{-1}$, corresponding to a Hubble time no larger than 15 Gyr. Assuming a constant spin-down rate, PSR J1024–0719 could only have lost 1.4 ms from its initial period in this time. Thus PSR J1024–0719 may have been born with an initial spin period of around 4 ms.

PSR J1024–0719 and two other sources in our list (J1744–1134 and J2124–3358) have been observed in X-rays (Becker & Trümper 1999). A deep observation with the VLT produced a possible detection of PSR J1024–0719 at optical wavelengths (Sutaria et al. 2003). Studies of high-energy emission benefit from accurate knowledge of the distance to the source, which these timing observations help to constrain.

We measure a small composite proper motion of $7 \pm 1$ mas yr$^{-1}$ for PSR J1603–7202, similar to the value of $8.7 \pm 0.7$ mas yr$^{-1}$ measured by Toscano et al. (1999b). This can be used, along with the NE2001 distance, to estimate the kinematic contribution to the observed value of $\dot{P}$, which we find to be $0.2\times10^{-20}$, or 14% of the observed value. We then compute a characteristic age upper limit of 16 Gyr, which is slightly older than the Hubble time, though smaller than the value of 21 Gyr determined by Toscano et al. (1999b). Assuming the measured intrinsic spin-down rate is constant, the pulsar must have been born with a period greater than 8 ms. NE2001 predicts that this pulsar is more than 1 kpc away and we cannot detect the signature of parallax.

The kinematic properties of PSR J1713+0747 were well measured by Splaver et al. (2005). Our parallax of $1.1 \pm 0.1$ mas is $2\sigma$ away from their measurement of $0.89 \pm 0.08$ mas and places the pulsar slightly closer to the Solar system, at a distance that corresponds very well to the NE2001 distance. The formal errors on our proper motion measurement are smaller but still consistent with the $1\sigma$ limit of Splaver et al. (2005).

Toscano et al. (1999a) measure the proper motion and parallax of PSR J1744–1134 and conclude that the pulsar is more than twice as far away as its small DM and the Taylor & Cordes (1993) galactic electron density model imply. The more recent NE2001 electron density model uses this pulsar as a calibrator, so there is no predictive power in the distance estimate. We measure a proper motion that is significantly different (about $3\sigma$ in both coordinates) to that previously published, and our parallax measurement of $2.1 \pm 0.4$ mas is smaller than the value of $2.8 \pm 0.3$ mas published by Toscano et al. (1999a), but within $2\sigma$. As our timing baseline increases, we will have greater confidence in these astrometric measurements.

We measure a parallax of $0.88 \pm 0.04$ mas for PSR J1909–3744, corresponding to a distance of $1140 \pm 50$ pc, more than twice that predicted by NE2001. We use Eq. 8.2 and its orbital period analogue to obtain upper limits that are close to the parallax distance, implying that a large fraction of the observed period and orbital period derivative are kinematic in origin. In fact, the characteristic age of this pulsar is slightly larger than the Hubble time, implying an initial spin period longer than 1.5 ms. Eq. 8.6 can be used in conjunction with our measurement of the inclination angle (from Shapiro delay) to place an upper limit of $3\times10^{-16}$ on $|\dot{x}|$. Fitting for this parameter yields a value of $1.5 \pm 1.4\times10^{-15}$, which is consistent with the predicted value given the large degree of uncertainty.

In summary, we find slight discrepancies in our measured proper motions when compared with previously published results, at the level of a few $\sigma$. This may indicate that the formal uncertainty returned by TEMPO is too small or that unmodelled systematic errors are present in our data. One possible explanation is our use of the DE405[2] Solar system ephemeris. Splaver et al. (2005) use very precisely determined pulse arrival times from PSR J1713+0747 to demonstrate that DE405 is significantly

[2]http://ssd.jpl.nasa.gov/iau-comm4/relateds.html

more accurate than its predecessor, DE200. Our timing of PSR J1909–3744 can be used to investigate the behaviour of these two ephemerides in a similar way. The total $\chi^2$ we obtain from a timing model fit (allowing all except the Shapiro delay parameters to vary) using DE405 is 2831. If we use the older DE200 ephemeris, the $\chi^2$ increases to 3108 and we obtain $\alpha = 19{:}09{:}47.4375075(8)$, $\delta = -37{:}44{:}14.31966(6)$, $\mu_\alpha = -9.78(2)$ and $\mu_\delta = -36.09(6)$. The difference of 277 $\chi^2$ units clearly indicates that DE405 is superior to DE200. The fitted values for position and $\mu_\alpha$ are inconsistent with the values listed in Table 8.3, the value for $\mu_\alpha$ in particular is more than $10\sigma$ away. However, the difference we measure corresponds to $\sim 0.3\,\mathrm{mas\,yr^{-1}}$, which is small compared to some of the discrepancies we find. For example, the proper motion in declination that we measure for PSR J0613–0200 is $3\,\mathrm{mas\,yr^{-1}}$ greater than that published by Toscano et al. (1999b). This is a factor of 10 greater than the differences we can attribute to the new reference frame of DE405. Indeed, changing the Solar system ephemeris and re-fitting for proper motion in the timing model for PSR J0613–0200 does not change the parameters significantly. In general, it is not advisable to compare astrometric parameters obtained with reference to DE200 and DE405, although the magnitude of the reference frame rotation is only $\sim 0.3\,\mathrm{mas\,yr^{-1}}$, which cannot explain some of the larger discrepancies we see.

We measure parallax in the timing residuals of PSR J2124–3358 for the first time. Although the significance of the detection is low, the corresponding distance estimate of 250 pc agrees with the NE2001 distance and both are well within the Shklovskii upper limit. Our proper motion measurements are consistent with those published by Toscano et al. (1999b). This pulsar has been detected in X-rays with several different instruments (Sakurai et al. 1999; Becker & Trümper 1999; Sakurai et al. 2001) and Gaensler, Jones & Stappers (2002) have identified a highly asymmetric optical bow shock in H$\alpha$ emission. Studies of this high energy emission may benefit from an improved distance estimate, which future timing will provide.

In contrast with Loehmer et al. (2004), we do not detect a significant parallax in the timing residuals of PSR J2145–0750, in the reference frame of either the DE200 or DE405 Solar system ephemeris. Loehmer et al. (2004) claim to detect a parallax of 2.0 mas, however if we assume this value and fit for position and proper motion, the $\chi^2$ increases by 15 compared to the no-parallax case. Given that our RMS timing residual is less than half that of Loehmer et al. (2004), and that the timing signature of parallax depends more on day-of-year coverage than the total span of observations (once there is a long enough baseline to break any covariance with other parameters), we should easily be able to detect a parallax as large as 2.0 mas. By computing a one-dimensional $\chi^2$ profile, we can rule out any parallax larger than 0.9 mas at the $2\sigma$ level, implying that the distance is greater than 1.1 kpc. Interestingly, the longer data span available to Loehmer et al. (2004) reveals secular evolution of the projected semi-major axis. However, if we adopt the value of $\dot{x} = 1.8 \pm 0.6 \times 10^{-14}$ lt-s s$^{-1}$ and fit for parallax, the result is $< 0$ and clearly unphysical. Given this paradox, we are reminded to treat astrometric timing measurements with some caution. Of course the Loehmer et al. (2004) parallax

| PSR | DM ($cm^{-3}pc$) | Distance (pc) | $n_e$ ($cm^{-3}$) | |
|---|---|---|---|---|
| | | | NE2001 | Measured |
| J0437–4715 | 2.645 | 159 | 0.019 | 0.0166 (5) |
| J0613–0200 | 38.779 | 480 | 0.023 | 0.08 (2) |
| J1022+1001 | 10.252 | 400 | 0.023 | 0.026 (9) |
| J1024–0719 | 6.485 | 520 | 0.017 | 0.012 (5) |
| J1713+0747 | 15.992 | 910 | 0.018 | 0.017 (2) |
| J1744–1134 | 3.139 | 470 | 0.008 | 0.007 (2) |
| J1909–3744 | 10.394 | 1140 | 0.023 | 0.0094 (4) |
| J2124–3358 | 4.596 | 250 | 0.017 | 0.018 (9) |

**Table 8.5: Estimated mean electron density along the line of sight to each pulsar for which we have a measurement of parallax. Uncertainties in DM and parallax can be found earlier in this chapter and are incorporated into the stated uncertainty in measured $n_e$.**

may be correct, as it agrees well with the DM distance and can be used to predict the measured scintillation velocity of the pulsar. Our distance limit implies a mean electron density of just $0.008\,cm^{-3}$ along the line of sight, which is small, but not unreasonable and very similar to that of PSR J1744–1134. Our density of timing points is low during the first year of observations and with time we may resolve this contradiction. Our Shapiro delay limits (see Section 8.7) on the system inclination and companion mass agree very well with those of Loehmer et al. (2004).

In all, we have independent distance measurements for 8 pulsars, which can be used to estimate the mean electron density along the line of sight to the source. Table 8.5 summarises our findings.

We see that two of the electron densities show significant deviations from the NE2001 model. Our parallax measurements for PSR J0613–0200 and PSR J1909–3744 are new and the corresponding electron densities are significantly at odds with the predictions of NE2001. In one case the density is under-estimated and in the other it is over-estimated. NE2001 was designed to be better than the Taylor & Cordes (1993) model, but even with greater complexity we see that it cannot always predict the correct distance. Using our parallax distances, we find that the average transverse velocity of our targeted MSPs is $87^{+31}_{-14}\,km\,s^{-1}$, in very good agreement with the value of $85 \pm 13\ km\,s^{-1}$ published by Toscano et al. (1999b).

In several of our millisecond pulsars, we have determined that the initial spin periods were near their current values and have therefore not changed much during their observable lifetimes, as suggested by Camilo, Thorsett & Kulkarni (1994). Bailes (submitted) recently showed that no recycled pulsars have spin periods less than $10\ m_c/m_p\,ms$, where $m_c$ is the companion mass and $m_p$ is the pulsar mass. Weak magnetic fields that lead to slow period evolution help to ensure this relation is not washed out by period evolution.

## 8.7 Orbital Parameters

### 8.7.1 Shapiro Delay

We can use either the detection or lack of observed Shapiro delay signature to constrain the orbital parameters of all the binary pulsars in our source list except PSR J1045–4509 and PSR J2129–5721. Shapiro delay is parameterised by its range and shape, which are represented in the DD timing model by two theory-independent parameters, $r$ and $s$. These correspond to companion mass, $m_c$, and $\sin(i)$ if we assume that General relativity is the correct theory of gravity. Measurement of both parameters and the pulsar mass function yields a value for the mass of the pulsar itself (again, assuming General relativity is correct). We construct two-dimensional $\chi^2$ maps to determine these two parameters, along with their uncertainty.

PSR J0437–4715 is only mildly relativistic, but its close proximity allows high-precision observations. van Straten et al. (2001) detect a very weak Shapiro delay and measure the rate of periastron precession with marginal confidence. Our measurement of the longitude of periastron at an epoch several years in advance of the previous measurement yields a rate of change equal to $0.02 \pm 0.01^o\,\mathrm{yr}^{-1}$, consistent with the measurement of $0.016 \pm 0.01^o\,\mathrm{yr}^{-1}$ made by van Straten et al. (2001). This rate of change is small and we are unlikely to obtain a better limit on the sum of the masses from General relativity for some time. van Straten et al. (2001) also measure the signature of annual orbital parallax. This breaks the degeneracy in our knowledge of the sign of the inclination angle, which cannot be determined from Shapiro delay alone. We allow the parameter $\Omega$ to vary freely while mapping the $i$–$m_c$ plane and can determine the inclination angle to high precision, but the companion mass is not as well constrained (see Fig. 8.5). We measure $i = 42.5 \pm 0.2^o$ and $m_c = 0.21 \pm 0.02\,M_\odot$, in good agreement with the values of $i = 42.75 \pm 0.09^o$ and $m_c = 0.236 \pm 0.017\,M_\odot$ published by van Straten et al. (2001).

The $1\sigma$ Shapiro delay contour for PSR J1022+1001 is quite extended, however if we assume a reasonable pulsar mass of $1.3\,M_\odot$, we find that the inclination angle should lie between $41^o$ and $53^o$ (the uncertainty in the sign of this angle is not stated explicitly but should be assumed hereafter unless otherwise stated). Assuming the same pulsar mass, the companion mass should lie somewhere in the range 0.8 to $1.3\,M_\odot$, making it almost certainly a heavy white dwarf. By similar arguments, we can constrain the inclination angle of PSR J0613–0200 to be between $59^o$ and $68^o$. This corresponds to a companion mass between 0.13 and $0.15\,M_\odot$. Even if we let the pulsar mass vary between 1.0 and $2.0\,M_\odot$, the companion mass remains between 0.1 and 0.2 $M_\odot$ (68% confidence). This makes it one of the most dissymmetric MSP binary systems suitable for tests of General relativity.

