The End of Accretion: The X-ray Binary/Millisecond Pulsar Transition Object PSR J1023+0038

Anne M. Archibald

Doctor of Philosophy

Physics

McGill University

Montreal, Quebec

2013-08-02

A thesis submitted to McGill University in partial fulfullment of the requirements of the degree of Doctor of Philosophy

DEDICATION

To my parents, who have been looking forward to this.

ACKNOWLEDGEMENTS

First of all I thank my supervisor, Vicky Kaspi, for taking a chance on a dissatisfied number theorist, and for having patience while I got up to speed. She taught me not just how to answer questions in astronomy but how to recognize questions whose answers would be interesting. I thank my parents, who supported me through some very difficult times. I also thank my friends Matt and Shannon for a good meal and good company every now and then. I particularly thank Sarah, who saw me through the home stretch.

This thesis contains work published elsewhere, in collaboration with other authors, and supported by other organizations. Acknowledgements from the original manuscripts follow.

- Acknowledgements for "A Radio Pulsar/X-ray Binary Link" (Chapter 3), Archibald et al. 2009: We thank J. Condon of the NRAO for pointing out that J1023 lay within our survey beam, B. Stappers for help with WSRT observations and analysis, and S. Bogdanov, J. Thorstensen, and Z. Wang for useful discussions. The Robert C. Byrd Green Bank Telescope (GBT) is operated by the National Radio Astronomy Observatory which is a facility of the U.S. National Science Foundation operated under cooperative agreement by Associated Universities, Inc. The Arecibo Observatory is part of the National Astronomy and Ionosphere Center, which is operated by Cornell University under a cooperative agreement with the National Science Foundation. The Parkes Observatory radio telescope is part of the Australia Telescope which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. The WSRT is operated by the Netherlands Institute for Radio Astronomy (ASTRON) with support from the Netherlands Foundation for Scientific Research (NWO). We thank the directors of all four facilities for their generous allocation of exploratory time to this project. ASP and GASP are collaborations funded by NSERC and the NSF; IHS and JvL thank their co-developers D. Backer, D. Nice, P. Demorest, R. Ferdman, I. Cognard and R. Ramachandran for their hardware and/or software contributions, and P. Demorest and W. van Straten for making the PSRCHIVE software compatible with ASP/GASP data. We further thank M. Bailes, S. Kulkarni and the others in the Caltech-Green-Bank-Swinburne-Recorder collaboration for the use of some of their computer nodes during GASP observations. JRB, VIK, DRL, and MAM are supported by a WV EPSCOR Research Challenge Grant. JWTH is an NWO Veni Fellow. VMK is supported via the Lorne Trottier Chair in Astrophysics and Cosmology, the Canada Research Chair program, an NSERC Discovery Grant, FQRNT, CIFAR, and CFI. IHS received support from NSERC, the CFI, the ATNF Distinguished Visitor program, and the Swinburne University of Technology Visiting Distinguished Researcher Scheme.

- Acknowledgements for "X-ray Variability and Evidence for Pulsations from the Unique Radio Pulsar/X-ray Binary Transition Object FIRST J102347.6+003841" (Chapter 4), Archibald et al. 2010: The authors thank Mallory S. E. Roberts for useful discussions. This work was based on observations obtained with *XMM-Newton*, an ESA science mission with instruments and contributions directly funded by ESA Member States and NASA. AMA is supported by a Schulich graduate fellowship. VMK acknowledges support from NSERC, FQRNT, CIFAR, and holds a Canada Research Chair and the Lorne Trottier Chair in Astrophysics and Cosmology. SB is a CIFAR Junior Fellow. JWTH is an NWO Veni Fellow. IHS received support from an NSERC Discovery Grant. MAM is supported by a WV EPSCOR Research Challenge Grant.

- Acknowledgements for "Long-Term Radio Timing Observations of the Transition Millisecond Pulsar PSR J1023+0038" (Chapter 5, Chapter 6, and Chapter 7): The authors would like to thank Chris Thompson and Joanna Rankin for productive discussions. The authors would also like to thank the Green Bank North Celestial Cap pulsar survey team for allowing us to observe J1023 at 350 MHz, doubling as validation scans for the survey. JWTH acknowledges funding through the Netherlands Foundation for Scientific research (NWO). The Westerbork Synthesis Radio Telescope is operated by the Netherlands Foundation for Research in Astronomy (ASTRON) with support from the NWO. Access to the Lovell Telescope is supported through an STFC consolidated grant.
- Acknowledgements for "No detectable radio emission from the magnetar-like pulsar in Kes 75" (Appendix A, Archibald et al. 2008): The authors thank Scott Ransom for helpful conversations during the writing of this paper. The authors are also grateful to the Parkes Multibeam Pulsar Survey team for providing the Parkes observations. Support for this work was provided to VMK by an NSERC Discovery Grant (Rgpin 228738-03), an R. Howard Webster Foundation Fellowship of the Canadian Institute for Advanced Research, Les Fonds de la Recherche sur la Nature et les Technologies, a Canada Research Chair and the Lorne Trottier Chair in Astrophysics and Cosmology. MAL is a Natural Sciences and Engineering Research Council (NSERC) PGS-D Fellow. MAM is supported by a WV EPSCoR grant.

STATEMENT OF ORIGINALITY AND CONTRIBUTION OF AUTHORS

The university guidelines for a manuscript-based thesis state that "PhD candidates have the option of including, as part of the thesis, the text of one or more papers submitted, or to be submitted, for publication, or the clearly duplicated text of one or more published papers."

In accordance with the above, the results presented in this thesis are original work that was, or will be, published in the following refereed articles:

- The manuscript "A Radio Pulsar/X-ray Binary Link" published in *Science* in 2009 October (Archibald et al., 2009) and its supporting material form Chapter 3. The paper describes the discovery of J1023 as a radio pulsar, discusses the optical history that makes it unique, and comments on the implications of the discovery for millisecond pulsar evolution. The discovery formed part of a larger survey, described in Boyles et al. (2013), in which I and all the other authors carried out observations, wrote and managed scripts to process the data on clusters, and examined processing output. I found the periodicity candidate corresponding to J1023 and brought it to the attention of the group. Dr. Ingrid Stairs noticed that it might correspond to the object FIRST J102347.67+003841.2, suggested by Thorstensen and Armstrong (2005) to be a neutron star LMXB. With help from Dr. Scott Ransom and Dr. Ingrid Stairs I carried out follow-up observations with the Arecibo and Green Bank telescopes. Dr. Ingrid Stairs carried out the dual-frequency observation with the Parkes telescope, whie Dr. Jason Hessels carried out observations with the Westerbork Synthesis Radio Telescope. Dr. Ingrid Stairs constructed the inital long-term timing solution, which I updated as new data arrived. Dr. Ron Remillard carried out an analysis of archival data from the RXTE all-sky monitor, while Drs. Rachel Rosen, Brad Barlow, and Bart Dunlap carried out and analyzed our optical observation of J1023. All authors participated in extensive discussions of the interpretation and implications of our discovery. I wrote the text of the article, with comments on each revision from all other authors.

- The manuscript "X-ray Variability and Evidence for Pulsations from the Unique Radio Pulsar/X-ray Binary Transition Object FIRST J102347.6+003841", published in *The Astrophysical Journal* in 2010 August (Archibald et al., 2010) and reproduced by permission of the AAS forms Chapter 4. The paper describes our investigation of the X-ray properties of J1023, and in particular the nature of its spectrum and variability. I, with help from Drs. Vicky Kaspi and Slavko Bogdanov, wrote the proposal to observe J1023 with XMM-Newton. I analyzed the X-ray data myself, using the long-term timing solution obtained in the above paper. All authors participated in extensive discussions of the implications of our observations and the relationship of J1023 to other similar systems. I wrote the text myself, with substantial comments on each revision from my coauthors.
- The manuscript "Long-Term Radio Timing Observations of the Transition Millisecond Pulsar PSR J1023+0038" by Anne M. Archibald, Victoria M.

Kaspi, Jason W. T. Hessels, Ben Stappers, Gemma Janssen, and Andrew G. Lyne, I plan to submit to the *Astrophysical Journal*. It forms Chapter 5, Chapter 6, and most of Chapter 7. The paper presents a detailed study of the timing and other radio phenomenology of J1023, a study of the γ-ray emission, and an discussion of the physical processes that give rise to these phenomena and their bearing on J1023's evolution. I carried out the radio observations using the Arecibo telescope and the long observations with the Green Bank Telescope. The shorter Green Bank Telescope observations were taken as validation scans for the Green Bank North Celestial Cap survey, of which I am part; I carried out all of these observations that were used in this thesis. The observations with the Westerbork Synthesis Radio Telescope were carried out by Drs. Jason Hessels and Gemma Janssen, while the observations with the Lovell telescope were carried out by Drs. Gemma Janssen and Ben Stappers. Dr. Andrew Lyne assembled a phase-coherent timing solution encompassing all the Lovell telescope data, which I used along with the timing solution obtained for Chapter 3 to build the full timing solution described here. I wrote the modifications for `tempo2`, I analyzed the radio and γ-ray observations, and I generated all the plots and diagrams. All the paper's authors participated in discussions of the implications of our results. I wrote the text, and all my coauthors provided commentary and suggestions on each draft.

