HIGH PRECISION DUAL FREQUENCY TIMING OF MILLISECOND PULSARS

Thesis by

Jagmit S. Sandhu

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California
2001
(Defended September 28, 2000)

The fact of the matter is that I do not really have a recipe for the corn cake I bake. Dona Alda, the wife of Mr. Renato of the Museum, told me how to do it, and in that way I learned, racking my brains until I got it right. (Was it not by loving that I learned to love ? Was it not by living that I learned to live ?) ...

Dona Flor in "Dona Flor and her two husbands" by Jorge Amado

Abstract

The science of precision millisecond pulsar timing can yield the most precise astrometric measurements ever made. This potential can only be realized through an extraordinary amount of investment in the experimental apparatus, the effort of many observers, and the close attention to details required to avoid the many pitfalls along the way. This thesis describes the work and results from a precision timing project aimed at monitoring the brightest millisecond pulsar, PSR J0437-4715. This pulsar is a very suitable target for such a study because of its small period (5.75 ms), low DM (2.69 $\mathrm{pc\,cm^{-3}}$), and high flux density ($\sim$ 90 mJy at 1.4 GHz). The initial work for this thesis involved completion and installation of the Fast Pulsar Timing Machine (FPTM) at Parkes observatory, Australia. With a bandwidth of 128 MHz and time resolution of 4 μs, this machine made a quantum jump in the time of arrival precision for PSR J0437-4715. The precision improved from $\sim$ 2 μs to $\sim$ 0.2 μs. In order to further enhance the signal-to-noise ratio achievable with this pulsar and probe the limits of precision pulsar timing, we have subsequently improved the FPTM significantly by doubling its bandwidth, so that it can record the pulsar radio emission over a 256 MHz bandwidth. This required us to double the IF processing hardware in the FPTM and implement numerous software modifications to control the observing apparatus, interface the FPTM with the observatory control computers, as well as process the data to produce final times of arrival.

Integration of just a few minutes with the 64 m Parkes radio telescope yield times of arrival for PSR J0437-4715 with a precision of 100 nanoseconds. The longer term residuals (3 years) show root-mean-square deviations of 500 nanoseconds. This excessively large scatter in the residuals has been traced to inaccurate polarization calibration, and a systematic quadratic trend of $\sim$ 5 μs in the times of arrival as a function of baseband frequency. The latter phenomenon has been simulated in software and shown to arise from the large dynamic range in the baseband spectrum.

Despite the systematic errors, our measurement of the pulsar's astrometric and binary parameters match the best obtained so far with other millisecond pulsars. This has allowed us to measure the pulsar's parallax, and the secular change in the binary's projected semi-major axis due to the system's proper motion. The latter effect restricts the inclination angle of the binary, $i < 43^\circ$. The parallax, along with the period derivative and orbital period derivative, enable us to constrain the distance of the pulsar, $162 < d < 205$ pc. A stability analysis of the pulsar's time of arrival residuals demonstrates that it matches the long term stability of the best studied millisecond pulsars, PSRs B1937+21 and B1855+09, at least on the time scale of the data available so far, 3 years. The precision in the pulsar position now matches the amplitude of the modulation of position expected from the pulsar's binary motion. Detection of this effect will require reduction of the calibration and spectral shape errors, as well as further refinements in the timing software used. Along with radio interferometric observations of the pulsar and optical detection of the white-dwarf companion, the pulsar timing position will provide the best contraints for frame tie between the ecliptic and extragalactic reference frames.

Acknowledgements

Blame it on Halley's comet. If it hadn't swung by the Earth's neighbourhood during my formative high school years ... If I hadn't broken the piggy bank to buy that 2.5 inch reflector ... If I hadn't found an old, dusty copy of Burnham's Celestial Handbook in the local library ... If my parents hadn't bought the Sky atlas 2000 on my insistence ... I wouldn't have learnt so many new things and had as much fun as I did during graduate school ! The journey from those amateur days to the ranks of aspiring professional astronomers has certainly been arduous, but would have been much more difficult had it not been for the help and encouragement of many individuals. Now it's time to thank those who helped and to say humble apologies to those I offended.

Professors Mathur and Mani at IIT Kanpur galvanized my enthusiasm with their fine teaching, and encouragement for graduate school. Dr. Ananthakrishnan at GMRT was instrumental in introducing me to Radio Astronomy, and for arranging that summer internship.

When I arrived at Caltech, I was a novice in research, quite ignorant of the current research activities and consequently unsure of what I wanted to do. Shri Kulkarni very quickly attracted me to his research group with his mesmerizing enthusiasm and boundless energy. Many of his students have waited for age to slow him down so that they can keep up with his pace. I don't think I have seen that happen yet. Shri has provided superb guidance, not just in words, but by setting the highest standards for others to follow. His greatest gift, which cannot easily be imitated, is his vision of the scientific possibilities looming over the horizon. The success of the FPTM is a tribute to his vision.

Jose Navarro displayed splendid patience in teaching me the ropes of the FPTM and answering my many questions. I hope some of his great skills in tackling hardware and software problems rubbed off on me. The upgrade to the FPTM might not have happened without his help and for that I have to thank him. John Yamasaki contributed to the completion of the FPTM upgrade both with his engineering skills and his beer bashes. Steve Padin provided invaluable technical guidance. Stuart Anderson de-mystified the inner workings of TEMPO.

Many others have helped in this intercontinental thesis project. Dick Manchester provided able guidance, support and hospitality in Australia. His efforts at pushing me to write that paper and this thesis are appreciated. Mathew Bailes helped by spreading good cheer all around and educating me about the intricacies of Aussie-Rules Football. The efforts of the entire Australian pulsar research group are highly appreciated. The large distance to Parkes precluded my attendance at most of the observing runs. They generously took up the slack. The Parkes support staff were of immense help at all times. I thank John Glowacki and Mike Kesteven for organizing those squash games.

Many thanks are due to some of the fellow denizens of Robinson who made the graduate school years more enjoyable. Gautam Vasisht and John Gizis injected just the right mixture of enthusiasm and cynicism. All the participants of the PJ trips helped with the much needed breaks, and by joining in on the stupid party games.

Jarrett Malone helped me outside graduate school in a new job with his generous help and his incessant proding to "go get that degree."

I don't know where I would be today without the love and support of my lovely wife, Pammi. She put up with my strange hours and my pessimistic dronings but always provided cheerful support. Now that this thesis is done, we can actually hope to lead a 'normal' life!

Last, but not least, my parents were ever present with constant encouragement and support. They seldom chided me for my lackadaisical efforts at writing home from half a world away. Astronomy was probably the last thing they expected one of their kids to study. Nevertheless, they taught us to make our own decisions and provided the freedom to do so. This thesis is dedicated to them, for being such fine shepherds, and to Shaan, to whom the future belongs.

Contents

List of Figures

List of Tables

Chapter 1 Pulsar Timing and its Applications

1.1 Introduction

1.1.1 Questions of Stability

More than 30 years have elapsed since the discovery of the first pulsar (Hewish et al. 1968). That event triggered a flurry of observational and theoretical activity aimed at uncovering the basic physics behind the pulsar phenomenon – their genesis, evolution, the origin of their radio emission and its ultimate cessation. Even though these efforts have greatly added to our knowledge base, many fundamental issues remain unclear and unsettled. These include the basic emisson process, the internal structure of neutron stars (whether it is ordinary nuclear material, meson condensates or the more exotic strange matter), the origin and evolution of their magnetic fields, etc. However, for the present, we can use pieces of the pulsar puzzle as valuable tools to conduct some interesting experiments. One such tool is the exceptional rotational stability of pulsars. In this respect pulsar astronomers have, reluctantly or otherwise, adopted a black box approach to the pulsar phenomenon. They may not understand how pulsars shine, but since their rotation periods have been observed to be extremely stable, they have decided to use that feature to make astrophysically interesting measurements.

Continuing in the above tradition, this thesis is centered on exploiting the exceptional temporal stability of pulsars, especially millisecond pulsars. Pulsar timing is akin to the study of clocks. Imagine being given a clock and asked to characterize its properties. The most fundamental parameter of a clock is its period. Any other information about the clock follows from knowledge about its period (or its inverse, frequency). Is it constant ? If not, how does it vary ? It should be immediately obvious that to study a clock, we need another clock whose properties we know very well. It then follows that we can establish the stability of the unknown clock only as well as our 'standard' clock. This is the essence of timing a pulsar. We model it as an oscillator which outputs a 'tick' once a period and we then compare the interval between these ticks to a standard clock, which in the usual case is an internationally established time scale like UTC.

The first few hundred pulsars discovered were observed to have period stabilities $(\delta P/P) \sim 10^{-12}$, compared to 10^{-13} for atomic clocks of that epoch. Impressive as this was, the really astonishing result in pulsar period metrology was obtained after 1982 with the discovery of PSR B1937+21 (Backer et al. 1982), the first millisecond pulsar. After a few years it was clear that its rotation was stable to $\sim 10^{-14}$, comparable to the stability of atomic clocks (whose performance is being continuously improved with better technology). We now had available a cosmic clock with which to compare the stability of our man made time scales. Many more millisecond pulsars have been discovered since then and a concerted effort undertaken by many researchers all over the world to study their long term spin stability. PSRs B1855+09 and B1937+21 have already demonstrated sub-microsecond timing stability over several years (Kaspi, Taylor, & Ryba 1994). These two pulsars represent the state of the art in millisecond pulsar timing and therefore it is instructive to look at the results obtained from their timing to help us chart the future for this field.

PSR B1937+21 has excellent short term precision because of its relatively high flux density ($\sim$ 10 mJy at 1400 MHz). Observations at Arecibo observatory typically produce 200 nanosecond residuals after a few minutes of integration. However, its long term timing residuals after 8 years of timing show excessive low frequency noise, in the form of a cubic polynomial with an amplitude of $\sim 7\ \mu$s. It is tempting to blame the pulsar for this behaviour. We can find some support for this conjecture from observations of long period pulsars. Arzoumanian et al. (1994) studied the long term stability of some slow pulsars and defined a stability parameter, Δ_8, equal to the base-10 logarithm of the cumulative pulsar phase error (in seconds) over a reference time interval of 10^8 seconds. They could fit their data with the relation:

$$\Delta_8 \equiv 6.6 + 0.6\, log\, \dot{P}, \tag{1.1}$$

where $\dot{P}$ is the pulsar period derivative. Although this defining relationship was derived from slow pulsars, the measured low frequency noise of PSR B1937+21 falls close to the value predicted by equation 1.1. Thus if the low frequency noise is indeed due to the pulsar, it would be consistent with the behaviour of slower pulsars.

However, another equally likely suspect is the interstellar medium (ISM). This is because the pulsar has a relatively high dispersion measure (DM) of 71.040 pc cm^{-3}, placing it at a nominal distance of 3.6 kpc. In fact, significant variations of the dispersion measure to PSR B1937+21 were reported with the timing results. Due to the turbulent nature of the ISM and relative motion between the pulsar and solar system, such variations are to be expected. It is unclear whether we completely understand the full range of effects that the ISM can have on the timing residuals. Another possible culprit is the solar system ephemeris used for transferring the topocentric times of arrival to the solar system barycenter. We lack sufficient precision in the masses of the outer planets and the asteriods to completely rule out their culpability in introducing long term biases in the ephemeris. Thus the nature of the low frequency noise of PSR B1937+21 is shrouded in doubt.

PSR B1855+09, on the other hand, is a low DM (13.3) pulsar, so we expect the ISM effects to be muted as compared to PSR B1937+21. Its period derivative is also almost an order of magnitude smaller. However, its flux density is low ($\sim$ 4 mJy at 1400 MHz). Consequently its timing precision is worse by about a factor of 4. Thus we probably lack the precision to test its low frequency noise, and indeed none has been observed so far. Clearly it also does not test our time keeping capability as well as PSR B1937+21.

It should now be clear that the main ingredient to push the current state of the art of pulsar timing is a bright, low DM, low period-derivative, short-period pulsar. This wish was fulfilled with the discovery of PSR J0437-4715 (Johnston et al. 1993). It has one of the lowest DMs (2.69), is very bright, $\sim$ 90 mJy at 1400 MHz, and has a low period derivative, $\sim 10^{-20}$ (compared to $\sim 10^{-19}$ for PSR B1937+21). Unfortunately, its presence at a high southern declination precluded observations from Arecibo observatory (which has been used to follow PSRs B1937+21 and B1855+09), and only allows observations with the smaller Parkes radio telescope. There was also an absence of observing equipment with the high time and bandwidth resolution required to fully face the timing challenges presented by PSR J0437-4715. This prompted an effort to build a new pulsar observing spectrometer by the Caltech pulsar group, which has been described in Navarro (1994). The work of this thesis included commissioning of this machine, enhancing its capabilities and analyzing the results for PSR J0437-4715 (and a few other millisecond pulsars). We next discuss some of the scientific motivation for this work, implementation and applications of pulsar timing (section 1.2 and 1.3), and justification for the technical enhancements carried out (section 1.4). Chapter two describes the hardware and software work accomplished during the course of this thesis. We subsequently present our results for PSR J0437-4715 in chapter three, followed by results for some other millisecond pulsars in chapter four.

1.1.2 Origin and Evolution of Millisecond Pulsars

Long term stability is not the only motivation for timing millisecond pulsars. This technique also furnishes many clues about their origin and evolution. It is necessary to outline the pertinent questions about these issues to aid us later on in interpreting the results obtained from pulsar timing.

The most physically plausible model for a pulsar is that it is a $1.4\,\mathrm{M}_{\odot}$ neutron star of radius about 10 km, having a predominantly dipolar magnetic field (which may or may not be aligned with the rotation axis). It is this magnetic field which causes the slow down of the pulsar through emission of magnetic dipole radiation and/or relativistic charged particles from the polar cap region (the exact mechanism is still unknown). From the measured braking torques of all pulsars, the deduced strength of the magnetic field is $\sim 10^{10} - 10^{12}$ G for slow pulsars and $10^8 - 10^9$ G for millisecond pulsars. This bimodal distribution of the magnetic field strengths of the two populations remains a mystery. Most arguments used to explain the deficit in the field strengths of millisecond pulsars start from the nature of their genesis. Two of the popular scenarios can be used to illustrate the divergence of thought about this issue.

The first maintains that millisecond pulsars were born as slow pulsars with normal periods ($\geq$ 10 ms) and magnetic field strengths via the supernova explosion of a high mass star, with the exception that they were in binary systems in which their companion was still on the main sequence. The pulsar is visible in the radio sky for a few million years and then fades away, spinning very slowly. Subsequently the companion evolves away from the main sequence and expands beyond its Roche lobe. Some or all of the matter outside the Roche lobe is accreted by the neutron star. The accretion adds angular momentum, thereby spinning up the neutron star to millisecond periods and it becomes visible again as a pulsar. The companion ends up as a white dwarf which cools away gradually by radiating away its residual thermal energy. This explains systems like PSR J0437-4715.

The second scenario holds that the millisecond pulsar is originally born as a white dwarf in a binary system. As in the previous hypothesis, its less-massive companion evolves sometime later, leading to mass transfer onto the white dwarf. If enough mass is transfered to tip the white dwarf mass over the Chandrasekhar limit ($1.4M_\odot$), it will implode to form a neutron star which may spin at millisecond periods if the white dwarf had sufficient angular momentum (and the neutron star may also be spun up by subsequent accretion). This is called the accretion-induced collapse model. In either scenario, it is quite possible that the accreted matter 'buries' the magnetic field, thus causing the observed deficit. Another possibility is that the magnetic fields of all pulsars decay gradually. Since millisecond pulsars are typically very old, $\sim$ 1–10 Gyr, we may simply be observing the end product of this decay.

There are important clues in the present day millisecond pulsar binaries to help us distinguish the correct evolutionary path. These include space velocities, masses of the neutron stars and their companions, orbital periods and eccentricities. Accretion-induced collapse models predict lower systemic velocities as compared to the model in which the pulsar is formed through a normal supernova explosion – a supernova occurs with a high probability of imparting a large recoil kick to the neutron star/binary system. The exceptional stability of millisecond pulsars permits us to measure their proper motion after just a few years of timing. A reasonable estimate of their distance can be obtained from their dispersion measure (if not measurable through parallax). Thus we can measure the transverse component of their velocity. A large data set of proper motions will therefore furnish an important clue to aid us in deducing the probable formation scenario. The few proper motion measurements available to date indicate that millisecond pulsars on the whole appear to have lower systemic velocities, ~ 100 km/s, than the younger pulsars, (~ 300 km/s). The newer millisecond pulsars discovered with the Parkes telescope are a good target for this purpose.

Neutron star mass measurements are crucial not only to unravel their formation mechanism but also to constrain the equation of state of the dense neutron star matter, about which we know very little. Pulsar timing measures five Keplerian parameters which by themselves are insufficient to measure the stellar masses. However, in some favourable cases – depending on the geometry of the system – a few post-Keplerian parameters can be measured. For instance, in the case of PSR B1855+09, a nearly edge-on circular low-mass binary, it has been possible to measure r and s, the range and shape of the Shapiro delay (Kaspi, Taylor, & Ryba 1994). Assumption of General Relativity theory then yields the orbital inclination and companion mass and the mass for the pulsar can be determined from the mass function. The amplitude of this signal is typically very small – a few microseconds – thus millisecond pulsar binaries are the most suitable candidates for this measurement. When the companion is also a neutron star and the orbit is small and eccentric, relativistic corrections provide enough information to test different theories of gravity and measure the pulsar masses independently of all theoretical assumptions. This has been the case for PSR B1913+16 (Taylor & Weisberg 1989).

Orbital periods, eccentricities and companion masses of millisecond pulsar binaries provide invaluable confirmation for the different formation and evolutionary scenarios. For instance, the fluctuation-dissipation theory of Phinney (1992) asserts a testable relationship between the orbital periods and eccentricities of such binaries (and has no adjustable parameters). This relationship is created by the occurrence of convective eddies in the envelope of the companion star towards the late stages of its evolution. These eddies assert time varying noncentral forces on the neutron star, thus generating significant eccentricity. The expected trend is:

$$\langle e\rangle^{1/2} = 1.5 \times 10^{-4}(P_b/(100\,\mathrm{d})) \tag{1.2}$$

for $P_b > P_{crit} \approx 25$ d. Despite the critical period cut-off, the trend has held up well for orbital periods ≥ 3 d. For shorter periods the expected eccentricities are $\leq 10^{-6}$ which have proven difficult to measure so far due to the lack of bright, short-period binaries. Among the newer millisecond pulsar binaries observable from Parkes, there are a few which are suitable for extending this test to short periods – PSR J0613-0200 (1.2 d), PSR J1911-1114 (2.7 d). There are also a number of longer period binaries whose eccentricities have not yet been well determined because of a lack of S/N. Pursuing a timing program for these systems will therefore be fruitful.

1.2 Brief Overview of Pulsar Timing

The details of pulsar timing can be found in several publications, for example, Manchester & Taylor (1977), Backer & Hellings (1986). I will outline some of the steps here in order to lay the groundwork for the explaining the scientific results accruing from this thesis.

As stressed earlier, in comparing the radio pulses of pulsars to clock ticks, we typically ignore the underlying mechanism which produces the tick – all that concerns us is that the tick be reproducible. There are good reasons to believe that this is the case. Ever since the discovery of pulsars, their pulse properties have been extensively studied. Even though individual pulses are known to vary considerably from each other, if we averge several thousand pulses and obtain a good signal to noise ratio, this average profile is seen to be stable over widely separated epochs. Averaging is important not just to overcome the stochastic noise, but also because some pulsars are known to show some systematic time variable trends in their pulsations - drifting subpulses, mode changing, etc., on short time scales. Fortunately these problems have not yet been seen in millisecond pulsars and therefore need not concern us here.

All pulsars can be modelled as oscillators in the following manner:

$$\phi = \phi_0 + \nu t + \dot{\nu} t^2/2 + ... \tag{1.3}$$

ϕ = pulse phase at time t

ϕ_0 = pulse phase at time t = 0

ν = pulsar spin frequency

$\dot{\nu}$ = pulsar spin frequency time derivative

Usually the frequency and its first derivative are sufficient to model the observed times of arrival. Frequency second derivatives have been measured for only a few pulsars (out of the $\sim$ 1000 known). These are generally so small that they are dominated by timing noise and cannot be used to predict phase reliably.

The basic measurement for pulsar timing analysis is the time of arrival (TOA) of a predetermined pulse phase at any given epoch. The fiducial phase most commonly used is the peak of the pulse profile for simple single pulse profiles, but it can be different for more complicated profiles. The TOA is determined in the following manner: Form a high signal to noise ratio profile by adding observations from several different epochs. This is called the standard profile. At any observation epoch for which a TOA is desired, the observed profile is correlated with this standard profile to determine the relative phase shift between the two. This shift is added to the start time of the observation to determine the TOA. In this discussion it is important to note that all observations must be total intensity profiles (Stokes I) since most pulsars are significantly polarized.

The observed TOAs are for a reference frame fixed to the observatory which is an accelerated frame – it is rotating around the center of the Earth, which itself is revolving around the Earth-Moon barycenter, which in turn is revolving around the Sun. The first step of TOA analysis involves transferring the observatory TOA to an inertial frame. The best realization of an inertial frame for this purpose has been a barycentric frame such as DE200 made available by JPL (Standish 1982) or PEP740R from the Center for Astrophysics. This transformation is as follows:

$$T_b = t_o + \mathbf{r}.\mathbf{n}/c + D/f^2 + \Delta t_{E\odot} + \Delta t_{S\odot} \tag{1.4}$$

T_b = TOA at the Solar system Barycenter
t_o = observatory TOA
$\mathbf{r}$ = barycentric radius vector to the observatory
$\mathbf{n}$ = barycentric unit vector towards the pulsar
c = speed of light
D = accounts for dispersion delay along the line of sight to the pulsar
f = is the observing frequency in the barycentric frame
$\Delta t_{E\odot}$ = Einstein variation of the Earth-based clock as it moves around the Sun
$\Delta t_{S\odot}$ = Shapiro delay of the pulsar signal in the gravitational field of the Sun

The dispersion term is required because the broadband radio frequency signal from the pulsar propagates through an ionized plasma which is frequency dispersive such that we receive the higher frequency components of the pulse earlier than the lower frequencies. The delay in arrival time as a function of frequency (with respect to infinite frequency) is:

$$\Delta t \;=\; \left(\frac{e^2}{2\pi mc}\int_0^d n_e dl\right)/f^2 \;\equiv\; \frac{D}{f^2} \;\equiv\; \frac{DM}{\alpha_d f^2}$$

$$DM \;=\; \int n_e dl \;\; \mathrm{pc\ cm^{-3}}\,, \qquad \alpha_d \;=\; 2.41\times 10^{-16}\ \mathrm{pc\ cm^{-3}\ Hz^{-1}} \tag{1.5}$$

Dispersion measure (DM) is the integrated electron density along the line of sight to the pulsar. α_d is an exact constant because of the definition of dispersion measure.

The above transformation is sufficient if the pulsar is isolated. In the case of a pulsar in a binary system, however, a further transformation is required to compensate for the binary motion:

$$T_{bin} = T_b + f(P_b, T_0, \omega, x, e) \tag{1.6}$$

T_{bin} = binary center of mass TOA
P_b = orbital period
T_0 = Time of periastron passage
ω = longitude of periastron passage
x = the projected semimajor axis $\equiv$ a sin i
e = eccentricity

$f(P_b, T_0, \omega, x, e)$ transforms the TOA to the binary center of mass using a Newtonian model and the five parameters required for that can be measured from the radial velocity variation of the pulsar period. When the pulsar has a heavy companion with an eccentric orbit, relativistic effects in the TOAs can be significant. Many timing models have been devised for such binaries and are used to measure the post-Newtonian effects and thereby provide a test for different theories of gravitation (Taylor & Weisberg 1989). Full exposition of $f(P_b, T_0, \omega, x, e)$ can be found in Backer and Hellings (1986).

Using the observed TOAs to derive a pulsar model is clearly a non-trivial process, and only a few software packages in the world are capable of processing pulsar timing data. The package used most widely (and the one used for this work) is TEMPO (http://pulsar.princeton.edu/tempo/). TEMPO takes the leading TOA in a series of TOAs as the starting pulse phase and then predicts all subsequent TOAs for each of the epochs corresponding to the observed TOAs, using the known pulsar parameters as a starting point. Typically the parameters are not known accurately initially. The errors between the observations and the predictions are known as residuals. These can span many periods when the parameters are known only roughly. Once the parameters have been narrowed down (by iteration) such that the residuals are less than half a period, they can be processed by TEMPO. As can be seen from equations 1.3, 1.4, 1.5 and 1.6, there are numerous factors which affect the residuals. Some of these are intrinsic to the pulsar while others are extrinsic and reflect our imperfect knowledge of the position, orientation and velocity of our observing platform (the solar system), as well as unknown perturbations by the interstellar medium. TEMPO attempts to modify these parameters so as to reduce the summed square of these errors, i.e., it performs a least-squares estimation of the pulsar parameters. Ideally, all the parameters would be calculated with equal precision and would have zero covariance. In practice, the different parameters have widely varying precisions and are covariant to some extent. It is easy to see that this is a result of different phase space coverage for each parameter at each observation epoch. Obviously as long a data set as possible (equally spread out over the possible phase space) is needed to reduce the covariance. TEMPO computes a covariance matrix for each solution and this can be used to judge the relative merit of each solution.

1.3 Applications of Precision Pulsar Timing

Most of the scientifically interesting results central to this thesis depend crucially on the precision of millisecond pulsar timing. It is therefore instructive to understand the inherent precision with which we can time tag a pulsar's TOA. The precision is roughly proportional to the pulse width divided by the signal to noise (S/N) of the profile. Taking a nominal duty cycle of 5%, and an S/N of 100, this implies an accuracy of:

$$\Delta t \;=\; 500\, P_{ms} \left(\frac{S/N}{100}\right)^{-1} \; ns \tag{1.7}$$

where P_{ms} is the pulsar period in milliseconds. Averaging over several epochs can then reduce the overall uncertainity to the level of 100 nanoseconds or less. The quantifiable implications of such accuracy can be appreciated if one realizes that converted into distance units, a precision of 500 ns is equal to a precision in the knowledge of the pulsar-Earth baseline of ~ 150 m (which is $\sim 10^{-17}$ of a distance of 1 kpc). Timing does not of course provide knowledge of any constant relative radial motion between the pulsar and the solar sytem – it is covariant with the pulsar period which is doppler shifted. Keeping the above precision in mind, let us now consider its metrological implications.

1.3.1 Astrometry and Kinematics

The least-squares fitting process in TEMPO solves for all the astrometric parameters for any pulsar – position, proper motion and, more rarely, parallax. The magnitude of the signature of these parameters in the timing residuals is given by:

$$\Delta TOA_{position,\lambda} \sim 2.42 \frac{\Delta\lambda}{1 \text{ mas}} \sin(\lambda)\ \mu s \tag{1.8}$$

$$\Delta TOA_{parallax} \sim 1.2 \frac{\cos^2(\lambda)}{d_{kpc}}\ \mu s \tag{1.9}$$

$$\Delta TOA_{proper\ motion} \sim 24.2 \frac{V_{100}}{d_{kpc}}\ \mu s\, yr^{-1} \tag{1.10}$$

where λ is the ecliptic latitude, $\Delta\lambda$ is the error in ecliptic latitude in mas, V_{100} is the pulsar's transverse speed in units of 100 km/s and d_{kpc} is the distance to the pulsar in kiloparsecs. The position term only describes the variation for ecliptic latitude. A similar term, without the sinusoidal dependence on λ, describes the variation for ecliptic longitude. Note that the proper motion term grows linearly with time while position and parallax have an annual and semi-annual periodic behaviour, respectively.

For most average-brightness millisecond pulsars positions can be measured to a few milliarcseconds, and for the brightest, for example, PSR J0437-4715, to a few tens of microarcseconds. These are among the most precise astrometric measurements ever made by any method, rivaling position determinations by VLBI. For example, the position of PSR B1937+21 has been determined to about 0.1 mas by timing (Kaspi, Taylor, & Ryba 1994), about 3 mas by VLBI (Bartel et al. 1996), and 30 mas with the VLA (Foster & Backer 1990). Such position determinations are crucial for several reasons. Timing positions are determined in the ecliptic reference frame which is based on orbital position of the Earth and hence on the orbits of the moon and planets. Radio interferometer positions are determined with respect to an equatorial reference frame tied to the rotation of the Earth and uses extragalactic radio sources (primarily quasars) as the definition points, and is insensitive to the orbital motion of the Earth. While the two reference frames have a high degree of internal consistency, the transformation between the two is not as good. For example, (Folkner et al. 1994) used lunar laser ranging measurements and VLBI measurements of extragalactic sources to derive a standard error of $\sim$ 3 mas between the two. Reducing this uncertainity is necessary for improvements in Earth orientation studies, better spacecraft navigation, and to uncover errors in modelling of solar system dynamics. Bright millisecond pulsars are some of the rare objects which are observable in both reference frames to high accuracy. Since there are four degrees of freedom in a frame tie (two in position offset and two in relative rotation), and each reference frame can have variable zonal errors, it is necessary to observe as many frame tie objects as possible. PSR J0437-4715, with its high flux density, promises to serve this role well. An added benefit of the frame tie is improved identification of sources in crowded galactic fields.

The small signature of parallax in the timing residuals implies that it can only be measured for nearby millisecond pulsars (d $\leq$ 1 kpc). Parallax obviously provides us with an absolute distance measure to a pulsar. Without parallax we have to rely on an uncertain distance derived from the dispersion measure, for example, the Galactic electron density model of Taylor and Cordes (1993). Though the statistical uncertainity in that model is claimed to be about 25%, individual distance estimates can be in error by as much as a factor of 2.

Measurement of the proper motions of millisecond pulsars is of interest for a variety of reasons. Proper motion measurements provide important confirmation of the dynamical interpretation of pulsar wind nebulae, for example in PSR B1957+20, where the proper motion agrees with the projected motion seen in the optical image. With an estimate of the pulsar distance (either through parallax or dispersion measure), proper motions can be used to derive their transverse velocities. These are considered to be important relics of millisecond pulsar birth events and are used to constrain their formation scenarios. Knowledge of the transverse speed is also necessary for determining better age estimates for pulsars. The characteristic ages of pulsars are defined as $P/2\dot{P}$ (for a braking index of 3). The measured period derivative (and binary period derivative when applicable) of pulsars contains contributions from its transverse velocity and radial acceleration:

$$\delta\dot{P} = \frac{P\,{v_{\perp}}^{2}}{cd} + \frac{P\,a_{||}}{c} \tag{1.11}$$

where $a_{||}$ is the radial acceleration,$v_{\perp}$ is the transverse speed and d is the distance to the pulsar (Shklovskii 1970, Camilo, Thorsett, & Kulkarni 1994). $a_{||}$ can be important for pulsars in globular clusters, or if they are far enough away that differential galactic rotation (between the pulsar and the solar system) is significant. For many pulsars the transverse velocity contribution is quite large and correspondingly decreases the characteristic age significantly (Camilo, Thorsett, & Kulkarni 1994). For most low-mass binary pulsars, their intrinsic $\dot{P_b}$ is negligible. The observed $\dot{P_b}$ can therefore be used to measure their distance (Bell & Bailes 1996).