PSR J1713+0747 can be timed with such high precision that both Shapiro delay and annual orbital parallax have been detected (Splaver et al. 2005). We therefore include $\Omega$ in the timing model and allow it to vary when computing the map over $i$–$m_c$ space (shown in Fig. 8.6). We measure $i = 76.6^{+1.5}_{-2.0}$ degrees and $m_c = 0.25 \pm$

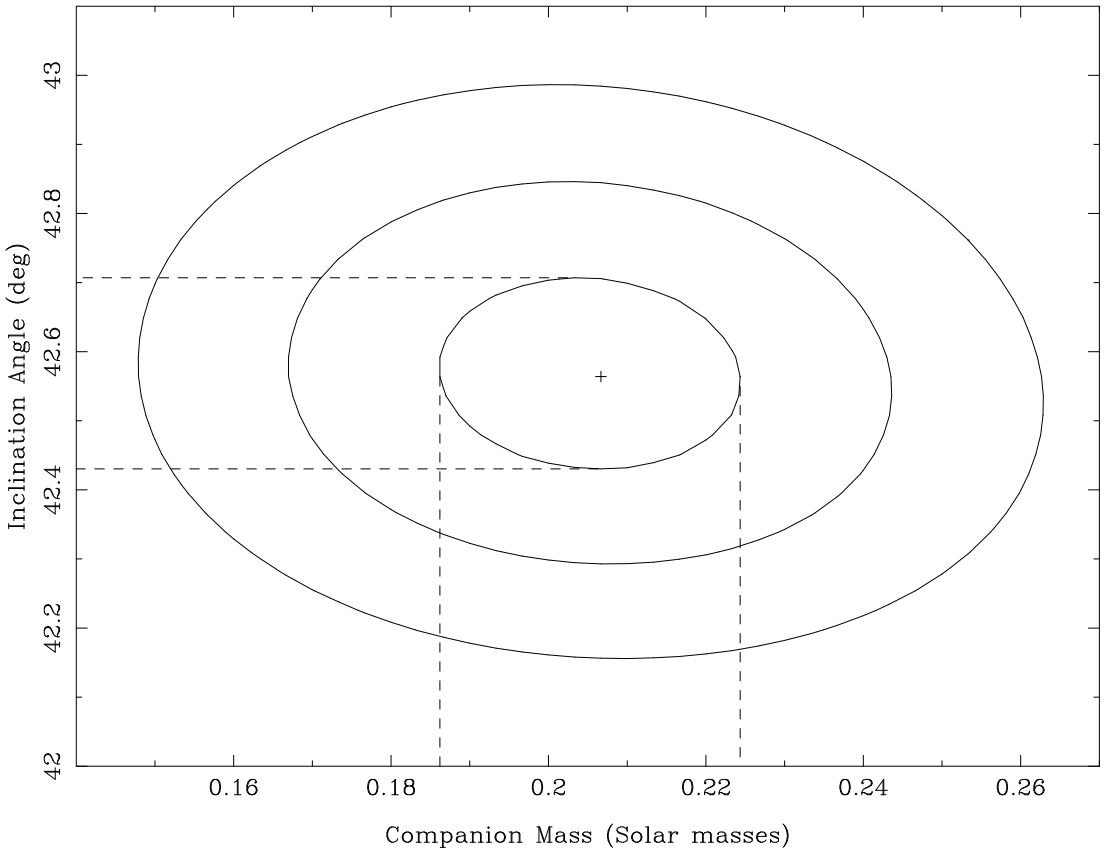


**Figure 8.5: $\chi^2$ map over $i$ and $m_c$ for 5 min integrations of PSR J0437–4715 in two bands centred on 1341 and 1405 MHz. The contours represent 1, 2 and 3$\sigma$ confidence levels.**

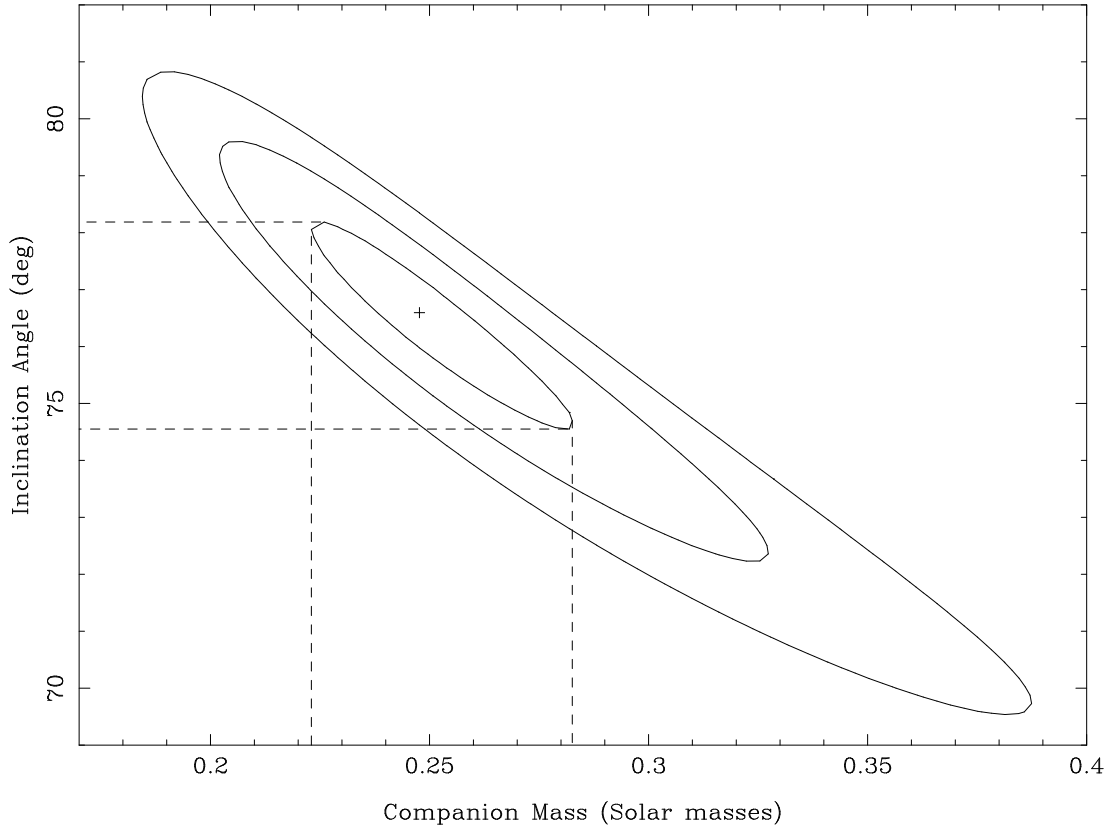


**Figure 8.6: $\chi^2$ map over $i$ and $m_c$ for 5 min integrations of PSR J1713+0747 in two bands centred on 1341 and 1405 MHz. The contours represent 1, 2 and 3$\sigma$ confidence levels.**

$0.03\,M_\odot$. The inclination angle is slightly larger and the companion mass slightly smaller than the values of $i = 71.9 \pm 1.8^o$ and $m_c = 0.28 \pm 0.03\,M_\odot$ that Splaver et al. (2005) measure, but both are consistent at the $2\sigma$ level.

For PSR B1855+09, we measure an inclination angle of $73^{+15}_{-13}$ degrees and $m_c = 0.21 \pm 0.07\,M_\odot$. Our results are less precise than, but consistent with, those presented by Kaspi, Taylor & Ryba (1994). Additional observations will increase the density of points in orbital phase and allow a more precise measurement of the Shapiro delay.

The large inclination and precise timing properties of PSR J1909–3744 make the signature of Shapiro delay highly significant. Fig. 8.7 shows that the inclination angle is constrained to lie between approximately $86.4^o$ and $86.7^o$ and the companion mass is $0.207 \pm 0.002\,M_\odot$. The orbital eccentricity is $1\times10^{-7}$, making this the most circular pulsar binary system ever observed, a record previously held by PSR J1012+5307 (Lange et al. 2001). We use the ELL1 binary model to better characterise the shape of the circular orbit, but the uncertainty in the derived longitude of periastron is still $\sim8^o$. It is therefore unlikely that the rate of periastron advance

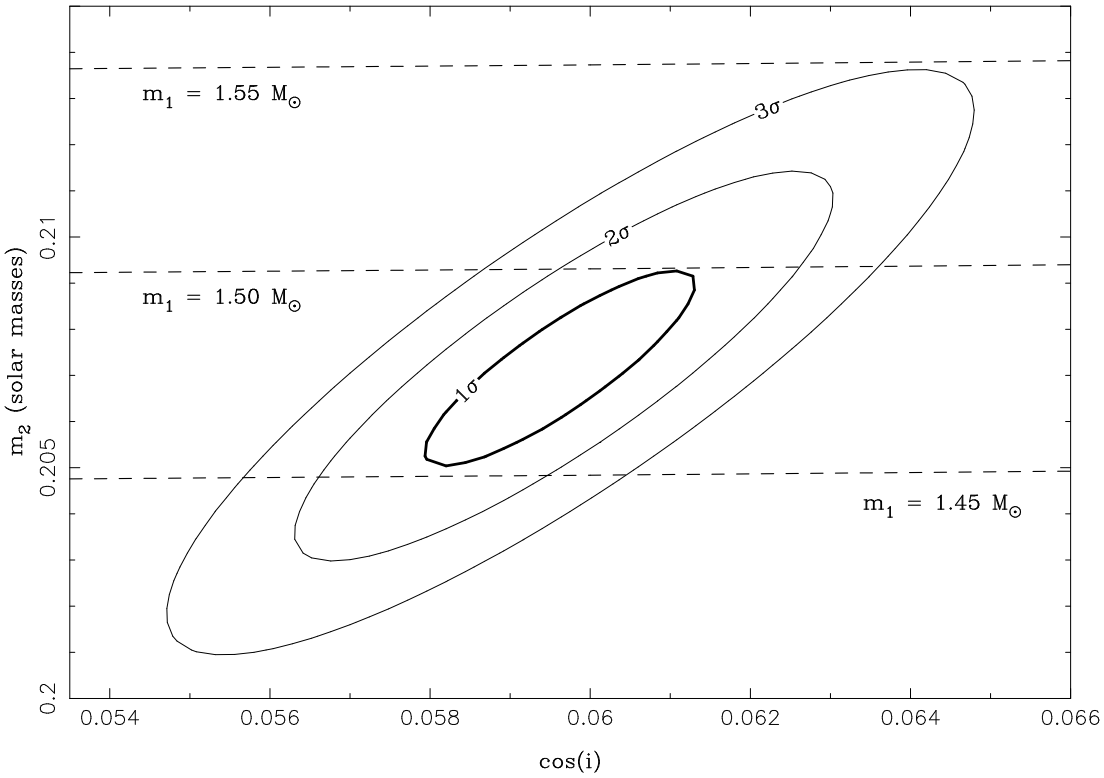


**Figure 8.7: $\chi^2$ map over $\cos(i)$ and $m_c$ for 5 min integrations of PSR J1909–3744 in two bands centred on 1341 and 1405 MHz. The contours represent 1, 2 and 3$\sigma$ confidence levels.**

will ever be measured. PSR J1909–3744 could provide a stringent test for the presence of gravitational dipole radiation, thanks to its low eccentricity and our precise knowledge of the component masses and inclination angle. Wex (2000) states that the best single-source limit is $\alpha_p < 0.02$, based on the timing of PSR J1012+5307. Unfortunately, the orbital period derivative of PSR J1909–3744 is dominated by kinematic effects and therefore tied to its distance from the Solar system. Using our parallax distance to correct for the kinematic contribution to $\dot{P}_b$, we can set a limit of $\alpha_p < 0.03$ (68% confidence), which is already close to the best constraint.

PSR J1909–3744 may also be a good system in which to test for violation of the strong equivalence principle (SEP) in the strong field regime (Damour & Schäfer 1991; Wex 2000; Stairs et al. 2005). Wex (1997) describes the characteristics a binary system must possess to provide a constraint. Firstly, the inequality ${P_b}^2/e > 10^7$ days$^2$ must hold. This is certainly true for PSR J1909–3744 thanks to its small eccentricity. Secondly, the binary system must be old enough for the orientation of the eccentricity vector to have been randomised with respect to the Galactic potential. After correcting for kinematic and galactic effects, the characteristic age of PSR J1909–3744 is at least 14 Gyr which is more than sufficient. Some amount of

statistical analysis will be necessary to derive a constraint on SEP violation due to the fact that we do not yet know the full three-dimensional orientation of the orbit. This is left for future work. With an orbital period of only 1.5 days, PSR J1909-3744 may have evolved differently to the wide-orbit binary millisecond pulsars studied by Stairs et al. (2005), offering an opportunity to test a new kind of system.

Loehmer et al. (2004) use their detection of $\dot{x}$ and the lack of Shapiro delay in the timing residuals of PSR J2145–0750 to place a limit of $0.5 \leq m_c \leq 1.0\,M_\odot$ on the companion mass. Although we do not detect any evolution of the projected semi-major axis within 3 years, assuming a $1.3\,M_\odot$ pulsar leads to a limit of $m_c > 0.5$ and $i < 45^o$, at the $2\sigma$ level.

Our limit on the mass of PSR J1713+0747 is $1.1 \pm 0.2\,M_\odot$, significantly lower than (but still consistent with) the $1\sigma$ limit of Splaver et al. (2005). Bailes et al. (2003) point out that most pulsar mass estimates come from members of DNS binary systems and show that the mass estimate for PSR J1141–6545 (which has a $\sim$1 $M_\odot$ companion) is slightly smaller than the average. Both PSR J1713+0747 and PSR J1909–3744 have much less massive companions, but the pulsar masses are also quite different. It will be necessary to study a wider range of dissymmetric pulsar binary systems in order to establish any mass relationship. Our binary timing models are only strictly valid when both stars can be treated as point masses. Pulsars with very low-mass companions (like PSR J0613–0200) will test this assumption.

### 8.7.2 Annual-Orbital Parallax

For a nearby pulsar with favourable orbital orientation, the heliocentric motion of the Earth can introduce periodic variations in two of the Keplerian binary parameters, as shown by Kopeikin (1995).

$$x_{\rm obs} = x_{\rm int}\left[1 + \frac{\cot(i)}{d}(\Delta_{\rm I_0}\sin(\Omega) - \Delta_{\rm J_0}\cos(\Omega))\right], \tag{8.7}$$

$$\omega_{\rm obs} = \omega_{\rm int} - \frac{\csc(i)}{d}(\Delta_{\rm I_0}\cos(\Omega) + \Delta_{\rm J_0}\sin(\Omega)). \tag{8.8}$$

Here, $\Delta_{\rm I_0} = \vec{r}\cdot\vec{I_0}$ and $\Delta_{\rm J_0} = \vec{r}\cdot\vec{J_0}$, where $\vec{r}$ is the position vector of the Earth in the Solar system (as a function of time) and $\vec{I_0}$, $\vec{J_0}$ and $\vec{K_0}$ are the set of basis vectors used to describe the pulsar binary system. The conventions used to describe a three-dimensional basis for the orientation of an orbit differ between Kopeikin (1995) and Splaver et al. (2005). The latter authors quantify these differences, most notably a rotation of $90^o$ about the line of sight. We adopt the basis used by Kopeikin (1995) and van Straten et al. (2001).

Depending on the orientation of the system, the timing signatures of Eq. 8.7 & 8.8 can have very different amplitudes. Because Eq. 8.7 introduces a dependence on $\tan(i)$, its detection breaks the usual degeneracy in the inclination angle, allowing us to tell whether the orbital angular momentum vector is pointing towards or away from the observer.