- The manuscript "No detectable radio emission from the magnetar-like pulsar in Kes 75", published in *The Astrophysical Journal* in 2008 November

(Archibald et al., 2008) and reproduced by permission of the AAS, forms Appendix A. This paper illustrates pulsar searching techniques similar to those used to discover J1023, as applied to a rather different problem: searching for radio pulsations from a known magnetar. I proposed for and carried out the Green Bank Telescope radio observation. Dr. Margaret Livingstone provided a long-term timing ephemeris for the pulsar in Kes 75, which I used to search for pulsations in the radio data. When I found strong dispersed single pulses in the data and realized they might be from RRAT 1846−02, I contacted Dr. Maura McLaughlin, who provided several observations of the same piece of sky that she had carried out. I analyzed these further observations as well, and computed upper limits on the possible brightness of the pulsar in Kes 75 as a radio pulsar. All authors participated in discussions about the implications. I wrote the text, and all my coauthors contributed comments and suggestions on each draft.

ABSTRACT

In this thesis, I describe the discovery and follow-up study of a transition object, PSR J1023+0038 (hereafter J1023). This object currently appears as a radio millisecond pulsar (MSP) but appears to have had an accretion disc in 2001. This places it at a poorly understood point along the pulsar evolutionary track, at the end of a low-mass X-ray binary (LMXB) phase and the beginning of a MSP phase. I discuss the evidence for this history and how it fits into our models of LMXB/MSP evolution. I also describe X-ray observations that show X-ray emission apparently coming partly from the pulsar, and partly from an intrabinary shock. I describe a long-term program of radio observations of the system, which show in detail the eclipses (ordered and random), dispersion measure variations, and orbital period variations. I also examine the system's γ-ray emission and find that it does not appear to be orbitally modulated but may be modulated at the pulsar period. I find that in its current state, radio emission from J1023 appears to be strongly affected by its Roche-lobe-filling companion, which seems to be magnetically active. This magnetic activity seems to support ionized material leaving the companion against the pulsar wind, and it also seems likely that it is responsible for the orbital period variations we observe. I attribute the companion's irradiation to heating by X-rays from a shock near the companion's surface. The observed variability in this shock may be related to activity within the companion. The question of whether or when J1023 will enter another active

episode remains open, although comparison with the newly discovered system M28I shows that it is possible.

ABRÉGÉ

Dans cette thèse, je décris la découverte et l'investigation d'un objet transitionel, PSR J1023+0038 (J1023 subséquemment). Cet objet à ce moment appert être un pulsar radio à période milliseconde (MSP) mais apparemment avaît un disque d'accrétion en 2001. Ce fait identifie J1023 comme étant en transision à une phase mal comprise dans la séquence evolutionnaire des pulsars, à la fin d'une phase binaire aux rayons X à faible masse (LMXB) et au début d'une phase MSP. Je discute les preuves pour cette interprétation et comment cette interprétation s'accorde avec nos modèles d'evolution LMXB/MSP. Je décris aussi les observations en rayons X de J1023 qui démontrent que l'émission des rayons X viens en partie du pulsar et en partie d'un choc entre le pulsar et son compagnon. Je décris un programme à long terme d'observations de J1023 dans les ondes radio, qui démontre en détail les éclipses (ordonnées et aléatoires), les variations en mesure de dispersion, et les variations en période orbitale. J'examine aussi l'émission du système en rayons γ et trouve qu'elle ne semble pas modulée à la période orbitale mais qu'elle pourrait être modulée à la periode du pulsar. Je trouve que dans son état actuel, l'émission radio de J1023 semble être fortement modifiée par son compagnon, qui remplit son lobe de Roche et qui semble être magnetiquement actif. Cette activité magnetique semble soutenir le matériel ionisé émis par le compagnon contre le vent du pulsar, et il semble aussi probable que cette activité est responsable pour les variations en période orbitale qui sont observées. J'attribue l'irradiation du compagnon au chauffage par rayons X venant d'un choc

proche de la surface du compagnon. La question de si ou quand J1023 se trouvera dans un autre épisode d'activité reste ouverte, mais la connection avec le système recemment découvert M28I démontre qu'il s'agit d'une possibilité.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Introduction

This thesis discusses a very unusual astronomical object, PSR J1023+0038 (also known as FIRST J102347.67+003841.2 or SDSS J102347.68+003841.0, and hereafter as J1023) and its role in clarifying the evolutionary sequence from low-mass X-ray binaries to millisecond pulsars.

I will begin with a description of these classes of object, their historical context, and some important puzzles (Chapter 1), then present a discussion of the key observing techniques and instruments (Chapter 2). The structure of the remainder of the thesis will be largely chronological, with Chapter 3 describing the discovery of J1023, Chapter 4 describing X-ray follow-up, Chapter 5 describing long-term radio monitoring, and Chapter 6 describing a γ-ray study of J1023. Chapter 7 will pull together all that we know of J1023 and try to describe the phenomena occurring in the system and relate them to the relationship between millisecond pulsars and their progenitors, and to the evolution of binary neutron stars in general.

Appendix A describes a study of two systems very different from J1023, the pulsar in the supernova remnant Kes 75 and the rotating radio transient RRAT J1846−02; this study shows the power of the search techniques that discovered J1023 and that are described in Section 2.3.

1.1 Pulsars

Hewish et al. (1968) reported a mysterious new kind of astronomical object. Their radio telescope had detected several astronomical sources, each of which emitted an extremely regular train of short pulses spaced by a characteristic period on the order of a second. The nature of these sources, which they named "pulsars," was puzzling, since obtaining such short pulses from such a distant object seemed to require tremendous density. Shortly after the discovery of the neutron, Baade and Zwicky (1934) had suggested that neutron stars might exist and be the end-results of supernovae. Although Hewish et al. (1968) had suggested neutron stars as possible candidates for pulsars, there were difficulties reconciling the regularity of the detected pulses with a radial pulsation mechanism. Gold (1968) first suggested that these pulsars might be *rotating* neutron stars. The discovery of regular pulsations, in radio (Comella et al., 1969), optical (Cocke et al., 1969), and then X-rays (Fritz et al., 1969) from a point source in the Crab Nebula provided convincing evidence that these objects were indeed neutron stars, and that the observed pulsations were the result of a relatively narrow beam sweeping past the Earth as the object rotated.

Careful timing of the pulsar in the Crab Nebula revealed that its 33.1-ms rotational period was increasing extremely gradually, by an amount of 36 ns day^{-1} (Richards and Comella, 1969). This gradual slowing was deduced to be the source of the power required to produce the observed radio, optical, and X-ray emission. Pacini (1967) had suggested that the Crab nebula might be powered by a neutron star, and Finzi and Wolf (1969) showed that the spin-down power of the pulsar was quite close to that required to produce the Crab Nebula, suggesting that the majority of this spin-down power emerges in the form of a particle wind. Other pulsars were also observed to be slowing gradually, and these objects as a class are taken to be powered by the loss of their rotational kinetic energy.

For most pulsars, very little about the object is measurable; in many cases the only physically meaningful observables are the period and period derivative, plus the hard-to-interpret pulse profile (average flux as a function of rotational phase). The pulsar astronomy community has therefore settled on a number of standard formulae (which I have drawn from Lorimer and Kramer, 2004) for calculating additional physical quantities describing a pulsar, based on certain assumed physical properties.

For the purposes of these calculations, pulsars are assumed to all have the same mass and radius. While this is known not to be true,[1] the goal of these formulae is to provide a basis for comparing pulsars rather than exact physical quantities. Models of electron degeneracy predict a maximum mass for white dwarfs, the Chandrasekhar limit, which is about $1.4M_{\odot}$ (Chandrasekhar, 1931, 1984). These formulae therefore assume that pulsar masses are all $1.4M_{\odot}$. Pulsar radii are assumed to be 10 km, a nice round number corresponding to a plausible density. Their moment of inertia, computed by assuming they are spheres of

[1] The technique of pulsar timing, discussed in Section 2.4, has allowed several neutron star masses to be inferred from light-bending effects (the Shapiro delay) or orbital decay due to gravitational-wave losses. These measured masses span a range from about 1.3 solar masses up to about 2.2 solar masses.