Recently, Kopeikin (1996) has pointed out some observable effects of proper motion on binary millisecond pulsars. For example, secular changes in the projected semi-major axis and the longitude of orbital periastron should be detectable with precision timing. These effects, when combined with the proper motion measurements, can be used to derive useful limits on some orbital parameters (which otherwise are poorly constrained).

1.3.2 Rate of Decay of Newton's Constant, $\dot{G}$

For low-mass low-eccentricity millisecond pulsar binaries, the orbital period decay rate due to gravitational radiation is expected to be very small. For example, the predicted $\dot{P}_b$ for PSR B1855+09 is $-9 \times 10^{-17}\mathrm{s\,s^{-1}}$. Present day observational accuracy in $\dot{P}_b$ is $\sim 10^{-11}\ \mathrm{s\,s^{-1}}$ after a few years of timing. Since this accuracy improves with time as $T^{5/2}$, it will take several hundred years to detect the effect of gravitational radiation in these systems! On the other hand, time variation of P_b would be a consequence of any time variability of the gravitational 'constant', G. A straightforward limit on $\dot{G}$ can be derived from binary pulsar period stability by using:

$$\frac{\dot{G}}{G} = -\frac{\dot{P}_b}{P_b} \tag{1.12}$$

(Damour, Gibbons, & Taylor 1988). The best limit from binary pulsar data so far is obtained from PSR B1913+16 (Taylor 1993):

$$\frac{\dot{G}}{G} = 4 \pm 5 \times 10^{-12}\ \mathrm{yr}^{-1}, \tag{1.13}$$

but this interpretation is complicated by the fact that both members of this system are neutron stars and the intrinsic $\dot{P}_b$ is significant $\sim -2 \times 10^{-12} ss^{-1}$. Therefore, low-mass binary millisecond pulsars are considered better candidates for this test since they are 'cleaner' and the intrinsic $\dot{P}_b$ is much lower. Kaspi, Taylor and Ryba (1994) used data for the binary millisecond pulsar B1855+09 to derive

$$\frac{\dot{G}}{G} = -9 \pm 18 \times 10^{-12}\ \mathrm{yr}^{-1}. \tag{1.14}$$

Similar limits using data from the millisecond pulsars in this thesis are unfortunately not much better because most of these pulsars have significant proper motion. As explained earlier, proper motion induces a significant $\dot{P}_b$. Unless the distance is accurately known, the proper motion contribution cannot be completely accounted for and the $\dot{G}/G$ limits will remain uninterestingly high. One exception is PSR J1713+0747, a bright millisecond pulsar with small proper motion. We have started observing this pulsar because it has already shown good long term stability from observations at Arecibo observatory (Camilo, Foster, & Wolszczan 1994). Within a few years of timing it should provide a limit on $\dot{G}/G \sim 3 \times 10^{-12}$ yr^{-1}.

1.3.3 The Gravitational Wave Background

Pulsar timing has important applications in the detection of gravitational waves. The Earth and the pulsar act as two free masses separated by a very long baseline, quite analogous to the two mirrors at the opposite ends of a laser interferometer. A gravitational wave passing the Earth (or the pulsar) perturbs the spacetime metric and 'stretches' or 'contracts' the distance between the two bodies. In terms of the dimensionless amplitude a of the gravitational wave, the strain produced is $\delta x/x \sim a$. On the time scale of a period of the wave, this will result in a displacement of $\delta x \sim \lambda a$, where $x = \lambda$ is the wavelength of the wave. This displacement is equivalent to a timing residual $\Delta t \sim \delta x/c \sim Ta$ (T = period of the wave). Therefore, the amplitude we can detect is $a \geq \Delta t/T$. Due to the sparse sampling of pulsar timing (days, weeks and months) we are only sensitive to T $\sim$ 1yr. Since $\Delta t \sim 1\mu$s, we can detect $a \geq 10^{-14}$. This is not sufficient to detect discrete events like binary mergers or supernovae with $a \leq 10^{-21}$ (Schafer 1991). These events also occur on a short timescale, to which pulsar timing is not sensitive. However, this level of sensitivity can be used to set limits on the energy density of a stochastic gravitational wave background which is predicted by many cosmological theories, much like the cosmic microwave background. Several authors (Detweiler 1979; Bertotti, Carr, & Rees 1983; Blandford, Narayan, & Romani 1984) have shown that the expected spectral density in timing residuals due to such a background is

$$P_g(f) = \frac{H_0^2}{8\pi^4}\,\Omega_g\,f^{-5} \qquad (1.15)$$

where Ω_g is the fractional energy density in gravitational waves per logarithmic frequency interval. With f measured in yr^{-1}, residuals in μs, and $H_0 = 100h$ km/s

Mpc^{-1}, $P_g(f)$ has the value

$$P_g(f) = 1.34 \times 10^4 \, \Omega_g \, h^2 \, f^{-5} \ \ \mu\text{s}^2\text{yr}. \tag{1.16}$$

Existing observations of pulsars therefore permit us to set an upper limit for $P_g(f)$. The best current results on $P_g(f)$ come from the timing residuals of PSR 1855+09 and PSR 1937+21 (Kaspi, Taylor, & Ryba 1994): $\Omega_g h^2 \leq 6 \times 10^{-8}$ (95% confidence limit).

Interpretation of timing residual noise in terms of a gravitational wave background requires care because there are other potential sources of noise such as clock and ephemeris error. It is possible to distinguish between the different noise sources because of their different angular signatures on the sky. Clock error will affect all pulsars equally while ephemeris error will have a dipolar signature. Intrinsic pulsar timing noise should be uncorrelated between different pulsars.

1.3.4 Millisecond Pulsar Timing as a Testbed for Time Scales and Ephemerides

We have already covered the present state of data on millisecond pulsar stability (section 1.1). Even though PSR B1937+21 exhibits low frequency noise in its timing residuals, the level of the noise – about 1 part in 10^{14} over several years – is such that atomic time scales cannot be ruled out as a possible source of all or part of this noisy behaviour. We require long term stability data on many more millisecond pulsars to be able to distinguish the real sources of low frequency noise – whether atomic time scales, solar system ephemrides or the millisecond pulsars themselves. This is feasable since each of them will have different signatures in the pulsar data and can be isolated by careful comparison. PSRs J0437-4715 and J1713+0747 are highly suitable for this purpose and we present the first few years of data for these pulsars in this thesis. So far they have not exhibited any low frequency noise but the data spans are still relatively small (3 yr).

1.4 Dual Frequency Timing

The potential for sub-microsecond timing from millisecond pulsars was apparent soon after the discovery of the first such pulsar, PSR B1937+21. Initial timing observations were limited in precision by a lack of observing equipment which could meet the challenges of fast sampling, and fine frequency resolution required to overcome interstellar dispersion. A steady improvement in equipment and data analysis led to impressive gains in precision to $\sim 0.2\ \mu$s (Kaspi, Taylor, & Ryba 1994). The high precision forced observers to pay close attention to all possible sources of timing error. These include time scales, barycentric correction (by using dynamical models of the solar system), and perturbations of the pulsar emission wavefront by the interstellar medium. Time scales and solar system ephemerides are monitored by a large community of scientists for many different applications and their errors are believed to be quite small. They have also gradually improved due to new technology (better atomic clocks) and new observations (spacecraft flybys). Interstellar propagation effects, on the other hand, have come to dominate the discussion about errors in precision timing. DM variations to PSR B1937+21 were detected by many observers (Rawley, Taylor, & Davis 1988, Cordes et al. 1990, Backer et al. 1993). The variations amounted to ~ 0.001 $\mathrm{pc\,cm^{-3}}\,yr^{-1}$, sufficient to induce delays of several microseconds in the arrival times.

The large impact of these changes motivated some observers to monitor the DM of other pulsars. Phillips & Woszczan (1991) detected both secular and random DM variations towards six of the seven pulsars that they observed. They could fit their results with $\sigma_{\mathrm{DM}} \propto \mathrm{DM}^{1.3\pm0.3}$, over a 10:1 range in DM. Backer et al. 1993 monitored a set of millisecond pulsars and also found secular and random variations. Their results were consistent with $\sigma_{\mathrm{DM}} \propto \mathrm{DM}^{0.5}$.

DM variations to pulsars should not come as a surprise. The interstellar medium (ISM) in the Galaxy is known to be inhomogenous on a wide range of length scales. This is true for all the different constituents of the ISM, including the ionized portion which is responsible for dispersion of radio signals. Large scale density fluctuations in the ISM can be easily created by energetic events like supernova explosions and stellar winds. Turbulence then cascades the energy at the larger scales to a wide range of smaller scales. The density irregularities are revealed through pulsar observations when the line of sight from the Earth to the pulsar sweeps through different regions of the ISM due to relative motion between the pulsar and Earth. This results in a host of observable effects – scintillation of the pulsar intensity due to destructive and constructive interference by different ray bundles, broadening of pulses, delays due to extra path lengths, a small wander of the apparent position of the pulsar in the sky, and variations in the dispersion measure.

Since pulsars are generally far away, $\sim 10^{20}$ cm, and the density irregularities are on a much smaller scale, $\sim 10^9 - 10^{15}$ cm, we can expect the ISM effects to be proportional to the pulsars' distance, and hence the DM (since DM $\propto$ distance). This is indeed the relationship deduced by the observers discussed above. For this reason, PSR J0437-4715, with its DM = 2.65 pc cm^{-3}, presents a very attractive target for precision timing. Using the observed DM variations as a yardstick, we would predict variations of only $\sim 10^{-5}$ pc cm^{-3} for this pulsar. The magnitude of the effect of a change in DM can be gauged from the following relationship:

$$\Delta t_{\Delta DM} = 2.06 \, \frac{\Delta \mathrm{DM}}{10^{-3}\,\mathrm{pc\,cm^{-3}}} \left(\frac{f}{1420\,\mathrm{MHz}} \right)^{-2} \mu s \tag{1.17}$$

where f is the observing frequency in MHz. Submicrosecond timing precision is therefore sufficient to detect DM changes $\sim 10^{-3}$ pc cm^{-3}. Since the DM variations towards PSR B1937+21 amounted to ~ 0.001 pc cm^{-3} yr^{-1}, they were readily revealed by precise timing. The observed variations prompted Cordes et al. (1990) to note that "the precision to which DM-induced timing perturbations can be removed from arrival times will have a large influence on the implied precision of PSR B1937+21 as a clock." Regular observations of PSR B1937+21 are now done at two widely different frequencies (1.4 and 2.8 GHz) so as to measure the DM variations and incorporate them into the timing model as an extra parameter (Kaspi, Taylor, & Ryba 1994). Foster and Cordes (1990) attempted a rigorous simulation of all known effects of ISM perturbations of arrival times for PSR B1937+21. They predicted that DM variations would dominate the perturbations with an amplitude of several microseconds. Even though these variations could be partially removed with dual frequency observations, there would be additional effects at the 100 – 300 nanosecond level because of geometrical path length differences and a wander of the pulsar position in the sky. Multiple frequency (3 or 4 widely separated frequencies) observations could be used to measure these effects, but the precision may never be better than the few hundred nanosecond level since the effects scale steeply with frequency (ν^{-2} and steeper) and may never be measurable with high accuracy.

The small DM and nearness of PSR J0437-4715 indicates that the anticipated ISM effects should be small and consequently the timing precision should be quite high. This prompted the construction of a new spectrometer – the Fast Pulsar Timing Machine (FPTM), (Navarro 1994) – to capitalize on this unprecedented opportunity. This machine used a total bandwidth of 128 MHz and used a fast correlator to record 256 lags for each polarization at a maximum time resolution of 4 μs. Observations of PSR J0437-4715 produced TOA residuals of 100 – 300 nanoseconds after 90 seconds of integration. Motivated by this big leap in precision (the previous spectrometer produced TOA residuals of $\sim$ microsecond), it was decided to expand the capabilities of the FPTM even further. With the addition of another analog section and reconfiguration of the correlator, the FPTM could be made to observe a total bandwidth of 256 MHz. Although this would be with half the frequency resolution, this was not of critical consequence for PSR J0437-4715 since its DM is so small. The bigger bandwidth would boost the S/N and potentially produce TOAs with residuals < 100 nanoseconds ! An additional advantage of this scheme was that the configuration could be used to observe the same pulsar at two widely different frequencies at the same time. In the 20-cm observing band these were usually 1.4 GHz and 1.7 GHz. The difference in arrival time between these frequencies due to a change in DM is given by

$$\Delta t_{\Delta DM} = 500 \, \frac{\Delta \mathrm{DM}}{10^{-3}\,\mathrm{pc\,cm^{-3}}} \, ns. \tag{1.18}$$

Since the TOA precision for PSR J0437-4715 is $\sim$ 100 – 300 ns, DM changes of

10^{-3} could be detected. It was also decided to add a larger set of analog filters with six different bandwidths (there were four originaly) to take advantage of interstellar scintillation (ISS). This is akin to an "if you can't beat'em, join'em" attitude. ISS typically enhances the pulsar signal in a narrow bandwidth whose size and location in the spectrum cannot be predicted. But these enhanced features typically are stable for a few minutes to hours, depending on the distance to the pulsar, its proper motion and the observing frequency. For the highest DM pulsars, for example, PSR B1937+21 with DM = 71.04 pc cm^{-3}, the ISS bandwidth is much smaller than the available bandwidth in the 20-cm observing band. This leads to a loss of S/N. It is therefore prudent to use the large bandwidth to search the available spectrum for the ISS enhancements and then zoom in on them with narrower bandwidths, thus reaping the benefits of better S/N and less dispersion smearing.

In chapter two we describe the design and fabrication of the new FPTM, along with a short description of the data analysis path. This is followed by results for PSR J0437-4715 in chapter three, and the remaining milllisecond pulsars in chapter four.

Chapter 2 Dual Band Upgrade

2.1 Introduction

A pulsar's radio signal can be modeled as an amplitude modulated noise process. The emission bandwidth is typically very large and can extend from frequencies as low as 25 MHz to as high as 30 GHz. The spectral index ranges between -1 and -3 for most pulsars. Using as large a bandwidth as possible results in increased signal-to-noise ratio (S/N) of the observed profile. However, the one major impediment in this path is dispersion smearing, which increases with bandwidth, leading to poorer timing resolution. Since dispersion $\propto 1/\nu^2$, observers are forced to observe at as high a frequency as possible, while preserving adequate S/N (since the spectral index is negative). A few different methods have traditionally been used to overcome dispersion:

1. Filter banks
2. Correlators
3. Direct recording of the baseband data on magnetic tape followed by software coherent dedispersion.

There have also been some limited implementations of real-time coherent dedispersion utilizing special purpose filters. Details about this and the above techniques can be found in (Hankins & Rickett 1975). The filter bank and correlator approach are equivalent in that both techniques break up the observing bandwidth into many smaller, usually equally-spaced, channels and the pulsar power in each frequency channel is then recorded. This limits the dispersion smearing to the bandwidth of the channel instead of the whole band. The smearing across one frequency channel (at a center frequency of 1420 MHz) is given by:

$$\Delta t_{DM} = 2.89\,(\mathrm{DM})\,\frac{\Delta f}{1\,MHz}\left(\frac{1420\,MHz}{f}\right)^{3}\;\mu s, \tag{2.1}$$

where DM is the dispersion to the pulsar in pc cm^{-3}, Δf is the channel width in MHz, and f is the center frequency of the channel in MHz. The channelization process leads to more complex hardware and a much higher data rate. The observer has to balance the need for larger S/N against dispersion smearing within a channel and the data recording capacity.

The approach adopted in the FPTM was to use a correlator as the basic building block. A correlator utilizes the Weiner-Khinchin relation to measure the power spectrum of the pulsar:

$$\int_{-\infty}^{\infty} r(\tau)\; e^{-j2\pi\nu\tau} d\tau = |H(\nu)|^{2} \tag{2.2}$$

where $r(\tau)$ is the autocorrelation function (ACF) of the radio signal, ν is the baseband frequency, and $|H(\nu)|^2$ is the frequency power spectrum of the radio signal. The correlator measures $r(\tau)$ and the power spectrum is then obtained simply by performing a Fourier transform of this measured ACF. This is in contrast to a filter bank in which the pulsar signal is split into individual frequency channels and then each is detected independently – no Fourier transformation is required.

In practice, the correlator is implemented with digital circuitry and hence measures the ACF only at discrete delays, called lags, and only up to a maximum lag. If the smallest lag is τ seconds, and there are N lags, the ACF is measured for delays equal to $0, \tau, 2\tau, ..., (N-1)\tau$ seconds. Since the ACF is symmetric with respect to the sign of the delay (negative and positive delays are the same), it is trivial to symmetrize the ACF before doing the FFT.

In its original configuration the FPTM (Navarro 1994) measured a 256 lag ACF in two independent polarizations for a given observing frequency. After the FFT this resulted in power values for 256 filter channels in each polarization, distributed contiguously over the observing bandwidth. The FPTM could also be configured to measure full Stokes parameters by devoting half the lags to measuring the cross-correlation between the two polarizations. This resulted in only half as many filter channels (128) per polarization.

In the FPTM the pulsar period is divided into a fixed number of intervals, 1024, or even submultiples of 1024, i.e., 512, 256, etc. These are called phase bins. During the first pass through a phase bin, the ACF is accumulated in a counter and at the end of the interval, stored in a specific memory location. During the next pulsar period for the phase bin, this stored sum is read back and added to the new accumulated sum and then written back to memory. At the end of the integration period, which typically lasts for 90 seconds, the ACF corresponding to each phase bin is written to a hard disk connected to the computer for processing offline and subsequent storage on magnetic tape.

We can get an estimate of the potential accuracy of a pulsar TOA derived from an FPTM observation. This is directly proportional to the time resolution of the 'Impulse Response' of the FPTM, which can be written as follows:

$$\Delta t = \left(\left(\max(4.0, (0.97 \frac{P}{1\,ms}) \right)^2 + \left(\frac{1}{\Delta f_{MHz}} \right)^2 + \left(\frac{2.89\, DM\, \Delta f_{MHz}}{f_{1420}^3} \right)^2 \right)^{1/2} \mu s \tag{2.3}$$

P is the pulsar period in milliseconds, Δf_{MHz} is the channel bandwidth in MHz, f_{1420} is the channel center frequency in units of 1420 MHz. The first term arises from dividing the pulsar period into 1024 or lesser number of phase bins. There is a practical limit to the minimum size of the pulsar phase bin, 4.0 μs. Thus any pulsar with $P < 4.096$ ms has a phase bin size $= 4$ μs, while those with $P > 4.096$ ms have phase bin size $= P/1024$. The second term is due to the finite rise time of a filter with bandwidth Δf and the third term represents dispersion smearing (equation 2.1). The three terms are added in quadrature since they are independent of each other. Figure 2.1 plots the different terms. For the highest DM pulsars, interstellar scattering overrides the above considerations and degrades the timing accuracy considerably. The best resolution (minimum Δt) is attained when the three terms are equal to each other. However, this is not possible for every period/DM combination because

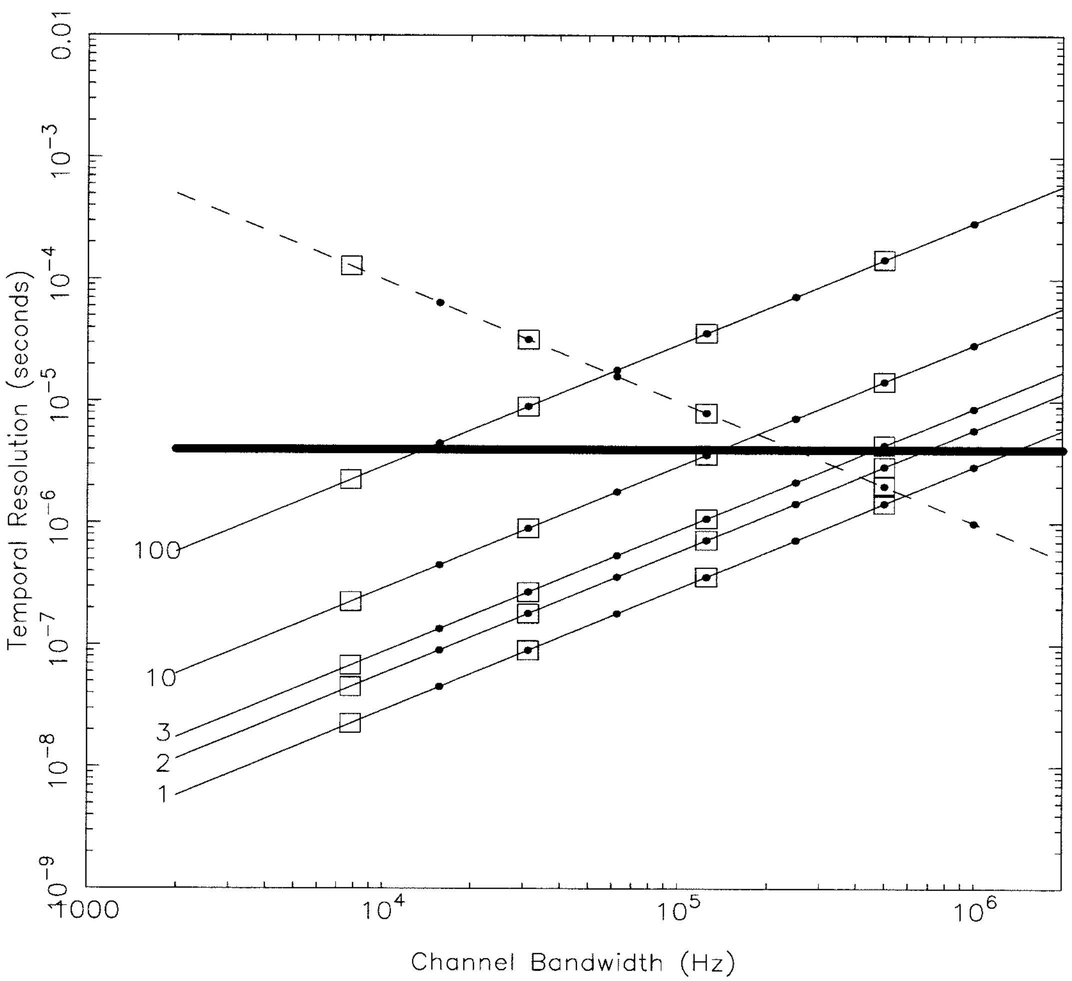

Figure 2.1: Time resolution in the FPTM. Open squares show bandwidths available in the original configuration of the FPTM. Filled circles show the bandwidths available after the dual band upgrade. The thick horizontal line marks the lower limit to the pulsar period phase bin. The dashed line shows the inverse bandwidth time constant while the solid lines indicate dispersion smearing (DM is shown at the left starting point of each line).

the FPTM only has a discrete set of bandwidths. For its original configuration the FPTM had 4 baseband filters with total bandwidths of 2, 8, 32 and 128 MHz. The corresponding channel bandwidths are obtained by dividing these values by 256, the number of lags in each polarization. These points are plotted as open squares in figure 2.1. It is clear from that figure that for low DM pulsars (DM ≤ 10) it is appropriate to use the maximum available bandwidth, but for moderate to high DM pulsars, it is better to reduce the bandwidth.

Figure 2.2 presents a general block diagram of the FPTM in its original configuration (from December 1993 to July 1995). Only one IF signal path is shown. There is another identical path for the second IF. The signal from the telescope, usually shifted to some intermediate frequency (IF) in the range 100 MHz - 700 MHz, is delivered to the first stage in the FPTM, the Upconverter. It translates the signal upwards so that the band of interest has a lower edge frequency of 1580 MHz. The Downconverter performs single sideband downconversion on this signal (always choosing the upper sideband) and shifts it to Baseband, 0 - 128 MHz. This Up/Down conversion scheme is needed for the following reason. The observatory IF is supplied at some frequency between 100 – 700 MHz and we need to translate this to baseband, 0 – 128 MHz, and have the flexibility to choose any 128 MHz (or smaller) band out of the total 600 MHz. This would require a local oscillator signal with the frequency range 100 – 572 MHz. It is difficult and expensive to build mixers which can mix a bandwidth (600 MHz in our case) which is quite large with respect to the local oscillator frequency. Most mixers work well when $\Delta f/f < 0.3$, where Δf is the required bandwidth, and f is the local oscillator frequency. We circumvent this problem by first mixing the signal UP to a higher frequency, 1580 MHz. This requires local oscillator frequencies in the range 1000 – 1400 MHz. The signal can then be conveniently downconverted to baseband. Another significant advantage in implementing this scheme was that the required downconverter had already been built by the Owens Valley Radio Observatory (OVRO) for their correlator. It is better to state that the FPTM adapted to the OVRO design and simply cloned the downconverter. Implementing the upconverter required only off-the-shelf components which were quickly available and relatively

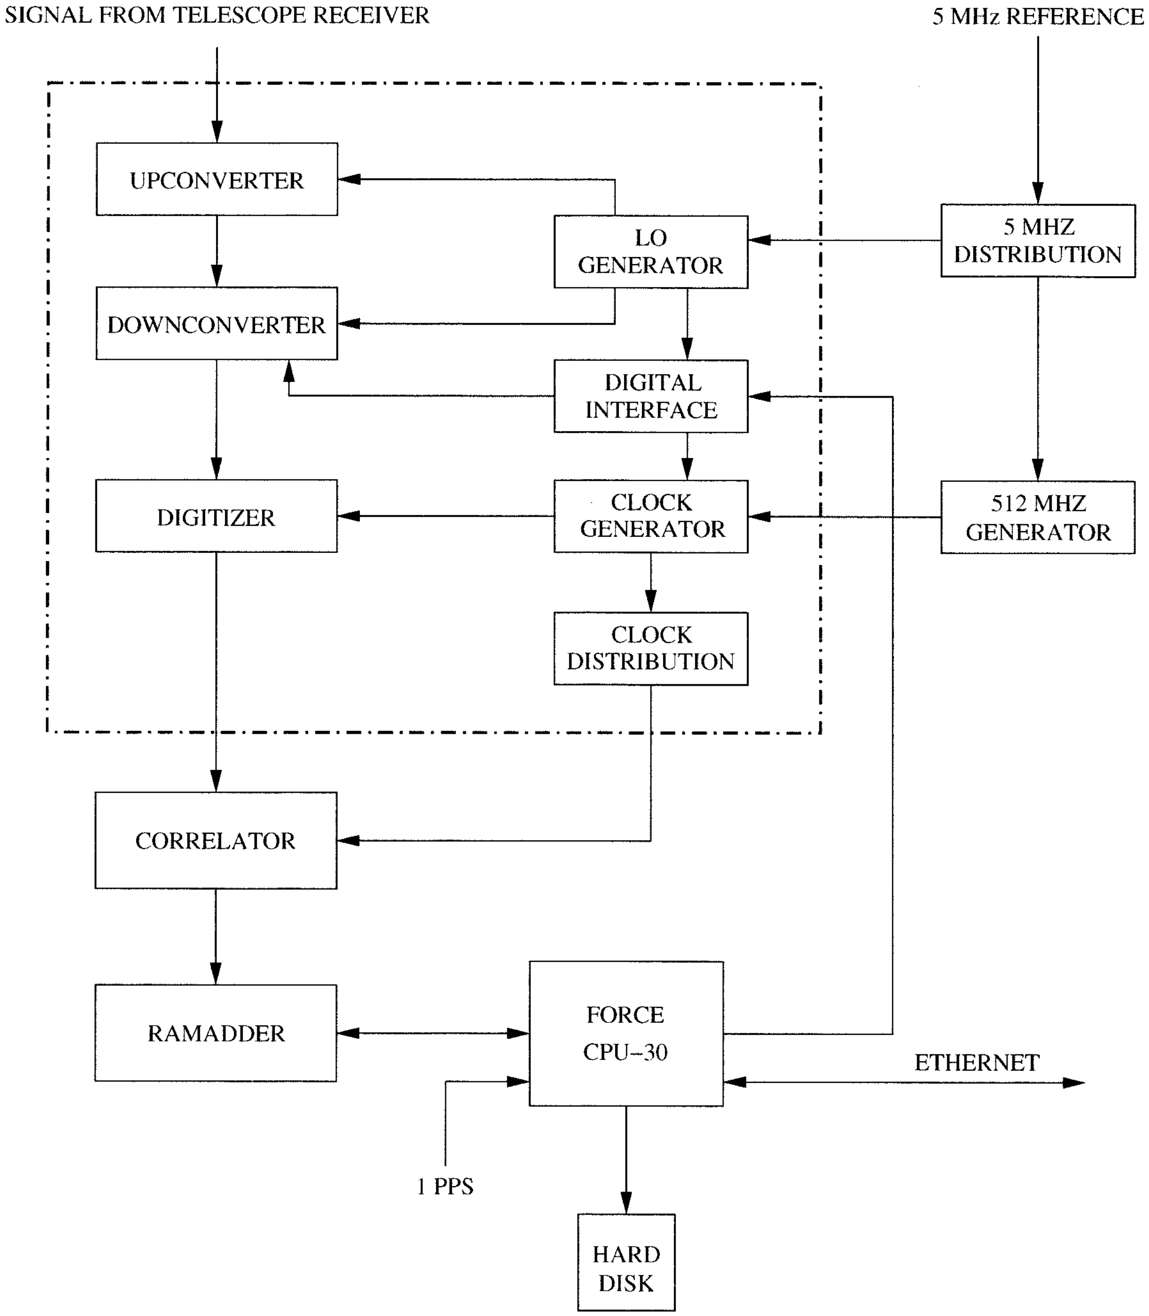

Figure 2.2: FPTM block diagram.

affordable.

After downconversion, one of four available baseband filters is selected and this filter determines the bandwidth of the signal delivered to the digitizer. This filter set is comprised of 2, 8, 32 and 128 MHz filters. The analog signal is sampled at a fixed frequency of 256 MHz in the Digitizer module. The clocks necessary for the digitizer are created by the Clock Generator card, which also generates the multiply and shift clocks necessary for the correlator boards. The card generates the clocks from a 512 MHz signal which is phase locked to a 5-MHz reference provided by the observatory. This reference signal is derived from a Hydrogen Maser, a stable frequency source.