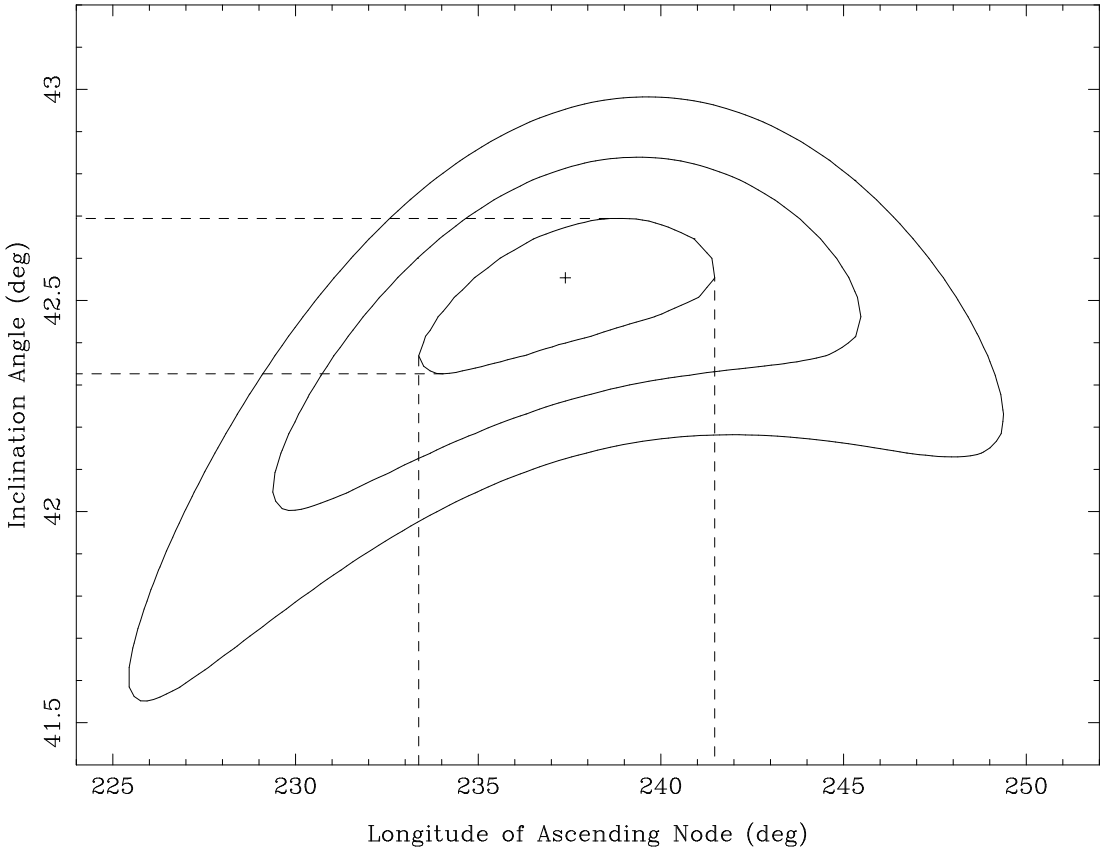


**Figure 8.8: $\chi^2$ map over $i$ and $\Omega$ for 5 min observations of PSR J0437–4715 in two bands centred on 1341 and 1405 MHz. The contours represent 1, 2 and 3$\sigma$ confidence levels.**

### PSR J0437–4715

Because the signature of Shapiro delay is very weak, including $r$ and $s$ in the timing model does not provide a good constraint. However, fitting for $i$ and $\Omega$ allows the inclination angle and orbital orientation to be measured. We obtain the values $\Omega = 237 \pm 4^o$ and $i = 42.5 \pm 0.2^o$. The $\chi^2$ map is shown in Fig. 8.8. van Straten et al. (2001) obtain a very similar value for $\Omega$ and a value of $42.75 \pm 0.09^o$ for $i$, which is consistent with our measurement, even though the data sets are completely independent.

It is interesting to note that the uncertainties in proper motion and $\dot{x}$ decrease as $t^{3/2}$. With more timing history, it would be possible to use Eq. 8.6 to obtain a very precise relationship between $i$ and $\Omega$. This would add a new constraint to Fig. 8.8. At the moment, our lack of precision in $\dot{x}$ means that we can only constrain the inclination angle to be less than $44.5^o$ using this method. The limit obtained from Fig. 8.8 is better. Presently we rely on trigonometric parallax to define the distance to the pulsar, which appears in the two equations (Eq. 8.7 and Eq. 8.8) that define

annual orbital parallax. In future, the Bell & Bailes (1996) distance estimate may prove to be even more precise than parallax and we will then be able to determine the orientation and inclination of the system to very high precision. This reduces the search for Shapiro delay to a one-dimensional problem and may offer a better constraint on the companion mass.

**PSR J1713+0747**

In contrast to PSR J0437–4715, Shapiro delay has a very significant effect on the timing residuals of PSR J1713+0747. Because Shapiro delay is highly covariant with the length of the projected semi-major axis, it is necessary to fit for the three parameters $i$, $\Omega$ and $m_c$ simultaneously to properly constrain the system. We also use Eq. 8.3 to compensate for secular evolution of the projected semi-major axis instead of fitting for an arbitrary $\dot{x}$. We map out a complete three-dimensional $\chi^2$ volume for a uniform grid of points in $i$, $\Omega$ and $m_c$ to determine the favoured combination. In order to determine uncertainties, we compute a two-dimensional map over $i$ and $\Omega$, leaving $m_c$ free (Fig. 8.9).
Although we identify several local minima in the $i$–$\Omega$ plane, there is a clear global minimum more than $3\sigma$ below the others. We measure $\Omega = 81^{+8}_{-9}$ degrees and $i = 76.6^{+1.5}_{-2.0}$ degrees. We note with interest that given the $90^o$ basis rotation with respect to Splaver et al. (2005), our value for $\Omega$ does not agree with previous measurements. The discrepancy may be due to our shorter timing history, which makes it more difficult to constrain the secular evolution of the binary parameters. If we compute $\dot{x}$ using Eq. 8.6 we obtain a value of $-6.2 \pm 1.5\times10^{-15}$. Fitting for $\dot{x}$ without $\sin(i)$ or $m_c$ yields a value of $-3.5 \pm 0.4\times10^{-14}$, but including the Shapiro delay parameters and our best estimate of $\Omega$ in the fit changes the value to $-5 \pm 4\times10^{-15}$, as shown in Table 8.3. This compares well with the expected value, though the uncertainty is large. The sign of the value is consistently negative, in contrast to the measurement of Splaver et al. (2005). Our measurements also allow an estimate of the pulsar mass, which we find to be $1.1 \pm 0.2\,M_\odot$. This is smaller than previous estimates, but still consistent with the typical value of $\sim 1.3\,M_\odot$.

## 8.8 Pulsar Stability

The timing presented in this chapter is some of the most precise ever published. 1 hr integrations of PSR J1909–3744 produce an RMS residual of 150 ns, which is only 0.4% of the pulse width. The effective sampling time (the reciprocal bandwidth of a single spectral channel) is $2\,\mu$s, an order of magnitude larger than the RMS residual. Traditionally, such precise timing has been plagued by large reduced $\chi^2$ values, indicative of systematic or underestimated errors. Observers were forced to ignore the problem or apply an arbitrary, uniform scaling to the TOA error estimates in order to compensate. We see in this chapter that CPSR2 data requires little or

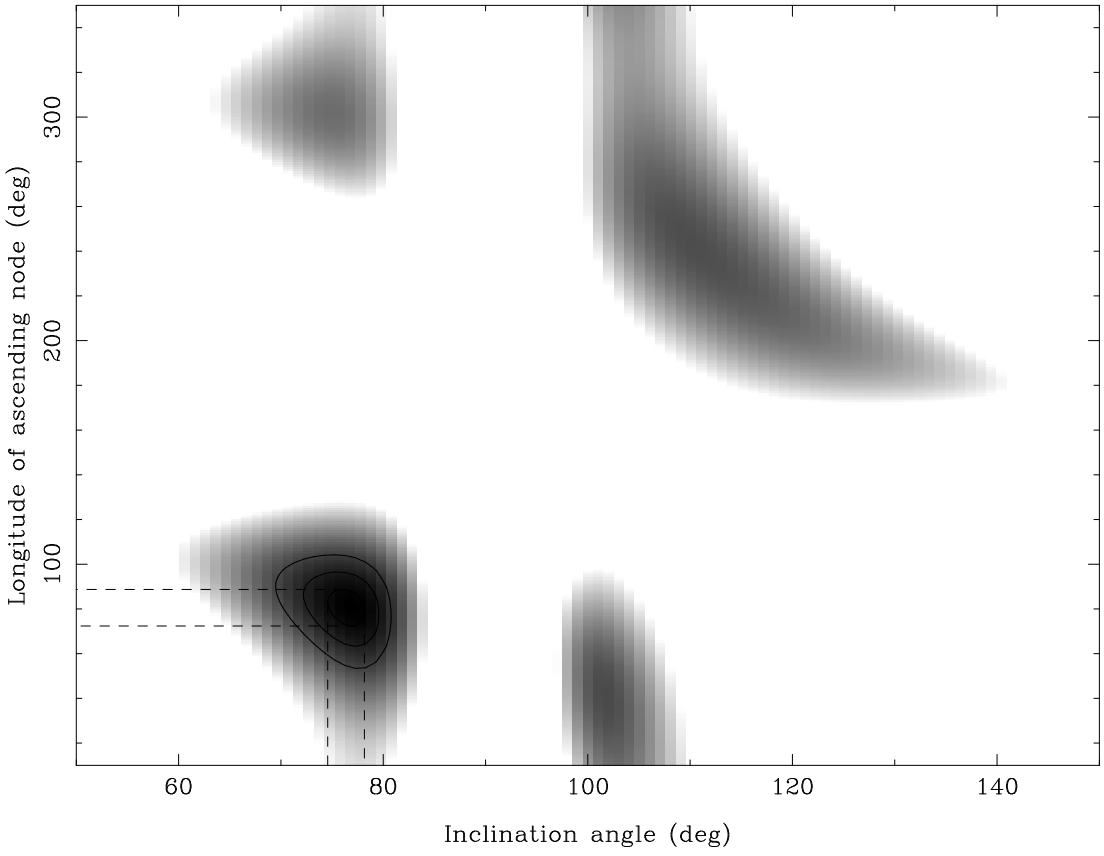


**Figure 8.9:** $\chi^2$ **map over** $i$ **and** $\Omega$ **for 5 min observations of PSR J1713+0747 in two bands centred on 1341 and 1405 MHz. The contours represent 1, 2 and 3**$\sigma$ **confidence levels. The colour map fades to white at a level of 6**$\sigma$**.**

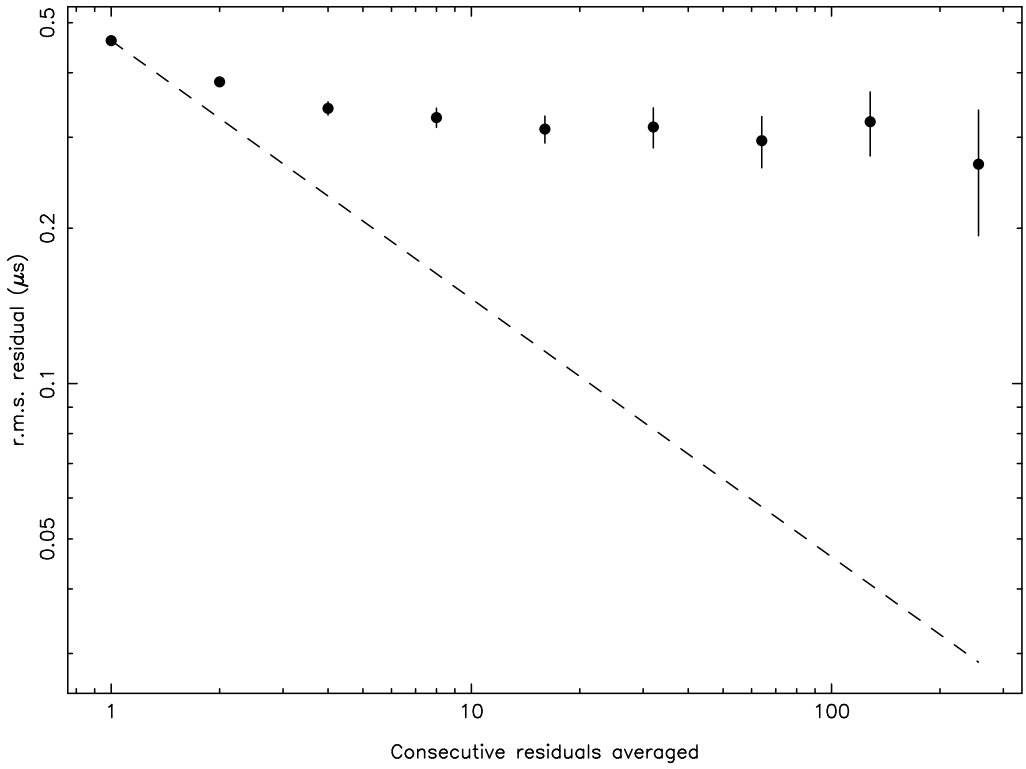


**Figure 8.10: Logarithmic plot of RMS timing residual against the number of residuals averaged. The dashed line represents the predicted trend for white-noise residuals. Clearly, observations of PSR J0437–4715 suffer from some unidentified systematic problem that reduces the timing stability.**

no error scaling as the reduced $\chi^2$ of our most precise timing sources is already close to unity.

### 8.8.1 Fundamental Timing Stability

One of the simplest ways to determine the presence of low-level systematic errors in any set of arrival times is to compute the way in which the RMS residual changes when successively larger numbers of consecutive timing residuals are averaged. For white-noise residuals, the RMS residual should decrease as the square root of the number of residuals averaged. Figures 8.10 and 8.11 show the results of this test when performed on PSR J0437–4715 and PSR J1713+0747.