constant density, is presumed to be 10^{45} g cm^2. These assumptions imply that the loss of spin-down energy

$$\dot{E} = 4\pi^2 I \dot{P} P^{-3} = 3.95 \times 10^{31} \text{ erg s}^{-1} \left(\frac{\dot{P}}{10^{-15}}\right) \left(\frac{P}{1 \text{ s}}\right)^{-3}, \tag{1.1}$$

where E is the spin-down energy, P is the pulsar spin period, and dots denote time derivatives. We know the spin-down process in pulsars is complicated — for example in the Crab we know most of the spin-down power emerges as a particle wind, and in fact it is not possible for sub-kilohertz electromagnetic waves to propagate through the interstellar medium — but it is customary for calculations to treat pulsars as if their spin-down were purely due to the radiation produced by a rotating magnetic dipole in a vacuum. This assumption leads to the formula

$$B = 3.2 \times 10^{19} \text{ G} \sqrt{(P/1 \text{ s})\dot{P}}, \tag{1.2}$$

where B is the magnetic field at the magnetic pole on the surface of the star. The assumption of pure magnetic dipole radiation also implies a particular functional form for the spin-down. In general, writing $\nu = 1/P$, we assume pulsars spin down according to

$$\dot{\nu} = -K\nu^n, \tag{1.3}$$

where K and n are presumed to be constant and n is called the "braking index." For magnetic dipole radiation, n is 3. Other radiation mechanisms have been modelled and often give roughly similar spin-down parameters but usually $n \neq 3$ (as seen in the pulsar described in Appendix A, among others). If we assume magnetic dipole radiation so that $n = 3$, and we assume further that the initial spin period P_0 was much smaller than the currently observed value P, we find the age of the pulsar is

$$\tau = P/2\dot{P}. \tag{1.4}$$

Since this is expected to differ from the actual age of real pulsars, it is usually called the "characteristic age." While all these quantities depend on a long list of assumptions known not to be exact, they nevertheless provide a standard basis for comparing pulsars.

One key derived quantity that does not depend on assumptions is the light cylinder radius

$$r_{LC} = cP/2\pi. \tag{1.5}$$

This is the distance from the pulsar's spin axis at which material would need to move at the speed of light to corotate; this condition defines a cylinder called the light cylinder. Magnetic field lines lying entirely within this cylinder might have fairly simple static configurations of plasma, but field lines crossing the light cylinder must undergo some process to couple them to the necessarily non-corotating fields outside. These "open" field lines are thought to be where most of the pulsar's magetospheric emission is generated, and their footprints on the pulsar surface are the "polar caps." The details of the pulsar mechanism remain poorly understood.

Given that the two principal pulsar observables are the period P and the period derivative $\dot{P}$, it is natural to plot these values for all known pulsars in a "P–$\dot{P}$" diagram like Figure 1–1. Such a diagram serves to classify pulsars based on characteristic age, spin-down power, or magnetic field.

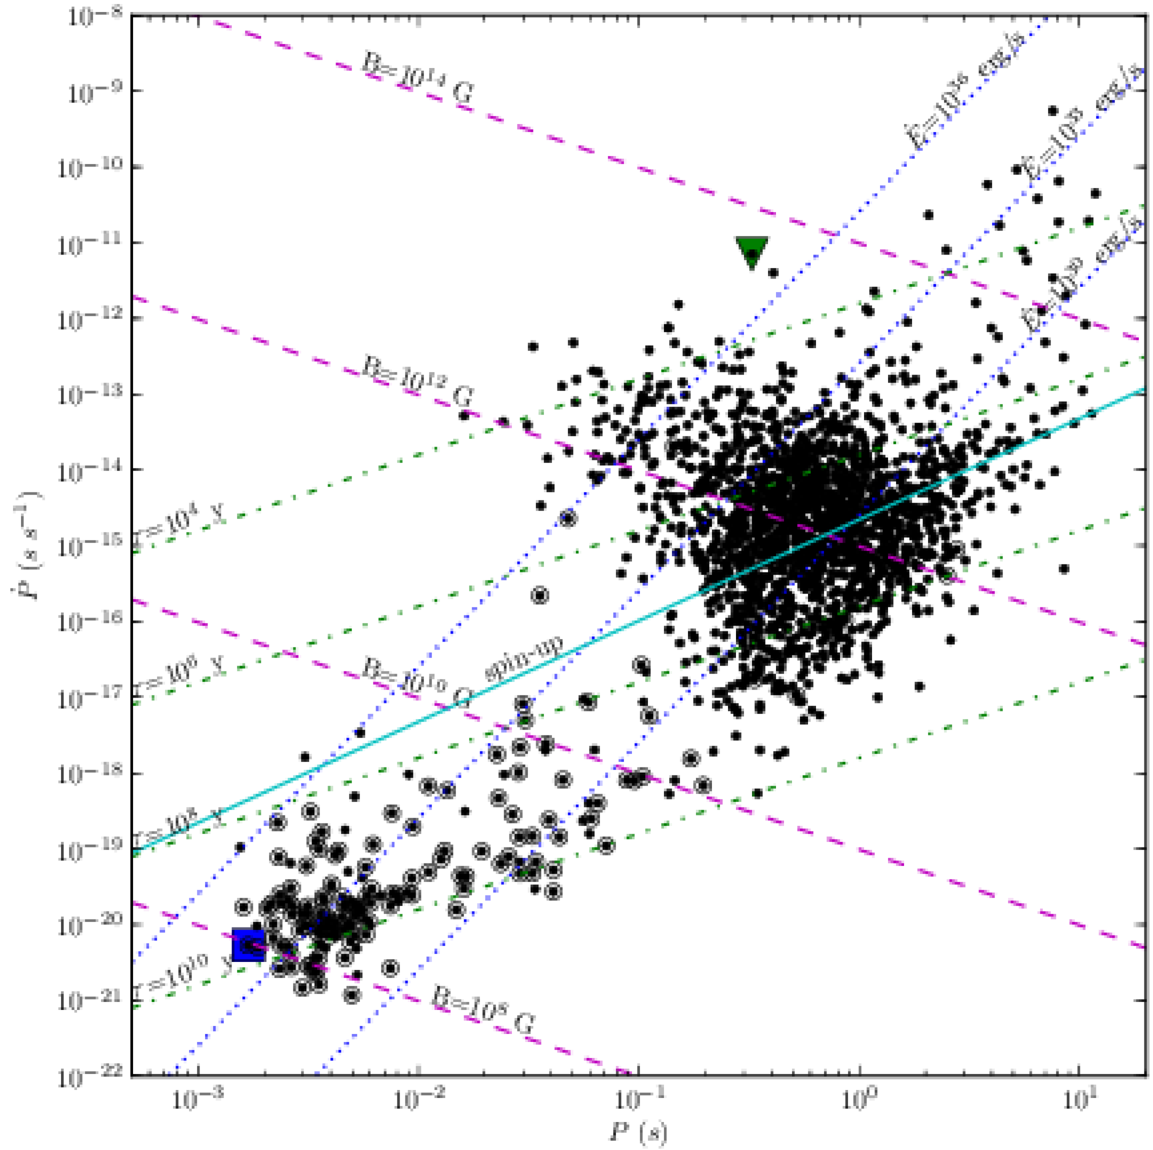

Figure 1–1 P–$\dot{P}$ diagram showing all pulsars in the ATNF pulsar catalog (Manchester et al., 2005). Circled dots indicate pulsars in binary systems. A blue square indicates J1023, the subject of the main body of this thesis, while a green triangle indicates the pulsar in the supernova remnant Kes 75, the subject of Appendix A. Dashed, dash-dotted, and dotted lines indicate lines of constant magnetic field (Equation 1.2), characteristic age (Equation 1.4), and spin-down power (Equation 1.1), respectively. The solid line indicates the maximum spin-up achievable by the recycling process described in Section 1.3 (and given by Equation 1.13); the fact that two millisecond pulsars are slightly above it is a reminder that many approximations enter into these models.

1.2 Millisecond pulsars

Backer et al. (1982) discovered a pulsar with the rotation period of 1.558 ms. This period, well below the 33.1-ms period of the very young Crab pulsar, was later found to be accompanied by a spin-down $\dot{P} = 1.2 \times 10^{-19}$, implying a characteristic age of 2×10^8 yr (Backer et al., 1983). The discovery of this "millisecond pulsar" (MSP) was soon followed by a number of other discoveries of pulsars with spin periods less than about 10 ms. Looking at Figure 1–1, there is clearly a cluster of MSPs now known having short periods, low magnetic fields, and large characteristic ages. These MSPs also, in contrast to pulsars generally, usually occur in binary systems. From the time of their discovery it seemed likely that MSPs shared an evolutionary history different from that of ordinary pulsars. In fact, shortly after the discovery of the first MSP, two papers suggested that it might have been formed by "recycling" an old pulsar through accretion of material from a companion (Alpar et al., 1982; Radhakrishnan and Srinivasan, 1982). This process would have spun the pulsar up, reducing the period to its current few-millisecond state. The magnetic field might have decayed during the pulsar's prior life, or the magnetic field might somehow have been reduced by the accretion process itself.