The digitized samples are fed into the correlator. The correlator consists of 8 boards each of which can process 64 lags = 512 lags total. In normal operation ('Autocorrelation Mode') these are split evenly between two independent polarizations. For full Stokes parameter measurements ('Polarization Mode'), the lags are divided into 4 sets of 128 lags each. Two sets measure autocorrelation while two measure cross-correlation. The correlator always performs multiplication at a fixed rate of 256 MHz but shifting data to the next lag is done by a different clock, called the Shift clock. This clock runs at the Nyquist rate (twice the bandwidth). The correlator boards perform limited accumulation of the correlation sums. The major accumulation is done in dedicated memory boards, called ramadder boards. There are 8 of these and each can accumulate 64 lags. There is thus a one-to-one mapping between correlator and ramadder boards. The ramadder boards are controlled by a FORCE CPU-30 computer through a master control board (MCB). All the parameters for a pulsar observation are programmed into the MCB at the start of an observation. These include the number of phase bins (normally 1024), an initial period and the number of periods for integration. The MCB derives its time standard from a 1 Pulse Per Second (1 PPS) signal provided by the observatory. This 1 PPS signal is derived from a stable frequency source (the hydrogen maser). Differences between the maser and UTC maintained at the NASA deep space station at Tidbinbilla (about 300 km) away are recorded daily to allow retrospective corrections to the maser time scale. The UTC at Tidbinbilla itself is compared daily with UTC(NIST) via GPS observations. Thus our final time scale is believed to be equivalent to UTC(NIST).

When the MCB is commanded to start observing, it starts accumulation on the next 10 second boundary. This is the time stamp to which all pulsar TOAs are tagged for that particular observation. During an observation, the FORCE CPU-30 computes the pulsar period every few seconds and updates the MCB with the new value. This is necessary since the apparent pulsar period can change significantly during the course of an observation. At the end of an observation the FORCE computer reads lag data from the ramadder boards and writes it to a hard disk. That data is then copied to magnetic tape for permanent storage. Full details of the architecture and operation of the FPTM can be found in Navarro 1994.

2.2 Dual Band Upgrade

The first results from the FPTM at Parkes were very promising. For PSR J0437-4715, the time of arrival (TOA) uncertainity after a 90 s integration was typically 100 – 300 ns. (at 1.4 GHz center frequency with 128 MHz bandwidth). Navarro (1994) reported that despite this excellent short term stability, the TOAs exhibited unusually large drifts and jumps over several hours and days. The drifts/jumps amounted to several μs. At the same time that the FPTM observes this pulsar, data is also collected with a filter bank system that was originally used to discover the pulsar. This filter bank has poorer timing resolution due to wider bandwidth filters and slower sampling. Navarro (1994) noted that both the FPTM and the filterbank showed similar drifts in their TOAs, leading them to speculate that the problem was most likely instrumental and not anything intrinsic to the pulsar. A factor in favour of this argument was that their observations were made using linear polarization feeds. Most pulsars have significant linearly polarized emission, leading to very different observed profiles in two independent linear polarization feed. If the relative calibration of the two feeds is incorrect, the total intensity profile (and hence the TOA) can be significantly different from the true profile. With the very high S/N obtainable for PSR J0437-4715, this was probably the cause of their observed drifts/jumps, especially since they did not have a calibration scheme in place for their data.

It is clear that instrumental problems remain to be resolved for further improvements in timing PSR J0437-4715. If these can be tamed, the possible long term precision of the timing could be < 100 ns, limited only by the pulsars inherent stability and our understanding of atomic scales and solar system ephemerides. As outlined in chapter one, interstellar effects on the TOAs could potentially amount to similar magnitude, if not higher. Therefore, it made sense to supplement the FPTM's capabilities to record two different frequency bands simultaneously, with maximum bandwidths of 128 Mhz each. The specific advantages of such a system would be:

1. With the 1.4 GHz receivers at Parkes, the available bandwidth is ≥ 500 MHz. With a dual band system the FPTM could observe 256 MHz, thereby improving the S/N.
2. Collect data at two significantly different frequencies with dual frequency feeds. This helps determine DM variations, while saving observing time.
3. Data in two different frequency bands would also help provide clues to solving the systematic drifts/jumps observed with the single band system, since the two bands would follow two physically independent IF/baseband paths.
4. With a greater variety of bandwidths available, 'scintillation hunting' could be used as a viable observing strategy for higher DM pulsars.

Keeping these motivations in mind, it was decided to make the following changes to the FPTM:

1. Add two more IF processing chains in the FPTM hardware.
2. Increase the number of observing bandwidths to 6. From figure 2.1 we can see that the minimum bandwidth of 4 MHz would adequately serve the highest DM pulsars. Accordingly, it was decided to use bandwidths of 4, 8, 16, 32, 64 and 128 MHz.

The hardware changes required for this purpose were confined to the dotted box in figure 2.2. In the following sections I describe the hardware and software modifications carried out to accomplish these aims.

2.2.1 Local Oscillator Generator

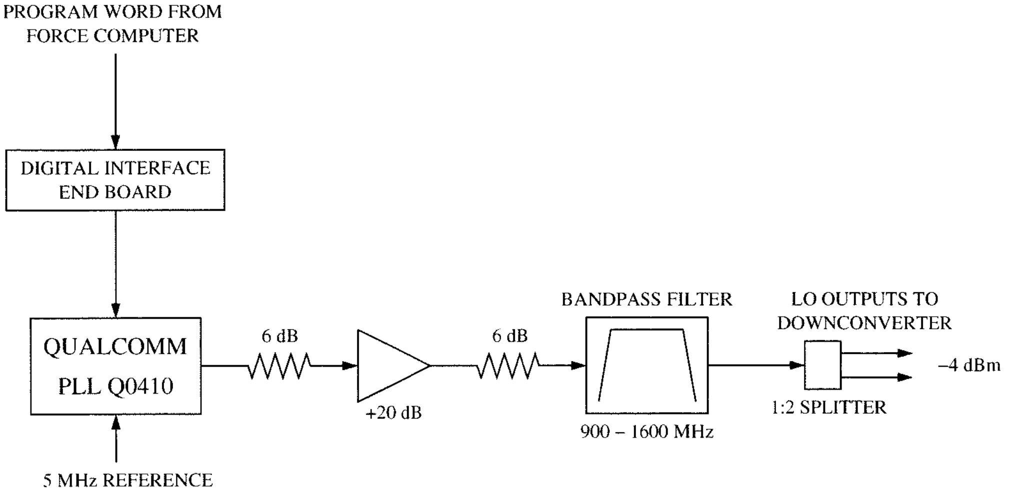

Figure 2.3: LO generator block diagram.

Each independent frequency band requires two local oscillator (LO) signals for the IF chain – one for the upconverter and one for the downconverter. This enables the FPTM to translate the observatory IF to baseband with flexibility. A dual band system would require four independent LOs for this scheme. However, it was possible to simplify the design (and cost) by adopting a common LO for the downconverter in the two observing bands. This way only three distinct frequencies were necessary – two for upconversion (one in each band) and one for downconversion. Since the FPTM could already generate two different LO frequencies, only one more LO generator was required. This was done by replicating the existing design, with the slight modification of a new bandpass filter at the output. Figure 2.3 shows a block diagram of the LO generator (a detailed schematic can be found in the FPTM documentation). The LO signal is generated from a Qualcomm Phase Locked Loop (PLL) synthesizer which can be programmed to produce frequencies in the range 900 – 1600 MHz, in 1.25 MHz steps. The bandpass filter is required to reject harmonics of the LO signal since the PLL is not a 'clean' generator. The amplification and attenuation have been chosen to produce two identical outputs at about -4 dBm power levels. Two outputs were designed so that this LO generator could be used as the common LO for downconversion in the two frequency bands. The PLL is phase locked to a 5-MHz reference derived from the observatory frequency standard (a hydrogen maser). The LO generator is programmed by the FORCE CPU-30 computer using a custom designed interface board which we dubbed the Digital Interface. Details about its operation can be found in Navarro 1994. The program **loset2** is used on the FORCE CPU-30 computer for this purpose.

2.2.2 Upconverter

The upconverter in the FPTM is designed to translate the IF signal from the telescope receiver from a lower IF frequency to a higher, fixed frequency range. The downconverter subsequently mixes the signal to baseband. The rationale for this scheme was explained in section 2.1 – basically it is meant to overcome the problem that the IF from different receivers is at widely different, and relatively low frequencies which cannot be directly mixed to baseband.

The original upconverter in the FPTM could process two IF signals, one in each polarization. Dual band operation requires four independent upconversion paths for two polarizations in two different frequency bands. Instead of adding a second upconverter, we decided to build a new one which could process four IF signals. This decision was also forced upon us by space limitations in the IF rack of the FPTM – there was no space for a second upconverter unit. The design of the new upconverter is essentially the same as the old one. The main changes are the use of a different mixer and the addition of an amplifier to boost the output signal strength. Figure 2.4 shows a block diagram of one IF frequency in the new upconverter. The other three are identical. The input IF signal can be in the range 10–680 MHz. It is mixed with an LO with a frequency such that the band of interest at the output has a start frequency of 1580 MHz. For example, if we want to observe the band 600–632 MHz

(IF frequency range), we would use a 980 MHz LO to shift the band to 1580–1612 MHz (980+600 = 1580 MHz). The output filter has a bandwidth of 256 MHz centered at 1580 MHz. Thus the output signal can have a maximum bandwidth of 256 MHz, even though we eventually only utilize 128 MHz. The amplification and attenuation have been so selected that there is a net amplification of about 3 dB, the exact value depending on the conversion loss of the mixer at the IF frequency in use. The DC block is required to prevent large impulsive interference from propagating through the IF chain and causing damage to the digitizer. The upconverter is a 'passive' device as far as the FORCE CPU-30 computer is concerned, i.e., it does not require any programming during observations since there are no programmable units inside it.

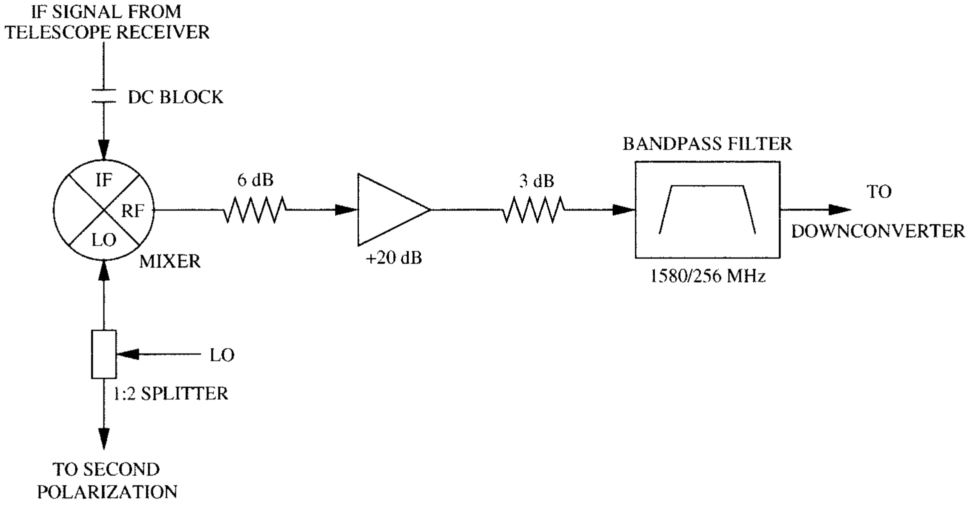

Figure 2.4: Upconverter block diagram.

2.2.3 Baseband Downconverter

As explained earlier, the downconverter always recieves an input signal in the frequency range 1580 ± 128 MHz. It uses an LO at 1580 MHz to move the upper 128 MHz to baseband, i.e., 0–128 MHz. The original design of the downconverter was cloned from the Owens Valley Radio Observatory millimeter array, with the addition of having a set of low pass filters at the output which were user selectable. As in the case of the upconverter, space limitations forced us to build a new downconverter unit which preserved the basic design of the original, but could process four IF signals.

Some component modifications were made to improve performance. Figure 2.5 shows a block diagram for one IF path. There is a variable attenuator at the input which can be varied from 0 to 31.5 dB in 0.5 dB steps. This attenuator is used to adjust the signal strength so that the digitizer is presented with a +13 dBm signal power. This is necessary because the thresholds in the digitizer are fixed and the correlator has optimum response when the signal strength is +13 dBm. The signal power changes from receiver to receiver and also from one part of the sky to another. We replaced the original variable attenuator, which had steps of one dB, with the new one which has attenuation steps of half dB, since one dB did not provide enough resolution for power control.

The detector following the attenuator is used to display the input signal strength on a front panel and serves as a visual aid for system health status. The single sideband mixer is designed to translate the upper sideband to baseband. Quadrature mixing is utilized to achieve this objective. The hybrid generates two copies of the input signal, 90 degrees out of phase with respect to each other. Each copy is then mixed down to baseband by the LO and sent through a chain of five amplifiers. The amplifiers are tuned to provide equal gain in the two chains but one set provides an extra 90 degrees of phase shift. When the two signals are then added the lower sideband product cancels, leaving only the upper sideband product at the output. The circuit containing the two sets of amplifiers is called the Quadrature filter. The difficulty in tuning this filter lies in achieving minimum sideband rejection throughout the band ($\sim$ 0.1 to 128 MHz) along with equal gain across the passband. The two requirements are antagonistic and have proved to be the most troublesome aspect of the downconverter. The filter contains roughly 50 independently tunable components – a sizeable challenge for meeting the two criteria. Despite many attempts, the performance has not met our original specifications. The sideband rejection is 20 dB, which is adequate, but the gain across the passband is disappointing. The gain at the higher end of the band is lower than that at the lower end by about 10 dB. This leads to significantly lower signal to noise in the upper 20–30% of the band. Thus we end up using only about 80 MHz of bandwidth instead of the designed 128

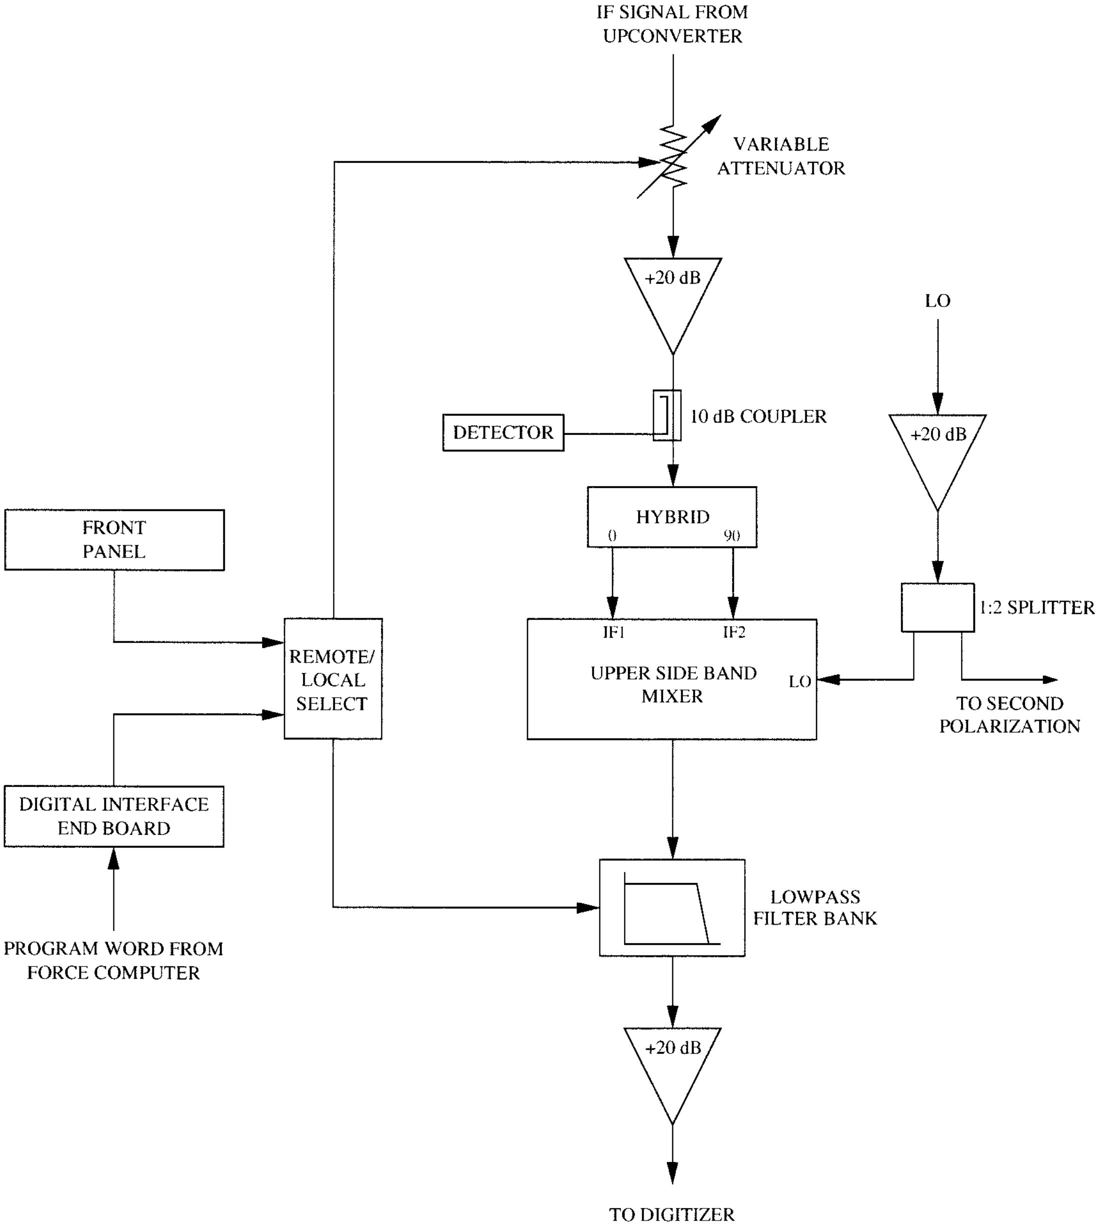

Figure 2.5: Downconverter block diagram.

MHz. Apart from the lower signal to noise, we believe the slope in the gain also leads to systematic modification of a pulsar profile across the band. Amelioration of this deficiency should therefore be a high priority in the near future.

The upper sideband mixer is followed by a bank of six-low pass filters. They have bandwidths of 4, 8, 16, 32, 64 and 128 MHz. These bandwidths were picked to provide more uniform coverage of the DM/period/ISS phase space which determines the time resolution for TOAs, as illustrated by 2.1. For a constant spectral power density across the band of interest, a bandwidth of 128 MHz has twice the total power compared to a 64 MHz bandwidth. However, the digitzer demands constant power for optimum performance. We are therefore required to insert extra attenuation in the path of the higher bandwidth filters in order to equalize the power no matter what the bandwidth in use. The attenuations used are 0, 3, 6, 9, 12, and 15 dB for the 4, 8, 16, 32, 64 and 128 MHz bandwidths, respectively. The filter bank was designed and built by Reactel Inc. In the specification for the bandwidths there was a misunderstanding about the nomenclature. Instead of specifying 128 MHz (for instance) as the 3 dB cutoff frequency of the filter, we stated 128 MHz as the passband, which meant that the cutoff frequency was shifted higher, to about 136 MHz. Since the sampling clock is at 256 MHz, this results in significant aliasing of the signal in the upper 15–20% of the bandwidth. We are therefore forced to reject the upper 20% of any one of the filters, leading to a corresponding decrease in signal to noise of about 1 dB. Caltech engineer John Yamasaki has designed and built a new set of filters with correct cutoff frequencies. These are now being used in the FPTM.

The downconverter requires two different controls for each IF signal path – the variable attenuator to set the signal strength and the low pass filter to be used. The FORCE CPU-30 controls these variables using the **bbc2** program which sets the appropriate bits in the digital interface end board inside the downconverter chassis. The chassis also has front panel controls for the same variables. Remote versus front panel control can be selected through a toggle switch on the front panel.

2.2.4 Clock Generator

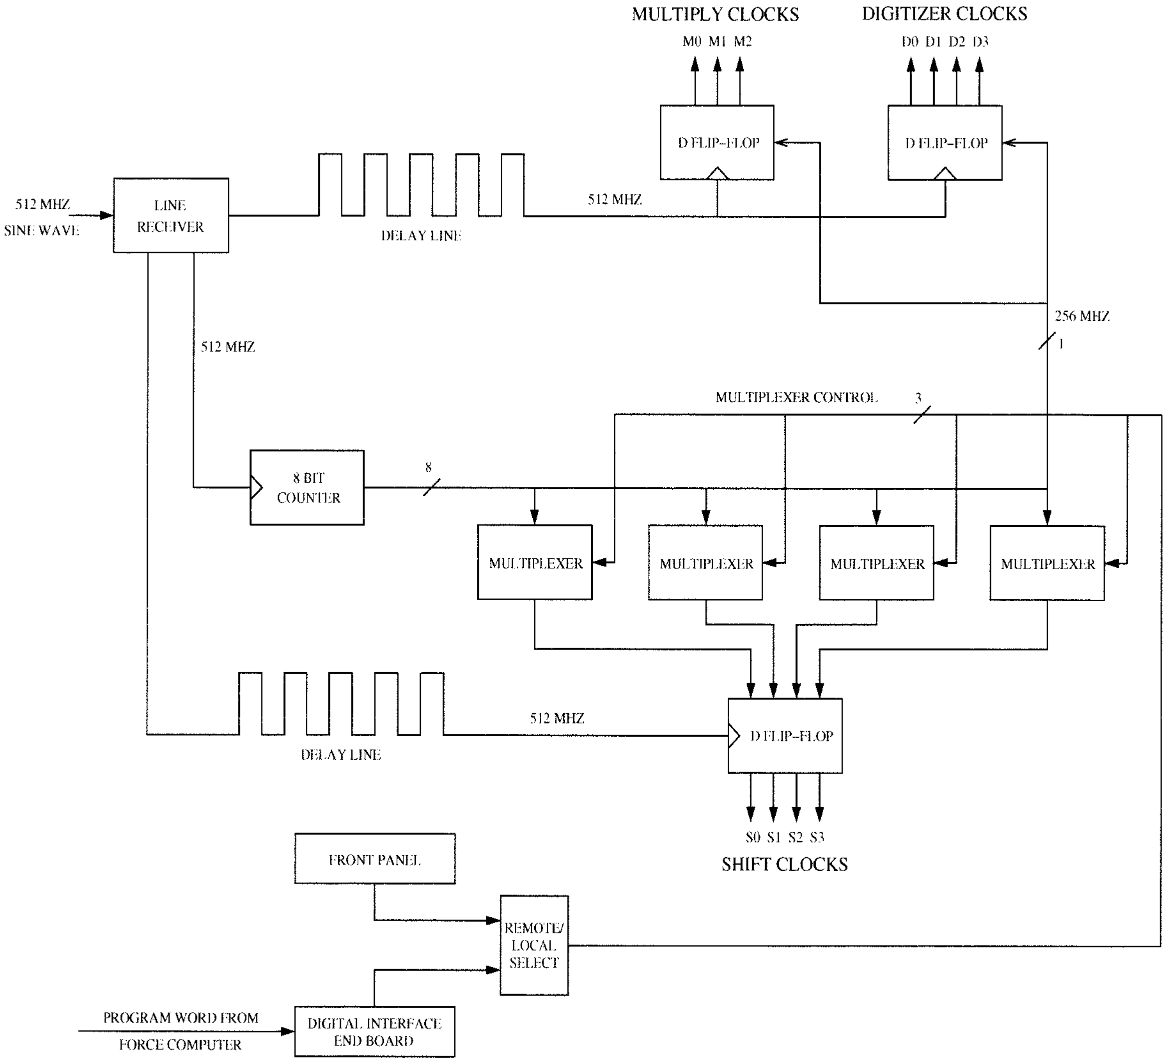

Figure 2.6: Clock generator block diagram.

The FPTM digital subsystem requires three clock signals: a fixed frequency (256 MHz) digitizing clock for the analog-to-digital converter, a fixed frequency clock (256 MHz) for the multiplier units in the correlator boards, and a variable frequency (8 - 256 MHz) for the shift blocks in the correlator boards. Out of these three, the digitizer and multiply clocks are common to both frequency bands but the shift clock depends on the bandwidth in use in each band. Thus in the most general case, four different clock frequencies are required. The generation of these clocks is the responsibility of the clock generator circuit. The original clock generator for the FPTM was borrowed from the Owens Valley Radio Observatory. This circuit used Gallium Arsenide chips because at that time they provided the high frequency performance required for the clocks. In the early nineties the manufacturer of the chips was acquired by another company which discontinued production of the chips used in the clock generator. At the same time comparable high frequency performance was emerging in ECLinPS technology offered by Motorola. We decided to build a new clock generator with this new techonology.

Designing a circuit at high frequencies ($\sim$ 1 GHz) requires attention to two main details – the signal trace has to be treated as a transmission line with associated path delays, and propagation delay through chips can significantly determine the timing skew between different signals. Transmission line analysis was done using standard formulations for microstrip lines, along with simulations using the Linecalc program. The printed circuit board manufacturer provided values for the dielectric constant which was then used to calculate the board thickness and trace width required to achieve a transmission line impedance of 50 Ohms. Figure 2.6 shows a block diagram of the clock generator circuit.

A 512 MHz sine wave, locked to the observatory frequency standard, is converted to a differential ECL signal by the line receiver. One copy of this signal drives an 8-bit binary counter. The upper six output bits of the binary counter have the desired frequencies of 256, 128, 64, 32, 16, and 8 MHz. These counter bits are used as inputs to four 8-bit multiplexers (making sure that the path lengths of all the signals are equal). Each multiplexer has 3 control bits which can be used to select the desired frequency for each shift clock. Even though only two distinct shift clocks are required for the dual band observations, we have provided for four in case more are needed in the future. The four shift clocks are sent to a 4-bit D flip-flop register for synchronization. The bit with the 256 MHz frequency is also sent to two 4-bit D flip-flop registers. The outputs of these registers represent the fixed frequency Digitzer and Multiply clocks.

It is important that all the output clocks be phase aligned to within a small fraction of the period (5%). This is necessary because the correlator boards are very sensitive to any relative delay between the clocks. For this purpose, it would be best if all the output D flip-flop chips were clocked by the same signal, ensuring that the output clocks would toggle at the same time. This could not be achieved with the board layout and so two different copies of the same clock are used to drive the 3 chips. The clock signals are derived from the 512 MHz sine wave using the line reciever and laid out to have the same path length from source to destination. This explains the different delay loops in the two clock paths. The output clocks also need to have the same path length from the chip pin to the output connector in the board housing. This was also achieved by inserting different delays in all the output clock paths. In case the calculated path lengths were found to be insufficient during testing, ground pads were placed next to the clock traces wherever possible so that the clock trace could be broken and extra delay inserted by using a loop of coaxial cable. As it turned out, testing revealed that all the clocks were adequately aligned and did not require such drastic modification. The board was laid out using ORCAD software and the board aluminum housing was designed with AUTOCAD.

2.2.5 Clock Distribution

Each of the eight correlator boards requires its own copy of Multiply and Shift clocks but the clock generator only produces one copy. The clock distribution circuit is designed to use the master clock to produce eight identical copies. This is done by using a single chip which is a 1-to-8 splitter. This ensures that all the clocks suffer nearly equal delays. The worst case differential delay is determined by the propagation delay skew of the chip, which is a few hundred picoseconds. The clock distribution circuit was also originally built using Gallium Arsenide technology, but OVRO had already redesigned and built it using ECLinPS chips. We borrowed the boards for our use but needed to do one major modification. Due to different requirements, the OVRO boards only produced **two** copies for a Shift clock and eight for the Multiply. Fortunately the board was laid out with such a symmetry that if we cut it into two pieces along one of the axes, one half could be used to split an input signal into eight. Using two such halves served our purpose of replicating both the Shift and Multiply clocks (one half for the multiply and one half for the shift). One such clock distribution circuit already existed in the FPTM. We needed to build only one more for dual band operation.

2.2.6 Digitizer and Threshold Monitors

The digitizer refers to both the analog to digital converter circuit as well as the chassis in which it is housed. The chassis also includes the clock generator box. The digitizer circuit operates at a constant 256 MHz sampling frequency. The Nyquist bandwidth is therefore 128 MHz. Each digitizer circuit can sample two baseband signals, which typically correspond to two independent polarizations in the same frequency band. Dual band operation therefore required one more digitizer circuit. A new, bigger digitizer chassis was needed to house two digitizers and the enhanced clock generator.

The digitizer produces two output bits (four level quantization) by comparing the input signal with three threshold levels. These thresholds are fixed at the time of testing the circuit at -0.9 V, 0 V, and $+0.9$ V. These values optimize the performance of the correlator for an input noise signal which has 1 V root mean square amplitude (and zero mean value), which is equivalent to $+13$ dBm power for a 50 Ohm impedance. For a Gaussian distributed noise signal we can calculate the probability of each sample falling in any of the quantization steps. Section 6.3 in Navarro 1994 contains a detailed discussion of these probabilities. Practically, each probability value provides us with a tool to monitor the performance of the digitizer. The digitizer has dedicated counters which keep track of the frequency with which the input signal exceeds each of the threshold levels. Given a sampling rate of 256 MHz, the probabilities can be used to compute the corresponding threshold crossing frequencies. Since each counter actually provides a divide by 16 output, the proper count rate to use is 16 MHz. For example, it is easy to show that a zero mean input signal should be greater than the middle threshold (0 V) half the time (the probability density of a zero mean gaussian is equally distributed about 0). This indicates that the Zero threshold counter should show an 8 MHz count rate. A different count rate would indicate that the input signal has a DC offset. Similarly it can be shown that the remaining two threshold counters should have count rates of 2.9 MHz each. A count rate less than this corresponds to a signal with insufficient power and vice versa. A count rate of zero or one which does not change when the input signal is increased or decreased might indicate that the digitizer is not working. The diagnostic value of these count rates encouraged us to build a dedicated frequency counter circuit which would display the count rates in real time. This circuit takes the threshold monitor outputs from the digitizer circuit and displays their frequencies on an LED display on the front panel. The FPTM operator can use them to quickly ascertain the health of the digitizer circuit, and it has indeed proven invaluable in debugging the FPTM time and again. The initial circuit for the single band FPTM was a wire wrap board. For the dual band upgrade, Caltech engineer John Yamasaki designed and laid out a printed circuit board which not only displays the count rates, but also shows the power supply levels.

2.2.7 Digital Interface

Many components of the FPTM need to be programmed with the correct parameters before starting any observation. These include the local oscillators for the Up and Down converters, the bandwidth selection in the Downconverter, the shift clocks in the clock generator, the attenuation settings for each IF channel, and the configuration of the correlator for either auto-correlation or cross-correlation. It is impractical to do these tasks manually through front panel controls before each observation. In order to automate the whole process, we designed a programming system called the Digital Interface. Appendix C in Navarro 1994 contains a detailed discussion of the Interface. The interface consists of one middle board and five end boards. Each end board resides in a destination module, for example the digitizer chassis. The FORCE CPU-30 sends a message to the middle board, consisting of data and an address corresponding to the end board for which the data is destined. The middle board then forwards this message to the end board. Each end board contains two 32 bit output registers. These registers are used to send data to whichever component needs to be programmed. The facility of the digital interface enormously simplified the observing procedure. Only a few seconds are spent to program all the components, compared to a few minutes without it. The middle board has been implemented with wire wrap and the end board was laid out for manufacture using ORCAD. One desirable feature for a future version of the interface would be the ability to read back the data in the end boards.