It is interesting to quantitatively consider the relative timing stability of the best pulsars in our sample using a figure of merit more descriptive than the RMS timing residual. To do this, we compute the $\sigma_z$ (Matsakis, Taylor & Eubanks 1997) statistic as a function of timescale (Fig. 8.12).
We must use 5 min average TOAs to maintain a large number of points, which

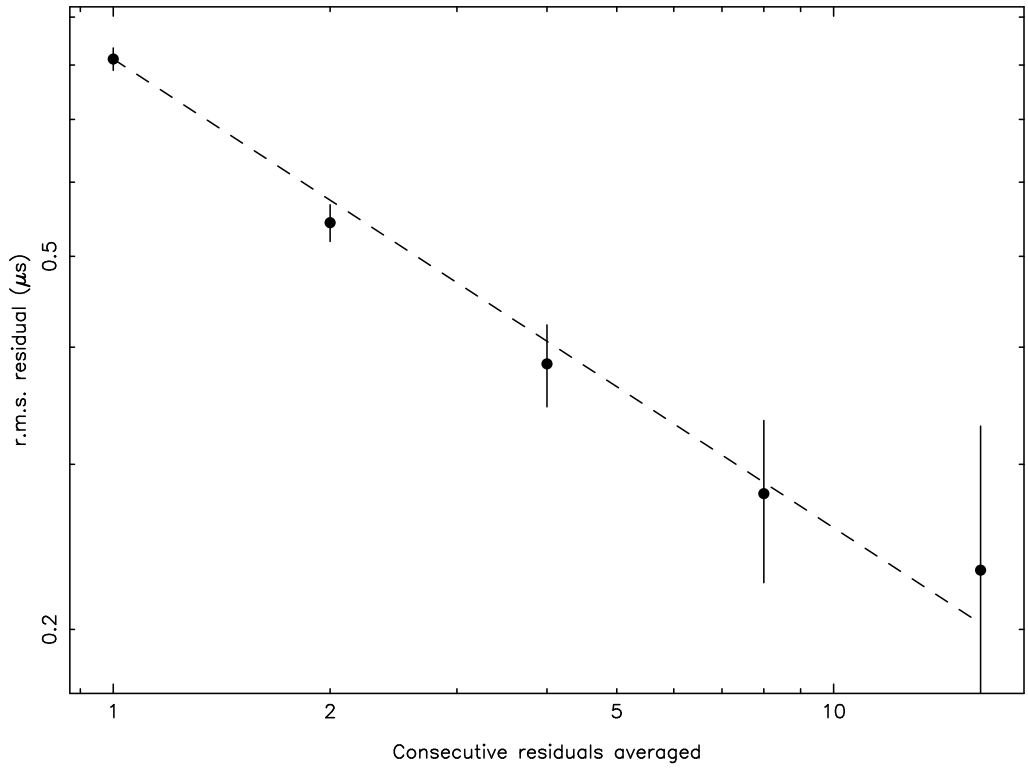


**Figure 8.11: As for Fig. 8.10. In contrast to PSR J0437–4715, PSR J1713+0747 is a very stable pulsar that can be timed very precisely given sufficient observing time.**

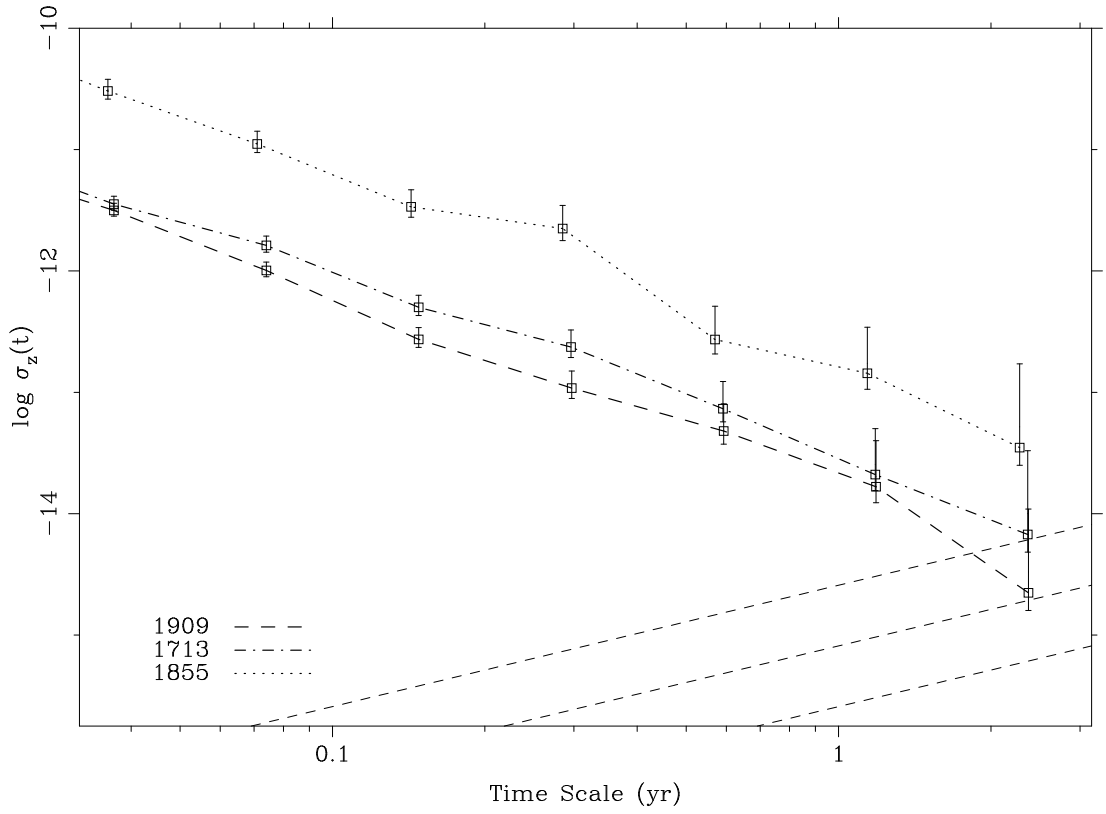


**Figure 8.12:** $\sigma_z$ **as a function of timescale for 5 min integrations of PSR J1909–3744, PSR J1713+0747 and PSR B1855+09. The diagonal dashed lines in the lower right of the plot represent limits associated with (from top)** $\Omega_g h^2 = 10^{-7}$, $10^{-8}$ **and** $10^{-9}$**. Although our total time span is small, we see that PSR J1909–3744 is significantly more stable than B1855+09.**

has the effect of increasing all $\sigma_z$ values relative to those computed with longer integrations (Kaspi, Taylor & Ryba 1994). PSR J1909–3744 and PSR J1713+0747 are the best pulsar clocks, far exceeding the stability of PSR B1855+09. With a few more years of observations we may be able to place more stringent limits on $\Omega_g$ than those already obtained (Kaspi, Taylor & Ryba 1994). We note that Splaver et al. (2005) require terms up to and including the eighth frequency derivative to "whiten" their timing residuals for PSR J1713+0747, over a time span of 12 years. We can obtain 125 ns RMS residuals over a time span of ∼3 years by fitting only for the first frequency derivative.

In general, it is difficult to obtain an RMS residual below 100 ns with current instruments and procedures. At this level of accuracy a large number of problems start to arise. These include imperfect calibration, instrumental distortions of the pulse profile, non-uniform bandpasses and uncertainties in the relationships between various timescales.

### 8.8.2 Dispersion Measure Variability

For 8 of our 15 pulsars, the measured DM is significantly different to other values that appear in the literature. In the case of PSR B1937+21, variations in DM measurements have already been studied and attributed to genuine changes in the observed DM (Kaspi, Taylor & Ryba 1994). Fitting for the first derivative of DM in the timing model we have constructed gives a value of $-0.0012 \pm 0.0004\,\mathrm{cm^{-3}pc\,yr^{-1}}$, which matches very well with the slope seen in Fig. 3 of Kaspi, Taylor & Ryba (1994).

In general, measurements of DM are subject to a number of systematic errors. Any misalignment of the standard template profiles used to produce TOAs at different observing frequencies will appear as a constant offset in the measured DM. Because pulse profiles can evolve with frequency, it is sometimes difficult to determine the best alignment and even more difficult to ensure that the offset does not change when new template profiles are constructed. A systematic offset could also be introduced by the observing system, which may have different signal path lengths at different frequencies. This means that the uncertainty in a DM measurement is often underestimated, which makes it difficult to compare measurements of DM taken with different telescopes and at different epochs.

In order to investigate the magnitude of any intrinsic and systematic DM variations in our data, we compare our measurements with a selection of previously published values and estimate the rate of change per year, assuming linear evolution. We also compute the time delay this would introduce if the incorrect DM were assumed in the timing model (Table 8.6).

The discrepancies between measured DM values correspond to systematic time delays of up to 5 $\mu$s per year, which would be catastrophic to any precision timing project. Monitoring the DM of a pulsar precisely requires regular observations at multiple frequencies over a long time span using the same hardware configuration.

| Pulsar | ΔDM ($cm^{-3}pc$) | ΔT (yrs) | ΔDM ($cm^{-3}pc\ yr^{-1}$) | Delay $\mu s$ | Reference |
|---|---|---|---|---|---|
| J0711–6830 | –0.0031 | 7.0 | –0.0004 | 0.8 | Toscano et al. (1999b) |
| J1024–0719 | –0.0027 | 5.4 | –0.0005 | 1.1 | Hobbs et al. (2004b) |
| J1045–4509 | 0.0138 | 7.6 | 0.0018 | 3.8 | Toscano et al. (1999b) |
| J1603–7202 | –0.0029 | 7.1 | –0.0004 | 0.8 | Toscano et al. (1999b) |
| J1713+0747 | 0.0009 | 6.0 | 0.0002 | 0.4 | Hobbs et al. (2004b) |
| B1855+09 | –0.0082 | 15.5 | –0.0005 | 1.1 | Kaspi, Taylor & Ryba (1994) |
| J2124–3358 | –0.0192 | 7.9 | –0.0024 | 5.1 | Toscano et al. (1999b) |
| J2145–0750 | –0.0012 | 6.2 | –0.0002 | 0.4 | Loehmer et al. (2004) |

**Table 8.6: Dispersion measure values and comparison with previous results. The second column shows the minimum difference between the DM value we measure and a previously published value for 8 pulsars where the change is outside the $1\sigma$ error margin. The third column shows the elapsed time between the two measurements and the fourth column shows the corresponding rate of change per year. The fifth column shows the offset that an uncorrected DM drift of this magnitude would introduce into the pulse arrival times each year. The final column lists the references from which the earlier DM values were obtained.**

In future, this will be necessary to avoid systematic errors and determine the true nature of the observed DM variations.

### 8.8.3 Pulsar-Based Time Standards

Long-term precision timing of millisecond pulsars has demonstrated stability comparable to the best terrestrial atomic clocks. It is therefore reasonable to test whether or not terrestrial time standards can be made redundant by using one pulsar to time another. The best reference source is the pulsar that has the smallest RMS residual, provided it does not exhibit any timing noise. We use PSR J1909–3744.

In practice, observatories require a stable local frequency standard implemented in hardware. At Parkes, the observatory time standard is a H-maser whose offset from UTC is monitored daily by comparison with time signals from GPS satellites. Observations of the pulsar we choose to use as a new time standard must be made using the local frequency reference. However, assuming our mathematical timing model is perfect, any systematic trend in the timing residuals (where all parameters are kept fixed and no clock corrections are applied) must correspond to errors in arrival time assignment, either due to profile corruption or a drifting local clock. These residuals can be used to correct the arrival times obtained from any other source. This mutual "referencing" of arrival times is fundamental to methods for detecting long-period gravitational waves.

We obtained our best timing model for PSR J1909–3744 using the clock corrections published by the observatory and then removed the GPS-PKS correction from TEMPO and performed a new fit, keeping all model parameters constant. The uncorrected residuals were then averaged within successive days, inverted in sign and

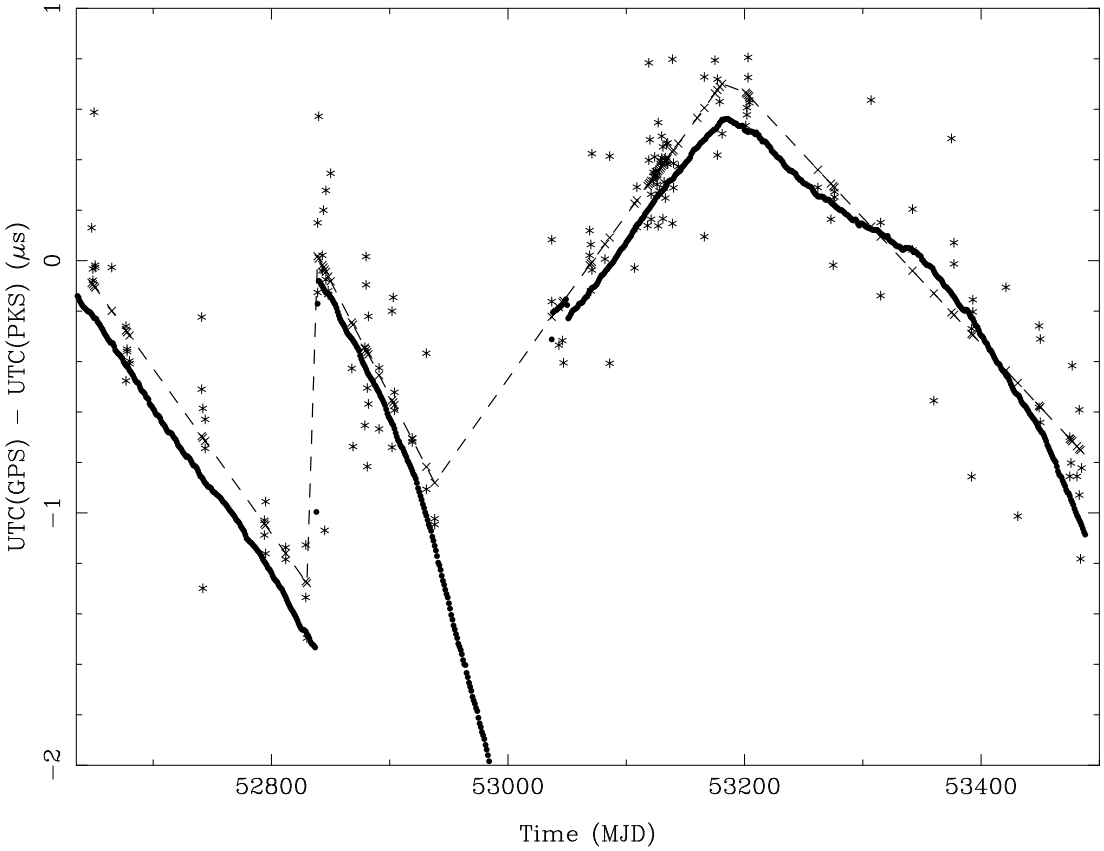


**Figure 8.13: The solid line represents measured Parkes clock offsets, obtained by comparing the maser-driven Parkes clock with GPS timing signals. The stars represent daily average clock corrections derived from the timing residuals of PSR J1909–3744 and the dashed line represents a best linear fit to the daily averages within the four regions bounded by adjustments to the maser. The linearised model is a very good approximation to the true Parkes clock drift.**

printed out in a format consistent with the TEMPO 'time.dat' clock correction file. These new corrections were used to time PSR J1713+0747 and PSR J1744–1134. Comparing the RMS timing residual obtained using the "official" clock corrections to that obtained with the corrections derived from PSR J1909–3744 indicates the effectiveness of the pulsar clock.