In a circular binary system, like those many MSPs are in and those in which almost[2] all are thought to have formed, one can form a coordinate system that

[2] The recently discovered PSR J0337+17 (Ransom et al., in preparation) is in a triple system, having two white dwarf companions. This strongly suggests that some MSPs can form in triple systems, possibly explaining the origin of

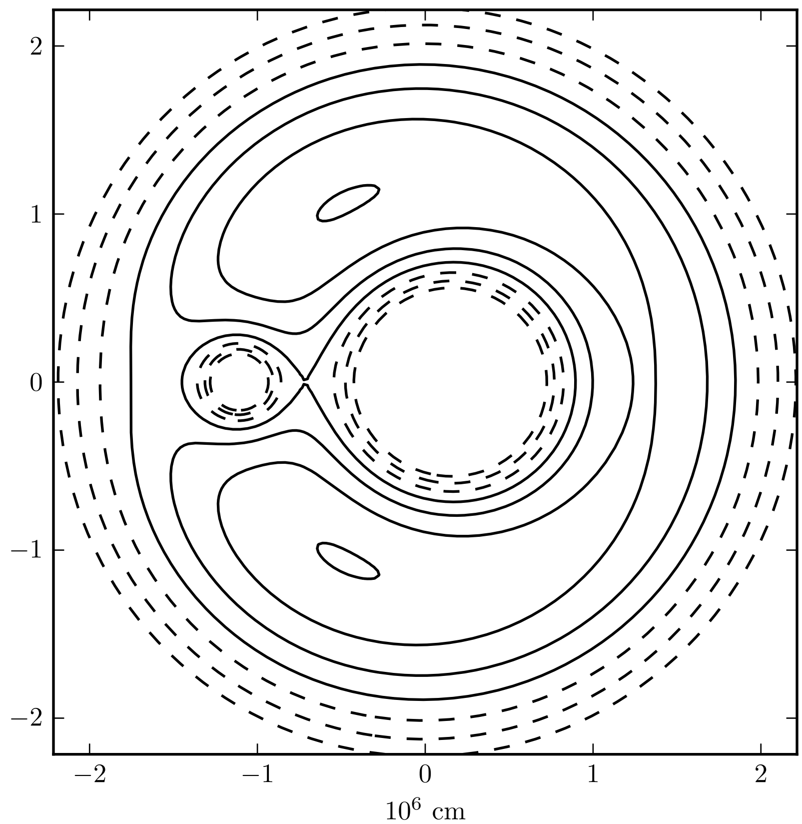

Figure 1–2 Geometry of the Roche potential as seen looking down on the orbital plane. This particular configuration is in fact that of J1023, according to our best estimate (see Chapter 5 and in particular Figure 5–6). Curves are equipotential lines, and L1 is the saddle point where the curves cross. Equipotentials below that of L1 are drawn dashed. The system centre of mass is at the origin of these coordinates.

rotates with the orbit. In such a coordinate system, the gravitational forces are augmented with a centrifugal force and a Coriolis force, but for non-moving objects the field can nevertheless be described by a potential (the Roche potential). This potential is pictured in Figure 1–2. Certain key features are of interest for binary systems in which accretion may have occurred. In particular, there are two wells in the potential, one around each object. Between the two objects is a point, the first Lagrange point (L1), where the forces are balanced. The equipotential line passing through this point forms a two-lobed hourglass-like figure; the lobe containing a mass is called that mass' Roche lobe. Because a star is composed of fluid, its surface will lie approximately along an equipotential surface. If a star is extensive enough to occupy a substantial part of its Roche lobe, the star will not be spherical. If the star grows large enough (or the Roche lobe shrinks enough) that the star overfills its Roche lobe, material will flow through L1 and fall towards the other body. (If both Roche lobes are full, the system will enter a highly dissipative and poorly understood common-envelope phase.) The basic Roche geometry may be evaluated numerically, but there are some useful standard approximations describing its shape. The size of the Roche lobe is expressed in terms of the radius r_R of a sphere with the same volume. For a mass ratio $q = M_1/M_2$, the radius r_R of the Roche lobe containing the first mass is given to

PSR J1903+0327, an MSP in an eccentric binary in the galactic field (Freire et al., 2011). Nevertheless a circular binary is thought to be the normal scenario for formation.

within about 2% by a pair of expressions by Paczyński (1971):

$$\frac{r_R}{a} = 0.38 + 0.2 \log_{10} q, \tag{1.6}$$

for $0.3 < q < 20$, or

$$\frac{r_R}{a} = 0.46224 \left(\frac{q}{1+q} \right)^{1/3}, \tag{1.7}$$

for $q < 0.8$, where a is the semi-major axis of the system. A smooth (but less analytically convenient) approximation valid for all q was given by Eggleton (1983):

$$\frac{r_R}{a} = \frac{0.49 q^{2/3}}{0.6 q^{2/3} + \log(1 + q^{1/3})}. \tag{1.8}$$

The fast rotation and gradual, regular spin-down have made MSPs valuable objects of study. Many applications simply use them as clocks, for testing general relativity (e.g. Antoniadis et al., 2013) or for measuring pulsar and companion masses (e.g. Demorest et al., 2010). One current project is an attempt to use an array of MSPs to detect many-year-timescale gravitational waves passing over the Earth (Hobbs et al., 2010). The emission physics of MSPs are complicated: for example, there are at least two observationally distinct sub-groups (based on the relation between their γ-ray and radio pulse profiles) that seem to require different emission mechanisms (Venter et al., 2009).

1.3 Low-mass X-ray binaries

The recycling process for MSPs is thought to involve accretion, heating material to high enough temperatures that it emits X-rays. Indeed, many accreting binary systems are observed as bright X-ray sources. For an excellent review article on X-ray binaries and their evolution, see Bhattacharya and van den

Heuvel (1991) or Tauris and van den Heuvel (2006). X-ray binaries are roughly classified by the mass of the donor star: in high-mass X-ray binaries the donor star is $\gtrsim 10M_\odot$, while in low-mass X-ray binaries (LMXBs) the donor is $\lesssim 1M_\odot$. (Intermediate-mass X-ray binaries occur but are rare.)

In a high-mass X-ray binary system, the companion is not much longer-lived than the star that gave rise to the compact object, so we expect stellar evolution to limit the time during which accretion can occur. Binary evolution processes also ensure that high-mass X-ray binaries are not long-lived. LMXBs, on the other hand, may undergo accretion for many millions of years. This allows sufficient time for a neutron-star companion to be spun up to a millisecond period.

While an LMXB is accreting, mass overflows the companion's Roche lobe and forms an accretion disc around the companion. For high mass transfer rates, this accretion disc carries material down to the neutron star surface, depositing it more or less uniformly. For lower mass transfer rates, the inflowing material will be channelled along the magnetic field lines to the magnetic poles. The radius at which the flow switches from Keplerian rotation to channeled flow down the field lines is the magnetospheric radius r_m; this is expressed in terms of the Alfvén radius r_A, the radius at which the ram pressure of the accretion flow (approximated as the kinetic energy density $\rho v(r)^2/2$) equals the static pressure of the magnetic field $(B(r)^2/8\pi)$; here $v(r)$ is the Keplerian velocity.[3] In terms of

[3] Some authors use the escape velocity rather than the Keplerian velocity here, but the difference is not great and can be absorbed into ξ.

pulsar parameters,

$$r_m = \xi r_A = \xi \left(\frac{B^4 R^{12}}{2GM\dot{M}} \right)^{1/7}, \tag{1.9}$$

(equation from Patruno and Watts, 2012; Lorimer and Kramer, 2004) where $\dot{M}$ is the rate at which material falls toward the neutron star and ξ is a correction factor, hopefully of order unity. If the magnetospheric radius is larger than the light cylinder radius, we expect the light cylinder to be free of plasma, and the pulsar mechanism should be active, producing radio emission and a strong particle wind.

A second key distance is the corotation radius r_{co}, the point at which the Keplerian rotation of the accretion disc matches the rotation of the neutron star (equation from Patruno and Watts, 2012):

$$r_{co} = \left(\frac{GMP^2}{4\pi^2} \right)^{1/3}. \tag{1.10}$$

If the magnetospheric radius is larger than the corotation radius (but smaller than the light cylinder radius), then the inner edge of the disc will be moving slower than the field lines, and any coupling should spin the pulsar down. If the magnetospheric radius becomes smaller than the cororation radius, the inner edge of the accretion disc will exert a spin-up torque on the neutron star. Thus for a given mass transfer rate, there is an equilibrium spin period P_{eq} for the neutron star (equation from Patruno and Watts, 2012):

$$P_{eq} = 1 \text{ ms} \left(\frac{B}{10^8 \text{ G}} \right)^{6/7} \left(\frac{M}{1.4 M_\odot} \right)^{-5/7} \left(\frac{\dot{M}}{10^{-9} M_\odot \text{ yr}^{-1}} \right)^{-3/7} \left(\frac{R}{10 \text{ km}} \right)^{16/7}. \tag{1.11}$$

Under most models of neutron star spin-up, the mass transfer rate is presumed to vary sufficiently slowly that the neutron star is always at or near this equilibrium. This equilibrium period can be written in terms of the pulsar period derivative, combined with the mass transfer rate. There is a natural maximum mass transfer rate, the Eddington rate, at which the radiation pressure is sufficient to support the infalling material against gravity (equation from Bhattacharya and van den Heuvel, 1991):

$$\dot{M}_{Edd} = \frac{R}{10^6 \text{ cm}} 1.5 \times 10^{-8} M_\odot \text{ yr}^{-1}. \tag{1.12}$$

We do not expect pulsars to be spun up beyond the equilibrium period corresponding to the Eddington rate; this defines a spin-up line in the $P - \dot{P}$ diagram above which we do not expect to find recycled pulsars:

$$P_{\text{spin-up}} = 1.89 \text{ ms} \left(\frac{B}{10^9 \text{ G}}\right)^{6/7}. \tag{1.13}$$

Although the exact position of this line depends on many assumptions (for example, the Eddington rate assumes a pure hydrogen plasma), it can be seen in Figure 1–1 that of the currently known MSPs, only two are above the line, and those two are close.