2.3 Software

Apart from the effort required to design and build all the hardware components for the single and dual band FPTM, a significant amount of time was spent in writing software for real time control as well as offline data analysis. In the following sections I describe the various software modules produced during these efforts.

2.3.1 Telescope Control

Implementing telescope control for observing with the FPTM at Parkes observatory was a unique challenge. This was because of the many different operating systems involved. OS9 is used in the FPTM control computer (a FORCE CPU-30), Unix was used in the observing computer, a Sun workstation, while the telescope was controlled by a VAX system running VMS. A complicating requirement for the observing scheme was that while the FPTM was observing a pulsar, another pulsar backend was also used for recording data. This was a multi-channel filterbank system controlled by the same VAX computer running the telescope. The following scheme was devised to meet the above constraints: The unix computer runs the master program (called 'Timer') which decides the observing sequence and duration. It sends commands to slave processes running on the other computers to coordinate their activities. In case of the FPTM this is done by running the observing program (**observe**) using a remote shell process (**rsh**). For the VAX computer a new program (**vtelcontrol**) was written which communicates with the unix program using sockets, a software communication implementation which utilizes TCP/IP. It receives the pulsar coordinates through the socket connection and sends the appropriate commands to the telescope control processes. **vtelcontrol** uses a shared memory flag with the filterbank observing program to indicate start and stop of observing.

2.3.2 FPTM Control Programs

Several significant modifications were necessary in the FPTM control software for doing dual band observations. The following programs can be used to control all aspects of observing. For each program, executing **program -h** on the FORCE CPU-30 displays usage information for that program.

bbc2 sets the attenuation and filter bandwidths for the downconverter, for both frequency bands or a single band only.

loset2 sets the upconverter and downconverter Local Oscillator frequencies.

sclk2 sets the shift clock in the two clock generators.

setl2 searches for the correct setting for the variable attenuators in the two downconverters for appropriate power level into the digitizer.

corset sets the correlator boards in auto-correlation or cross-correlation mode.

observe The main observing program. It runs all the above programs before starting an observation.

2.3.3 Data Analysis Programs

Typically, the FPTM is programmed to observe a millisecond pulsar for 25 minutes. This duration is broken up into 16 observing intervals of about 90 seconds each, separated by short gaps of about 2 seconds. 90 seconds is close to the maximum observing duration before the memory in the ramadder boards is filled to capacity.

Hence any longer observations have to be done in multiples of 90 seconds. The 2 second gap is required to read out all the memory boards. The 90 and 2 second intervals are called long and short integrations, respectively. The transition from a long to a short (and vice versa) is always done on a pulsar period boundary so that successive long integrations can be added phase coherently. Data for each long integration is written to disk and given a unique filename. This filename is of the form YYMMDD_HHMMSS_XXX.rfd, where YY,MM,DD,HH,MM and SS stand for the year, month, day, hour, minutes and seconds of the day, respectively. XXX is an increasing sequence of integers to keep track of the long integrations. For example, an observation made on 31st March, 1997, starting at approximately 09:39:05 (UTC), consisting of four long integrations will produce the following four files:

970331_093905_000.rfd
970331_093905_001.rfd
970331_093905_002.rfd
970331_093905_003.rfd

The **.rfd** suffix is short for **raw folded data**. In the offline processing the very first step would be a Fourier transform of the ACFs in the rfd file to obtain the frequency power spectrum for each pulsar phase bin. Before this step though, we need to carefully analyze each ACF in order to locate bad lag data, which can occur either due to known systematic causes or from hardware irregularities. The data must also be calibrated and searched for interference, which must be removed. The calibration scheme used is as follows: Once a day during an observing session the telescope is pointed towards Hydra (3C218), a radio source with a calibrated flux density. The Parkes telescope has a noise diode which can be turned on to inject noise in the received signal. It is controlled by a TTL signal from the FPTM. A three minute observation is done on Hydra with the noise diode controlled by a square waveform (a 50% duty cycle 'pulsar'). This is followed by a similar three minute observation two degrees away from Hydra. The two data sets are combined later on to calculate the noise diode absolute flux density. The noise diode is thus established as a secondary calibrator for pulsar observations. Before any millisecond pulsar observation during the rest of the day, a three minute pulsed diode observation is done a few degrees away from the pulsar position (this is only to avoid any contribution to the received flux from strong pulsars). This observation and the calculated noise diode density yields the gain in each channel and also the system noise level in Janskies.

The following programs are used to convert the recorded ACFs to TOAs:

rfd2obs This program processes an rfd file. It searches for bad lags and phase bins and replaces them with appropriate mean values. It applies a correction to the ACFs to compensate for two-bit quantization. The ACFs are then Fourier transformed to produce frequency domain data. If the observation was in single frequency mode or polarization mode, one output file would be produced. If it was a dual band observation, two output files would be produced. The output files are given a **.obs** suffix and a **c** prefix for the first frequency band and a **d** prefix for the second file for a dual band observation. For example, if 970331_093905_000.rfd were dual band observation data, two output files would be produced with the names:
c970331_093905_000.obs
d970331_093905_000.obs

The program can also be used to display data in the ACF, phase bin or frequency domain.

gethdr Uses the observation parameters in the YYMMDD_HHMMSS_000.rfd header to produces two header files:
cYYMMDD_HHMMSS.hdr
dYYMMDD_HHMMSS.hdr

corresponding to the two frequency bands. For single band or polarization observations, only the cYYMMDD_HHMMSS.hdr file is produced. These header files are used by subsequent programs to process the data.

getcal Given any .obs file (produced by rfd2obs), it searches for the nearest (in time) noise diode calibration file that can be used to calibrate that particular .obs file.

Pulsar off-source noise diode observations are given the special suffix **.cal**, while noise diode absolute flux density files are given the special suffix **.scal**.

tcedar This program is run immediately after getcal. It uses the data in .cal and .scal files to calculate the gain in each channel. These gains are applied to the data in a .obs file and written out to a file with the suffix .obs0.

archiver This is the program that puts many long integrations together into one data file, after they have been calibrated, i.e., converted to .obs0 form. The final data file is called an archive and has the suffix **ar**. For each pulsar, archiver averages the data in frequency and time according to parameters stored in a pulsar catalogue. For example, a pulsar might have the frequency channels summed from 128 to 32 and the phase bins reduced from 1028 to 256. The parameters are set individually for each pulsar and are aimed to strike a balance between preserving maximum frequency and time resolution and saving storage space.

ar2scal This program uses the data in the primary calibrator noise diode observations ('Hydra observations') to produce absolute flux density values for the noise diode.

rfd2ar This is a script that runs all the above mentioned programs in the correct order to convert raw data from any observing session to final archive format. The only intervention required from the operator is production of the appropriate calibration file (.scal) using the ar2scal program (see above).

treduce This is a general purpose program which can display data in an archive file and produce TOAs.

getscint If the observer wants to narrow his observing bandwidth onto a scintillation island, he needs to compute the center frequency that can be used, since the LOs in the FPTM cannot be set to any frequency (their resolution is 1.25 MHz). This program will compute the best center frequency, given three inputs: the telescope receiver LO in use, the desired center frequency for the scintillation band and the desired bandwidth.

plk This program plots the residuals produced by TEMPO. Pre-fit or Post-fit residuals can be plotted vesus date, orbital phase, serially, day of the year, observation frequency, and input error. A verbose option can be used to display the residual values in ASCII format.

Chapter 3 PSR 0437-4715

3.1 Introduction[1]

PSR J0437-4715 was discovered in the Parkes southern pulsar survey (Johnston et al. 1993; Manchester et al. 1996). Among the known millisecond pulsars, it is the nearest and brightest. It has a period of 5.75 milliseconds and is in a binary system with a 5.74 day orbital period. A number of observers detected its companion through optical observations and identified it as a low-mass white-dwarf (Bell, Bailes, & Bessell 1993, Danziger, Baade, & Della Valle 1993, Bailyn 1993). Timing (Bell et al. 1995) and optical (Danziger et al. 1993) measurements show that the system has a large proper motion, $\sim$ 140 mas yr^{-1}, corresponding to a transverse space velocity of $\sim$ 100 km s^{-1} at the dispersion measure (DM) derived distance of 140 pc. The low DM, high flux density ($\sim$ 90 mJy at 1.4 GHz), and short spin period make PSR J0437-4715 a very attractive target for precision timing. The average time of arrival (TOA) uncertainity after a few minutes of observation at 20-cm wavelength is about 100 nanoseconds. This not only allows us to measure the astrometric and kinematic parameters with high precision, but also establishes PSR J0437-4715 as a suitable candidate for comparing the stabilities of pulsars with man made atomic clocks.

In this chapter, we report on timing measurements for PSR J0437-4715 made at Parkes Observatory, Australia, using the FPTM. This instrument has significantly better timing resolution compared to the filterbank-based systems which were used to report the earlier timing results (Johnston et al. 1993; Bell et al. 1993; Bell et al. 1996). This has allowed greatly improved determinations of the basic pulsar and

---

[1] A substantial part of the results presented in this chapter (section 3.3) were published in Sandhu et al. 1997. Here we present a more thorough analysis of the systematic errors, as well as a measurement of the long-term stability of PSR J0437-4715.

orbital parameters, and significant measurements of both the annual parallax and the secular change in the projected orbital semi-major axis.

3.2 Observations and Analysis

Observations were made using the 64 m Parkes radio telescope with three different dual-polarization receiver systems, centered at 0.66, 1.4 and 1.6 GHz from 1994 January 21 to 1996 August 23. The 660 MHz system used orthogonal linear polarizations with a bandwidth of 32 MHz and 256 lags per polarization. Some early 1994 observations at 1.4 GHz were made with orthogonal linear polarizations. The rest of the observations at 1.4 and 1.6 GHz used orthogonal circular polarizations. Initially, we recorded a single band of 128 MHz bandwidth with 256 lags for each polarization, yielding 1/2 MHz channel bandwidth. In July 1995 we upgraded the system for dual frequency observations. Since the dispersion smearing for PSR J0437-4715 at 1.41 GHz is only $\sim 7\ \mu s$ across 1 MHz, we observed the pulsar at two different frequencies, each with 128 lags and 128 MHz bandwidth. The two frequencies were usually centered at 1.41 and 1.66 GHz.

The observation method has been described in chapter two. For PSR J0437-4715 we typically devoted 24 minutes of observing time at each epoch, divided into 16 90 s integrations. 1024 phase bins were used for each lag. Time synchronization was provided by starting the first of a series of 90 s integrations coincident with a 10 s pulse from the Observatory time system. The Observatory timescale was referred to that of the NASA Deep Space Station at Tidbinbilla using a microwave link, and this in turn was referred to UTC (NIST) using common-view GPS observations. The final timescale is believed to agree with UTC (NIST) to better than 100 ns.

The raw data were then further processed as follows. First, the ACF data were corrected for the non-linear effects of the two-bit digitization. The ACFs were multiplied by a triangular window function and then Fourier transformed to yield power spectra for each of the 90 s integrations. The triangular window reduces the sidelobe response for each channel from 13.5 dB to 27.5 dB. This helps to reduce the effect of narrow band interference on adjacent channels. There is a degradation in the channel width, but the trade-off for PSR J0437-4715 in terms of enhanced dispersion smearing is minimal. The spectra were calibrated on a flux density scale using the method described in chapter two. Frequency channels affected by narrow-band interference were also deleted. Next, the spectra were dedispersed, compressed into 32 channels and stored on disk. Standard total-intensity profile templates of high signal-to-noise ratio were formed for each of the observing systems by adding many (~ 100) separate observations. Finally, pulse TOAs were obtained by cross-correlating these standard templates with fully dedispersed 90 s integrations. A total of 4970 TOAs were used from the 2.6 yr of observations.

3.2.1 Sources of TOA Error

As shown in Navarro 1994, and confirmed by further observations, the TOA residuals for PSR J0437-4715 exhibit unusually large drifts and jumps over the course of a few hours and days. Here we present some examples of this behaviour and discuss different reasons which might give rise to systematic deviations.

Polarization Calibration Issues

The TOA calculation algorithm assumes that the observed pulse profile in each of the 90 s integrations is identical to the standard profile, except for additive radiometer noise. This condition is difficult to satisfy in practice because the pulsar emission is significantly polarized. Unless calibration of the two received polarizations is 100% accurate, there are bound to be subtle changes in the total intensity profile for each observation, depending on the paralactic angle during the observation and the relative gain error between the two polarizations. These profile changes translate directly into TOA errors. Our calibration scheme was not in place for most of 1994 and the data collected during that period show the drastic effect of mis-calibration on the TOAs.

Figure 3.1 shows the full pulse profile as well as a zoomed in section around the pulse peak, illustrating the pulse shape difference in the two polarizations. This pro-

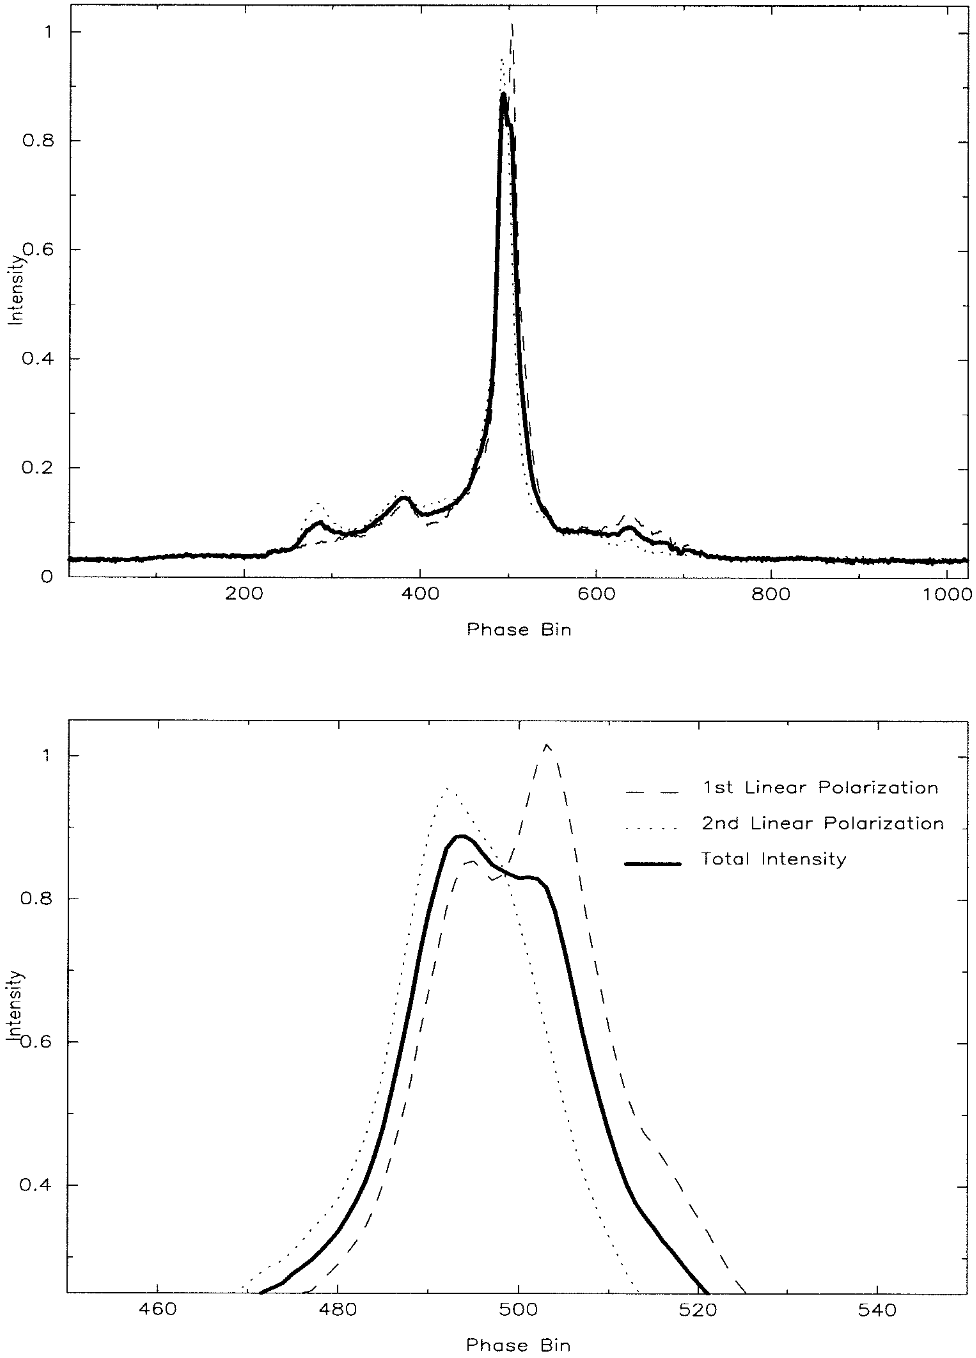

Figure 3.1: Pulse profile of PSR J0437-4715 at 1.4 GHz, in total intensity and each individual linear polarization. The top panel shows the pulse profile for a full period, divided into 1024 phase bins. The lower panel displays a zoomed in view of the pulse peak. The intensity scale is arbitrary and the linear polarization intensities have been scaled up by a factor of two for convenient comparison with the total intensity. This data was collected on 4th July, 1994, and used 128 MHz bandwidth.

file data was collected on 4th July 1994 using independent linear polarizations at 1.4 GHz center frequency and 128 MHz bandwidth. Figure 3.2 shows TOAs obtained from the same observing system. Each TOA was obtained from a 90 s observation and the whole observing session lasted about 6 hours, starting 3.5 hours before meridian crossing and ending 5 hours after. The pulsar crossed the local meridian at fractional MJD 0.99. The middle points are the TOA residuals for the total intensity profile for each 90 s observation, and the other two sets are the corresponding residuals in each linear polarization. The extreme swing in TOA for each polarization is explained by the change in the corresponding pulse profile with parallactic angle. The more gradual, but significant, drift in the total intensity TOAs occurs because we have simply added the two linear polarizations (in the absence of calibration). The opposite swings of the two polarizations cancel each other to some extent, but not completely, leading to the observed drift in the total intensity TOA. Using this same data set and applying variable gain to the two polarizations, we can demonstrate that a 1% calibration error produces a shift in TOA of about 100 ns. With formal TOA uncertainity from the template fitting of typically 100 ns, our data are clearly susceptible to such errors.

Since accurate calibrations were not obtained for the 1994 data, we used two different methods to reduce the systematic trends seen in figure 3.2. In the first, the pulse energies in the two orthogonal polarizations were equalized before forming the total-intensity profile. In the second, the relative amplitudes of the two polarizations were varied before adding to form the total-intensity profile. A TOA was calculated for each ratio, along with the associated TOA error. The ratio which minimized the TOA error was selected as the correct 'calibration' and the corresponding TOA used in the final solution. This approach can be likened to a matched filter in which we attempt to optimize the fit of the observed profile to that of the standard profile. Figure 3.3 illustrates the pre-fit residuals from the two methods, along with the uncalibrated pre-fit TOA residuals. The data set is the same as shown in figure 3.2. Even though the two approaches reduce the systematic trend considerably, the correction is not complete. The post-fit residuals in the two solutions show significant dependence

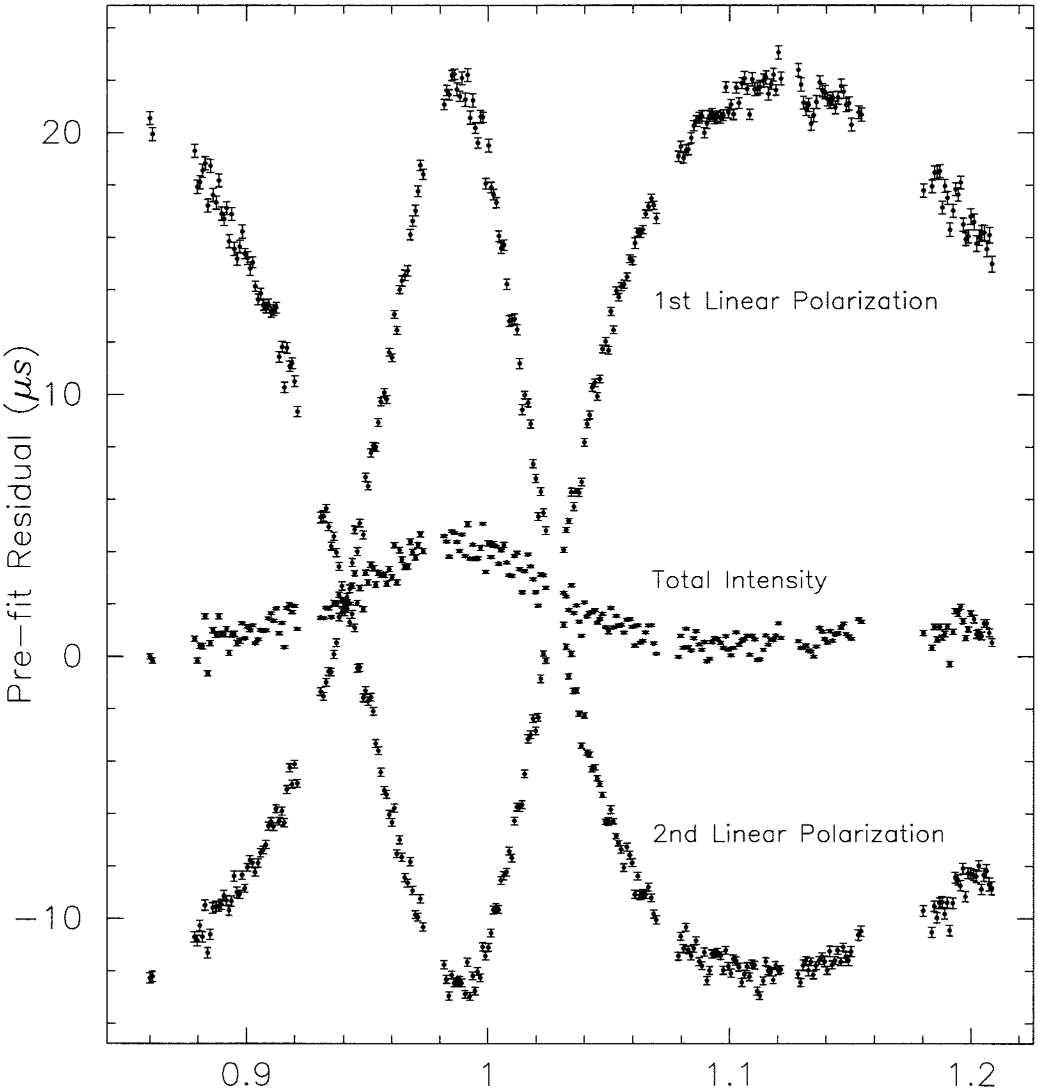

Figure 3.2: Time of Arrival (TOA) residuals for 1.4 GHz center frequency observations, with 128 MHz bandwidth. TOA residuals are plotted every 90 s, in total intensity as well as for each linear polarization. These observations were made on 4th July 1994.

on hour angle (especially in the case of linear polarization data). The remaining systematics in the two methods are comparable in magnitude, though the equal-fluxes approach appears slightly better.

Calibrated Data

Our calibration scheme was implemented in late 1994 and is explained in detail in chapter two. It consists of using a pulsed noise diode as a secondary calibrator. The noise diode is calibrated in absolute flux density once a day by observing the extragalactic radio source Hydra. Just before observing a pulsar, we make a three minute observation, slightly off-source, using the noise diode pulsed at 90 milliseconds (an artificial pulsar). This observation is used off-line to calculate the gains in the two polarizations. The pulsar data are then processed using these gains. In figure 3.4 we show pre-fit TOA residuals for calibrated data obtained on 20th April, 1996, using two orthogonal circular polarizations. As in figure 3.2, the total-intensity TOAs are plotted along with the TOAs for each polarization, for each 90 s observation. After July 1995, the dual band system was in place. Therefore, we are able to display data from two different frequency bands in figure 3.4. The two bands were centered at 1426 MHz and 1665 MHz and utilized total bandwidths of 128 MHz each. It is clear that the parallactic angle dependence of each polarization is still quite strong and is different in the two observing bands. The calibration appears to have worked in this case since the total-intensity residuals do not exhibit the deviations shown by the individual polarizations. There is a secular drift of the residuals through the day of about 1 microsecond. Since this is shared by all the polarizations, we can attribute it to an error in either the pulsar model or the time keeping, but not due to calibration errors.

Figure 3.5 illustrates pre-fit TOA residuals for calibrated data obtained on 4th June, 1997, using orthogonal linear polarizations. The same parallactic angle dependence is observed. However, in this case, the calibration does not seem to have done as good a job since we can clearly see the total intensity TOA residuals in the two different frequency bands diverging and then converging through the course of the

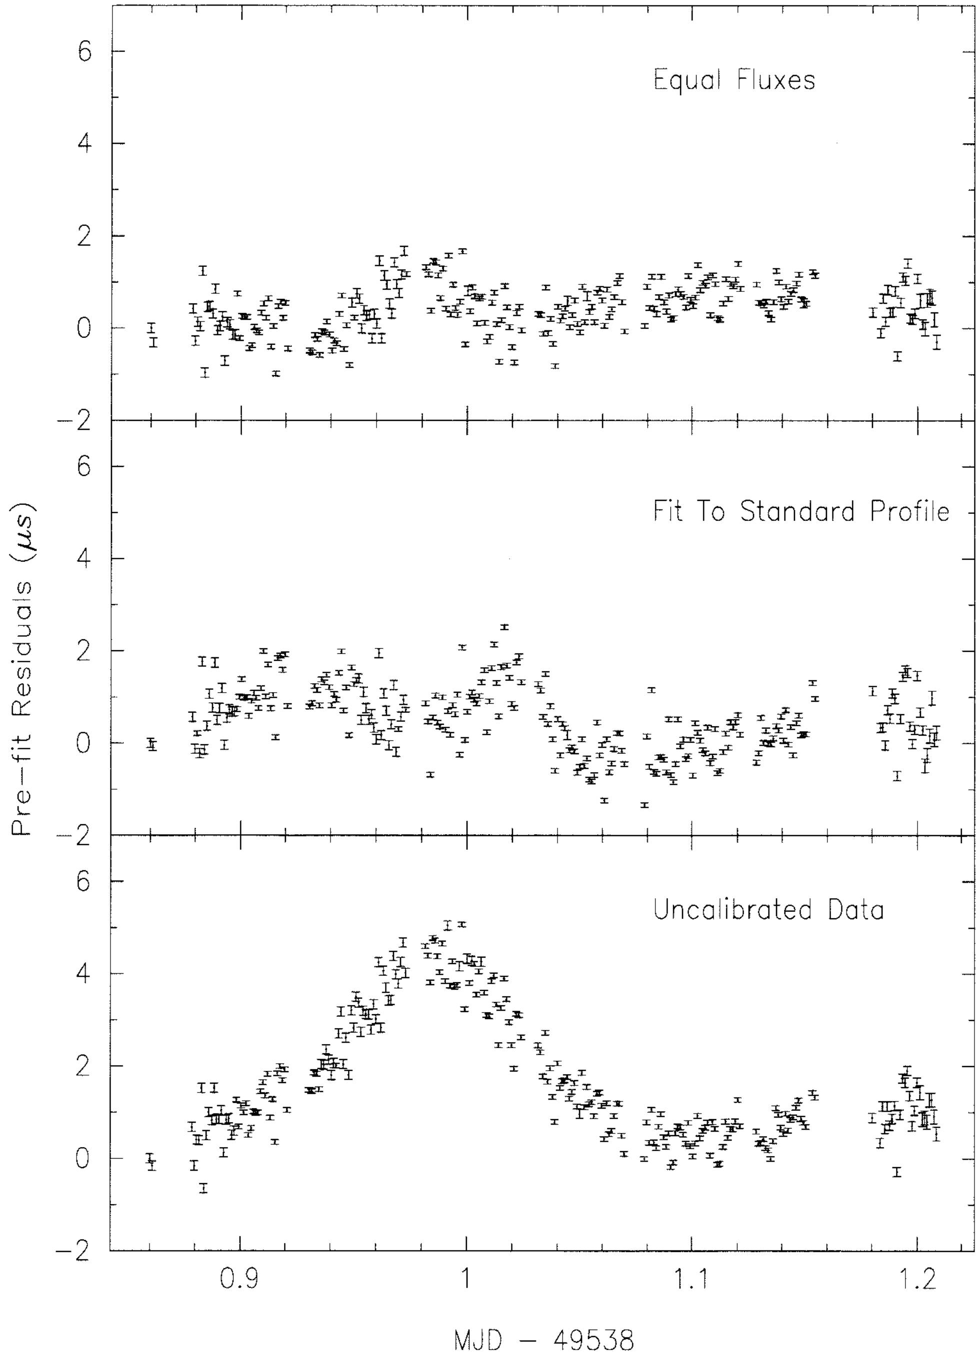

Figure 3.3: The top two panels show the pre-fit TOA residuals obtained with the methods outlined in the text for data obtained on 4th July 1994. The bottom panel shows the raw pre-fit residuals.