Daily average residuals from PSR J1909–3744 contain significant scatter at the level of ∼100 ns, which we smooth by using a process of linear interpolation. The new clock corrections are broken into 4 segments, bounded by points where the first derivative is undefined due to hardware adjustment of the maser rate. We fit a straight line to each segment, and obtain smoothed clock corrections from the linearised model. This is a reasonable approximation; to first order, most clock errors consist of a steady drift at some small rate. Fig. 8.13 summarises our approach.

Using the linearised PSR J1909–3744 clock corrections, we obtain an RMS resid-

ual of 133 ns for 1 hr integrations of PSR J1713+0747 and 882 ns for 5 min integrations of PSR J1744–1134. This corresponds to an increase of 8 ns for PSR J1713+0747 and a decrease of 8 ns for PSR J1744–1134. The figure of 133 ns represents one of the smallest RMS residuals ever obtained.

Unfortunately, pulsar time standards suffer from a number of limitations at the present time. Fig. 8.13 illustrates that a very high density of observations is required to accurately monitor clock drifts; the gap in observations around MJD 53000 leaves a significant discrepancy between the linearised model and the official corrections. For this reason, the RMS timing residual of PSR J0437–4715 increases by 100 ns when timed against PSR J1909–3744. Observations of PSR J1713+0747 and PSR J1744–1134 were always made within days of observing PSR J1909–3744, making them ideal test cases. Our method also involves an element of "bootstrapping", in that the official clock corrections were used to obtain the best timing model for PSR J1909–3744. A truly independent timescale would rely on a free-floating reference, which amounts to accepting that the timing model we use for the reference pulsar is perfect at this epoch and will remain so in the future. Such faith in current models is not warranted.

## 8.9 Conclusions

We have systematically studied pulse arrival times from 15 millisecond pulsars and demonstrated the remarkable stability of the CPSR2 baseband recording and coherent dedispersion system. We obtained some of the smallest RMS timing residuals ever seen. This allowed several new parallax distance estimates, two of which were at odds with the predictions of the NE2001 electron density model. Shapiro delay or its non-detection were used to constrain all but two of the binary pulsars in our source list. We measured annual orbital parallax in PSR J0437–4715 and PSR J1713+0747. Future observations of these systems will be necessary to confirm the orbital alignments and obtain more precise measurements of secular evolution of the projected semi-major axes. In addition, we have highlighted that the DE405 Solar system ephemeris is superior to the earlier DE200 version and found that PSR J1909–3744 can be used as an independent time standard with stability similar to the best atomic clocks.

Decadal timing programs with sub-100 ns RMS precision will probe cosmologically interesting limits on the gravitational wave background (Jenet et al. 2005). At this level of precision, sources of systematic error abound. Future improvements in accuracy will be limited by scintillation, dispersion measure variations, digitisation artifacts, polarisation purity, ionospheric variations, Solar system ephemerides and the long-term stability of telescope and instrumental hardware. It will be necessary to eradicate or compensate for all these if we are to improve below 100 ns RMS.

# Chapter 9

# Conclusion

"If you want to make an apple pie from scratch, you must first create the universe."

*Carl Sagan*

Here, we review the findings described in previous chapters and discuss their implications for precision pulsar timing. We also outline ways in which this research could be extended to yield new results in the future. Finally, we discuss the evolution of radio astronomical instrumentation in light of the rapid pace of technological advance in the modern era.

## 9.1 Summary of Results

- We have shown in Chapter 5 that the integrated profile of PSR J1022+1001 is stable. CPSR2 observations of this pulsar were timed to a level of precision an order of magnitude better than any previously obtained. We note that given its 16.5 ms period, our level of timing precision for PSR J1022+1001 is typical of the millisecond pulsars.

- In Chapter 6, we showed that the mean profile of PSR J1141–6545 has grown dramatically over time, most likely as a consequence of geodetic precession. This phenomenon will one day allow us to map the overall shape and polarimetric structure of a pulsar emission cone, providing additional insight into the nature of the emission mechanism.

- CPSR2 observations of the 22.7 ms pulsar in the double pulsar binary system (PSR J0737–3039A) yielded high S/N profiles at two previously unpublished radio frequencies. We calibrated the data to an absolute flux scale and corrected for instrumental polarisation in order to produce two mean profiles that

were presented in Chapter 7. We also timed this pulsar over a period of 15 days and used the detection of Shapiro delay to constrain the inclination angle of the orbit to $88.5^{+0.8}_{-1.1}$ degrees, in good agreement with other timing and scintillation studies.

- Chapter 8 described high precision timing observations of 15 millisecond pulsars over a period of 3 years. CPSR2 was shown to produce high quality data, from which we detected parallax in 8 systems, constrained the masses and inclination angles of several binary pulsars and detected annual orbital parallax in two nearby sources. We measured several RMS timing residuals of order 200 ns, amongst the smallest ever obtained. We used the timing residuals of PSR J1909–3744 as a reference clock against which we obtained an RMS residual of 133 ns for PSR J1713+0747. Linking the timing of two pulsars in this way is central to the concept of a pulsar timing array.

### 9.1.1 Limiting Factors

The motivation for this thesis was the need to push pulsar timing precision to new levels. Although we can obtain ∼200 ns RMS residuals with stable instrumentation, it is not clear that we are capable of pushing RMS residuals below 100 ns with current technology. At the levels of precision we are beginning to reach, there are a number of potential sources of systematic error above and beyond those that are routinely compensated for:

- Digitisation artifacts (including scattered power) may corrupt the folded profile, especially if the signal is strong.
- Polarisation impurities in the observing system (and improper calibration) may corrupt the folded profile.
- Dispersion measure variations due to relative motion between the pulsar and the ISM and fluctuations in the Solar wind and the ionosphere may introduce unmodelled time delays. These DM variations are very difficult to disentangle from time delays introduced by systematic errors in the method used to align standard template profiles from different frequency bands, or even changes in the relative signal path length when different observing systems are used.
- Narrow-band scintillation may randomly alter the effective observing frequency.
- Timing noise may begin to dominate the residuals.
- Solar system ephemerides may not allow sufficiently precise (or consistent) barycentric corrections.
- The use of different receivers and instruments causes discontinuities in the timing history.

These effects, and undoubtedly others that we are currently unaware of, make it difficult to achieve sub-100 ns RMS timing residuals. In future, practitioners of pulsar timing will have to employ wide-band detectors with multi-bit digitisers and real-time monitoring capabilities to ensure the level of data quality required to fully investigate these and other sources of uncertainty.

Telescopes are continuously upgraded to maximise their scientific output, but this makes it difficult to obtain an unbroken timing series spanning more than a few years. Although it is possible to combine the data from distinct instruments, phase jumps are always necessary to cross the boundaries, reducing the power of a long timing baseline.

## 9.2 Future Work

We must diligently monitor PSR J1141–6545, PSR B1534+12 and PSR B1913+16, which already exhibit signs of geodetic precession, and PSR J0737–3039A, in case its mean profile also begins to change. The geodetic precession timescale for this pulsar is only ∼70 years; provided the geometry of the system is favourable, we should expect to see rapid profile evolution. The lack of any detectable change in the first 12 months of data (Manchester et al. 2005b) is mysterious. In addition, we must assess the timing of PSR J0437–4715 at 10 cm, to see whether higher precision can be obtained at this relatively short wavelength.

If General relativity is the correct theory of gravity, we expect there to be a background of cosmological gravitational radiation associated with super-massive binary black holes at the centres of galaxies. To detect this radiation with a significance greater than $4\sigma$ would require residual-whitening methods and 20 pulsars timed to 100 ns RMS over the course of 5 years (Jenet et al. 2005). The significance scales with the number of pulsars in the ensemble and the detectable background amplitude scales with the RMS timing residual. At the moment, we have sufficient timing precision to detect the expected background level, but far too few sources to make the detection significant. The best way forward is to increase the number of MSPs that can be timed with the precision we achieve in this thesis. Unfortunately, this is not only a matter of increasing the S/N of the observed profile. Although we study 15 MSPs in Chapter 8, only 3 are likely to achieve 100 ns RMS over 5 years. Only the fastest MSPs have the intrinsic timing precision required. There are currently 84 pulsars with periods less than 5 ms, but many of these reside in globular clusters and are therefore difficult to detect. Their timing is also often detrimentally affected by acceleration in the gravitational field of the cluster.

While pulsars like J0711–6830 may benefit from more sensitive observing systems, PSR J1045–4509 has a mean S/N of 55 and appears to time poorly for fundamental reasons. Of the 84 short-period MSPs, we might expect roughly 1/4 to be suitable for timing at the 100 ns level, provided a sufficiently sensitive instrument exists. Solving the problem with integration time would tax the resources of a

single telescope (an ensemble of 20 pulsars observed for 8 hours each would take a week of telescope time). The sensitivity and multibeam capabilities of the proposed Square Kilometre Array[1] (SKA) should make the task trivial, but with new wide-band timing instruments it may be possible to detect the cosmological gravitational wave background with existing telescopes some 10 years before the SKA could be completed.

Jenet, Creighton & Lommen (2005) derive an expression for the expected timing signature caused by the presence of an intermediate-mass binary black hole within a globular cluster. With amplitudes ranging from 5 to 500 ns (depending on how close the binary black hole is to the line of sight), the signature of a 10–1000 $M_\odot$ black hole binary pair may be detectable with current timing methods. Instrumental sensitivity would ideally have to increase by an order of magnitude to make a search for this type of timing signature in globular cluster pulsars feasible.

While the requirements of the Australian Pulsar Timing Array may seem difficult to meet at present, I believe we are only one or two instrumental iterations (perhaps an additional factor of two or three in sensitivity) away from new scientific capabilities.

## 9.3 Future Technology Development

New developments in software, algorithms, analytical methods and instrumentation always have broad-ranging implications. In the case of pulsar timing, higher precision will lead to better constraints on the parameters of individual systems and therefore to new and more stringent tests of gravitational theories. Developments will also lead, via the emerging concept of the pulsar timing array, to new insight on cosmological models and the gravitational wave background.

The scientific instruments of tomorrow will not consist of complex, hard-wired single-use systems that cost disproportionately large amounts of money and are quickly made obsolete. The scientific community is no longer an isolated body operating far in advance of consumer demand. It is certainly true that scientific applications offer some of the most strenuous tests of general-purpose digital hardware, but at least in radio astronomy, the gap is closing. During the course of this thesis we have proven that within the limitations of 2-bit digitisation, CPSR2 is a robust and versatile system that can perform the arduous task of precision pulsar timing and polarimetry as well as any other instrument in the world. By using widely-available general purpose processors as the central engine, it was immediately possible to take advantage of the wealth of existing software and experience surrounding the Linux operating system and the higher level programming languages (in particular, C++). The modular design philosophy of CPSR2 and PSRCHIVE represents an innovation that was long overdue in the pulsar community and has been used in professional engineering circles for many years. Modularity is a concept so simple and pow-

[1] http://www.skatelescope.org

erful that it can vastly reduce development and testing time and the unnecessary duplication of effort, if applied with common sense.

As re-programmable devices like FPGAs become cost effective and widely available, the constraints imposed by desktop computing standards will start to evaporate. Our ultimate goal should include adaptive instruments that can automatically be configured for a specific experiment by allocating dynamic logic units built up around core components capable of detection, filtering, spectrum analysis, and so on. The interface between software and hardware will continue to blur. There will of course be significant challenges; primarily the need for world-wide standards. It is likely that some sort of meta-data approach (similar to the International Virtual Observatory Alliance standard[2]) may be the ultimate solution, where a block of information must carry around its own instruction manual. This generality, coupled with wide-area network bandwidths several orders of magnitude greater than we have today, will coincide with a new era in global communications. Scientists in general and radio astronomers in particular will supply the drive and the applications that will test the first examples of this new technology. An instrument such as the SKA would require unprecedented self-diagnosis and intelligent re-configuration capabilities simply to handle the necessary data rates and wide-ranging analytical tasks. Similarly, a world-wide electronic VLBI network will test intelligent data routing schemes under high loads, while simultaneously providing unprecedented spatial resolution. Active cooperation between industry and the academic community will be essential to achieve the goals of both parties.