The actual behaviour of LMXBs is extremely complicated, often involving episodic accretion, thermonuclear bursts, and varying accretion rates (see Bhattacharya and van den Heuvel, 1991, for a review). In a few objects, the accretion rate is sometimes low enough that the accretion is channeled onto the magnetic poles, and one can observe coherent pulsations in the X-ray data. Such objects are called accreting millisecond X-ray pulsars (AMXPs), and the presence of coherent pulsations is a powerful tool to study their behaviour; Patruno and Watts (2012) present a recent review.

1.4 Black widows and redbacks

As part of a survey for MSPs, Fruchter et al. (1988b) discovered a MSP, PSR B1957+20, in a binary system that exhibited radio eclipses. It was already clear that the eclipsing region was larger than the companion's Roche lobe, so these eclipses were not the result of the companion blocking our line of sight; rather, there was some ionized material in the system, mostly near the companion, that was blocking radio signals. The authors therefore suggested that the pulsar might be evaporating its companion; such a phenomenon could explain the existence of non-binary MSPs. As a result PSR B1957+20, and the class of similar pulsars, came to be called "black widow" pulsars, after a species of spider famous for[4] the female's consuming the male during mating.

In addition to the eclipses observed in the black widow pulsar, long-term timing showed that the orbital period was varying in a quasi-periodic way over several years. Applegate and Shaham (1994) proposed a mechanism, gravitational quadrupole coupling, in which magnetic activity in the companion combined with

[4] In fact, sexual cannibalism is common among spiders, but the black widow spider was named for the phenomenon.

differential rotation varies the companion's gravitational quadrupole, which has the effect of varying the orbital period.[5]

Since the discovery of PSR B1957+20, several other black widow systems have been discovered (PSR J0610−2100, PSR J2051−0827; Burgay et al., 2006; Stappers et al., 1996, respectively). They are recognized by their characteristic frequency-dependent radio eclipses and other signatures of ionized material in and around the system.

Another class of systems has recently been recognized; these systems show all the same signs of interaction between the pulsar and its companions, but the companion is more massive and seems to be non-degenerate. If stellar evolution is still going on in these companions, then we would expect it to drive mass transfer as the companion evolves off the main sequence. Indeed, the discovery of J1023, described in Chapter 3, suggests strongly that these systems are at the end of an LMXB phase and the beginning of a MSP phase. These systems have been dubbed "redbacks."[6] There are now known several redbacks in the Galactic field (e.g.

---

[5] For example, if the companion becomes more oblate, the gravitational field in its equatorial plane becomes stronger; to maintain a stable orbit the orbital velocity must increase. Since the orbital angular momentum remains unchanged, the orbital separation must decrease and the orbital period must also decrease. This is explained and the formulae are worked out in detail in Applegate (1992).

[6] Redbacks are named after an Australian species of spider in which the male actively assists the female in sexual cannibalism (Andrade, 1996). The implication is that these systems are like the black widows except that in this case the companion is driving binary evolution and therefore its own ablation by Roche lobe overflow.

PSR J2215+51, PSR J1723−28; Roberts, 2011). Several previously known systems (e.g. 47 Tuc W, PSR 1740−5340) in globular clusters are now also classified as redbacks (although given the high rate of stellar interaction in clusters their evolutionary histories are less clear).

1.5 The puzzling origin of millisecond pulsars

The general model of LMXBs "recycling" old pulsars into MSPs is well accepted, but there remain many questions about how the process actually occurs. From a binary-evolution point of view, Deloye (2008) compares the population of currently known LMXBs and MSPs, looking at binary periods and companion masses, and finds it difficult to explain how evolution can turn the first population into the second.

The end of accretion in an LMXB is poorly understood. In particular, models of binary evolution often predict that mass transfer should continue indefinitely (Deloye, 2008), but ionized material entering a neutron star's (NS) magnetosphere is expected to "short out" magnetospheric emission (Stella et al., 1994). How then can we explain the observed radio MSPs?

J1023, the subject of this thesis, is an interesting system for studying this question. As we describe in Chapter 3, this system seems to have had an accretion disc in 2001, but it is now active as a radio MSP. This suggests that we are observing the end of an LMXB phase and the birth of a radio MSP, during which the system undergoes transient LMXB phases. This is consistent with a scenario put forward by Burderi et al. (2001) in which LMXBs, late in their evolution, may have episodes of mass transfer but spend most of their time in “radio ejection” phases, during which the wind from an active radio MSP sweeps material overflowing the companion out of the system.

In its current quiescent state, J1023 does not appear to contain a disc (see Chapter 3), but it does show signs of substantial ionized material in the system (Chapter 5), an X-ray shock possibly due to interacting winds (Chapter 4), and orbital period variations possibly due to motions in the companion (Chapter 5). These phenomena have also been observed in the black widow and redback systems. It is my hope that the study of J1023 and these phenomena it exhibits will clarify how LMXB phases end and MSPs are born.

CHAPTER 2
Observing Pulsars

In this section I will summarize key observing techniques and instruments used for the research in this thesis. For an excellent overview of pulsar observing techniques, see Lorimer and Kramer (2004).

2.1 The radio signal from pulsars

The radio signal from a pulsar is normally an extremely regular pulse whose intensity and polarization profiles are characteristic of the pulsar. This pulse is broadband, with flux density typically falling off as a power law with observing frequency: $F_f \propto f^\alpha$. The spectral index α has a wide range, with population average -1.8 and standard deviation 0.2 (Lorimer and Kramer, 2004).

The emission time of the pulse from a pulsar is roughly independent of observing frequency, depending only on the moment the pulsar's beam sweeps past us. The pulse is then modified as it propagates through the interstellar medium. Though tenuous, the ionized parts of the interstellar medium introduce a frequency-dependent delay characteristic of the propagation of waves above the plasma frequency. To a good approximation (plasma frequency approximated as zero), this delay is

$$\Delta t_{\mathrm{DM}} = 4.15\ \mathrm{ms}\ (f/1\ \mathrm{GHz})^{-2}\ (\mathrm{DM}/\mathrm{pc\ cm}^{-3}), \tag{2.1}$$

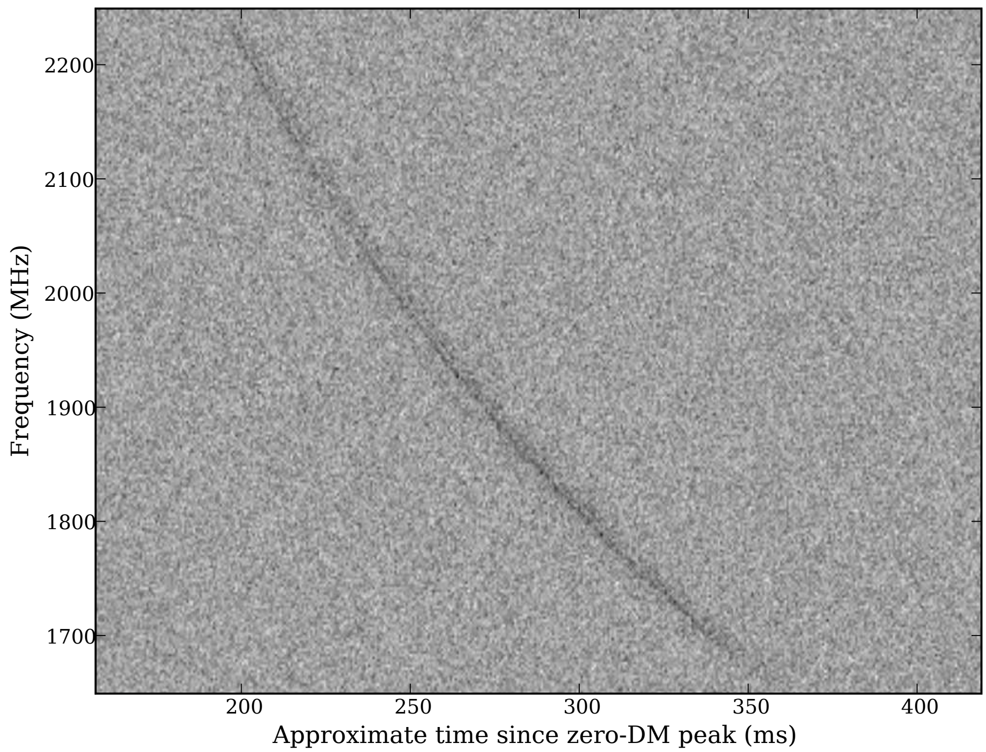

Figure 2–1 Dispersion delay of a single pulse from RRAT 1856−02. This figure shows flux density as a function of time and observing frequency, taken from the long observation described in Appendix A. The rotating radio transient (RRAT), also described in Appendix A, emits pulses very infrequently, though always at the same rotational phase. This is a single pulse, presumably emitted as a broadband pulse roughly simultaneous at all frequencies; propagation through the interstellar medium at a DM of 237 pc cm^{-3} has delayed the lower frequencies relative to the higher, showing the characteristic $1/f^2$ curve implied by Equation 2.1.

where DM is the integrated column density of electrons, known as the dispersion measure (hereafter abbreviated to DM; Lorimer and Kramer, 2004). Typical pulsars have DMs of tens to hundreds of pc cm^{-3}. Given a painstakingly constructed map of the free electron density in the Galaxy (e.g. Cordes and Lazio, 2002), a DM measurement can be used to infer a rough estimate of the distance to a pulsar.