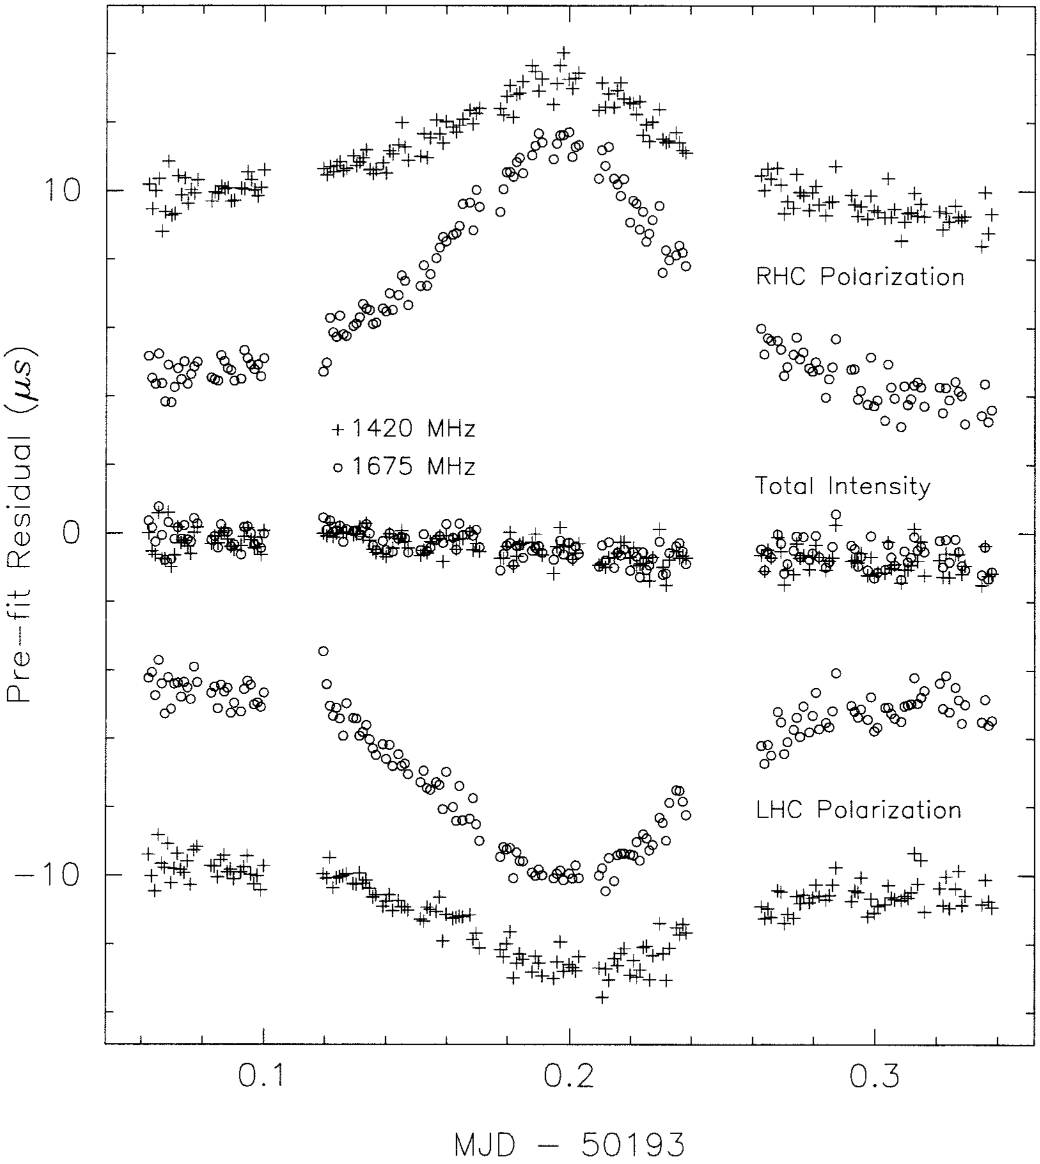

Figure 3.4: Time of arrival (TOA) residuals for two different frequency bands observed with orthogonal circular polarization feeds on 20th April 1996. Each TOA is from a 90 s integration, using 128 MHz bandwidth. Error bars (not shown) are about the size of the plotted symbols.

day. The divergence amounts to about 2 microseconds at its peak. DM variations would cause both frequency bands to shift in the same direction, albeit with a small difference. Since we only see one band shifting significantly, we can safely rule out DM variation as a source of the divergence. Each TOA is derived from a 128 MHz bandwidth with a 90 s observation duration. The relative gain variation between each polarization should therefore only be $\sim 1/(B\tau)^{1/2} \sim 10^{-6}$. A residual shift of 2 μs implies a gain variation of $\sim$ 20%. This discrepancy suggests that our calibration scheme is not reliable at the sub-microsecond level of precision.

We calibrate our secondary calibration source (the noise diode) once a day and collect noise diode data about once every half an hour. It is quite possible that a change in the elevation angle of the telescope, as it tracks the pulsar, distorts the telescope and feed support such that the calibration data is no longer valid at the 1% level after just a few minutes. This area requires careful study and is crucial to realize the full potential of the TOAs for PSR J0437-4715 . In the meantime, it is important to make observation using circular polarization feeds since that type of data appears to exhibit smaller systematics.

Frequency Structure of the TOAs

Apart from calibration problems, another source of concern about the TOAs for PSR J0437-4715 has been the frequency dependence of the residuals. Figure 3.6 shows a good example of this phenomenon. The plot displays pre-fit TOA residuals as a function of frequency for a 25 minute observation at a center frequency of 1410 MHz. The channel spacing for the data is 1 MHz (128 channels across 128 MHz total bandwidth). For this receiving system the FPTM was observing the lower-sideband sky frequency. Hence the frequency scale is inverted, though it does represent the correct baseband frequency order processed by the correlator, i.e., increasing frequency from left to right.

Two specific characteristics stand out from figure 3.6:

- Timing accuracy degrades towards higher baseband frequencies.

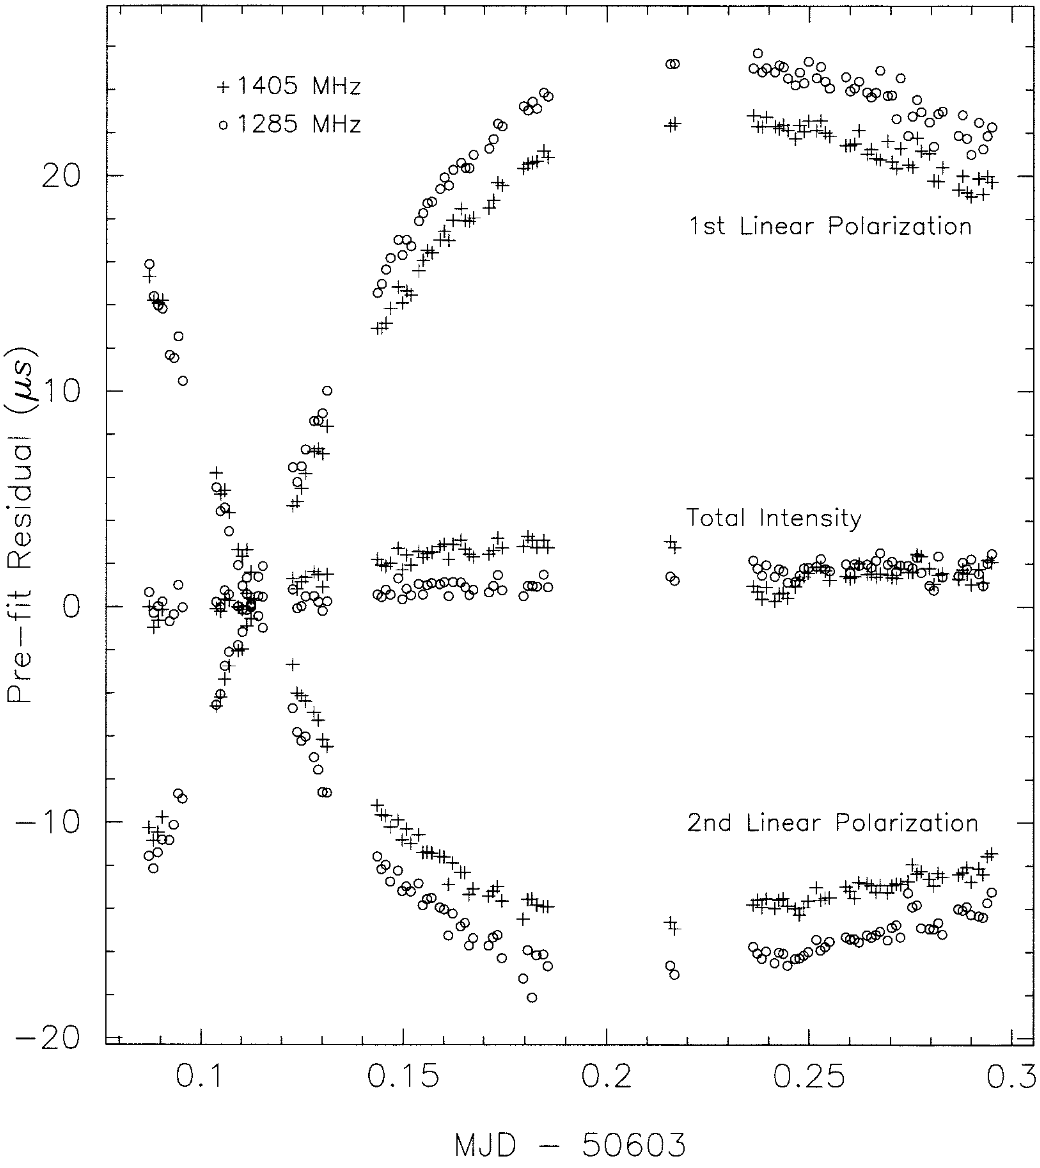

Figure 3.5: Time of arrival residuals for two different frequency bands observed with orthogonal linear polarization feeds on 4th June 1997. The total intensity TOA residuals in the two bands can be seen to diverge and then converge through the course of the day. Each TOA is from a 90 s integration, using 128 MHz bandwidth. Error bars (not shown) are about the size of the plotted symbols.

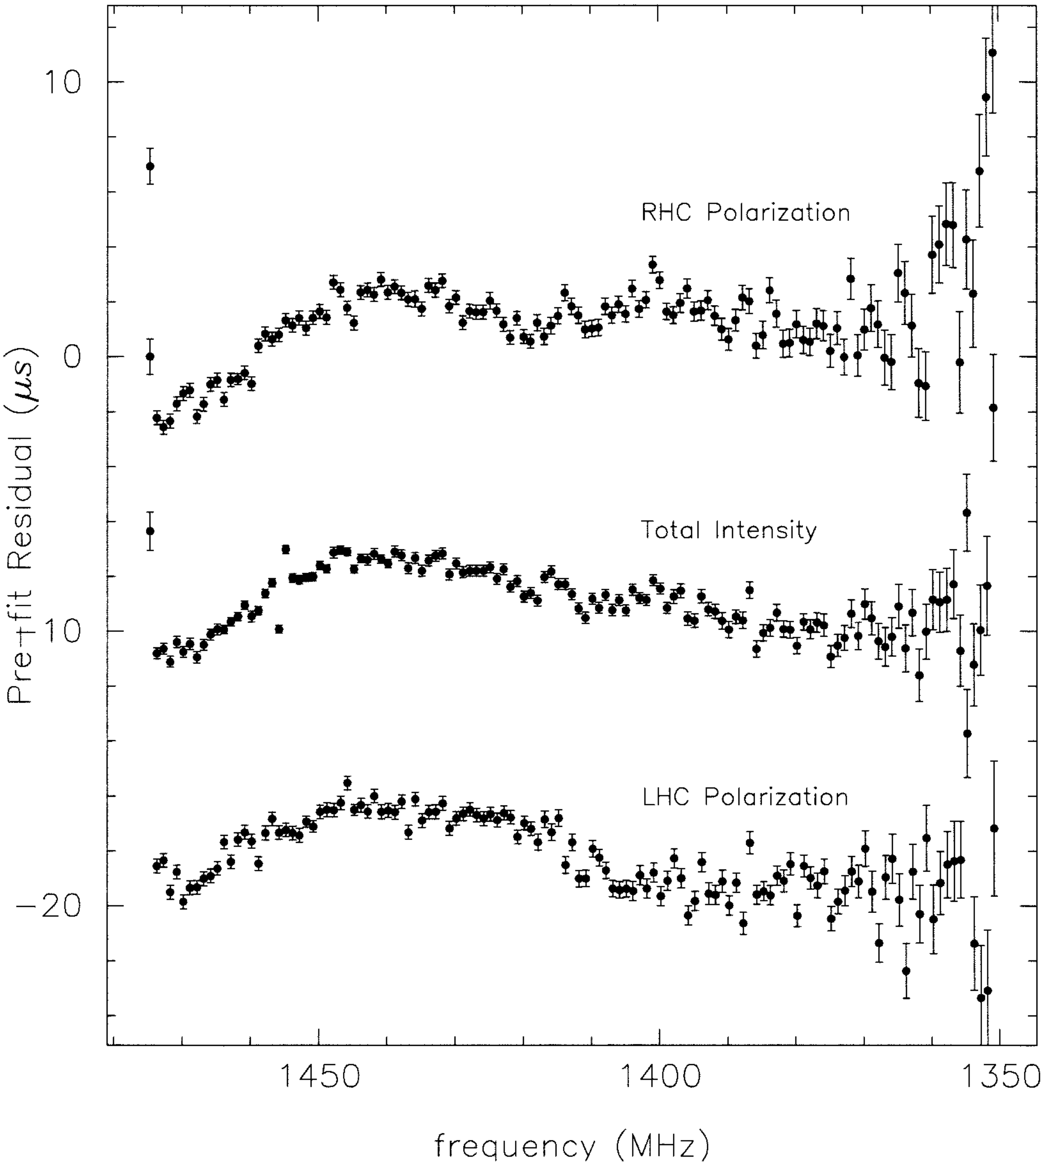

Figure 3.6: Time of arrival residuals as a function of frequency for an observation made on 20th April 1996 and lasting 25 minutes. Residuals are plotted as a function of channel frequency in two orthogonal circular polarizations as well as total intensity. The offset in the two polarizations is due to different pulse shapes.

- The residuals display a systematic curvature of $\sim 2\ \mu$s amplitude.

Though the two polarizations do not exhibit exactly the same systematic wiggle, the overall trend is the same in both of them and shows up in the total-intensity residuals as well. The same pattern can be seen in adjacent observing bands, i.e., the pattern is independent of the center frequency being used, at least within a range of a few hundred MHz. Since the timing accuracy is worse at lower frequencies, the pattern is less noticeable in the 70-cm and 50-cm observing bands.

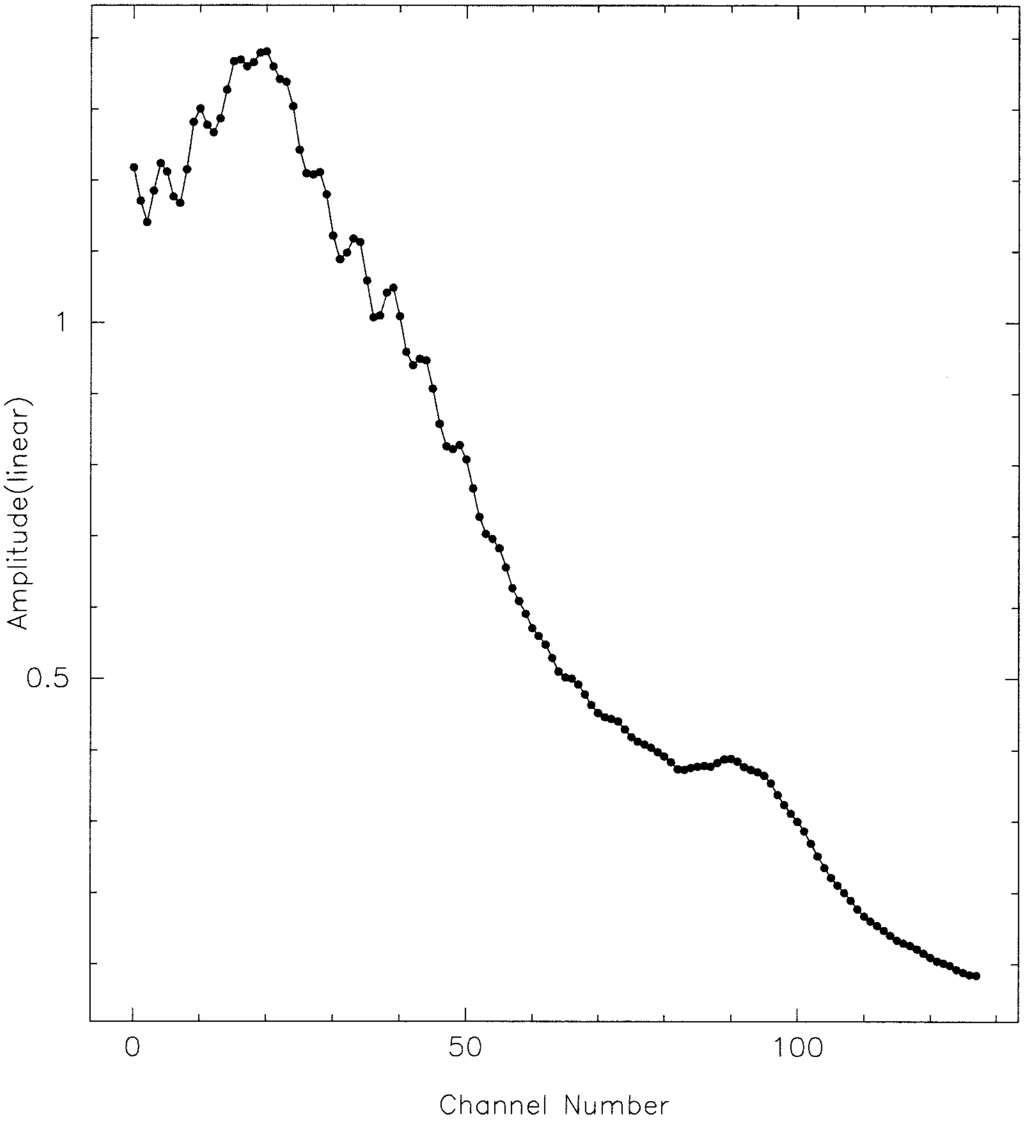

Figure 3.7: Observed bandpass at 1420 MHz using 128 MHz bandwidth. The plot shows the power spectral density recorded by the correlator in the FPTM on 20th April 1996.

We believe that the reason for both of the above features lies with the shape of the input spectral density to the digitizer. Figure 3.7 shows a typical spectral density plot as recorded after processing by the correlator. An ideal bandpass would have equal power distributed throughout the bandwidth. However, as discussed in section 2.2.3, the downconverter suffers from poor gain at the higher frequencies, resulting in the bandpass shown in figure 3.7. One effect of the distorted bandpass is uneven S/N response across the bandwidth. The quantization noise from the two bit digitizer is spread across the whole bandwidth, even though most of the signal is in the lower half of the bandwidth. Consequently, the S/N at the output of the correlator is worse at the higher baseband frequencies (corresponding to lower frequencies in figure 3.6). This phenomenon explains the degradation of the timing accuracy for the higher frequencies.

The 'wiggle' in the residuals is harder to explain. We can discount a DM error since that would produce a linear trend in the residuals. That leaves the baseband shape as the major suspect. Establishing the connection is difficult because the two phenomenon are in different dimensions - the bandpass is in amplitude while the 'wiggle' is in time. To properly uncover any possible connection, we implemented a full simulation of the FPTM signal processing in software. This is a substantial undertaking since the processing power required amounts to several gigaflops. Fortunately we had access to a supercomputer at Caltech.

End-to-End Simulation of the FPTM

The simulation involved the following steps:

- Generate a gaussian distributed random noise signal at the Nyquist rate (128 MHz).
- Modulate the noise signal at the pulsar period with a high S/N pulse profile with an amplitude consistent with that seen by the FPTM.
- Disperse the noise signal to simulate interstellar dispersion.

- Apply a bandpass shape filter to match the real bandpass shape of the down-converter.
- Digitize the noise signal using the same scheme as the real digitizer.
- Correlate the digitized signal in the same manner as the correlator chips.
- Accumulate the correlator products in memory to build a sufficient S/N pulse profile.

The above scheme was implemented in a C language program and run on the 512 node NCUBE machine in the Caltech Supercomputer center. This computer afforded a reduction in processing time by about a factor of 500, as compared to a SparcStation 20. The program was run with the following different options, to study their effect on the TOA residuals:

- Use a flat bandpass and DM = zero.
- Use a flat bandpass and use the DM for PSR J0437-4715.
- Use the observed bandpass but set the DM to zero.
- Use the observed bandpass and the DM for PSR J0437-4715 (the 'real life' situation).

The output of the simulation was processed in the manner identical to that for a regular FPTM observation. The TOAs were derived with respect to the modulating profile used in the simulation (the 'standard profile'). The results for the zero DM simulations and the flat bandpass simulations revealed no systematic trends in the TOA residuals as a function of frequency. The results using the real bandpass and the DM for PSR J0437-4715, on the other hand, were very interesting and are plotted in figure 3.8, along with the results from the simulation which used a flat bandpass. Compared with figure 3.6, there is a striking similarity in the overall pattern, if not in the exact magnitude for each channel. It is gratifying to note that the degradation

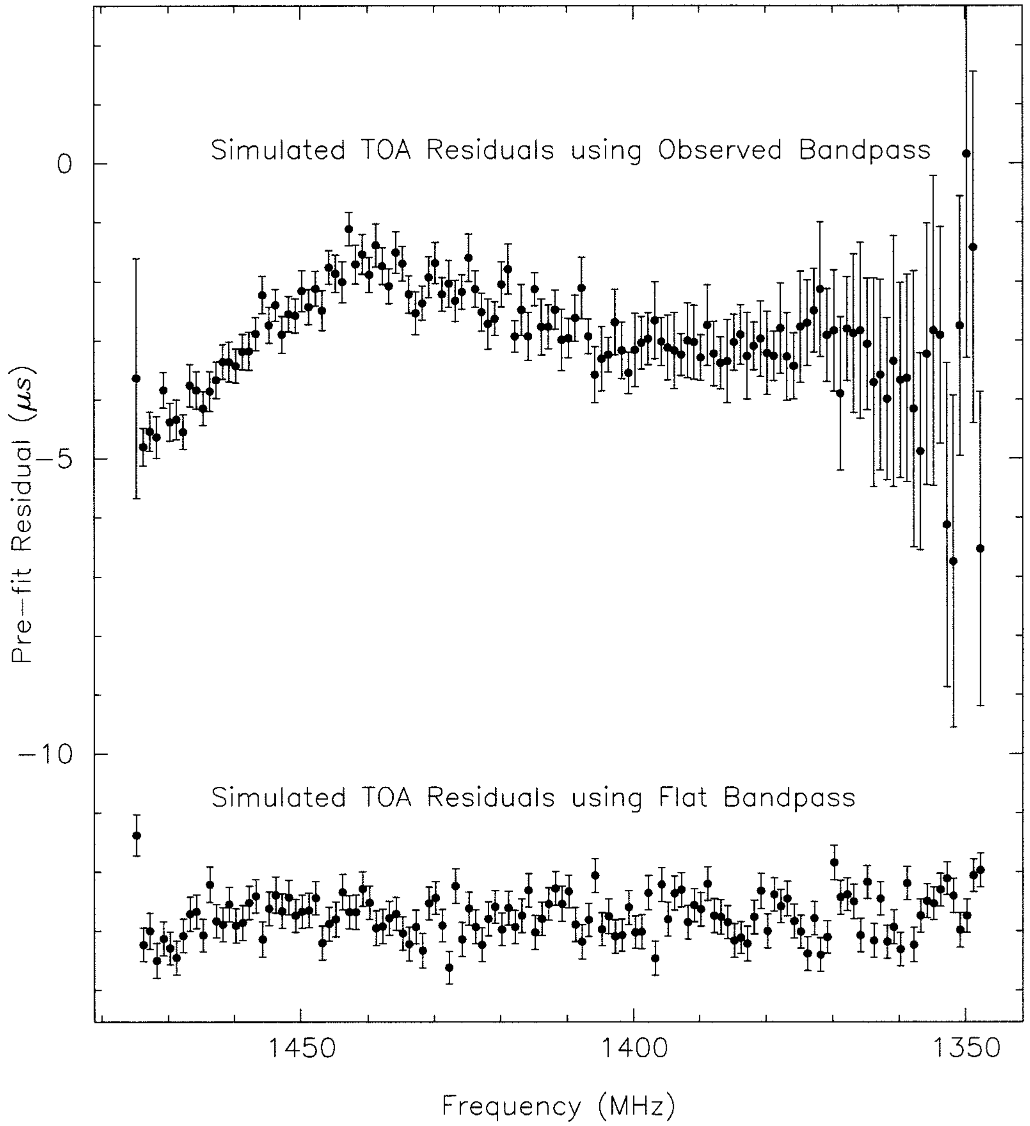

Figure 3.8: Simulated TOA residuals as a function of frequency and the bandpass shape used in the simulation. The two data sets have been plotted with an offset of 10 μs to avoid overlap.

in accuracy at the higher channels is reproduced, leading to some confidence in our simulations.

It is difficult to pinpoint the mechanism by which the TOA residuals are distorted. The amplitude of the effect – a few microseconds over the full bandwidth – is less than the size of one phase bin used for averaging – 5.7 microseconds for PSR J0437-4715. We believe, and as the simulations demonstrate, that the main culprit is the bandpass shape. A correlator does not offer good channel-to-channel isolation of the signal power, i.e., the sidelobe level for each channel is much higher compared to a well designed filter bank. When there is significant variation of power from one channel to another, enough power can leak from the higher power channel into the weaker channel, leading to distortion of the profile. For a flat window function in the autocorrelation domain, the highest sidelobe level in the frequency domain is only 13.5 dB below the main peak. In contrast, the triangular window ('Parzen Window') offers a 27 dB sidelobe level, but also leads to a main lobe which is wider by about a factor of 2. The residuals in figures 3.6 and 3.8 were obtained with a flat window function to avoid smoothing out the systematic trend by using a triangular window. We have actually processed all of our pulsar data using the triangular window in order to reduce the influence of narrow band interference on adjacent channels. The trend is still present, but with a slightly smaller amplitude.

The only way to permanently remove the frequency domain systematic trends is to correct the bandpass shape so that the power is equally distributed throughout the band. As discussed in section 2.2.3, this is a difficult undertaking. An attempt has been made by Caltech engineer John Yamasaki, but the bandpass remains more or less the same. A radical redesign of the downconverter system may be required to meet the requirements.

3.3 Results

The TOAs for PSR J0437-4715 were analysed using the program TEMPO (Taylor & Weisberg 1989) with the JPL ephemeris DE200 (Standish 1982) to give the pulsar and binary parameters. We actually processed two sets of TOAs. The two sets differed only in the methods used to 'calibrate' the 1994 data, as described in the previous section. In one set we equalized the flux in each polarization before summing to form the total intensity profile. In the second set we varied the relative gain for the two polarizations till the least TOA error was obtained. The two methods produced rms residuals of 490 and 550 ns, respectively, with significant differences in some of the astrometric and binary parameters. Figure 3.9 shows post-fit residuals for one of the two solutions (the other set of residuals show no significant difference and hence have not been plotted). For each of the two solutions we also varied the time offset between different observing systems, and different center frequency observations. This is necessary because the pulse profile changes significantly with observing frequency, bandwidth and the receiver system. Three time offsets were required and were of the order of a microsecond.

Astrometric and binary parameters for PSR J0437-4715 obtained from TOAs for frequencies above 1 GHz are given in Table 3.1. Parameter values are the mean of those given by the two solutions. Errors (given in parentheses) are in the last quoted digit and are the sum of half the difference between the two solutions and the mean rms error for that parameter, as given by TEMPO. Typically, the difference between the two solutions was $2 - 10$ times the formal rms error.

The DM was obtained by including 660 MHz TOAs in the data set and solving for DM. Fractional DM variations were $< 10^{-4}$, consistent with DM fluctuation found by Phillips & Wolszczan (1991) for several pulsars. They reported mean $\delta\mathrm{DM} = 1.5 \times 10^{-5}\mathrm{DM}^{1.3}$, which predicts $\delta\mathrm{DM} = 5.3 \times 10^{-5}$ for PSR J0437-4715 .

3.3.1 Position and Proper Motion

The pulsar position (relative to the DE200 reference frame) is determined to extremely high precision ~ 50 micro-arcseconds, allowing us to determine the pulsar proper motion very accurately and to make a significant measurement of the annual parallax. The positional accuracy matches that of PSR B1937+21 (Kaspi, Taylor, & Ryba

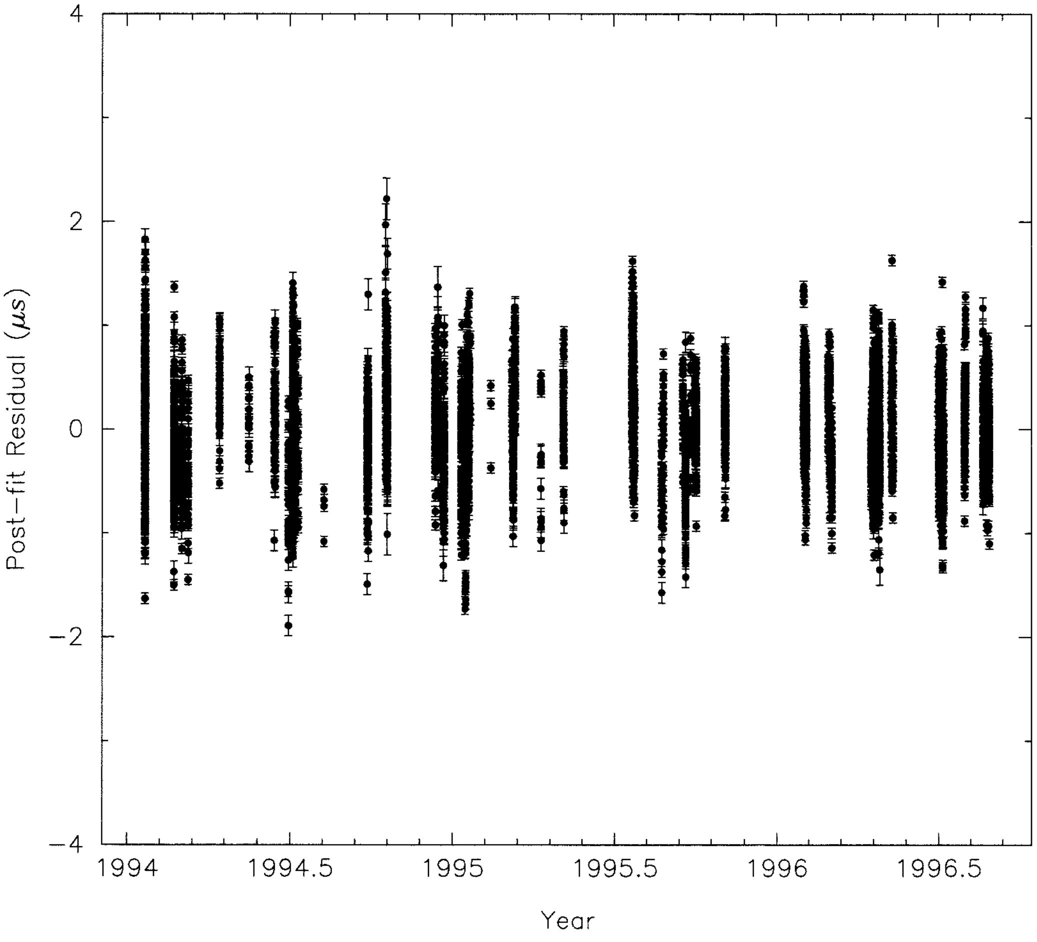

Figure 3.9: Post-fit residuals for PSR J0437-4715 for the full 2.6 year data span

Observed and Derived Parameters for PSR J0437-4715

Parameter	Value
R. A. (J2000)	$04^{\rm h}\ 37^{\rm m}\ 15^{\rm s}.748182(4)$
Dec. (J2000)	$-47^\circ\ 15'\ 08''.23145(5)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	121.34(6)
μ_δ (mas y^{-1})	-72.50(3)
Annual parallax (mas)	5.6(8)
Period, P (ms)	5.75745182525633(6)
Period derivative $\dot{P}$ (10^{-20}) (s s^{-1})	5.7295(9)
Period second derivative d^2P/dt^2 (10^{-30}) (s s^{-2})	0.80(40)
Epoch of period and position (MJD)	50019.00
Dispersion Measure (cm^{-3} pc)	2.6469(1)
Binary period, P_b (d)	5.741042353935(350)
$x = a_p \sin i$ (s)	3.36668528(4)
Eccentricity	0.00001920(2)
Long. of periastron, ω (°)	1.793148(20000)
Epoch of periastron (MJD)	50000.49656856(40000)
$\dot{x}$ (10^{-12}) (lt-s s^{-1})	0.082(4)
$\dot{\omega}$ (deg y^{-1})	0.0(11)
$\dot{P}_b$ (10^{-12}) (s s^{-1})	5.5(15)
$\dot{e}$ (s^{-1})	0.0003(14)
Timing data span (MJD)	49373 – 50323
RMS timing residual (μs)	0.49, 0.53
Parallax distance (pc)	178(26)
DM Distance (pc)	140
$\dot{P}$ distance limit (pc)	205
Transverse Velocity (km s^{-1})	120
Position angle of proper motion (°)	120.5(1)

Table 3.1: The two RMS timing residual values are for the two different solutions obtained from the 1994 data (as described in the text). Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

1994), which was obtained after 8 yrs of timing, compared 2.6 yrs for PSR J0437-4715. An interferometric determination of the pulsar position will greatly assist in measuring the offset (in position and rate) between the ecliptic reference frame used for timing and the equatorial reference frame used for interferometry (determined from positions of extragalactic radio sources). The position of PSR J0437-4715 in the southern sky is fortuitous for this purpose, since the other millisecond pulsars for which precise position determinations have been made through timing, PSRs B1937+21, B1855+05 (Kaspi et al. 1994) and J1713+0747 (Camilo et al. 1994), are all clustered together in the northern sky. The position of PSR B1937+21 has been determined to about 3 mas by VLBI (Bartel et al. 1996), and 30 mas with the VLA (Foster & Backer 1990). Since PSR J0437-4715 is brighter, we can expect a more precise determination of its position with interferometry, thus allowing for a more accurate frame tie.