[2]http://www.ivoa.net

# References

Anderson S. B., Gorham P. W., Kulkarni S. R., Prince T. A., 1990, Nature, 346, 42

Arzoumanian Z., Nice D. J., Taylor J. H., Thorsett S. E., 1994, ApJ, 422, 671

Baade W., Zwicky F., 1934, Proc. Nat. Acad. Sci., 20, 254

Backer D. C., Sallmen S., 1997, AJ, 114, 1539

Backer D. C., Kulkarni S. R., Heiles C., Davis M. M., Goss W. M., 1982, Nature, 300, 615

Bailes M. et al., 1994, ApJ, 425, L41

Bailes M. et al., 1997, ApJ, 481, 386

Bailes M., Ord S. M., Knight H. S., Hotan A. W., 2003, ApJ, 595, L49

Bailes M., 1988, A&A, 202, 109

Baring M. G., 2004, Adv. Space Res., 33, 552

Barker B. M., O'Connell R. F., 1975a, Phys. Rev. D, 12, 329

Barker B. M., O'Connell R. F., 1975b, ApJ, 199, L25

Bartel N., Ratner M. I., Shapiro I. I., Cappallo R. J., Rogers A. E. E., Whitney A. R., 1985, AJ, 90, 318

Becker W., Trümper J., 1999, A&A, 341, 803

Bell J. F., Bailes M., 1996, ApJ, 456, L33

Bhattacharya D., van den Heuvel E. P. J., 1991, Phys. Rep., 203, 1

Biggs J. D., 1990, MNRAS, 245, 514

Bignall H. E. et al., 2003, ApJ, 585, 653

Blandford R., Teukolsky S. A., 1976, ApJ, 205, 580

Blandford R. D., Romani R. W., Applegate J. H., 1987, MNRAS, 225, 51P

Blumenthal G. R., Tucker W. H., 1974, Ann. Rev. Astr. Ap., 12, 23

Bracewell R., 1965, The Fourier Transform and its Applications. McGraw–Hill, New York

Britton M. C., 2000, ApJ, 532, 1240

Bureau International des Poids et Mesures, 1998, The International System of Units

(SI). Organisation Intergouvernementale de la Convention du Metre, Pavillion de Breteuil, France

Burgay M. et al., 2003, Nature, 426, 531

Camilo F., Nice D. J., Shrauner J. A., Taylor J. H., 1996, ApJ, 469, 819

Camilo F., Kaspi V. M., Lyne A. G., Manchester R. N., Bell J. F., D'Amico N., McKay N. P. F., Crawford F., 2000a, ApJ, 541, 367

Camilo F., Lorimer D. R., Freire P., Lyne A. G., Manchester R. N., 2000b, ApJ, 535, 975

Camilo F., Thorsett S. E., Kulkarni S. R., 1994, ApJ, 421, L15

Cheng K. S., Ho C., Ruderman M., 1986, ApJ, 300, 500

Cheng K. S., 2004, Adv. Space Res., 33, 561

Clifton T. R., Lyne A. G., Jones A. W., McKenna J., Ashworth M., 1992, MNRAS, 254, 177

Cognard I., Backer D. C., 2004, ApJ, 612, L125

Coles W. A., McLaughlin M. A., Rickett B. J., Lyne A. G., Bhat N. D. R., 2005, ApJ, 623, 392

Cooley J. W., Tukey J. W., 1965, Math. Comput., 19, 297

Cordes J. M., Lazio T. J. W., 2002, astro-ph/0207156

Cordes J. M., Lazio T. J. W., 2003, astro-ph/0301598

Cordes J. M., Wolszczan A., Dewey R. J., Blaskiewicz M., Stinebring D. R., 1990, ApJ, 349, 245

Cordes J. M., Wasserman I., Blaskiewicz M., 1990, ApJ, 349, 546

Crawford F., Kaspi V. M., Manchester R. N., Lyne A. G., Camilo F., D'Amico N., 2001, ApJ, 553, 367

D'Amico N., Stappers B. W., Bailes M., Martin C. E., Bell J. F., Lyne A. G., Manchester R. N., 1998, MNRAS, 297, 28

D'Amico N. et al., 2001, in Palumbo G., White N., eds, AIP Conference Proceedings 599: Stellar Endpoints, AGN and the X-ray Background. Gordon & Breach, Singapore, p. 598

Damour T., Deruelle N., 1985, Ann. Inst. H. Poincaré (Physique Théorique), 43, 107

Damour T., Deruelle N., 1986, Ann. Inst. H. Poincaré (Physique Théorique), 44, 263

Damour T., Ruffini R., 1974, Academie des Sciences Paris Comptes Rendus Ser. Scie. Math., 279, 971

Damour T., Schäfer G., 1991, Phys. Rev. Lett., 66, 2549

Damour T., Taylor J. H., 1992, Phys. Rev. D, 45, 1840

Davies J. G., Hunt G. C., Smith F. G., 1969, Nature, 221, 27

Deich W. T. S., Kulkarni S. R., 1996, in van Paradijs J., van del Heuvel E. P. J., Kuulkers E., eds, Compact Stars in Binaries: IAU Symposium 165. Kluwer, Dordrecht, p. 279

Demorest P., Ramachandran R., Backer D. C., Ransom S. M., Kaspi V., Arons J., Spitkovsky A., 2004, ApJ, 615, L137

Dennett-Thorpe J.; de Bruyn A. G., 2000, ApJ, 529, 65

Dewey R. J., Taylor J. H., Weisberg J. M., Stokes G. H., 1985, ApJ, 294, L25

Dodson R., Legge D., Reynolds J. E., McCulloch P. M., 2003, ApJ, 596, 1137

Dodson R. G., McCulloch P. M., Lewis D. R., 2002, ApJ, 564, L85

Edwards R. T., Bailes M., 2001a, ApJ, 547, L37

Edwards R. T., Bailes M., 2001b, ApJ, 553, 801

Edwards R. T., Bailes M., van Straten W., Britton M. C., 2001, MNRAS, 326, 358

Esposito L. W., Harrison E. R., 1975, ApJ, 196, L1

Faulkner A. J. et al., 2004, MNRAS, 355, 147

Faulkner A. J. et al., 2005, ApJ, 618, L119

Foster R. S., Backer D. C., 1990, ApJ, 361, 300

Foster R. S., Wolszczan A., Camilo F., 1993, ApJ, 410, L91

Fruchter A. S., Stinebring D. R., Taylor J. H., 1988, Nature, 333, 237

Gaensler B. M., Jones D. H., Stappers B. W., 2002, ApJ, 580, L137

Gavriil F. P., Kaspi V. M., Roberts M. S. E., 2004, Adv. in Space Res., 33, 592

Gold T., 1968, Nature, 218, 731

Goldreich P., Julian W. H., 1969, ApJ, 157, 869

Griffiths D. J., 1999, Introduction to Electrodynamics. Prentice-Hall International, Inc.

Gunn J. E., Ostriker J. P., 1969, Nature, 221, 454

Hamilton P. A., Lyne A. G., 1987, MNRAS, 224, 1073

Hanisch R. J., Farris A., Greisen E. W., Pence W. D., Schlesinger B. M., Teuben P. J., Thompson R. W., Warnock A. I., 2001, A&A, 376, 359

Hankins T. H., Rickett B. J., 1975, in Methods in Computational Physics Volume 14 — Radio Astronomy. Academic Press, New York, p. 55

Hankins T. H., Stinebring D. R., Rawley L. A., 1987, ApJ, 315, 149

Hari Dass N. D., Radhakrishnan V., 1975, Astrophys. Lett., 16, 61

Heiselberg H., Pandharipande V., 2000, Ann. Rev. Nucl. Part. Sci., 50, 481

Helfand D. J., Manchester R. N., Taylor J. H., 1975, ApJ, 198, 661

Hellings R. W., Downs G. S., 1983, ApJ, 265, L39

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

Hobbs G. et al., 2004a, MNRAS, 352, 1439

Hobbs G., Lyne A. G., Kramer M., Martin C. E., Jordan C., 2004b, MNRAS, 353, 1311

Hotan A. W., Bailes M., Ord S. M., 2004, MNRAS, 355, 941

Hotan A. W., Bailes M., Ord S. M., 2005, MNRAS, 362, 1267

Hotan A. W., Bailes M., Ord S. M., 2005, ApJ, 624, 906

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

Hulse R. A., Taylor J. H., 1974, ApJ, 191, L59

Hulse R. A., Taylor J. H., 1975, ApJ, 195, L51

Ilyasov Y. P., Imae M., Hanado Y., Oreshko V. V., Potapov V. A., Rodin A. E., Sekido M., 2005, Astronomy Letters, 31, 30

International Telecommunication Union, 2002, ITU-R Recommendation TF.460-6

Jacoby B. A., Bailes M., van Kerkwijk M. H., Ord S., Hotan A., Kulkarni S. R., Anderson S. B., 2003, ApJ, 599, L99

Jacoby B. A., Hotan A., Bailes M., Ord S., Kuklarni S. R., 2005, ApJ, 629, L113

Jenet F. A., Anderson S. B., 1998, PASP, 110, 1467

Jenet F. A., Creighton T., Lommen A., 2005, ApJ, 627, 125

Jenet F. A., Hobbs G. B., Lee K. J., Manchester R. N., 2005, ApJ, 625, L123

Johnston S., Galloway D., 1999, MNRAS, 306, L50

Johnston S., Lyne A. G., Manchester R. N., Kniffen D. A., D'Amico N., Lim J., Ashworth M., 1992, MNRAS, 255, 401

Johnston S. et al., 1993, Nature, 361, 613

Kaspi V. M., Manchester R. N., Siegman B., Johnston S., Lyne A. G., 1994, ApJ, 422, L83

Kaspi V. M., Lackey J. R., Mattox J., Manchester R. N., Bailes M., Pace R., 2000a, ApJ, 528, 445

Kaspi V. M. et al., 2000b, ApJ, 543, 321

Kaspi V. M., Taylor J. H., Ryba M., 1994, ApJ, 428, 713

Kedziora-Chudczer L., Jauncey D. L., Wieringa M. H., Walker M. A., Nicolson G. D., Reynolds J. E., Tzioumis A. K., 1997, ApJ, 490, 9

Kirkpatrick S., Stoll E., 1981, J. Comp. Phys., 40, 517

Konacki M., Wolszczan A., Stairs I. H., 2003, ApJ, 589, 495

Kopeikin S. M., 1995, ApJ, 439, L5

Kopeikin S. M., 1996, ApJ, 467, L93

Kramer M., Xilouris K. M., Camilo F., Nice D., Lange C., Backer D. C., Doroshenko O., 1999, ApJ, 520, 324

Kramer M. et al., 2003, MNRAS, 342, 1299

Kramer M., 1998, ApJ, 509, 856

Lai D., Rafikov R. R., 2005, ApJ, 621, 41

Lange C., Camilo F., Wex N., Kramer M., Backer D., Lyne A., Doroshenko O., 2001, MNRAS, 326, 274

Large M. I., Vaughan A. E., Wielebinski R., 1968, Nature, 220, 753

Levenberg K., 1944, Quart. Appl. Math., 2, 164

Livingstone M. A., Kaspi V. M., Gavriil F. P., Manchester R. N., 2005, ApJ, 619, 1046

Livingstone M. A., Kaspi V. M., Gavriil F. P., 2005, ApJ, 633, 1095

Loehmer O., Kramer M., Driebe T., Jessner A., Mitra D., Lyne A. G., 2004, A&A, 426, 631

Lorimer D. R., Bailes M., Dewey R. J., Harrison P. A., 1993, MNRAS, 263, 403

Lorimer D. R. et al., 1995a, ApJ, 439, 933

Lorimer D. R., Yates J. A., Lyne A. G., Gould D. M., 1995b, MNRAS, 273, 411

Lorimer D. R., Lyne A. G., Bailes M., Manchester R. N., D'Amico N., Stappers B. W., Johnston S., Camilo F., 1996, MNRAS, 283, 1383

Lyne A. G., Smith F. G., 2004, Pulsar Astronomy, 3rd ed. Cambridge University Press, Cambridge

Lyne A. G., Pritchard R. S., Graham-Smith F., Camilo F., 1996, Nature, 381, 497

Lyne A. G. et al., 1998, MNRAS, 295, 743

Lyne A. G. et al., 2004, Science, 303, 1153

Lyne A. G., Manchester R. N., Taylor J. H., 1985, MNRAS, 213, 613

Lyne A. G., Pritchard R. S., Smith F. G., 1988, MNRAS, 233, 667

Lyne A. G., 1971, MNRAS, 153, 27P

Manchester R. N., Peters W. L., 1972, ApJ, 172, 221

Manchester R. N., Peterson B. A., 1989, ApJ, 342, L23

Manchester R. N., Taylor J. H., 1977, Pulsars. Freeman, San Francisco

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

Manchester R. N., Lyne A. G., Robinson C., D'Amico N., Bailes M., Lim J., 1991, Nature, 352, 219

Manchester R. N. et al., 1996, MNRAS, 279, 1235

Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, AJ, 129, 1993

Manchester R. N. et al., 2005, ApJ, 621, L49

Marquardt D., 1963, SIAM J. Appl. Math., 11, 431

Matsakis D. N., Taylor J. H., Eubanks T. M., 1997, A&A, 326, 924

McLaughlin M. A. et al., 2004, ApJ, 613, L57

Morris D. J. et al., 2002, MNRAS, 335, 275

Narayan R., Vivekanand M., 1983, A&A, 122, 45

Narayan R., 1992, Philos. Trans. Roy. Soc. London A, 341, 151

Navarro J., 1994, PhD thesis, California Institute of Technology

Ord S. M., Bailes M., van Straten W., 2002, MNRAS, 337, 409

Ord S. M., Bailes M., van Straten W., 2002, ApJ, 574, L75

Ord S. M., van Straten W., Hotan A. W., Bailes M., 2004, MNRAS, 352, 804

Ord S. M., Jacoby B. A., Hotan A. W., Bailes M., 2006, MNRAS, accepted for publication

Ostriker J. P., Gunn J. E., 1969, ApJ, 157, 1395

Pacini F., 1968, Nature, 219, 145

Paczyński B., 1990, ApJ, 348, 485

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

Radhakrishnan V., Shukre C. S., 1985, in Srinivasan G., Radhakrishnan V., eds, Supernovae, Their Progenitors and Remnants. Indian Academy of Sciences, Bangalore, p. 155

Ramachandran R., Kramer M., 2003, A&A, 407, 1085

Ransom S. M., Kaspi V. M., Ramachandran R., Demorest P., Backer D. C., Pfahl E. D., Ghigo F. D., Kaplan D. L., 2004, ApJ, 609, L71

Ransom S. M., Hessels J. W. T., Stairs I. H., Freire P. C. C., Camilo F., Kaspi V. M., Kaplan D. L., 2005, Science, 307, 892

Reichley P. E., Downs G. S., Morris G. A., 1970, ApJ, 159, 35

Riess A. G. et al., 2005, ApJ, 627, 579

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

Ryba M. F., Taylor J. H., 1991, ApJ, 371, 739

Sakurai I., Kawai N., Negoro H., Torii K., Saito Y., Shibata S., 1999, Astronomische Nachrichten, 320, 339

Sakurai I., Kawai N., Torii K., Negoro H., Nagase F., Shibata S., Becker W., 2001, PASJ, 53, 535

Sandhu J. S., Bailes M., Manchester R. N., Navarro J., Kulkarni S. R., Ander-

son S. B., 1997, ApJ, 478, L95

Sayer R. W., Nice D. J., Taylor J. H., 1997, ApJ, 474, 426

Segelstein D. J., Rawley L. A., Stinebring D. R., Fruchter A. S., Taylor J. H., 1986, Nature, 322, 714

Shabanova T. V., Lyne A. G., Urama J. O., 2001, ApJ, 552, 321

Shemar S. L., Lyne A. G., 1996, MNRAS, 282, 677

Shibazaki N., Murakami T., Shaham J., Nomoto K., 1989, Nature, 342, 656

Shklovskii I. S., 1970, Sov. Astron., 13, 562

Smith F. G., 1968, Nature, 218, 325

Splaver E. M., Nice D. J., Stairs I. H., Lommen A. N., Backer D. C., 2005, ApJ, 620, 405

Stairs I. H., Arzoumanian Z., Camilo F., Lyne A. G., Nice D. J., Taylor J. H., Thorsett S. E., Wolszczan A., 1998, ApJ, 505, 352