In addition to the bulk frequency-dependent delay summarized by DM, the interstellar medium can introduce more complex frequency- and time-dependent changes to the pulsar signal. Inhomogeneities in the interstellar medium produce scintillation: refraction and/or diffraction by the interstellar medium near the line of sight increase or decrease the flux density at each observing frequency. As the line of sight sweeps through the interstellar medium, these intensities change as a function of time. In addition, the ensemble of lines of sight that contribute to the detected pulsar signal produces a range of delays, so that observed pulsar profiles acquire roughly-exponential "scattering tails" that are longest at low frequencies. The complex phenomena of scintillation and scattering are described in detail in Lorimer and Kramer (2004). For some kinds of observing it is sufficient to summarize the scintillation with a bandwidth Δf and timescale Δt over which features appear and disappear within a given observing band.

2.2 Radio telescopes

In the course of the research described in this thesis I have made use of observations taken with a number of different radio telescopes. Here I will give a brief description of each.

Figure 2–2 The Robert C. Byrd telescope at Green Bank, West Virginia. This photo was taken during the mid-2007 azimuth track replacement, so that the telescope was constrained to operate as a transit instrument.

The Robert C. Byrd telescope at Green Bank, pictured in Figure 2–2, is near the town of Green Bank, West Virginia. Operated by the National Radio Astronomy Observatory in the United States, this telescope is in a National Radio Quiet Zone, within which regulations restrict radio-frequency emissions and observatory staff attempt to limit the production of radio-frequency interference. The telescope is a 100-m by 110-m elliptical off-axis reflector with an unobstructed aperture. Frequencies above about 900 MHz are focused by a subreflector to the scondary focus, where several receivers are mounted on a turret, while 350 MHz and 820 MHz receivers can be mounted on an arm that extends into the prime focus. The primary reflector consists of panels roughly 1 m by 2 m with linear actuators at each corner, allowing the telescope surface to be adjusted to compensate for distortions

Figure 2–3 The Arecibo telescope. The Gregorian subreflector and azimuth arm are visible here; the opening in the center of the dish is used during maintenance to lift equipment up to the receiver room.

due to the weight of the telescope or heating from the Sun; this active surface is generally used only above 2 GHz.

The William E. Gordon telescope at Arecibo, operated by SRI international, USRA, and UMET under cooperative agreement with the National Science Foundation, is in the municipality of Arecibo in Puerto Rico. The telescope, pictured in Figure 2–3, is a 300-m spherical reflector built into a roughly circular valley in the local karst. The telescope is pointed by moving a Gregorian subreflector mounted on an azimuth arm. This configuration allows the telescope to observe out to a zenith angle of 20 degrees; the telescope can observe sources at declinations from 0 to +40°, though near these limits the sources are visible only for a short time each day. The telescope mounts single-beam receivers for a variety of bands including 327 MHz and 1400 MHz, and it also mounts a seven-beam receiver for 1400 MHz.

The Parkes Radio Telescope, operated by the Australia Telescope National Facility, is a 64-m fully-steerable dish near the town of Parkes, in Australia. The receiver cabin is at the prime focus, and contains several detectors, including a 13-beam receiver for 1400 MHz and a dual-band receiver capable of simultaneously observing at 600 and 3000 MHz.

The Lovell Telescope, operated by the University of Manchester, is a 76-m fully-steerable dish at the Jodrell Bank Observatory in the north-west of England. The altitude bearings are 15-inch gun turret bearings from the HMS Revenge and Royal Sovereign. Receivers are mounted in the "focus box" at the prime focus of the dish.

The Westerbork Synthesis Radio Telescope, operated by ASTRON, is near the village of Westerbork in the northern Netherlands. The telescope consists of a linear East-West array of 14 25-m dishes spaced over 2.7 km. This telescope was designed to act as an interferometer, computing pairwise correlations to allow high-resolution synthesis imaging (Burke and Graham-Smith, 2009). When observing point sources such as pulsars, though, it is operated in a much simpler "tied array" mode. In this mode, an accurate pulsar position is used to compute the delays between the arrival time of the signal at each telescope, and the sampled voltage data streams are delayed by this amount and added. This effectively synthesizes a single narrow beam — the same beam whose averaged flux forms the central pixel of an interferometric image — directed at the pulsar, with the same sensitivity as a single dish with the same total collecting area — effectively a 94-m single dish.

2.3 Finding pulsars in radio surveys

Most of the known pulsars have been found as part of a planned pulsar survey. In fact, most of the known pulsars were found in a single survey, the Parkes Multibeam Pulsar Survey (Manchester et al., 2001, and many later papers). Lorimer and Kramer (2004) provide an overview of how such surveys are generally conducted. For definiteness, though, I will give an overview of the survey process here. I will provide examples of quantities taken from Boyles et al. (2013) and Lynch et al. (2013), which describe a specific modern survey, the Green Bank Telescope 350 MHz drift-scan pulsar survey; in fact this is the survey that discovered J1023.

The process of finding new pulsars is normally a search for periodic radio signals in each of a grid of pointings on the sky. The pointings may be discrete telescope pointings, where the telescope tracks a piece of sky for the (typically a few minutes) duration of the observation, then slews to the next pointing. Alternatively, the survey may be conducted in "drift scan" mode, where the telescope's azimuth and elevation are held fixed while the Earth turns; a "pointing" is then the time from when a point on the sky enters the telescope beam to when that point leaves the telescope beam (e.g. $\sim$140 seconds). Figure 2–4 shows the sky coverage of the GBT350 drift-scan survey. Observations may also be carried out with a multi-beam receiver, observing multiple points in the sky simultaneously.

The observing frequency chosen for pulsar searching depends on many factors, including the availability of suitable receivers, the beam size (particularly for drift scan surveys), the Galactic background noise, and the expected spectral index of

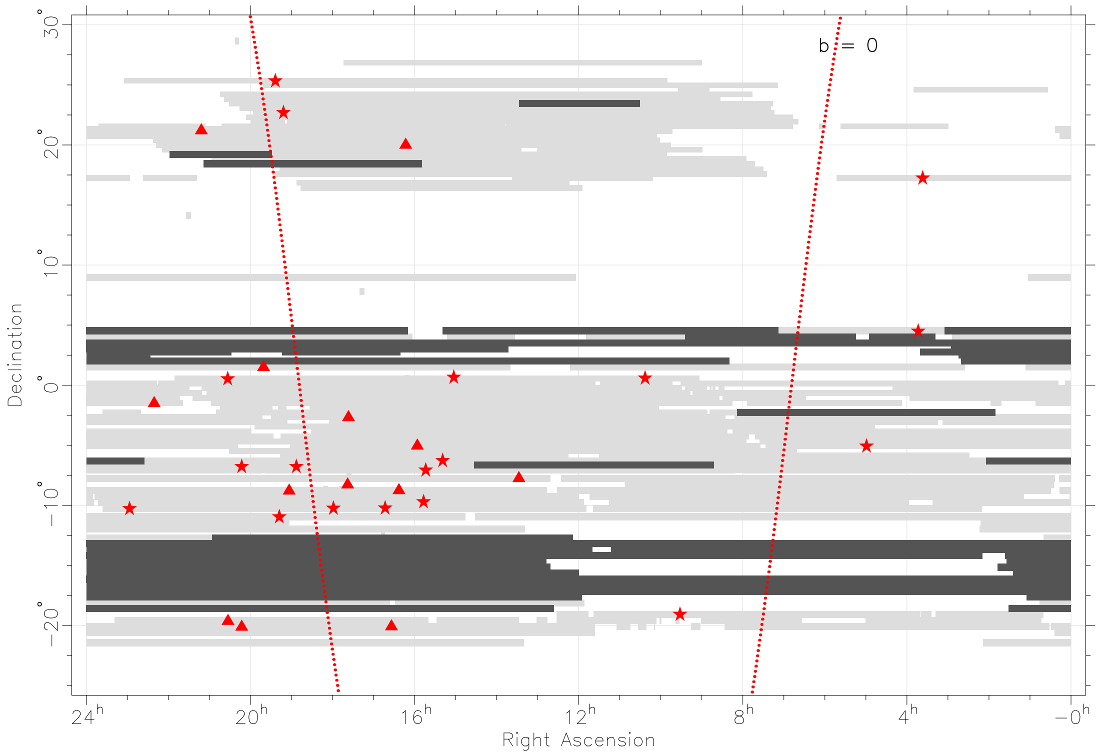

Figure 2–4 The GBT 350 MHz drift-scan survey coverage. Grey regions indicate data examined by the survey team, while black regions indicate data examined by the Pulsar Search Collaboratory, an outreach organization. Red stars and triangles denote new pulsars discovered by the survey. The dotted red line indicates the plane of the Galaxy. Figure from Boyles et al. (2013).

the pulsars of interest. Most recent surveys have been carried out at 350 MHz or 1400 MHz.