It is instructive to consider future improvements of the pulsar position. The minimum size of the binary orbit is $c\,x = c\,(a_p \sin i/c) = 10^{11}$ cm. At a distance of 178 pc (see Distance, section 3.3.2), this is equal to an angular size $\theta = 40\,\mu$as. Since we have determined $i < 43^\circ$ (section 3.3.3), this limit actually increases to $\theta = 55\,\mu$as. The binary motion will therefore modulate the position of the pulsar with this amplitude (with the binary period). This amplitude is equivalent to our current precision and hence will prevent any further improvement in the position determination. Overcoming this obstacle will require adding to TEMPO's capabilities so that it can model the binary modulation of position (and other relevant parameters).

Our timing proper motion is consistent with the measured optical proper motion (Danziger et al. 1993) and with previous timing values, but our new distance estimate increases the transverse speed slightly, to 120 ± 20 km s^{-1}. This is consistent with the speed inferred from scintillation observations by Nicastro & Johnston (1995) if allowance is made for the revised distance estimate. The velocity is moderate compared with many of the longer period pulsars for which we have accurate velocities, and is in accord with the velocities of other millisecond pulsars (Toscano et al. 1999).

3.3.2 Distance

The Taylor and Cordes model (Taylor & Cordes 1993) for the galactic electron density distribution gives a distance estimate for PSR J0437-4715 of 140 pc with an uncertainty of about 25%. Our parallax determination places the pulsar at a distance of 178 ± 26 pc, indicating that the electron density along this line of sight has been overestimated by about 30%. We obtain an average electron density of 0.015 cm^{-3}, which is somewhat higher than the average $n_e \sim 0.005$ cm^{-3} believed to exist in the local interstellar bubble (Cox & Reynolds 1987). This indicates that the bubble probably does not extend all the way to the pulsar and provides an additional constraint in determining the bubble's morphology.

The proper motion of any object induces an apparent period derivative, in addition to its intrinsic value. This effect was first described in the context of pulsars by Shklovskii (1970) and its implications for millisecond pulsar timing were discussed by Camilo, Thorsett, & Kulkarni (1994). Since PSR J0437-4715 has a large proper motion, this effect contributes $\sim$ 70% of the observed period derivative ($\dot{P}$). This is significant, since $\dot{P}$ is a fundamental parameter which gives us estimates of the pulsar age, magnetic field and spin-down luminosity. The kinematic contribution to the observed period derivative is given by:

$$\dot{P}_{PM} = \frac{v^2 P}{cd} \tag{3.1}$$

where P is the pulsar period, v is the transverse velocity and d is the distance to the pulsar (Shklovskii 1970). Since the actual observable is proper motion, μ, we can recast equation 3.1 into a more meaningful form by using $v = \mu d$. Hence $\dot{P}_{PM} = \mu^2 P d / c$. Using our parallax distance, we obtain $\dot{P}_{PM} = (4.9 \pm 0.7) \times 10^{-20}\,\mathrm{s\,s^{-1}}$. Since $\dot{P}_{obs} = 5.729 \times 10^{-20}\,\mathrm{s\,s^{-1}}$, we obtain $\dot{P}_{int} = (0.8 \pm 0.7) \times 10^{-20}\,\mathrm{s\,s^{-1}}$, where $\dot{P}_{int}$ is the intrinsic period derivative. If we assume that $\dot{P}_{int}$ is positive, $\dot{P}_{obs}$ sets an upper limit on $\dot{P}_{PM}$ and hence on the distance. Using $\dot{P}_{PM} \leq \dot{P}_{obs}$, we obtain $d \leq 205$ pc. This is just inside our 1σ upper parallax distance limit of 208 pc.

The proper motion of PSR J0437-4715 also makes a significant change in the orbital period derivative. Bell & Bailes (1996) show that the expected kinematic contribution to $\dot{P}_b$ is very much greater than the general relativistic $\dot{P}_b$. For an inclination angle $i = 30°$, general relativity predicts $\dot{P}_b = -3.5 \times 10^{-16}\,\mathrm{s\,s^{-1}}$, whereas the measured $\dot{P}_b = 5.0(15) \times 10^{-12}\,\mathrm{s\,s^{-1}}$. The observed $\dot{P}_b$ can therefore be used in the same manner as the pulse period derivative to yield an independent distance estimate. We obtain $d = 227 \pm 62$ pc. This distance is consistent with the parallax and period derivative estimates. Since the relative error in $\dot{P}_b$ scales as $t^{-5/2}$, this term should ultimately provide the most accurate distance to the object.

3.3.3 Orbital Inclination

Kopeikin (1996) discussed several observable effects of proper motion for binary pulsars. For PSR J0437-4715 the most significant is a secular variation in $x \equiv a_p \sin i$ (the projected semi-major axis of the pulsar orbit). This occurs because the inclination i of the orbit to our line of sight changes as the system moves across the sky. The expected secular increase in x for PSR J0437-4715 is about $1\mu\mathrm{s\ yr^{-1}}$. There is also a secular variation in ω, the longitude of periastron but for small eccentricity (as in PSR J0437-4715), any secular change in ω is absorbed into a redefinition of the orbital period. The effect on x is given by

$$\delta\dot{x} = x \cot i \ (-\mu_\alpha \ \sin\,\Omega \ + \ \mu_\delta \cos\,\Omega) \ \mathrm{s\,s^{-1}}, \tag{3.2}$$

where $\Omega\,(0 \leq \Omega < 2\pi)$ is the position angle of the ascending node of the binary orbit, and $\mu = (\mu_\alpha,\ \mu_\delta)$ are the components of the proper motion. Since orientation of the orbit with respect to proper motion is unknown, we have to solve equation 3.2 for i using the permissible range of Ω. This yields an upper limit on $i < 43°$. The likelihood of $\dot{x}$ being caused by a real change in a_p can be safely ruled out because the observed $\dot{P}_b$ is very small (a change in a_p would be accompanied by a change in $\dot{P}_b$). The observed $\dot{P}_b$ implies that any intrinsic $\dot{x} \leq 3.4 \times 10^{-17}\,\mathrm{s\,s^{-1}}$, many orders of magnitude smaller than the observed value.

3.3.4 Masses and Age of the Binary System

Apart from the astrometric and kinematic measurements, a fundamental goal of pulsar timing is to measure the masses of the pulsars and their companions in binary systems. These provide the pivotal points for testing theories of neutron star formation and binary evolution. The five Newtonian parameters of the timing solution only provide us with the mass function which has three unknowns - pulsar and companion masses and the orbital inclination. Fortunately, for PSR J0437-4715 the limit on orbital inclination narrows down the parameter space for the masses. Figure 3.10 illustrates the mass limits implied by $i \leq 43°$, showing that $M_{WD} \geq 0.22$ $\mathrm{M}_\odot$(for $M_{NS} = 1.4$ $\mathrm{M}_\odot$).

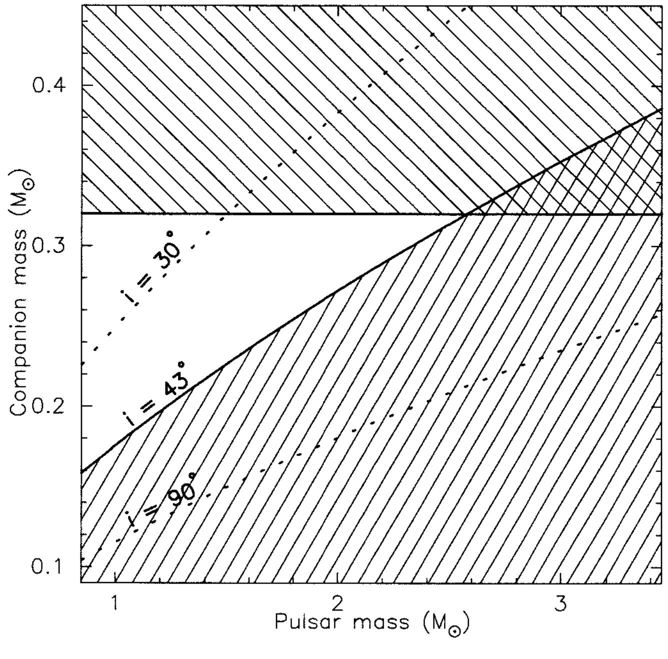

Figure 3.10: Mass constraints for the PSR J0437-4715 system. The lower hatched area is excluded by the limit on inclination angle obtained from the secular change in the projected orbital semi-major axis. The upper hatched area is excluded by the limit derived using the parallax distance and the cooling models of Hansen (1996). Dotted curves indicate the mass function for the labelled inclination angle.

Recently, Hansen (1996) computed detailed cooling models for low-mass helium white dwarfs with hydrogen atmospheres; the companion to PSR J0437-4715 is believed to be a star of this type. Following Hansen, we assume that $T_{\rm eff} = 3950 \pm 550\,{\rm K}$, and $m_I = 19.5 \pm 0.2$. Using his cooling curves and our parallax distance limits, we obtain an upper limit $M_{WD} \leq 0.32\ {\rm M}_\odot$, which is also indicated in figure 3.10. The mass range $0.22\ {\rm M}_\odot \leq M_{WD} \leq 0.32\ {\rm M}_\odot$ lies slightly higher than the value predicted by Phinney & Kulkarni (1994), although we note that our lower limit can be decreased by assuming a smaller pulsar mass. From the derived minimum white-dwarf mass of $0.22\,{\rm M}_\odot$, we can also put an upper limit on the distance to the binary and check the consistency of our parallax determination. This can be done because the minimum mass implies a maximum radius for the white dwarf, and for a given range of $T_{\rm eff}$, the observed apparent magnitude implies an upper limit on the distance. Using Hansen's cooling curves, we obtain $d \leq 183$ pc, just larger than the parallax distance and consistent with our other limits. In figure 3.11, we have collected together the different distance limits to PSR J0437-4715. These include the parallax determination, the period derivative, the orbital period derivative, and cooling model limits. The overlap range from the figure is $162 < d < 205$ pc, slightly narrower than the parallax distance determination.

The cooling curves of Hansen (1996) also provide an age estimate for the white dwarf: $3.1 \leq t_{\rm cool} \leq 6.0\,{\rm Gyr}$. The timing age of the pulsar is defined as $\tau = P/\,(n-1)\,\dot{P}$, where n is the braking index. For $n = 3$, our $\dot{P}_{int}$ limits imply that $\tau \leq 6.0\,{\rm Gyr}$, consistent with the range of $t_{\rm cool}$. Our derived $\dot{P}_{int}$ also puts an upper limit on the pulsar magnetic field $B \leq 3 \times 10^8$G.

3.4 Long Term Stability

3.4.1 Millisecond Pulsars as Stable Clocks

In view of the stability of millisecond pulsars, many authors have suggested their use as long term phase references to supplement atomic time scales. These time

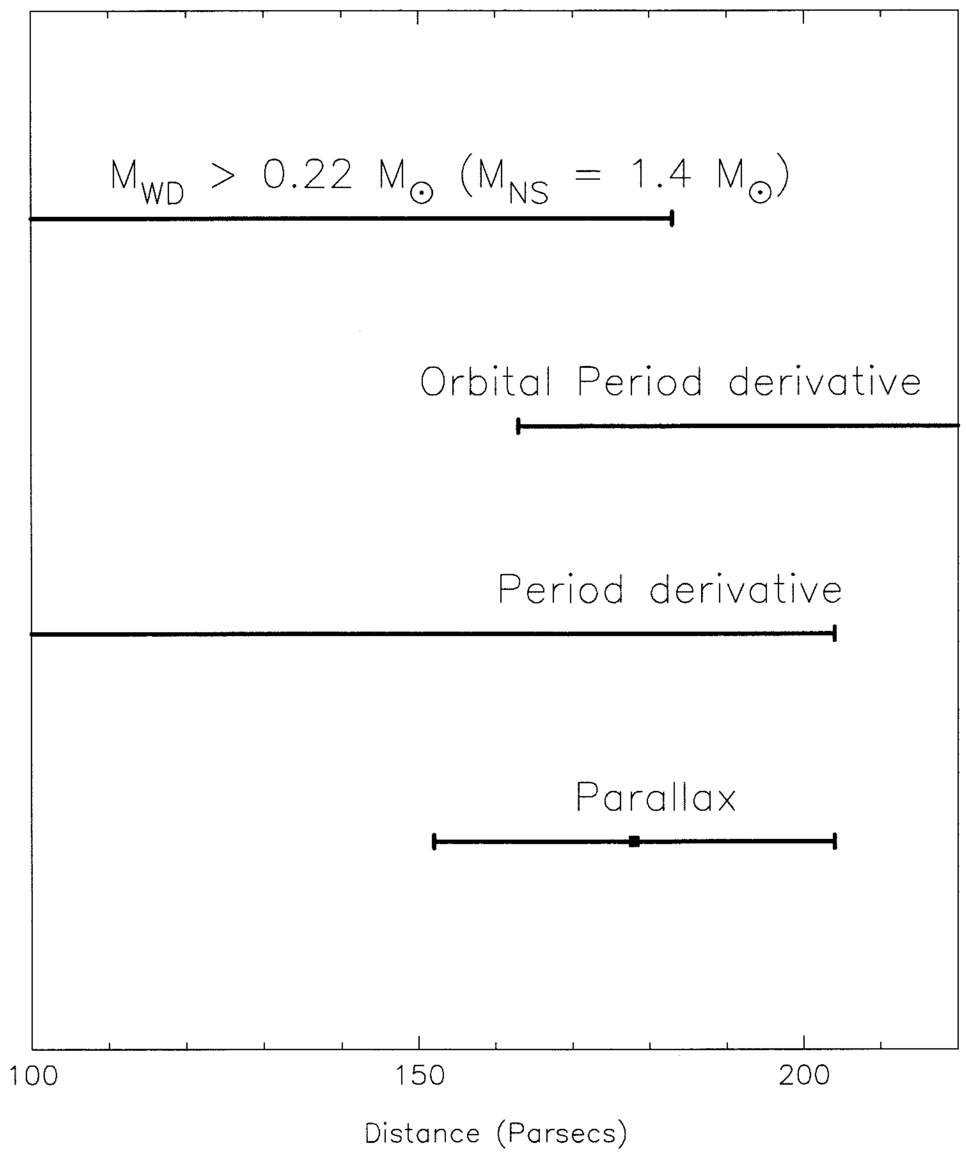

Figure 3.11: Various distance limits for PSR J0437-4715 as derived in the text. Error bars represent 1-σ uncertainties (where applicable).

scales are constructed from an ensemble of several hundred atomic clocks all over the world. Their data are combined with appropriate weighting proportional to their stability. In order to use millisecond pulsars for time keeping, we would also need to have such an ensemble with well studied properties. It could then be used to isolate errors in the atomic time scale as well as the solar system ephemeris used to produce barycentric pulsar TOAs. This is possible since a clock error will produce an indentical signature for all pulsars while an ephemeris error will produce a dipolar signature on the sky. The data on millisecond pulsar stability is still meager and the tools used to quantify it are in evolution. Here we present our stability analysis for PSR J0437-4715, comparing it with atomic time scales and other closely studied millisecond pulsars.

The best time scales on Earth are based on atomic clocks and exhibit stabilities of 10^{-14}, i.e., in three years (roughly equal to the length of the data span for PSR J0437-4715) the clock used to measure the TOAs will be in error by $\sim 10^{-14} \times 10^{8} = 10^{-6}$ s. Since the TOA precision for PSR J0437-4715 is a fraction of this, the error in the time scale is easily absorbed by the fit for period and it is not surprising that the period for PSR J0437-4715 is now known to a precision of about 10^{-14} (Table 3.1). Since pulsar timing is essentially a clock comparison experiment between the pulsar and an atomic time scale, it is easy to absorb any instability in the time scale into the pulsar parameters. For example, Guinot & Petit (1991) state that TAI (International Atomic Time) exhibits an annual variation with a peak-to-peak amplitude of 300 ns. This happens because TAI is comprised of many industrial Caesium clocks which suffer seasonal frequency variations due to the effects of humidity. Some of this variation therefore shows up in TAI. If the timing precision for a pulsar is high, some of this annual variation will be absorbed by parameters like position or parallax (those with an annual or semi-annual period).

Clearly, it is difficult to isolate the cause of any observed instability in millisecond pulsar timing data – whether it is intrinsically due to the pulsar, or because of some external factors. Kaspi et al. (1994) demonstrate this issue very clearly with their data on the millisecond pulsars PSRs B1937+21 and B1855+09. Both pulsars are very stable and have sub-microsecond rms timing residuals on short time scales (less than a year). However, their long term behaviour is quite different. After several years of timing, the residuals for PSR B1937+21 show large low-frequency noise in the form of a cubic polynomial variation, while PSR B1855+09 residuals have negligible power in the cubic term. The cubic polynomial is expected to be the dominant term in the residuals because the fit for period and period derivative absorbs any power in the linear and quadratic coefficients. The absence of low frequency noise in PSR B1855+09 could, however, be masked by higher radiometer noise since it has a lower flux density than PSR B1937+21. The latter pulsar is also much further away, leading to significant variations in its arrival times to do inhomogenieties in the interstellar medium (ISM). The first order effect is variations in its DM. The authors (Kaspi, Taylor, & Ryba 1994) have measured the DM variations and accounted for them in the timing model and claim that further ISM effects are negligible, and consequently blame the pulsar for the observed instability. Others authors (Foster & Cordes 1990) have suggested that the DM correction may not be complete at the sub-microsecond level and additional perturbations (position drift of the pulsar, geometrical path length differences) could account for comparable ($\sim$ 500 ns) effects in the arrival times. The source of PSR B1937+21's observed instability is therefore not resolved.

The results for PSR J0437-4715 promise to be free of such doubtful issues because of several factors. It is very bright ($\sim$ 90 mJy at 1.4 GHz), leading to high S/N, and its low DM (2.69) ensures that ISM effects will be small and negligible. Its period derivative is also smaller than that of PSR B1937+21 by at least a factor of 10. This parameter has been seen to correlate well with the level of timing noise observed in slower pulsars (Arzoumanian et al. 1994). The timing noise of PSR B1937+21 is consistent with this correlation (equation 1.1). Long term observations of PSR J0437-4715 should therefore be very helpful in answering the question of whether millisecond pulsars are the most stable clocks known to us.

3.4.2 Measuring (In)stability

Clock instabilities have traditionaly been measured by comparing the square root of the Allan variance, σ_y, of the second differences of clock offset measurements. Second differences are used because even though atomic clocks may have frequency offsets between each other, their frequency drifts are much smaller. On the other hand, pulsars have non-zero frequency derivatives because of their spin down torques and a fit for the period derivative removes the second differences. Consequently the Allan variance is inapplicable to pulsars.

Recently Matsakis, Taylor and Eubanks (1997) have defined a new statistical measure to allow easy comparison between atomic clocks and pulsars. Analogous to the Allan variance, they define $\sigma_z(\tau)$ as the weighted root-mean-square of the coefficients of the cubic polynomial terms fitted over intervals of length τ. The polynomial approach is suitable since pulsar observations are not made at regular intervals, yet this definition produces equivalent results when used for equally spaced data (for example atomic clock offset measurements). They applied this statistic to the data for PSRs B1937+21 and B1855+09 and demonstrated its utility in making meaningful comparisons with atomic clock data. We have adopted their approach and used the program *sigmaz* developed by the above authors and made available freely for all pulsar observers. Figure 3.12 displays our results for PSR J0437-4715. Results for PSRs B1937+21 and B1855+09 are shown as well (the data for these pulsars was provided by the authors of *sigmaz*). Also plotted are results for a cesium clock with superior performance, maintained at the U.S. Naval Observatory (USNO) (Matsakis 2000, personal communication). TT96 is a retrospective time scale created by the Bureau International des Poids et Mesures (BIPM), taking into account any corrections necessary to make UTC more stable.

Like the other two pulsars, PSR J0437-4715 displays increasing stability (decreasing σ_z) up to its maximum data span of about three years. This range corresponds to the dominance of white phase noise in the residuals. Beyond 2-3 years, both PSR B1937+21 and PSR B1855+09 display higher instability corresponding to low fre-

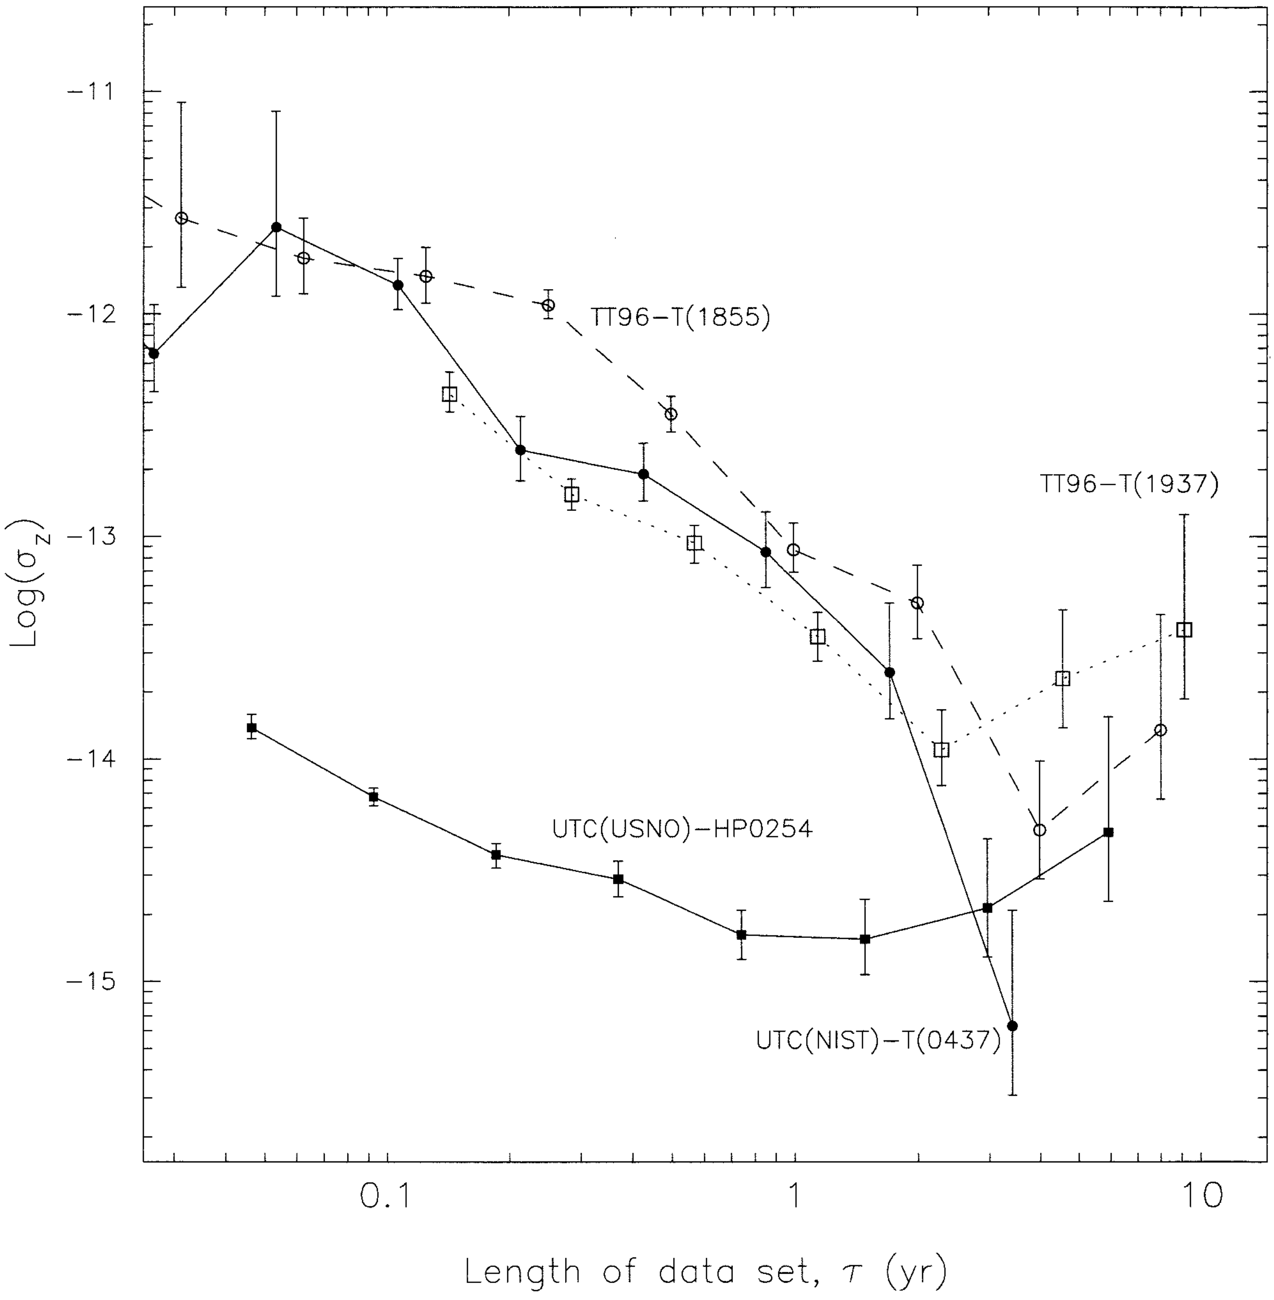

Figure 3.12: Instability measurements for PSR J0437-4715 and two other millisecond pulsars, according to the definition of σ_z in the text. Also shown is data for a cesium atomic clock maintained at the USNO.

quency noise caused, perhaps, by frequency random walk (the cause of low frequency noise in pulsars and atomic clocks, is an unanswered question). It is too early to say anything conclusive about the long term behaviour for PSR J0437-4715, but at least in the short term its stability is equivalent, if not better, than the millisecond pulsars studied so far.

3.5 Future

As figure 3.9 shows, our currently achieved timing precision of 500 ns is worse than the TOA accuracy from template fitting. The principal reasons for this discrepancy appear to be errors in calibration and systematic deviations in frequency due to an uneven bandpass. For short time scales of up to a few minutes we have been able to achieve 100 ns rms noise, and it will be interesting to see whether the long term precision can be pushed down to this unprecedented level (or maybe even lower).

Recently, Britton et al. (2000) have studied the influence of miscalibration on the TOAs for PSR J0437-4715. Using the Stokes parameters, $S = (I, Q, U, V)$, representation of the pulsar radiation, they demonstrate that the transformations of these parameters during propagation and detection of the radiation can be represented by the Lorentz transformations. Thus the Stokes parameters behave like a Lorentz 4-vector. The total intensity, I, is the timelike component and the other three parameters, which represent the polarized flux, form the spacelike component. Consequently, the total intensity can mix into the polarized flux and viceversa, depending on the cross-coupling between the two detected polarization states. Since the degree of polarization of the pulsar radiation varies through the pulse phase, such mixing will distort the total intensity pulse shape, thereby leading to an error in the TOA.

Their analysis offers one possible solution to this problem. They show that the quantity $(I^2 - Q^2 - U^2 - V^2)$ is invariant under Lorentz transformations (up to a multiplicative constant), i.e., it is the polarimetric analogue of the Lorentz invariant. Interpreting the Stokes parameters in the usual sense, this quantity is just the unpolarized component of the pulsar radiation. A viable timing strategy could therefore be to time this invariant profile instead of the total intensity. They tested this proposition with data from the FPTM as well as a coherent dedispersion tape recorder system and demonstrated that the calibration systematics in the TOAs are removed by timing the invariant profile. This has prompted us to change the FPTM observing strategy such that all Stokes parameters are recorded for PSR J0437-4715. The fact that the FPTM can be configured for measuring cross-correlation is invaluable in this regard. The only remaining systematic effects are due to the uneven bandpass.

It is gratifying to note that, despite the systematic errors in our data set, our various independent distance estimates agree well with one another, and that we are able to achieve a level of timing accuracy comparable to the best now available.

Chapter 4 A Menagerie of Millisecond Pulsars

4.1 Introduction

PSR J0437−4715 has been the focus of study for this thesis because of its short period, high flux density and low dispersion measure. The resulting high precision timing offers invaluable information for studying the long term stability of atomic clocks, frame tie between the interferometric, ecliptic and optical reference frames, and putting limits on the level of any gravitational wave background. Though many other millisecond pulsars are observable with the FPTM from Parkes, they do not offer the same high quality of timing precision as PSR J0437−4715 . Nevertheless, observing and monitoring these pulsars on a longer term is a worthwhile scientific pursuit for many reasons. We review some of these motivations in this section and present results subsequently.