Stairs I. H., Thorsett S. E., Taylor J. H., Wolszczan A., 2002, ApJ, 581, 501

Stairs I. H. et al., 2005, ApJ, 632, 1060

Stairs I. H., Lyne A. G., Shemar S., 2000, Nature, 406, 484

Stairs I. H., Thorsett S. E., Arzoumanian Z., 2004, Phys. Rev. Lett., 93(14), 141101

Stairs I. H., Thorsett S. E., Camilo F., 1999, ApJS, 123, 627

Stairs I. H., 2004, Science, 304, 547

Staveley-Smith L. et al., 1996, PASA, 13, 243

Stinebring D. R., Ryba M. F., Taylor J. H., Romani R. W., 1990, Phys. Rev. Lett., 65, 285

Stinebring D. R., Smirnova T. V., Hankins T. H., Hovis J., Kaspi V., Kempner J., Meyers E., Nice D. J., 2000, ApJ, 539, 300

Stone E. H., 1925, Man, 25, 69

Sutaria F. K., Ray A., Reisenegger A., Hertling G., Quintana H., Minniti D., 2003, A&A, 406, 245

Taylor J. H., Cordes J. M., 1993, ApJ, 411, 674

Taylor J. H., Weisberg J. M., 1982, ApJ, 253, 908

Taylor J. H., Weisberg J. M., 1989, ApJ, 345, 434

Taylor J. H., 1992, Philos. Trans. Roy. Soc. London A, 341, 117

Toscano M., Britton M. C., Manchester R. N., Bailes M., Sandhu J. S., Kulkarni S. R., Anderson S. B., 1999, ApJ, 523, L171

Toscano M., Sandhu J. S., Bailes M., Manchester R. N., Britton M. C., Kulkarni S. R., Anderson S. B., Stappers B. W., 1999, MNRAS, 307, 925

van Straten W., Bailes M., Britton M., Kulkarni S. R., Anderson S. B., Manchester R. N., Sarkissian J., 2001, Nature, 412, 158

van Straten W., 2003, PhD thesis, Swinburne University of Technology

van Straten W., 2004, ApJS, 152, 129, preprint, astro-ph/0401536

van Straten W., 2006, astro-ph/0510334

Walker M. A., 1998, MNRAS, 294, 307

Weisberg J. M., Taylor J. H., 2002, ApJ, 576, 942

Weisberg J. M., Romani R. W., Taylor J. H., 1989, ApJ, 347, 1030

Wex N., Kalogera V., Kramer M., 2000, ApJ, 528, 401

Wex N., 1997, A&A, 317, 976

Wex N., 2000, in Kramer M., Wex N., Wielebinski R., eds, Pulsar Astronomy - 2000 and Beyond, IAU Colloquium 177. Astronomical Society of the Pacific, San Francisco, p. 113

Wolszczan A., 1990. IAU Circ. No. 5073

Xilouris K. M., Kramer M., Jessner A., von Hoensbroech A., Lorimer D., Wielebinski R., Wolszczan A., Camilo F., 1998, ApJ, 501, 286

# Appendix A

# Computing Arrival Times

Computing the arrival time of a pulse involves determining the relative phase shift between the observed profile and a standard template profile (the phase of which provides the fiducial point). Taylor (1992) describes this process in detail and outlines a method for measuring the relative shift between two pulse profiles in the frequency-domain. This method takes advantage of the Fourier shift theorem (Bracewell 1965), which states that any time-domain shift $\tau$ will manifest itself in the frequency-domain as the slope of a linear phase gradient. Taylor (1992) presents two equations (A7 and A8) that can be solved iteratively to determine $\tau$. These equations include a summation over all frequency components up to the Nyquist limit, which leads to a number of practical problems (described below). The mathematical expression derived for $\tau$ does not in any way depend on the original sampling time, meaning that under optimal conditions, the algorithm can be used to determine the relative phase shift to less than a single bin of precision. Taylor (1992) noted that in practice it is not possible to measure the relative shift to better than approximately 10% of a single phase bin. In recent years, RMS timing residuals as small as 3% of a phase bin have been achieved (van Straten et al. 2001) using baseband recording, however it remains the case that few pulsars can be timed to this level of precision.

Data used for pulsar timing vary greatly in S/N according to the specifications of the observing system and the intrinsic flux density of the source. In addition, each pulsar has a different characteristic profile shape. Timing algorithms must therefore be able to function within a wide parameter space. To contrast with the frequency-domain method, we developed a purely time-domain algorithm that computes the discrete cross-correlation function of two profiles and fits a Gaussian curve to a phase range (0.0125 units by default) surrounding the peak. The centroid of the fitted Gaussian is taken to be the relative phase shift, sub-bin precision is effectively obtained by interpolation. Fitting is performed using generalised code based on the Levenberg-Marquardt (Levenberg 1944; Marquardt 1963) scheme. The error associated with the shift measurement is the distance the centroid of the Gaussian must be displaced to double the $\chi^2$ of the fit. Note that the error estimate obtained using this method is not a conventional $1\sigma$ error, the required scale factor

was determined by experiment.

We note that the idea of time-domain fitting is by no means new to pulsar timing. Our method differs from the usual approach of fitting a parabola to the peak bin and its single nearest neighbour on either side. By interpolating over a larger number of bins (typically of order 25), we decrease the sensitivity of the algorithm to random amplitude variations, making it more robust and suitable for use even when the S/N of the profile is low.

PSRCHIVE (Hotan, van Straten & Manchester 2004) contains implementations of Taylor's frequency domain algorithm and the Gaussian interpolation scheme described above. We find that the frequency-domain method is best used to fine-tune a shift estimate. If the relative shift is large, the phase gradient in frequency space may not be monotonic, making determination of its slope more difficult. It is therefore necessary to perform a discrete cross correlation of the two profiles and rotate one by the coarse shift so obtained to ensure that both are approximately aligned before moving into the frequency-domain. PSRCHIVE implements this algorithm with an additional loop that successively lowers the cut-off frequency used in the summation from the Nyquist frequency down to the $32^{nd}$ spectral component. The best fit from this set of trials is returned as the final shift estimate.

## A.1 Systematic Errors

It is well-known that RMS pulse timing residuals are often inconsistent with the individual pulse arrival time uncertainties. This is usually attributed to "systematic errors" of unknown origin that somehow contaminate the mean pulse profile. To combat this, an arbitrary scaling of errors is often used to bring the reduced $\chi^2$ of any subsequent model fit closer to unity. We bring to light two hitherto unknown regimes where standard analyses would cause incorrect estimates of both arrival times and derived parameters. The first is when the S/N of an individual observation is low or the pulse duty cycle is large, and the second is when a standard profile is created from the sum of several low S/N observations that are then timed against this summation.

Although both the time- and frequency-domains contain the same information, the two methods described above have quite different susceptibility to systematic errors. Using simulated data, we tested the behaviour of these two shift algorithms in different S/N regimes. For simplicity, we simulate a Gaussian profile morphology; a standard template with a centroid of 0.5 phase units was constructed synthetically and given a small amount of unique, random noise, obtained using the R250 random number generator (Kirkpatrick & Stoll 1981). We also constructed an ensemble of test profiles that were given unrelated noise (according to the desired S/N level) and shifted randomly about phase 0.5 by changing the centroid of the Gaussian function.

Our initial tests were performed using a pulse duty cycle of 0.2, roughly similar to that of the two components present in the mean profile of PSR J0737–3039A. We

find that the frequency-domain algorithm under-estimates the reported error in the shift quite significantly. In all cases, the estimate converges to the true error as the S/N increases; however even at a S/N of 50, the mean reported error is a factor of 4 smaller than the mean true error. In contrast, the mean error estimate obtained by fitting in the time-domain is correct to within a factor of 0.5 even when the S/N is as low as 10. The performance of the frequency-domain algorithm improves as the duty cycle of the synthetic Gaussian profiles is reduced. If the pulse is less than 0.05 phase units wide, the mean error estimate returned by this method is also within a factor of 0.5 of the true mean error when the S/N is of order 10.

When the noise level becomes a significant fraction of the total amplitude in any given pulse phase bin, the frequency-domain algorithm performs poorly. Since the Fourier transform of a Gaussian profile with width $w$ is a Gaussian with width $1/w$, it is clear that a wide Gaussian profile contains no high-frequency spectral components. Noise is inherently "white" and contains spectral power at all frequencies, thus the high-frequency spectral components of a wide Gaussian profile are dominated by its baseline noise. This is detrimental to the performance of the frequency-domain algorithm, which uses high-frequency spectral components to extend the lever-arm that defines the slope of the phase gradient. The time-domain algorithm gives reliable shift estimates even when there is no narrow structure in the pulse profile (or the S/N is low) but offers less precision when the S/N is high. If the frequency spectra of the template and the observed profile are dominated by the structure of the pulse itself, the frequency-domain algorithm can provide very precise shift estimates. Any narrow features in the mean profile contribute additional power to the high-frequency spectral components, increasing the effectiveness of the algorithm.

We also investigated a scenario in which the standard template profile is constructed from a summation of the profiles whose arrival times are being measured, as is often the case in practice. While this is the best way to accurately characterise the shape of the mean pulse, it has the side effect that noise in an observed profile may correlate with its own contribution to the noise in the template baseline. We investigated the problem by constructing a template from half the simulated profiles in a given ensemble. The simulated profiles were constructed with a S/N of 14 and the first half of the ensemble was copied and integrated to form a high S/N template. To best illustrate the danger of "self-standarding", we analytically shifted the centroid of each Gaussian profile in the ensemble (leaving the baseline noise fixed), using a linear ramp with a slope of 1/256 bins per unit ensemble index. We then used the frequency-domain algorithm to estimate the shift and its uncertainty for each profile in the ensemble.

Figure A.1 shows the analytic shift applied to each profile in the 512 member ensemble and the shift returned by the frequency-domain algorithm, with the corresponding error estimate. The shifts reported in the first half of the ensemble cluster around zero instead of the analytic linear ramp. Profiles in the latter half of the ensemble return to the linear ramp. It is clear that the "scatter" about the mean is smaller in the first half of the ensemble, showing that the RMS deviation is

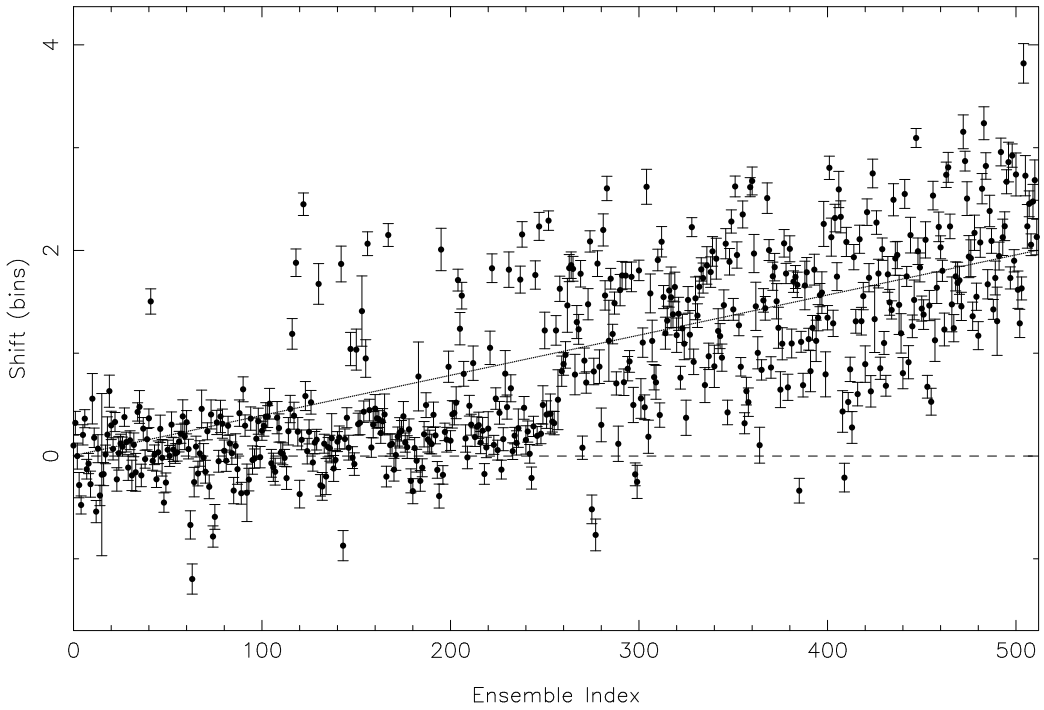


**Figure A.1: Results of a single simulated timing experiment. Each profile in the ensemble of 512 was constructed with a S/N of 14. The first 256 profiles were summed to form a standard template. The diagonal dashed line represents the analytic shift given to each profile and the data points represent the shift reported by the frequency-domain algorithm. This algorithm does not properly characterise the true analytic shift when the ensemble profile has made a significant contribution to the template.**

artificially reduced when "self-standarding" occurs. This problem diminishes with increasing S/N and is insignificant once each of the observed profiles has a S/N of $\sim$25. The time domain algorithm does not suffer from the same problem, even when the observed profiles have S/N $< 10$. The noise present in a given observed profile imprints itself on the standard template; small amounts of high-frequency information are preserved even after the addition of many other observed profiles. The time-domain algorithm is less sensitive to high-frequency structure. In general, only the first $\sim$10% of the spectral components present in a given profile are dominated by the structure of the pulse. Including higher frequency components when fitting for the relative shift in the frequency-domain can clearly lead to systematic errors.

In summary, the mean error estimate returned by the frequency-domain shift algorithm is only physically valid if the spectra of both the standard template and the observed profiles are dominated by structural information. It is also possible to measure an artificially low RMS timing residual if the standard template profile was constructed by summing together all the observed profiles. The correlation between baseline noise in an observed profile and its own imprint on the standard template can be strong enough to overwhelm the true shift. Care should be taken to ensure that reported RMS residuals are not dominated by this effect. "Self-standarding" can be alleviated by reducing the number of spectral components used when fitting for the linear phase gradient in the frequency-domain. The same effect can be achieved by applying a low-pass filter to the spectrum of the standard template. Using a synthetic, analytically constructed standard template would also prevent spurious correlations, but may not provide a good match to the morphology of the mean pulse.

# Appendix B

# CPSR2 Software Architecture

The CPSR2 data acquisition and communications code was written by Willem van Straten and Steve Ord in the months before August 2002. It resides in the Swinburne pulsar software repository, in the directory soft_swin/online/cpsr2/das. Software for monitoring the cluster, reporting its status back to the observer and controlling the operation of the recorder resides in soft_swin/online/cpsr2/GUI.