During a pointing, the receiver amplifies a range of frequencies (e.g. 325–375 MHz) and mixes them down to a frequency range centered on zero (using quadrature mixing[1] to produce the real and imaginary parts of a complex signal; Burke and Graham-Smith, 2009). This complex signal is then sampled by an analog-to-digital converter, channelized, and the average power in each channel is recorded frequently. The width of each channel is selected so that the expected smearing due to dispersion within that channel is small enough to permit detection of all pulsars of interest (e.g. 2048 channels of 24.4 kHz). Similarly, the sampling time[2] for each channel is chosen small enough to avoid smearing out pulsars of interest (e.g. 81.92 μs). While radio telescopes are unavoidably sensitive to polarization, this is generally not used in pulsar searching, so typically one records only the summed power coming from both polarization channels. Nevertheless,

[1] Quadrature mixing is a technique for frequency down-conversion where the input signal is mixed with (multiplied by) a pair of signals in quadrature: cosine and sine waves at some frequency f. The two outputs are then treated as a single complex signal, and an input signal at frequency f' produces a complex output signal at frequency $f' - f$. Since the output signal is complex, positive and negative frequencies can be distinguished, and if the output channels each have a bandwidth B, any input frequency from $f - B$ to $f + B$ can be accurately represented.

[2] Each sample measures the average power within a channel since the last sample. This is distinct from the notion of sampling generally used when discussing Fourier transforms and aliasing, and it results in a loss of sensitivity to pulsars with spin frequencies near or above the sampling rate; see Ransom et al. (2002).

the necessity of recording the power in thousands of channels tens of thousands of times a second produces large data rates (e.g. 90 GB per hour).

Once the data have been recorded, the analysis requires substantial computation. Since the DM is completely unknown and the signal-to-noise in a single channel is minuscule, one must choose a list of trial DMs up to some maximum (e.g. ten thousand DMs up to 1000 pc cm^{-3}), dedisperse the data at each DM — that is, time-shift each channel appropriately and then add them — and analyze the resulting time series.

Analyzing each time series for periodicity is largely a process of taking a fast Fourier transform and searching for peaks, but there are a number of subtleties (see Ransom et al., 2002, for much more detail). First of all, although a normal pulsar is extremely stable over survey observing times, a pulsar will not typically be exactly at a Fourier frequency, so its energy will be spread over several frequencies in the power spectrum (typically as a $\sin x/x$ profile). One therefore typically uses some form of interpolation to search for peaks between Fourier frequencies. Pulsar pulse profiles are also typically relatively narrow peaks, so much of the power appears in higher harmonics. One must therefore somehow combine the power from several harmonics; this is typically done incoherently (ignoring their relative phase) to limit the computational burden.

For pulsars in binaries, the binary motion affects the observed signal. The orbital velocity introduces a simple shift in the observed frequency, not a problem since we search all plausible frequencies anyway. The orbital acceleration, on the other hand, can introduce frequency shifts within the observation time. This can smear the pulsar's power over several Fourier frequencies. Fortunately there are efficient techniques for searching for accelerated pulsars as well.

There are several suites of software tools for carrying out this searching operation; in this thesis we used PRESTO (Ransom, 2001a). Such tools are assembled into a "pipeline" whose input is raw observation data and whose output is a list of candidate periodic signals.

Unfortunately, the vast majority of periodic signals detected by the above process are of terrestrial (or near-Earth) origin. These signals may be due to internal "birdies" in the receiver, interference generated by telescope hardware, sparks from car engines, stray signals from cell phones or microwave ovens, radar transmissions, satellite transmissions, or other unidentifiable sources. All these non-pulsar signals are called "radio-frequency interference" (RFI). Most telescope installations attempt to minimize the levels of RFI in their neighbourhoods, and survey pipelines carry out a certain amount of automated RFI excision. Nevertheless, an essential part of pulsar searching is somehow sorting through the many candidates due to RFI looking for the few genuine pulsars.

2.4 Timing pulsars with radio telescopes

Once a radio pulsar has been found, its period, DM, and approximate location are known. For a measurement of the period derivative, or if the pulsar is in a binary system, the binary parameters, more work is needed. The technique of phase-coherent timing permits extremely precise measurement of many pulsar parameters. Lorimer and Kramer (2004) provide a description of this technique, but I will provide an overview here.

The key idea of phase-coherent timing is that a single observation, or a single part of an observation, can yield the arrival time of a single pulse. If one knows the number of turns between this and another pulse whose arrival time we know, one has a strong constraint on the pulsar's apparent rotation. One can then use this constraint to predict the number of turns to the next pulse whose arrival time is measured; if this number of turns is correct, the residual phase error sharpens the constraints on the pulsar's apparent rotation. This ability to track the exact number of turns between observations relies on observations with sufficient cadence and quality to maintain a good prediction of the pulsar's rotation.

A pulsar's apparent rotation contains many pieces of information, often more easily understood in terms of delays in pulse arrival times relative to the pulsar's intrinsic spin-down. For example, the Earth's motion around the Sun introduces delays in the pulse arrival times as the distance from the observatory to the pulsar changes; these delays depend on the pulsar's position in the sky (the effect of an error in assumed position on the delays is shown in Figure 2–5). If the pulsar is in a binary system, the binary motions will also induce delays. Accurate measurement of these pulse arrival times can permit fitting for these and many other pulsar parameters, which much nevertheless be chosen carefully to avoid unmeasurable or highly-covariant combinations (for example, a constant radial velocity of the pulsar is conventionally absorbed into the pulsar and orbital periods through a redefinition of the second to accommodate the Doppler effect).

Observations for the purpose of phase-coherent timing resemble pulsar search observations in certain ways. One uses the same receivers, and one channelizes the