The period and period derivative of a pulsar are some of the parameters that can be measured easily soon after its discovery. We then use them to derive certain fundamental parameters for each pulsar. For example, the characteristic age of any pulsar is defined as $\tau = P/(n-1)\dot{P}$, where P is the period, $\dot{P}$ is the period derivative and n is the braking index, which equals 3 for magnetic dipole radiation. τ is an upper limit to the true age for the pulsar – the time since it started spinning down. However, as discussed in chapter one, we know that the measured period derivative also contains contributions from its transverse velocity and any radial acceleration:

$$\delta\dot{P} = \frac{P{v_\perp}^2}{cd} + \frac{P a_{||}}{c} \tag{4.1}$$

where $a_{||}$ is the radial acceleration,$v_\perp$ is the transverse speed and d is the distance to the pulsar (Shklovskii 1970, Camilo, Thorsett, & Kulkarni 1994). $a_{||}$ can be important for pulsars in globular clusters, or if they are far enough away that differential galactic rotation (between the pulsar and the solar system) is significant. Since the pulsars being observed with the FPTM are all relatively nearby and none of them is in a globular cluster, we can safely ignore this term. For many pulsars the transverse velocity contribution is quite large and correspondingly changes the characteristic age significantly: Our measurements of the proper motion of PSR J0437−4715 modifies the derived characteristic age from 1.6 Gyr to 11 Gyr!

The interpretation of τ as an age estimate for the pulsar assumes special importance in pulsar binary systems in which the companion is a white dwarf. We believe that millisecond pulsars are formed when slowly spinning pulsars are spun up by accretion of material from their companions which have expanded beyond their Roche lobe (as part of their evolutionary cycle). The end of the accretion process leaves behind a millisecond pulsar and a white dwarf. This epoch marks the start of two clocks – one advanced by the gradual slowing down of the pulsar and the second by the cooling of the white dwarf as it radiates away its residual thermal energy. Consistent 'readings' of these two clocks will be a vote of confidence in our current understanding of the origin and evolution of such systems. Having an accurate measurement of the system age is also useful in deriving the original spin period of the pulsar. The true equation of state of degenerate nuclear matter (assumed to be the constituent matter of neutron stars) still remains unresolved – there are many proposed candidates. Either the discovery of a pulsar with period $\leq$ 1 ms, or deriving an original spin period $\leq$ 1 ms for a millisecond pulsar will root out the incorrect equations (since the different equations predict different minimum periods at which the neutron star will break apart due to centrifugal forces).

Much work has been done in recent years to advance our knowledge of the cooling of white dwarfs (Hansen & Phinney 1998a), especially low-mass Helium-core white dwarfs which predominate in pulsar binaries. We are now in a position to obtain a reasonable assessment of a white dwarf's mass and cooling age from a measurement of its surface temperature and absolute luminosity. While the former is obtained from photometric measurements, the latter can only be derived from the apparent magnitude by using some distance estimate. With about a dozen pulsar companions having been observed with optical telescopes, distance measurements remain one of the biggest sources of uncertainity in deriving masses and ages (Hansen & Phinney 1998b). This is another field in which the timing of millisecond pulsars is valuable. While the pulsars are too weak to allow measurement of parallax, Bell & Bailes (1996) pointed out that measuring the orbital period derivative of circular binary pulsars gives us an indirect way to derive their distance. This is because the measured orbital period derivatives can be entirely attributed to the effect of proper motion (the expected contribution from gravitational radiation is negligible). The measurement of orbital period derivative is helped by the fact that its timing signature grows as $T^{2.5}$, where T is the span of observations. This provides a major motivation for our timing program in which we are monitoring several binary millisecond pulsars.

Precision timing demands close attention to a number of details, and human error in analysing the timing residuals is not unknown. Therefore, it makes sense for different observers to keep tabs on the same pulsars so that they can verify their results with each other. Towards this end we have also started a program of regularly observing some millisecond pulsars which have the potential for precision timing and that are also being monitored by other pulsar astronomers. These include PSRs J1713+0747 and J1939+2134. We also report on our observations of these pulsars in this chapter.

4.2 PSR J1713+0747

PSR J1713+0747 is a 4.5 ms pulsar in a 67-day binary. It is relatively bright, the DM is not large (15.99 pc cm^{-3}), and has shown quite stable timing behaviour (Camilo, Foster, & Wolszczan 1994), making it a promising target for long term monitoring. Camilo, Foster & Wolszczan (1994) also reported a significant detection of parallax, and the general relativistic Shapiro delay (the delay suffered by the pulsar's signal as it travels through the gravitational field of its companion). Two post-Keplerian parameters are used to identify the signature of Shapiro delay: $r = (Gm_2/c^3)$, and $s = sin\, i$, where G is Newton's contant, m_2 is the mass of the companion and i is the inclination of the orbit to our line of sight. These are also called the range and shape, respectively. Detection of the delay is clearly favoured by large inclination angles. The χ^2 minimization procedure in TEMPO produces assymetric confidence intervals in the $m_2 - sin\, i$ plane, as illustrated by figure 1 of Camilo, Foster & Wolszczan (1994). Their 99% confidence region is relatively large implying that the inclination angle is not very favourable. The global minimum in their χ^2 plot implies extremely high masses for both the pulsar and the companion, a result inconsistent with our current understanding of how low-mass binary pulsar systems are formed. Consequently they used astrophysical arguments to put upper limits on their mass measurements: $1.2 \leq M_{ns} \leq 2.0$ $\mathrm{M}_\odot$ and $0.27 \leq m_2 \leq 0.4$ $\mathrm{M}_\odot$, the lower limits being derived from their 2 σ confidence intervals.

Starting January 1994, we have been observing PSR J1713+0747 about once a month at Parkes. Most of the observations are at 20-cm wavelength, with some infrequent observations done at 50-cm to monitor any DM variations. With integrations of 90 s, typical TOA uncertainties are 1−−2 μs. Table 4.1 presents our results for the pulsar and binary parameters obtained from a least squares fit from TEMPO using 448 TOAs over 3 years of data, using the relativistic timing model of Damour and Deurelle (1986). Since the shape of the χ^2 surface as a function of m_2 and $sin\, i$ is not symmetric, we followed the fitting procedure of Ryba & Taylor (1991) and mapped out the χ^2 statistic over the relevant range in the two parameters. Figure 4.1 shows our results. The contours enclose regions of 68.3%, 95.4% and 99.73% confidence for m_2 and $sini$. We have also plotted solutions for the measured mass funtion for various assumed pulsar masses. Though our confidence intervals cover a relatively wide range, it is gratifying to note that the canonical neutron star mass of 1.4 $\mathrm{M}_\odot$ falls inside our 1 σ region and close to the global minimum.

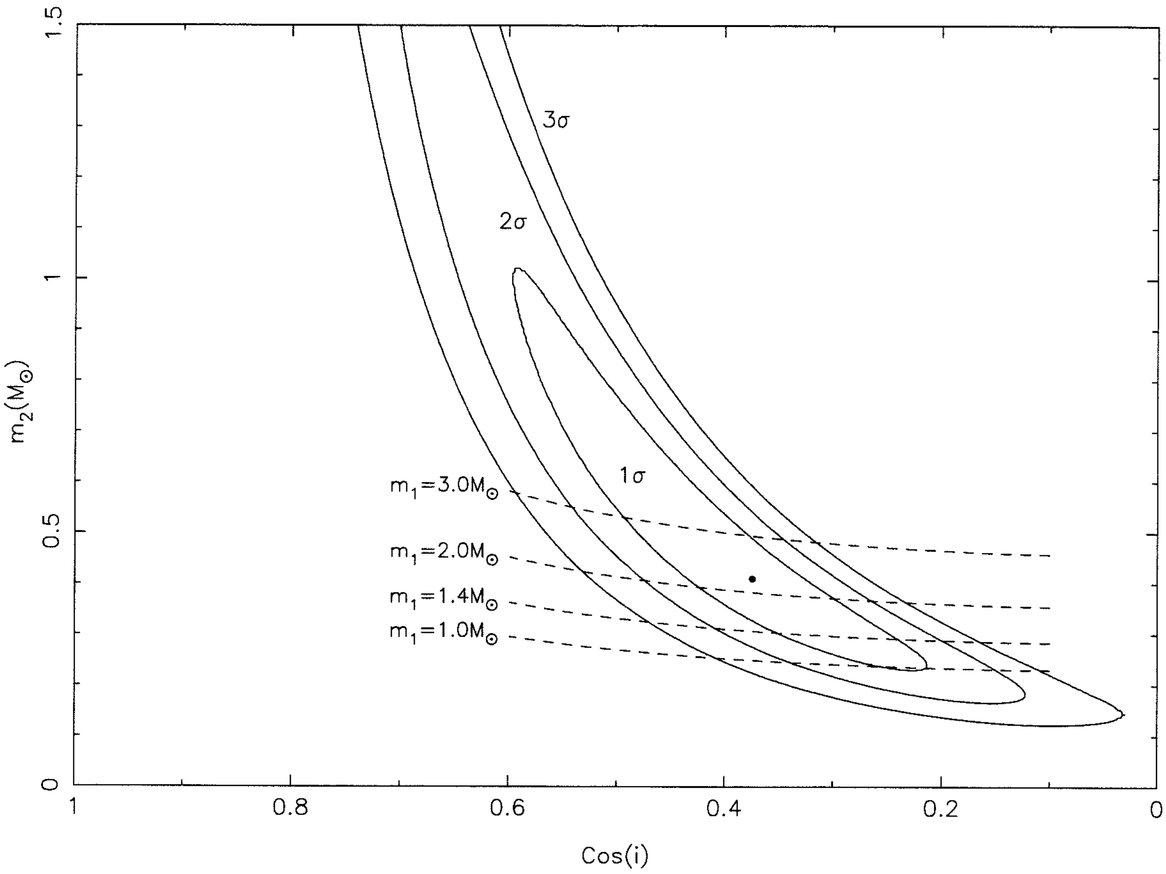

Figure 4.1: χ^2 surface in the m_2–$sin\, i$ plane for PSR J1713+0747. The solid dot marks the global minimum while the contours identify the confidence intervals. The dotted lines are constraints on the companion mass obtained from the mass function by assuming different pulsar masses.

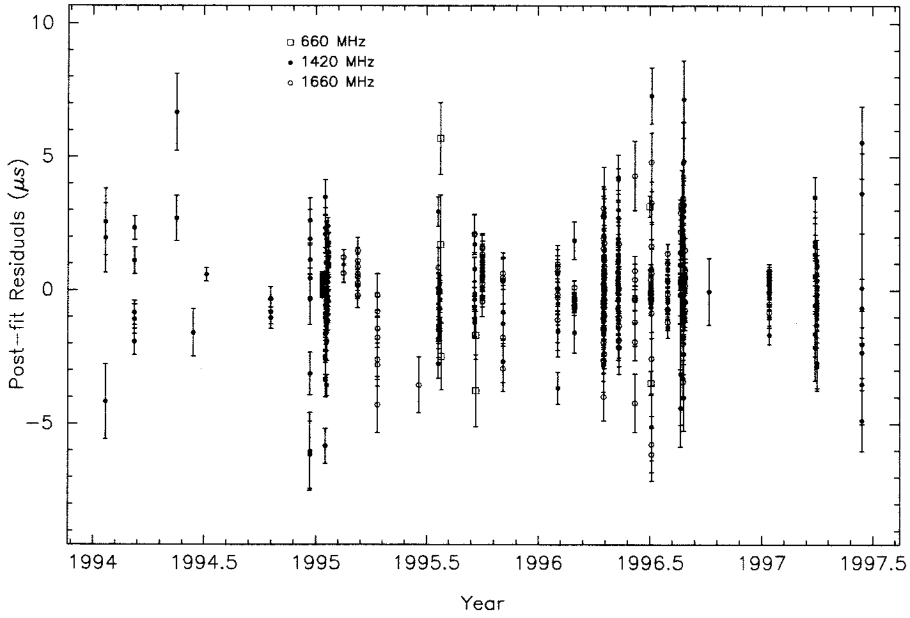

Figure 4.2: Post-fit timing residuals for PSR J1713+0747.

In general, our results are consistent with Camilo, Foster & Wolszczan (1994), but with better precision – by at least an order of magnitude in almost all parameters. Our measurement of parallax: $\pi = 1.2 \pm 0.2$ mas implies a distance d $= 0.83^{+0.2}_{-0.1}$ kpc, slightly lower than and in a narrower range from the d $= 1.1^{+0.5}_{-0.3}$ kpc measurement derived by them from parallax. The Taylor-Cordes model (Taylor & Cordes 1993) for the Galactic electron density distribution yields another distance estimate d $=$ $0.89^{+0.11}_{-0.07}$ kpc, in good agreement with our parallax measurement. Combined with the proper motion, our lower distance estimate implies a transverse speed of 24^{+5}_{-4} km/s. This speed and distance decreases the observed period derivative by only about 4%. The characteristic age is correspondingly increased slightly to 8.8×10^9 yr.

Lundgren, Camilo, & Foster 1996 claimed detection of the companion of PSR J1713+0747 using the Hubble Space Telescope and reported $m_V = 26.0(1)$ and $m_I =$ $24.1(1)$. Hansen & Phinney (1998) used these measurements and the parallax distance of Camilo, Foster & Wolszczan (1994), along with their cooling models of low-mass Helium core white dwarfs (both with thick and thin Hydrogen envelopes) to compute the likely cooling age and mass of the white dwarf. They found acceptable solutions

Observed and Derived Parameters for PSR J1713+0747

Parameter	Value
R. A. (J2000)	$17^{\mathrm{h}}\ 13^{\mathrm{m}}\ 49\overset{\mathrm{s}}{.}527807(2)$
Dec. (J2000)	$07^{\circ}\ 47'\ 37\overset{\prime\prime}{.}54571(9)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	4.74(5)
μ_δ (mas y^{-1})	-3.8(1)
Annual parallax (mas)	1.2(2)
Period, P (ms)	4.57013652362754(2)
Period derivative $\dot{P}$ (10^{-21}) ($\mathrm{s\,s}^{-1}$)	8.531(2)
Epoch of period and position (MJD)	50019.00
Dispersion Measure (cm^{-3} pc)	15.99267(6)
Binary period, P_b (d)	67.825129882(8)
$x = a_p \sin i$ (s)	32.3424208(4)
Eccentricity	0.000074946(6)
Long. of periastron, ω (°)	176.192687(7590)
Epoch of periastron (MJD)	50030.649871(1430)
$\dot{x}$ (10^{-15}) (lt-s s^{-1})	0.1(30)[a]
$\dot{\omega}$ (10^{-6} deg y^{-1})	0.0(1)[a]
$\dot{P}_b$ (10^{-11}) ($\mathrm{s\,s}^{-1}$)	−0.3(23)[a]
$\dot{e}$ (10^{-15} s^{-1})	0.1(2)[a]
Timing data span (MJD)	49374 – 50614
RMS timing residual (μs)	0.85
Parallax distance (pc)	833
DM Distance (pc)	890
Transverse Velocity (km s^{-1})	24^{+5}_{-4}
Composite proper motion (mas y^{-1})	6.09(9)
Position angle (celestial) of proper motion (°)	128(1)
Position angle (galactic) of proper motion (°)	192(1)

Table 4.1: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

[a] These parameters were held at zero while all other parameters were fit. Subsequently all parameters were held fixed at their listed values while each of these parameters was fit individually.

for $0.15 \leq m_2 \leq 0.31$ M$_\odot$ for thick Hydrogen envelopes, and $m_2 \leq 0.27$ M$_\odot$ for thin envelopes. They noted that they could find solutions up to 0.45 M$_\odot$ if a smaller distance, comparable to the dispersion measure derived distance, is used. Our parallax measurement thus widens the acceptable range of companion masses. A more accurate distance estimate is therefore required. The proper motion of PSR J1713+0747 will produce an apparent orbital period derivative of about $4.4 \times 10^{-13}\mathrm{s\,s^{-1}}$. With our present accuracy of $3 \times 10^{-11}\mathrm{s\,s^{-1}}$, it will take about 15 years of timing to detect this perturbation in $\dot{P}_b$ (since the accuracy of $\dot{P}_b$ improves as $T^{2.5}$). Ultimately $\dot{P}_b$ will provide the most accurate distance measure. We note that even though our Shapiro delay determination does not put significant constraints on the constituent masses at this time, future timing should shrink the confidence intervals such that the lower limit on M_{ns} becomes interesting (see figure 4.1). Another parameter that should be measurable is the secular change in the projected semi-major axis, similar to our detection of that effect for PSR J0437−4715. Using $i = 70$ ° (suggested by the Shapiro delay) and the measured proper motion, we can estimate the magnitude of the effect by using equation 3.2: $\dot{x} \sim 1.1 \times 10^{-14}$. Given the present precision of 3.0×10^{-14}, this could be achievable in about 10 years of timing, yielding a useful limit on the inclination.

4.3 PSR B1937+21

PSR B1937+21 is the first millisecond pulsar discovered and the fastest to date. It spins at 1.55 ms and is isolated. Despite its high DM (71.040 pc cm^{-3}), it is a good candidate for long term monitoring because of its large flux density. Kaspi, Taylor & Ryba (1994) reported long term timing results from 8 yrs of data using the Arecibo 300 m radio telescope. They obtained timing residuals of 0.2 μs after fitting for spin and astrometric parameters, including parallax ($\leq$ 0.28 mas) and proper motion ($\sim$ 0.5 mas/yr). They also confirmed earlier reports of significant variations of the dispersion measure of PSR B1937+21. They observed at two widely spaced frequencies at regular intervals and were able to measure and thereby account for any changes in the DM in their timing solution. Despite their careful analysis of the TOAs, the post-fit residuals have an unmistakable signature of a cubic term indicating the presence of some long term noise source. In the presence of low frequency noise the cubic term is expected to be the dominant feature in the post-fit residuals since the fit for period and period derivative absorbs any power in the linear and quadratic coefficients. By careful comparison of the residuals with different atomic clock data, solar system ephemeris and the post-fit residuals for PSR B1855+09 (which appear to not display any similar low frequency noise), they came to the conclusion that the culprit is most likely the pulsar. However, the large distance to the pulsar makes it likely that interstellar propagation effects could account for some, if not all, of the observed noise. The possibility of solving this mystery and repeating their high precision astrometry encouraged us to pursue timing this pulsar on a long term basis.

We have been observing PSR B1937+21 since January 1995. Due to its high DM we are forced to adopt a strategy of using the widest available bandwidth (128 MHz) to hunt for interstellar scintillation induced maxima and then zooming in on them with narrower bandwidths. All the TOAs used have been obtained from observations which used either 32 or 4 MHz of bandwidth. Typical 90 s TOA uncertainity for data at 20-cm observing wavelength was 1–3 μs. Though we did some observations at 50-cm wavelength, in order to monitor DM changes, they have not been of a sufficiently good quality to serve their intended purpose. Part of the problem is that the scintillation maxima at 50-cm have a smaller bandwidth and more ephemeral, making the 'hunt' process more difficult and time consuming. For the future we need to devote more time to obtain good data at 50-cm at regular intervals, preferably at every epoch.

Table 4.2 presents our latest fit parameters for PSR B1937+21 utilizing 215 TOAs from about 2 years of data. The post-fit timing error is 1 μs. Though our position and period determination precision almost equals that of Ryba & Taylor 1991, we have still some way to go to determine the proper motion and parallax. The data length span is also not adequate to measure the frequency second derivative reported by them. The current results certainly encourage us to continue timing this object.

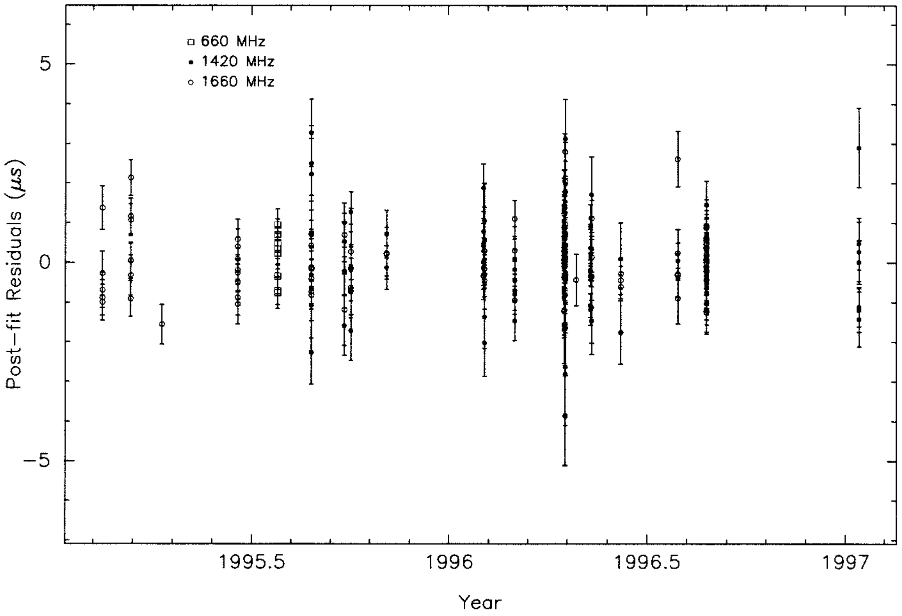

Figure 4.3: Post-fit timing residuals for PSR J1939+2134.

Observed and Derived Parameters for PSR J1939+2134

Parameter	Value
R. A. (J2000)	$19^{\mathrm{h}}\ 39^{\mathrm{m}}\ 38\overset{\mathrm{s}}{.}560128(3)$
Dec. (J2000)	$21^\circ\ 34'\ 59\overset{\prime\prime}{.}13883(9)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	0.1(2)
μ_δ (mas y^{-1}).......................	-0.7(2)
Annual parallax (mas)	0.1(4)
Period, P (ms)	1.557806492432726(8)
Period derivative $\dot{P}$ (10^{-19}) ($\mathrm{s\,s^{-1}}$) ..	1.05111(1)
Epoch of period and position (MJD)	50100.00
Dispersion Measure (cm^{-3} pc)	71.03668(6)
Timing data span (MJD)	49925 - 50322
RMS timing residual (μs)	1.00

Table 4.2: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

4.4 PSR J1744-1134

After PSR J0437−4715, this millisecond pulsar is the most promising of the new pulsars discovered in the Parkes survey. It has a low DM (3.1 pc cm^{-3}) and a short duty cycle. Typical TOA precision is 0.5−−1.0 μs in a 90 s integration. Our timing solution (Table 4.3) has vastly improved parameters than the previously published solution, including a determination of the parallax. Our parallax $\pi = 3.2(4)$ mas places the pulsar at a distance of 312^{+45}_{-34} kpc, compared to the dispersion measure distance of 170^{+20}_{-20} kpc (using the model of Taylor & Cordes (1993)). As with our parallax measurement of PSR J0437−4715, this indicates that the electron density has been overestimated along the line of sight to this pulsar. We obtain an average electron density of 0.01 cm^{-3} whereas the model assumes 0.019 cm^{-3}. Our precise determination of proper motion permits us to subtract the kinematic contribution to $\dot{P}$. This amounts to a 16% correction in $\dot{P}$ and we get $\dot{P} = 7.51 \times 10^{-21}\mathrm{s\,s}^{-1}$. This is one of the lowest period derivatives among all the millisecond pulsars, and consequently the activity parameter of Arzoumanian (1994) is also one of the lowest. This parameter has been observed to indicate good timing stability for pulsars and therefore motivates us to continue the timing observations for the long haul.

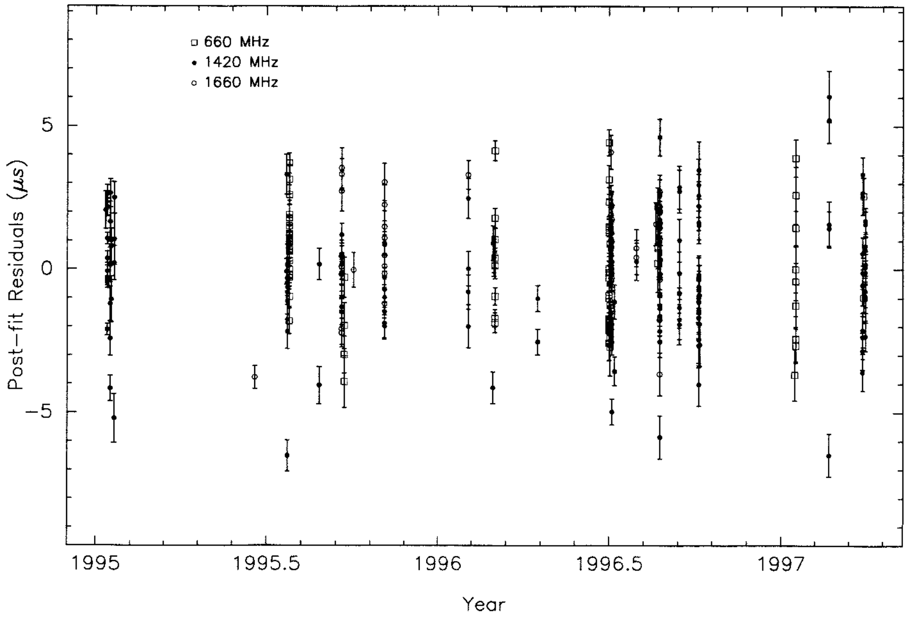

Figure 4.4: Post-fit residuals for PSR J1744-1134.

Observed and Derived Parameters for PSR J1744-1134

Parameter	Value
R. A. (J2000)	$17^{\mathrm{h}}\ 44^{\mathrm{m}}\ 29^{\mathrm{s}}.389921(8)$
Dec. (J2000)	$-11^\circ\ 34'\ 54''.5661(6)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	18.4(2)
μ_δ (mas y^{-1})	-11.0(9)
Annual parallax (mas)	3.2(4)
Period, P (ms)	4.07454587489526(3)
Period derivative $\dot{P}$ (10^{-21}) ($\mathrm{s\,s^{-1}}$)	8.931(4)
Epoch of period and position (MJD)	50134.00
Dispersion Measure (cm^{-3} pc)	3.13949(4)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	2.02
Composite proper motion (mas y^{-1})	21.4(5)
Position angle (celestial) of proper motion (°)	120(2)
Position angle (galactic) of proper motion (°)	108(2)

Table 4.3: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

4.5 The Rest

The following pulsars have low flux densities and do not offer good prospects for long term timing observations. Their solutions are presented here because in most cases we have made significant measurements of proper motions and orbital eccentricities, compared to their previously published solutions where only upper limits or very wide error bars were presented. The eccentricities can now be compared with the models of Phinney (1992), with which they are in good agreement, though the eccentricity of PSR J0613-0200 (see Table 4.4) is still not very precise. Furthur timing is therefore desirable to improve the proper motion and eccentricity measurements, and in the rare case, measure secular changes in the orbital parameters (due to proper motion). We have attempted to estimate the magnitude of $\dot{P}_b$ and $\dot{x}$ in each case by using our measured proper motions with the dispersion measure distances (for $\dot{P}_b$) and by assuming an inclination angle = 60°, the median expected inclination angle, for $\dot{x}$. PSR J1455-3330 and PSR J1643-1223 hold the best hope for realizing these measurements[1].

---

[1] Since these solutions were calculated, updated solutions using a larger data set from the same observing system have been presented in Toscano *et al.* (1999). For easy reference, we present those results in Appendix A.

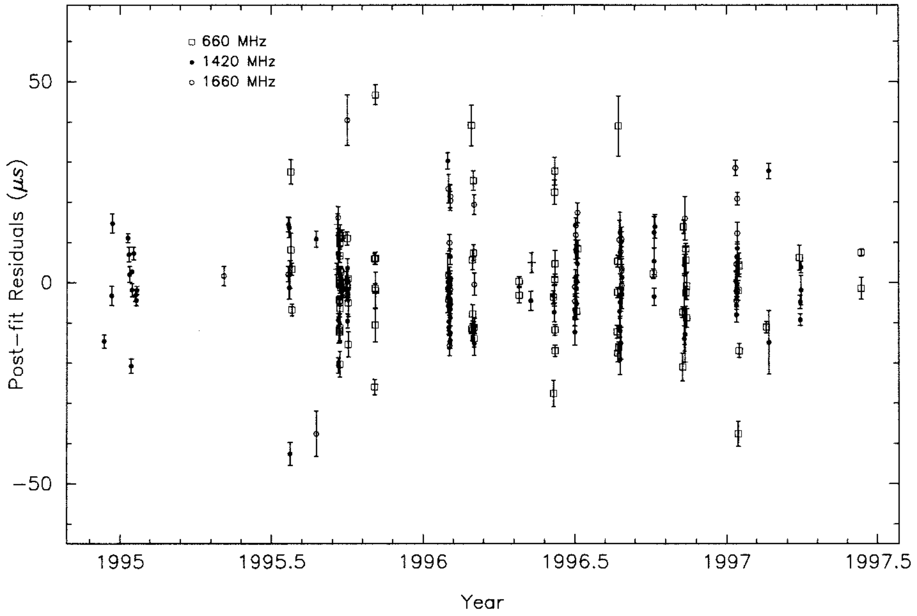

Figure 4.5: Post-fit residuals for PSR J0613-0200.

Observed and Derived Parameters for PSR J0613-0200

Parameter	Value
R. A. (J2000)	$06^{\mathrm{h}}\ 13^{\mathrm{m}}\ 43^{\mathrm{s}}.9729(1)$
Dec. (J2000)	$-02^\circ\ 00'\ 47''.064(5)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	2(1)
μ_δ (mas y^{-1})	-15(2)
Period, P (ms)	3.0618440360792(9)
Period derivative $\dot{P}$ (10^{-21}) ($\mathrm{s\,s^{-1}}$)	9.57(2)
Epoch of period and position (MJD)	49512.00
Dispersion Measure ($\mathrm{cm^{-3}}$ pc)	38.7924(2)
Binary period, P_b (d)	1.1985125578(12)
$x = a_p \sin i$ (s)	1.0914434(14)
Eccentricity	0.0000065(23)
Long. of periastron, ω (°)	50.5696(213838)
Epoch of periastron (MJD)	49512.43432(7119)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	13
Composite proper motion (mas y^{-1})	15(2)
Position angle (celestial) of proper motion (°)	171(5)
Position angle (galactic) of proper motion (°)	108(5)
Expected $\dot{x}$ (10^{-14}) (lt-s s^{-1})	0.1
Expected $\dot{P}_b$ (10^{-13}) ($\mathrm{s\,s^{-1}}$)	3

Table 4.4: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

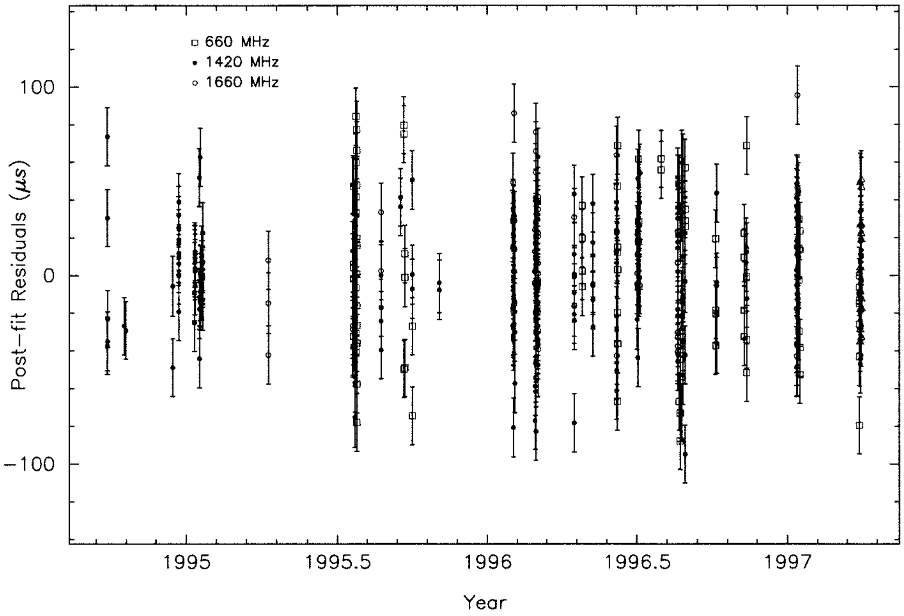

Figure 4.6: Post-fit residuals for PSR J1045-4509.