## B.1 Control System Components

Many of the essential control system tasks run as "daemons"; jobs that have no controlling terminal and perform their tasks in the background. Often, these daemons can only be accessed via network sockets. Figure B.1 shows the software hierarchy in more detail.

CPSR2 was designed to be highly modular; each hardware device had one or more associated control daemons which could be queried or issued commands via a network socket. This reduced the complexity of any individual component, but also demanded a complex network of interprocess communication. The central control daemon, cpsr2d, ran on the gateway machine in the Parkes control room, with network access to the cluster and the outside world. It accepted user input from a number of sources and propagated the correct sequence of commands throughout the rest of the system in order to start or stop an observation, reset the system, alter a header parameter and so on. The EDT–60 DMA cards were controlled by the cpsr2_dmad daemon that ran on both primary nodes. In addition, the second primary node ran a daemon called cpsr2_ffdd that communicated via a serial interface to the FFD. The secondary nodes each ran two data-recording processes, cpsr2_recv and cpsr2_dbdisk, either in the foreground or the background. Their purpose was to accept data over Gb Ethernet, buffer it in RAM and then write it to disk. All machines ran the ekg_daemon, which gathered system information and reported it to any other program via a network connection across the 100 Mb Ethernet layer. The secondary nodes also ran the dm_daemon, which controlled data reduction and catalogued the files for quick reference. All daemons ran in the background, the user

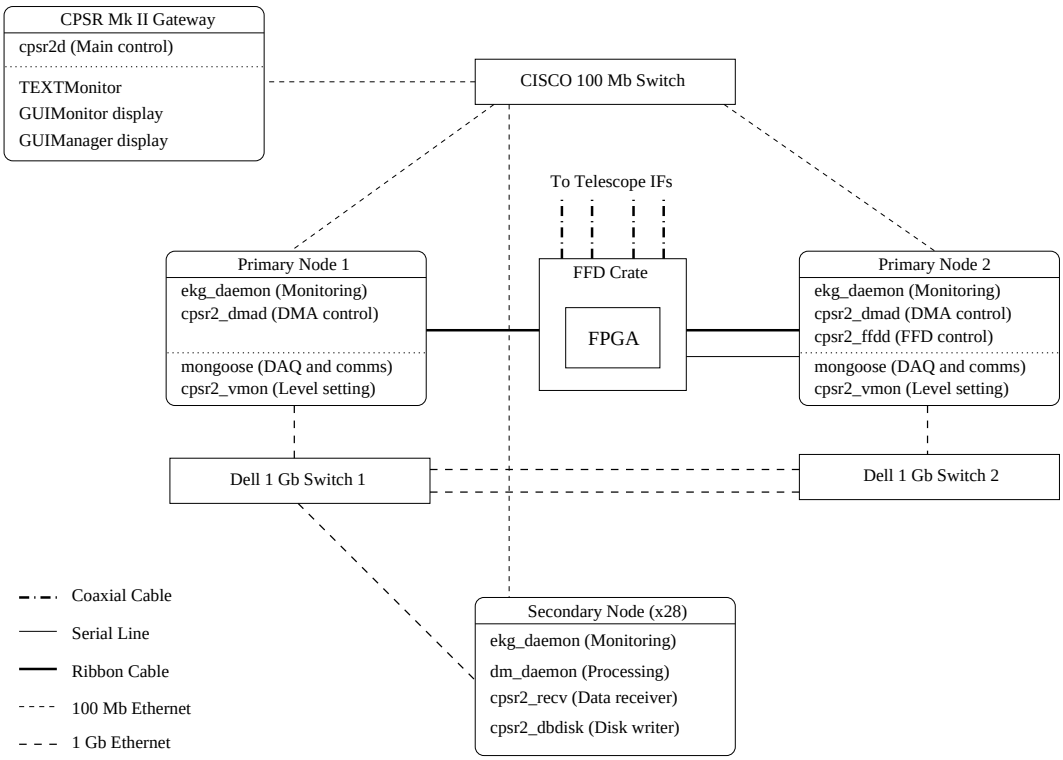


**Figure B.1: This diagram shows the primary hardware components of the CPSR2 instrument and lists the software programs that were run on each machine. Programs listed above the dotted line ran in daemon mode, hidden from the user. Programs below this line required terminal I/O and generally provided some sort of feedback to the observer. The gateway ran a daemon called `cpsr2d` that was responsible for interpreting high-level user input, either directly via a specially designed graphical user interface or through `TCS`. Commands sent to `cpsr2d` were propagated to the appropriate network sockets on the primary nodes, which initiated data recording and network communications. Each machine ran the `ekg_daemon` monitoring process, which could be queried over the 100 Mb Ethernet layer.**

did not normally need to interact with them in any way. However, each daemon also had an ASCII command line interface that could be accessed via a telnet session, for low-level debugging and control.

The primary data acquisition system was named mongoose, and ran within a terminal on each of the primary nodes, so that status and debugging information were visible to the observer. The data path made extensive use of shared memory buffers. These buffers could be a significant fraction of the total extended memory on an individual machine and had to be allocated within the the operating system before the control programs were started. The sequence of commands used to allocate shared memory buffers and run the programs that attach to them was stored in a script that the ekg_daemons could access. Because a mongoose crash would leave the shared memory status flags (semaphores) in an inconsistent state, it was often necessary to rebuild these memory buffers as part of the crash recovery process.

To begin an observation, it was first necessary to ensure that each machine in the cluster (excluding the gateway) was running an ekg_daemon. The control process cpsr2d was run on the gateway machine. Once started, these daemons helped automate the remaining setup tasks. The manual control interface, GUIMonitor, was run on one of the secondary nodes, from a terminal opened on the gateway. While running the monitor on the gateway itself would simplify the process, this was impossible because it directly tests the health and latency of the Gb Ethernet layer, which was not physically connected to the gateway machine. Once started, the GUIMonitor would spend about 10 seconds polling the cluster nodes and then report back their status. To ensure a clean start, it was usually necessary to perform a "master reset", using the supplied button on the GUI. This ran through a scripted sequence of commands, shutting down any lingering jobs on the secondary nodes, removing the shared memory blocks, re-allocating them and finally re-launching all of the communications software. The monitor was not able to start up any program requiring terminal access, so the user had to start a terminal session on both of the primary nodes and run the mongoose. This checked the availability of secondary nodes on the network, the consistency of the internal clock and the status of the primary shared memory buffers and EDT buffers. If all was ready, the mongoose waited on a condition variable that was set when the FFD started recording. The mongoose had a "target list" that the user could edit via the GUIMonitor. The target list specified which machines each primary node would attempt to send data to. This was useful when observing in two widely separated frequency bands, as the processing load associated with coherent dedispersion was a strong function of wavelength (due to the fact that longer wavelengths were dispersed to a greater degree and required longer FFTs). This meant that more cluster nodes had to be allocated to the primary node whose wavelength was longest. When an observation was started, the mongoose terminal displayed an updating time since start and any error messages that were generated. It also briefly displayed the current file header, which could be visually scanned to verify the configuration of the system. The mongoose was also responsible for launching cpsr2_vmon, which extracted a small

number of samples from the data stream to monitor the digitiser statistics. This program opened a graphical display window that plotted the digitiser statistics in real-time.

## B.2 Real-Time Monitoring

Each ekg_daemon maintained an internal database of useful information, including the 5-minute load average of the host machine, the amount of free space on the scratch disk, the number of free buffers in the main shared memory block and the current network activity status. These parameters were made available as simple ASCII strings on request, via a telnet session connected to port 10123 of the host machine. This communications socket was 2-way, the ekg_daemon could parse certain commands and perform tasks on the host machine. These tasks included rebuilding the shared memory space and launching the cpsr2_recv and cpsr2_dbdisk programs. The ekg_daemon also used a set of logical conditions to determine the overall status of the host machine, which could be one of several pre-defined states, including OFFLINE, ONLINE, RECEIVING and so on. This status variable was queried by the mongoose before initiating a data connection, to ensure that the machine was capable of receiving data.

The GUIMonitor was built using the Trolltech Qt[1] widget set. It had a number of tabbed screens that reported the essential parameters of the cluster (obtained from the ekg_daemons) and could also be used to start and stop observations or manually enter the header parameters that cpsr2d inserted into each file. Alternatively, TCS could be used to control CPSR2 in schedule mode, reducing the amount of use intervention required. However, the communications link between TCS and cpsr2d was not bi-directional. If CPSR2 crashed, TCS had to be interrupted manually or it would continue sending commands.

Most serious errors in the recording process resulted in immediate termination of the mongeese. The code was written in this way to ensure the integrity of any data ending up on disk, by stopping the recorder at the first sign of trouble. The GUIMonitor was compatible with the University of Edinburgh's Festival[2] Speech Synthesis System. The gateway hosted a Festival server which spoke any text string sent to it by the GUIMonitor. Numerous tests on the status of the recording system were performed at regular intervals and if a problem was detected, the user was warned with an informative spoken message.

Data processing was managed via a graphical user interface called GUIManager. This was analogous to (and based on the same code as) the GUIMonitor. The dm_daemon process was run on each of the secondary nodes, to construct a record of the files that had been written to disk. The information could be accessed directly via a telnet session to port 10125 on the host machine, but in most cases all

[1] http://www.trolltech.com/products/qt
[2] http://www.cstr.ed.ac.uk/projects/festival/

communication was handled by the GUIManager. The dm_daemons were capable of launching PSRDISP inside a thread that was monitored until completion, at which point a set of scripted commands were run to place the resulting folded archive in the correct location in a central storage area. The data partition on each secondary node was exported via a network file system (NFS) link, as was the central archive storage partition (which physically resided in the first primary node). All the other machines could then automount whatever disks they required on demand.

The number of PSRDISP processing threads allowed to co-exist simultaneously was user-configurable, though running any more than two simultaneously would overload the secondary nodes. By selecting the observation of interest from the list presented within GUIManager, it was possible to obtain a list of individual data files still on the cluster, their distribution across the nodes and a measure of the disk space they occupied. PSRDISP could then be run on the selected observation. The user could directly specify the desired spectral resolution and number of phase bins produced. They could also customise the length of the first FFT in order to optimise processing speed. An ASCII database containing the best processing parameters for each regularly observed pulsar was kept on the first primary node and parsed by the GUIManager. This information was used to present a set of default parameters whenever a new job was initiated. If the current source was not in the database, a conservative fall-back strategy was suggested.

## B.3 Current Status

The tools described above were in use until the end of 2004. Subsequent developments intended to simplify the system for new observers made some of these programs redundant. The CPSR2 control system continues to evolve and observers searching for up-to-date instructions should refer to the Swinburne on-line manual which can be found at the URL: http://astronomy.swin.edu.au/pulsar/observing/cpsr2.

# Appendix C

# Publications

The following publications appear as chapters in this thesis:

**PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis,** Hotan, A. W.; van Straten, W.; Manchester, R. N. *Publications of the Astronomical Society of Australia,* Volume 21, January 2004, Issue 3, pp. 302-309.

**PSR J1022+1001: Profile Stability and Precision Timing,** Hotan, A. W.; Bailes, M.; Ord, S. M. *Monthly Notices of the Royal Astronomical Society,* Volume 355, December 2004, Issue 3, pp. 941-949.

**Geodetic Precession in PSR J1141-6545,** Hotan, A. W.; Bailes, M.; Ord, S. M. *The Astrophysical Journal,* Volume 624, May 2005, Issue 2, pp. 906-913.

**PSR J0737–3039A: Baseband Timing and Polarimetry,** Hotan, A. W.; Bailes, M.; Ord, S. M. *Monthly Notices of the Royal Astronomical Society,* Volume 362, October 2005, Issue 4, pp. 1267-1272.

**High-Precision Baseband Timing of 15 Millisecond Pulsars,** Hotan, A. W.; Bailes, M.; Ord, S. M. *Monthly Notices of the Royal Astronomical Society,* Volume 369, July 2006, Issue 3, pp. 1502-1520.

The following papers were prepared with the assistance of the candidate during the course of this thesis:

**Self-Consistency of Relativistic Observables with General Relativity in the White Dwarf-Neutron Star Binary PSR J1141-6545,** Bailes, M.; Ord, S. M.; Knight, H. S.; Hotan, A. W. *The Astrophysical Journal,* Volume 595, September 2003, Issue 1, pp. L49-L52.

**PSR J1909-3744: A Binary Millisecond Pulsar with a Very Small Duty Cycle,** Jacoby, B. A.; Bailes, M.; van Kerkwijk, M. H.; Ord, S.; Hotan, A.; Kulkarni, S. R.; Anderson, S. B. *The Astrophysical Journal,* Volume 599, December 2003, Issue 2, pp. L99-L102.

**Polarimetric profiles of 27 millisecond pulsars,** Ord, S. M.; van Straten, W.; Hotan, A. W.; Bailes, M. *Monthly Notices of the Royal Astronomical Society,* Volume 352, August 2004, Issue 3, pp. 804-814.

**The Mean Pulse Profile of PSR J0737-3039A,** Manchester, R. N.; Kramer, M.; Possenti, A.; Lyne, A. G.; Burgay, M.; Stairs, I. H.; Hotan, A. W.; McLaughlin, M. A.; Lorimer, D. R.; Hobbs, G. B.; Sarkissian, J. M.; D'Amico, N.; Camilo, F.; Joshi, B. C.; Freire, P. C. C. *The Astrophysical Journal,* Volume 621, March 2005, Issue 1, pp. L49-L52.

**Long-Term Variations in the Pulse Emission from PSR J0737-3039B,** Burgay, M.; Possenti, A.; Manchester, R. N.; Kramer, M.; McLaughlin, M. A.; Lorimer, D. R.; Stairs, I. H.; Joshi, B. C.; Lyne, A. G.; Camilo, F.; D'Amico, N.; Freire, P. C. C.; Sarkissian, J. M.; Hotan, A. W.; Hobbs, G. B. *The Astrophysical Journal,* Volume 624, May 2005, Issue 2, pp. L113-L116.

**The Mass of a Millisecond Pulsar,** Jacoby, B. A.; Hotan, A.; Bailes, M.; Ord, S.; Kulkarni, S. R. *The Astrophysical Journal Letters,* Volume 629, August 2005, Issue 2, pp. L113-L116.

**High-Precision Timing of PSR J1600-3053,** Ord, S. M.; Jacoby, B. A.; Hotan, A. W.; Bailes, M. *Monthly Notices of the Royal Astronomical Society,* Accepted for publication.