Observed and Derived Parameters for PSR J1045-4509

Parameter	Value
R. A. (J2000)	$10^{h}\ 45^{m}\ 50^{s}\!.1930(1)$
Dec. (J2000)	$-45^\circ\ 09'\ 54''\!.194(1)$
$\mu_\alpha \cos\delta$ (mas y^{-1}).........................	-5(2)
μ_δ (mas y^{-1})................................	10(2)
Period, P (ms)	7.4742241059192(6)
Period derivative $\dot{P}$ (10^{-20}) ($s\,s^{-1}$)	1.764(7)
Epoch of period and position (MJD)	50160.00
Dispersion Measure (cm^{-3} pc)	58.1712(5)
Binary period, P_b (d)	4.083529171(9)
$x = a_p \sin i$ (s)	3.015131(2)
Eccentricity	0.000021(1)
Long. of periastron, ω (°)	238(4)
Epoch of periastron (MJD)	50161.87442(4686)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	32
Composite proper motion (mas y^{-1})	11(2)
Position angle (celestial) of proper motion (°)	333(13)
Position angle (galactic) of proper motion (°)	306(13)
Expected $\dot{x}$ (10^{-14}) (lt-s s^{-1})	0.3
Expected $\dot{P}_b$ (10^{-13}) ($s\,s^{-1}$)	3

Table 4.5: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

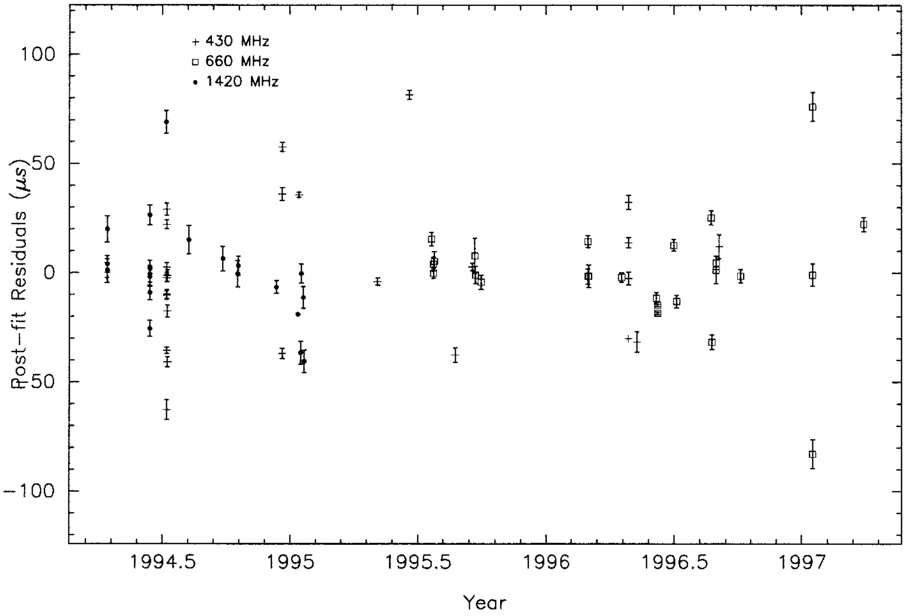

Figure 4.7: Post-fit residuals for PSR J1455-3330.

Observed and Derived Parameters for PSR J1455-3330

Parameter	Value
R. A. (J2000)	$14^{\mathrm{h}}\ 55^{\mathrm{m}}\ 47.^{\mathrm{s}}9609(6)$
Dec. (J2000)	$-33^{\circ}\ 30'\ 46.''36(2)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	0(7)
μ_δ (mas y^{-1})	-4(14)
Period, P (ms)	7.987204793955(7)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s}^{-1}$)	2.41(2)
Epoch of period and position (MJD)	49500.00
Dispersion Measure (cm^{-3} pc)	13.5727(6)
Binary period, P_b (d)	76.1745667(8)
$x = a_p \sin i$ (s)	32.362222(9)
Eccentricity	0.0001689(3)
Long. of periastron, ω (°)	223.0(2)
Epoch of periastron (MJD)	49513.369568(3079)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	28
Composite proper motion (mas y^{-1})	4(14)
Position angle (celestial) of proper motion (°)	172(91)
Position angle (galactic) of proper motion (°)	202(91)
Expected $\dot{x}$ (10^{-14}) (lt-s s^{-1})	1
Expected $\dot{P}_b$ (10^{-13}) ($\mathrm{s\,s}^{-1}$)	60

Table 4.6: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

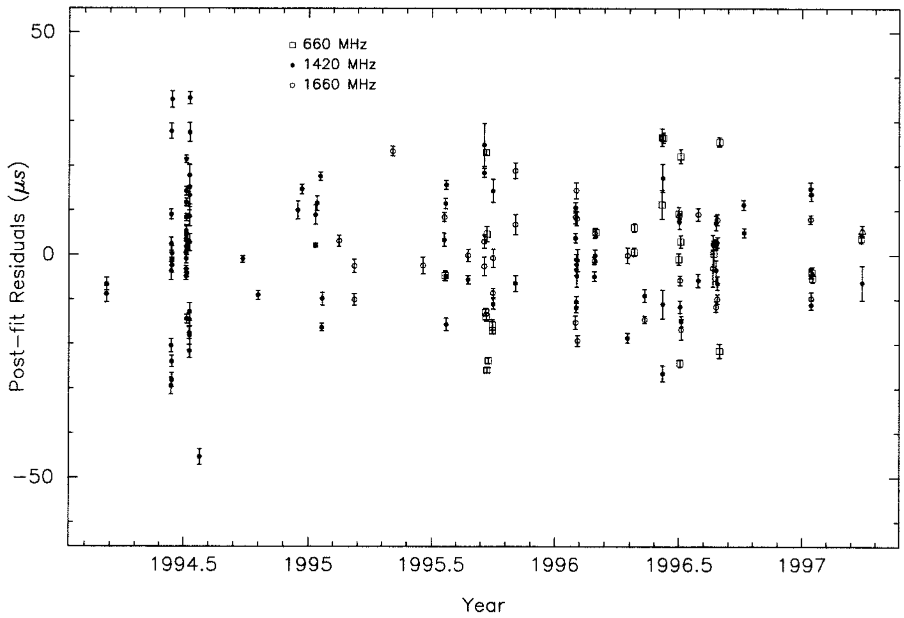

Figure 4.8: Post-fit residuals for PSR J1643-1224.

Observed and Derived Parameters for PSR J1643-1224

Parameter	Value
R. A. (J2000)	$16^{\mathrm{h}}\ 43^{\mathrm{m}}\ 38^{\mathrm{s}}\!.15556(8)$
Dec. (J2000)	$-12^{\circ}\ 24'\ 58''\!.727(6)$
μ_α cos δ (mas y^{-1})	5(2)
μ_δ (mas y^{-1})	11(11)
Period, P (ms)	4.6216414461348(3)
Period derivative $\dot{P}$ (10^{-20}) (s s^{-1})	1.846(3)
Epoch of period and position (MJD)	50019.00
Dispersion Measure (cm^{-3} pc)	62.4110(4)
Binary period, P_b (d)	147.017395(1)
$x = a_p \sin i$ (s)	25.072615(2)
Eccentricity	0.0005058(2)
Long. of periastron, ω (°)	321.84(2)
Epoch of periastron (MJD)	50019.019540(7435)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	14
Composite proper motion (mas y^{-1})	13(10)
Position angle (celestial) of proper motion (°)	22(21)
Position angle (galactic) of proper motion (°)	76(21)
Expected $\dot{x}$ (10^{-14}) (lt-s s^{-1})	2
Expected $\dot{P}_b$ (10^{-13}) (s s^{-1})	200

Table 4.7: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

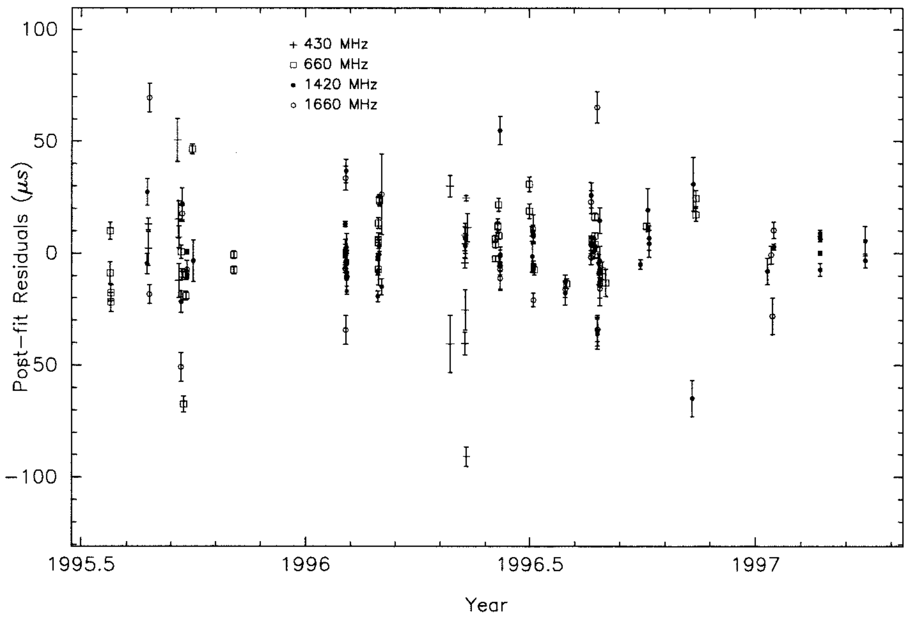

Figure 4.9: Post-fit residuals for PSR J2129-5721.

Observed and Derived Parameters for PSR J2129-5721

Parameter	Value
R. A. (J2000)	$21^{\mathrm{h}}\ 29^{\mathrm{m}}\ 22\overset{\mathrm{s}}{.}7535(3)$
Dec. (J2000)	$-57^\circ\ 21'\ 14\overset{''}{.}090(2)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	8(5)
μ_δ (mas y^{-1})	-6(11)
Period, P (ms)	3.7263484183906(7)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s^{-1}}$)	2.08(2)
Epoch of period and position (MJD)	50228.30
Dispersion Measure ($\mathrm{cm^{-3}}$ pc)	31.8498(4)
Binary period, P_b (d)	6.62549299(4)
$x = a_p \sin i$ (s)	3.500567(3)
Eccentricity	0.000011(2)
Long. of periastron, ω (°)	190(5)
Epoch of periastron (MJD)	50227.50986(9694)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	22
Composite proper motion (mas y^{-1})	10(9)
Position angle (celestial) of proper motion (°)	133(51)
Position angle (galactic) of proper motion (°)	242(51)
Expected $\dot{x}$ (10^{-14}) (lt-s s^{-1})	0.3
Expected $\dot{P}_b$ (10^{-13}) ($\mathrm{s\,s^{-1}}$)	2

Table 4.8: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

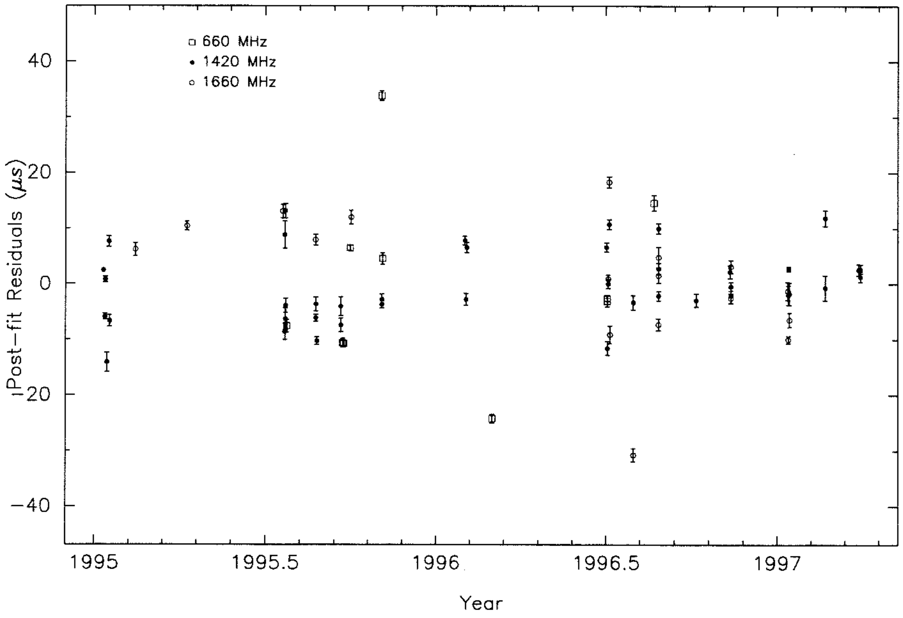

Figure 4.10: Post-fit residuals for PSR J0711-6830.

Observed and Derived Parameters for PSR J0711-6830

Parameter	Value
R. A. (J2000)	$07^{\mathrm{h}}\ 11^{\mathrm{m}}\ 54\overset{\mathrm{s}}{.}2208(2)$
Dec. (J2000)	$-68^\circ\ 30'\ 47\overset{''}{.}5923(9)$
μ_α cos δ (mas y^{-1})..........................	-17(1)
μ_δ (mas y^{-1})...............................	15(1)
Period, P (ms)	5.4909684153884(4)
Period derivative $\dot{P}$ (10^{-20}) (s s^{-1})	1.493(5)
Epoch of period and position (MJD)	50100.00
Dispersion Measure (cm^{-3} pc)	18.4099(5)
Timing data span (MJD)	49925 – 50540
RMS timing residual (μs)	10
Composite proper motion (mas y^{-1})	23(2)
Position angle (celestial) of proper motion (°)	312(4)
Position angle (galactic) of proper motion (°)	238(4)

Table 4.9: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

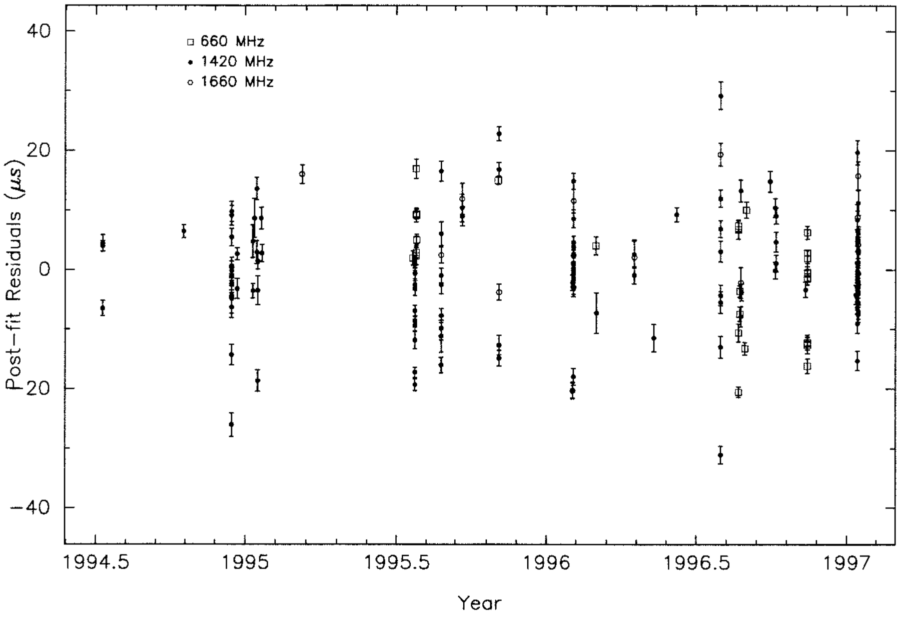

Figure 4.11: Post-fit residuals for PSR J2124-3358.

Observed and Derived Parameters for PSR J2124-3358

Parameter	Value
R. A. (J2000)	$21^{\mathrm{h}}\ 24^{\mathrm{m}}\ 43^{\mathrm{s}}.86279(7)$
Dec. (J2000)	$-33^{\circ}\ 58'\ 44''.221(1)$
μ_α cos δ (mas y^{-1}).........................	-16(1)
μ_δ (mas y^{-1})...............................	-51(2)
Period, P (ms)	4.9311148586700(2)
Period derivative $\dot{P}$ (10^{-20}) (s s^{-1})	2.055(2)
Epoch of period and position (MJD)	50019.00
Dispersion Measure (cm^{-3} pc)	4.5975(3)
Timing data span (MJD)	49500 – 50464
RMS timing residual (μs)	10
Composite proper motion (mas y^{-1})	54(2)
Position angle (celestial) of proper motion (°)	197(1)
Position angle (galactic) of proper motion (°)	282(1)

Table 4.10: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Appendix A Updated Solutions

Updated solutions for some of the millisecond pulsars discussed in chapter four were published in Toscano et al. 1999. We present them here for convenience and easy reference for other observers. These solutions have improved positions and proper motions.

Observed and Derived Parameters for PSR J1744-1134

Parameter	Value
R. A. (J2000)	$17^{\rm h}\ 44^{\rm m}\ 29^{\rm s}\!.390963(5)$
Dec. (J2000)	$-11^\circ\ 34'\ 54''\!.5746(5)$
$\mu_\alpha \cos\delta$ (mas $\rm y^{-1}$)	18.72(6)
μ_δ (mas $\rm y^{-1}$)	-9.5(4)
Annual parallax (mas)	2.8(3)
Period, P (ms)	4.07454587512695(3)
Period derivative $\dot{P}$ (10^{-21}) $(\rm s\,s^{-1})$	8.9405(9)
Epoch of period and position (MJD)	50434.00
Dispersion Measure ($\rm cm^{-3}$ pc)	3.1388(3)
Timing data span (MJD)	49730 – 51209
RMS timing residual (μs)	0.47
Composite proper motion (mas $\rm y^{-1}$)	21.0(2)
Position angle (celestial) of proper motion (°)	116.8(9)

Table A.1: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J0613-0200

Parameter	Value
R. A. (J2000)	$06^{\rm h}\ 13^{\rm m}\ 43^{\rm s}\!.97336(2)$
Dec. (J2000)	$-02^\circ\ 00'\ 47''\!.0970(7)$
$\mu_\alpha \cos\delta$ (mas $\rm y^{-1}$)	2.0(4)
μ_δ (mas $\rm y^{-1}$)	-7(1)
Period, P (ms)	3.06184403674401(5)
Period derivative $\dot{P}$ (10^{-21}) $(\rm s\,s^{-1})$	9.572(5)
Epoch of period and position (MJD)	50315.00
Dispersion Measure ($\rm cm^{-3}$ pc)	38.7792(1)
Binary period, P_b (d)	1.1985125566(4)
$x = a_p \sin i$ (s)	1.0914441(5)
Eccentricity	0.0000038(10)
Long. of periastron, ω (°)	34(15)
Epoch of periastron (MJD)	50315.38(5)
$g = e\, sin\, \omega (10^{-6})$	2.1(9)
$h = e\, cos\, \omega (10^{-6})$	3.1(10)
Time of ascending node (MJD)	50315.2693799(1)
Timing data span (MJD)	49700 – 50931
RMS timing residual (μs)	3.5

Table A.2: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J1045-4509

Parameter	Value
R. A. (J2000)	$10^{\mathrm{h}}\ 45^{\mathrm{m}}\ 50\overset{\mathrm{s}}{.}1927(1)$
Dec. (J2000)	$-45^\circ\ 09'\ 54\overset{''}{.}195(1)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	-5(2)
μ_δ (mas y^{-1})	6(1)
Period, P (ms)	7.4742241060982(5)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s^{-1}}$) ..	1.766(3)
Epoch of period and position (MJD)	50277.00
Dispersion Measure ($\mathrm{cm^{-3}}$ pc)	58.147(5)
Binary period, P_b (d)	4.083529200(6)
$x = a_p \sin i$ (s)	3.015134(2)
Eccentricity	0.0000197(13)
Long. of periastron, ω (°)	243(4)
Epoch of periastron (MJD)	50276.26(5)
$g = e\,sin\,\omega(10^{-6})$	-17.5(14)
$h = e\,cos\,\omega(10^{-6})$	-8.9(15)
Time of ascending node (MJD)	50273.5068944(5)
Timing data span (MJD)	49623 – 50931
RMS timing residual (μs)	16.2

Table A.3: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J1455-3330

Parameter	Value
R. A. (J2000)	$14^{\mathrm{h}}\ 55^{\mathrm{m}}\ 47\overset{\mathrm{s}}{.}9618(3)$
Dec. (J2000)	$-33^\circ\ 30'\ 46\overset{''}{.}350(7)$
μ_α cos δ (mas y^{-1})	5(6)
μ_δ (mas y^{-1})	24(12)
Period, P (ms)	7.987204795504(3)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s^{-1}}$)	2.42(1)
Epoch of period and position (MJD)	50239.00
Dispersion Measure (cm^{-3} pc)	13.578(7)
Binary period, P_b (d)	76.1745675(6)
$x = a_p \sin i$ (s)	32.362233(5)
Eccentricity	0.0001697(3)
Long. of periastron, ω (°)	223.8(1)
Epoch of periastron (MJD)	50275.17(2)
$g = e\,sin\,\omega(10^{-6})$	-117.4(2)
$h = e\,cos\,\omega(10^{-6})$	-122.5(4)
Time of ascending node (MJD)	50227.817383(2)
Timing data span (MJD)	49700 – 50779
RMS timing residual (μs)	1.4

Table A.4: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J1643-1224

Parameter	Value
R. A. (J2000)	$16^{\mathrm{h}}\ 43^{\mathrm{m}}\ 38\overset{\mathrm{s}}{.}15583(4)$
Dec. (J2000)	$-12^\circ\ 24'\ 58\overset{\prime\prime}{.}720(4)$
$\mu_\alpha \cos \delta$ (mas y^{-1})	3(1)
μ_δ (mas y^{-1})........................	-8(5)
Period, P (ms)	4.6216414465636(1)
Period derivative $\dot{P}$ (10^{-20}) (s s^{-1}) ..	1.849(1)
Epoch of period and position (MJD)	50288.00
Dispersion Measure (cm^{-3} pc)	62.4121(2)
Binary period, P_b (d)	147.0173943(7)
$x = a_p \sin i$ (s)	25.072613(1)
Eccentricity	0.0005058(1)
Long. of periastron, ω (°)	321.81(1)
Epoch of periastron (MJD)	50313.040(5)
$g = e\, sin\, \omega (10^{-6})$	-312.71(9)
$h = e\, cos\, \omega (10^{-6})$	397.5(1)
Time of ascending node (MJD)	50181.618211(1)
Timing data span (MJD)	49645 – 50932
RMS timing residual (μs)	6.8
Composite proper motion (mas y^{-1}).	8(5)

Table A.5: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J2129-5721

Parameter	Value
R. A. (J2000)	$21^{\mathrm{h}}\ 29^{\mathrm{m}}\ 22^{\mathrm{s}}.7539(3)$
Dec. (J2000)	$-57^\circ\ 21'\ 14''.093(2)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	7(2)
μ_δ (mas y^{-1})	-4(3)
Period, P (ms)	3.7263484187800(5)
Period derivative $\dot{P}$ (10^{-20}) $(\mathrm{s\,s^{-1}})$	2.074(4)
Epoch of period and position (MJD)	50444.00
Dispersion Measure (cm^{-3} pc)	31.855(4)
Binary period, P_b (d)	6.62549308(5)
$x = a_p \sin i$ (s)	3.500559(3)
Eccentricity	0.000068(22)
Long. of periastron, ω (°)	178(12)
Epoch of periastron (MJD)	50445.9(2)
Timing data span (MJD)	49983 – 50931
$g = e\, sin\, \omega (10^{-6})$	0.3(1.4)
$h = e\, cos\, \omega (10^{-6})$	-7(2)
Time of ascending node (MJD)	50442.643011(1)
RMS timing residual (μs)	1.2
Composite proper motion (mas y^{-1})	8(4)

Table A.6: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J0711-6830

Parameter	Value
R. A. (J2000)	$07^{\mathrm{h}}\ 11^{\mathrm{m}}\ 54\overset{\mathrm{s}}{.}21813(6)$
Dec. (J2000)	$-68^\circ\ 30'\ 47\overset{''}{.}5793(4)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	-15.7(5)
μ_δ (mas y^{-1})	15.3(6)
Period, P (ms)	5.4909684158061(1)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s^{-1}}$)	1.4900(1)
Epoch of period and position (MJD)	50425.00
Dispersion Measure ($\mathrm{cm^{-3}}$ pc)	18.4106(3)
Timing data span (MJD)	49922 – 50929
RMS timing residual (μs)	2.1
Composite proper motion (mas y^{-1})	22(6)

Table A.7: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Observed and Derived Parameters for PSR J2124-3358

Parameter	Value
R. A. (J2000)	$21^{\mathrm{h}}\ 24^{\mathrm{m}}\ 43\overset{\mathrm{s}}{.}86194(6)$
Dec. (J2000)	$-33^\circ\ 58'\ 44\overset{''}{.}257(1)$
$\mu_\alpha \cos\delta$ (mas y^{-1})	-16(1)
μ_δ (mas y^{-1})	-14(1)
Period, P (ms)	4.9311148591481(1)
Period derivative $\dot{P}$ (10^{-20}) ($\mathrm{s\,s^{-1}}$)	2.054(1)
Epoch of period and position (MJD)	50288.00
Dispersion Measure ($\mathrm{cm^{-3}}$ pc)	4.6152(2)
Timing data span (MJD)	49644 – 50933
RMS timing residual (μs)	5.2
composite proper motion (mas y^{-1})	49(2)

Table A.8: Uncertainties quoted are in the last digit and represent 1 σ estimates, inclusive of systematic effects.

Bibliography

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

Backer, D. C., Hama, S., Van Hook, S., & Foster, R. S. 1993, ApJ, 404, 636.

Backer, D. C. & Hellings, R. W. 1986, ARA&A, 24, 537.

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

Bailyn, C. D. 1993, ApJ, 411, L83.

Bartel, N., Chandler, J. F., Ratner, M. I., Shapiro, I. I., Pan, R., & Cappallo, R. J. 1996, AJ, 112, 1690.

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

Bell, J. F., Bailes, M., & Bessell, M. S. 1993, Nature, 364, 603.

Bell, J. F., Bailes, M., Manchester, R. N., Lyne, A. G., Camilo, F., & Sandhu, J. S. 1996, MNRAS, 286, 463.

Bell, J. F., Bailes, M., Manchester, R. N., Weisberg, J. M., & Lyne, A. G. 1995, ApJ, 440, L81.

Bertotti, B., Carr, B. J., & Rees, M. J. 1983, MNRAS, 203, 945.

Blandford, R. D., Narayan, R., & Romani, R. W. 1984, J. Astrophys. Astron., 5, 369.

Britton, M. C., Van Straten, W., Bailes, M., & Toscano, M. 2000, in Pulsar Astronomy - 2000 And Beyond, IAU Colloquium 177, ed. M. Kramer, N. Wex, & R. Wielebinski, Astronomical Society of the Pacific, p. 73.

Camilo, F., Foster, R. S., & Wolszczan, A. 1994, ApJ, 437, L39.

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

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

Cox, D. P. & Reynolds, R. J. 1987, ARA&A, 25, 303.

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

Damour, T., Gibbons, G. W., & Taylor, J. H. 1988, Phys. Rev. Lett., 61, 1151.

Danziger, I. J., Baade, D., & Della Valle, M. 1993, A&A, 276, 382.

Detweiler, S. 1979, ApJ, 234, 1100.

Folkner, W. M., Charlot, P., Finger, M. H., Williams, J. G., Sovers, O. J., Newhall, X. X., & Standish, E. M. 1994, A&A, 287, 279.

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

Foster, R. S. & Cordes, J. M. 1990, ApJ, 364, 123.

Guinot, B. & Petit, G. 1991, A&A, 248, 292.

Hankins, T. H. & Rickett, B. J. 1975, Meth. Comp. Phys., 14, 55.

Hansen, B. 1996. PhD thesis, California Institute of Technology.

Hansen, B. & Phinney, E. 1998a, MNRAS, 294, 557.

Hansen, B. & Phinney, E. 1998b, MNRAS, 294, 569.

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

Johnston, S. et al. 1993, Nature, 361, 613.

Kaspi, V. M., Taylor, J. H., & Ryba, M. 1994, ApJ, 428, 713.

Kopeikin, S. M. 1996, ApJ, 467, L93.

Lundgren, S. C., Camilo, F., & Foster, R. S. 1996, in Pulsars: Problems and Progress, IAU Colloqium 160, ed. M. Bailes, S. Johnston, & M. Walker, Astronomical Society of the Pacific, In Press.

Manchester, R. N. et al. 1996, MNRAS, 279, 1235.

Manchester, R. N. & Taylor, J. H. 1977, Pulsars (San Francisco: Freeman).

Matsakis, D. N., Taylor, J. H., & Eubanks, T. M. 1997, A&A, 326, 924.

Navarro, J. 1994. PhD thesis, California Institute of Technology.

Nicastro, L. & Johnston, S. 1995, MNRAS, 273, 122.

Phinney, E. S. & Kulkarni, S. R. 1994, ARA&A, 32, 591.

Rawley, L. A., Taylor, J. H., & Davis, M. M. 1988, ApJ, 326, 947.

Ryba, M. F. & Taylor, J. H. 1991, ApJ, 371, 739.

Sandhu, J. S., Bailes, M., Manchester, R. N., Navarro, J., Kulkarni, S. R., & Anderson, S. B. 1997, ApJ, 478, L95.

Shklovskii, I. S. 1970, Soviet Astron., 13, 562.

Standish, E. M. 1982, A&A, 114, 297.

Taylor, J. H. 1993, in Particle Astrophysics, IVth Rencontres de Blois, ed. G. Fontaine & J. Trân Thanh Vân (Gif-sur-Yvette, France: Editions Frontieres), p. 367.

Taylor, J. H. & Cordes, J. M. 1993, ApJ, 411, 674.

Taylor, J. H. & Weisberg, J. M. 1989, ApJ, 345, 434.

Toscano, M., Sandhu, J., Bailes, M., Manchester, R., Britton, M., Kulkarni, S., Anderson, S., & Stappers, B. 1999, MNRAS, 307, 925.