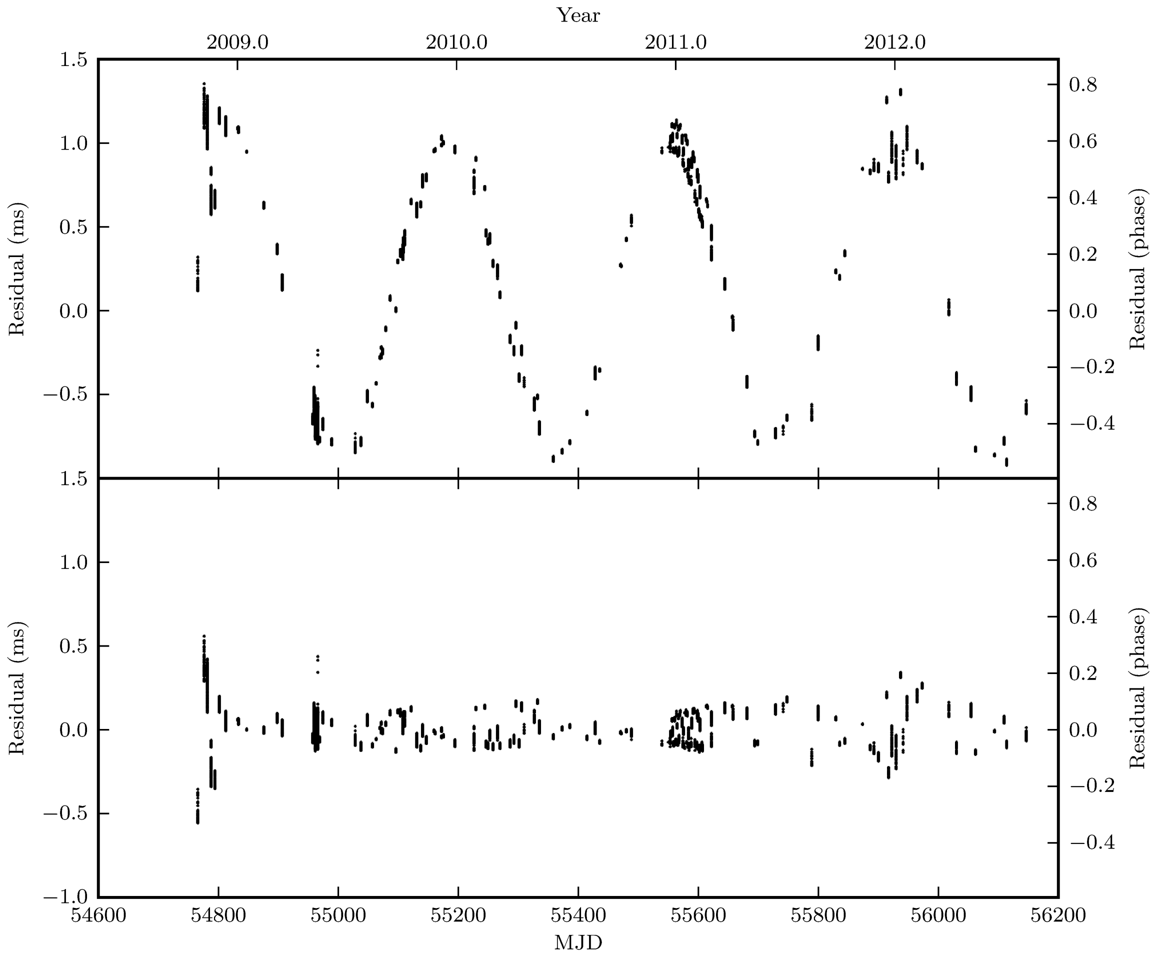

Figure 2–5 The effect on pulsar timing of a position error. These plots show the timing residuals, that is, how early or how late pulses arrive compared to a model of the pulsar's rotation. The top panel uses a model with a one arcsecond error in declination, while the bottom panel uses the best-fit model for this pulsar. The pulsar in this case is J1023, and the data are those described in Chapter 5; note that J1023 suffers from a variety of complicated timing behaviours that make the residuals of the best-fit model substantially larger than one would expect from their microsecond-level scatter (the error bars corresponding to this scatter are too small to be seen on this plot; the apparent vertical extent of points is actually the result of many measurements within an orbit being spread vertically by imperfect modelling of the changing orbital parameters). Fortunately Deller et al. (2012) measured, among other things, the position of J1023 to sub-milliarcsecond accuracy, so it is not necessary to fit for the astrometric parameters of J1023.

incoming signal into many channels. In some observing systems, it is possible to use the known DM and signal-processing techniques to coherently dedisperse the signal within each channel (Hankins and Rickett, 1975); in others one simply uses enough channels to minimize smearing due to within-channel dispersion. In either case one typically also uses a model of the pulsar's rotation to "fold" the incoming data, averaging the power in each of many (hundreds to thousands of) phase bins. This produces an average pulse profile for each channel; these are accumulated for a subintegration time (tens of seconds to minutes) and then dumped to disk to begin a new subintegration. The result is a data (hyper-)cube, power as a function of frequency channel, phase bin, subintegration number, and polarization (either Stokes I, Q, U, and V or some other representation). Several standard formats exist for storing such data cubes; the emerging standard is PSRFITS (Hotan et al., 2004), a specific form of FITS file, but PRESTO has its own format, as do several telescope backends.

Given such a data cube, one typically averages together enough frequency channels and subintegrations to obtain an adequate signal-to-noise ratio. Each profile in the resulting data cube is then compared to a reference profile to determine its phase shift, and therefore the time of arrival of a fiducial point on the pulse within that frequency band and subintegration. Each such measurement is called a time of arrival (TOA). TOAs are usually expressed as a modified Julian date (MJDs), the number of days since 1858 Nov 17 at midnight. At this point the data volume is usually modest enough to express in one of several text formats.

Given a collection of TOAs from a pulsar and a model for its rotation, one will typically use one of the programs `tempo`[3] or `tempo2` (Hobbs et al., 2006). These programs read in the TOAs, use the rotation model (see Edwards et al., 2006, for a detailed decription) to assign a turn number to each pulse, and compute the phase residuals. They then treat the problem as a linear one, using the derivative of the timing model with respect to each parameter, and carry out a least-squares fit (when iterated, this is known as the Gauss-Newton algorithm for solving nonlinear least-squares problems). If the turn numbers are correct, one or two steps of this process typically result in a model that is very nearly a least-squares fit to the TOAs. If one or more of the turn numbers are incorrect, this process will not improve the model, and these programs rely on human intervention to adjust turn numbers until the fit can be carried out.

2.5 X-ray observations of pulsars

Pulsars can emit X-rays in a variety of ways. They are sufficiently small and massive that they emit thermal X-rays for quite a long time after their birth, or if heated by accretion or by particles streaming down from their magnetosphere. This magnetosphere may produce X-rays itself, through curvature radiation or inverse Compton scattering. Pulsars also produce strong particle winds, which may create X-ray-emitting shocks when these winds meet the interstellar medium, or a companion, or that companion's wind. A few pulsars, known as magnetars, with high magnetic fields can emit X-rays in the form of months-long outbursts,

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

Figure 2–6 Artist's impression of the *XMM-Newton* telescope. Image thanks to ESA/Ducros.

milliseconds-long bursts, or persistent emission, through processes that remain somewhat mysterious. For an overview of the ways in which pulsars emit X-rays, see Kaspi et al. (2006), or Mereghetti (2008) for more on magnetars.

Observing X-rays must be done from outside the atmosphere. There are currently a number of X-ray telescopes in orbit, but the most relevant ones operate on similar principles. A few serve as all-sky monitors, including the All-Sky Monitor aboard the now-deorbited *Rossi X-ray Timing Explorer* (which we use in Chapter 3). Most of the currently active X-ray telescopes are imaging telescopes similar to the *XMM-Newton* telescope (pictured in Figure 2–6), which we used for Chapter 4 and which I will describe here.

The *XMM-Newton* satellite, 10-m long and with a mass of 4 tonnes, was launched in 1999 into a high elliptical orbit having apogee of 114000 km and perigee of 7000 km. The lower part of this orbit is within the radiation belts, and is not used for science observations. The satellite carries three coaxial X-ray telescopes and a small optical monitor instrument. Redundant star trackers are able to maintain attitude stability of one arcsecond. The downlink was designed to handle 64 kbit s^{-1}, so telemetry volume can restrict available time resolution. For a more detailed overview of the satellite see Jansen et al. (2001).

The three X-ray telescopes aboard *XMM-Newton* use identical mirror modules. Each mirror module consists of 58 nested mirrors using Wolter I (grazing-incidence) geometry. Their focal length is 7.5 m, and the half-energy width of the point spread function (PSF) is 15 arcseconds. In the focal plane assembly are five instruments: the two reflection grating spectrometers (which I will not discuss further as they are only useful for bright sources with narrow spectral features), two EPIC MOS cameras, and the EPIC PN camera.

The EPIC MOS cameras, described in detail in Turner et al. (2001), provide 1.1 arcsecond pixels using front-illuminated CCDs based on metal-oxide-silicon (MOS) technology. The EPIC PN camera, described in detail in Strüder et al. (2001), provides 4.4 arcsecond pixels using back-illuminated CCDs based on PN (p-doped-n-doped junction) technology (which is less vulnerable to radiation damage). Both types of camera can operate in several modes, including a full-frame readout mode for normal imaging, modes with reduced window sizes for faster readout, and one-dimensional imaging modes in which very fast readout is achieved by averaging along one CCD dimension. The PN camera provides faster readout in general, and in particular, its one-dimensional timing mode has a sample time of 29 μs, while the timing mode of the MOS cameras has a sample time of 1.75 ms. Spectral calibration of timing modes is generally poor, but the PN timing mode can be used for high-time-resolution observations of pulsars. The timing capabilities of *XMM-Newton* are described in detail in Kirsch et al. (2004), but in particular, the absolute event timing has an estimated accuracy of 100 μs.

Both types of EPIC camera are capable of recording the energy of incoming photons, and the ability to compare X-ray spectra with models is a key capability of *XMM-Newton*. According to the XMM users' handbook[4] , the MOS and PN cameras have spectral resolutions of ~ 70 and ~ 100 eV at 1 keV, respectively, rising to ~ 200 and ~ 180 eV at 10 keV, respectively. The current status of spectral (and timing) calibration is described in the *XMM-Newton* Calibration Technical Note XMM-SOC-CAL-TN-0018.[5] In general terms, the absolute effective area is accurate to $\pm 10\%$, the relative effective area is accurate to $\pm 5\%$, and the absolute energy scale is accurate to ± 10 eV. For particular observations, it is important to obtain up-to-date calibration data from the *XMM* Science Operations Center, as calibration becomes better understood and as radiation damage to the instrument affects its calibration. For this reason, although observations downloaded from the *XMM* archive are accompanied with preprocessed data, users are advised to

[4] The XMM users' handbook is available from `http://xmm.esac.esa.int/external/xmm_user_support/documentation/uhb_frame.shtml`

[5] This note is available from `http://xmm.vilspa.esa.es/external/xmm_sw_cal/calib/documentation.shtml`.

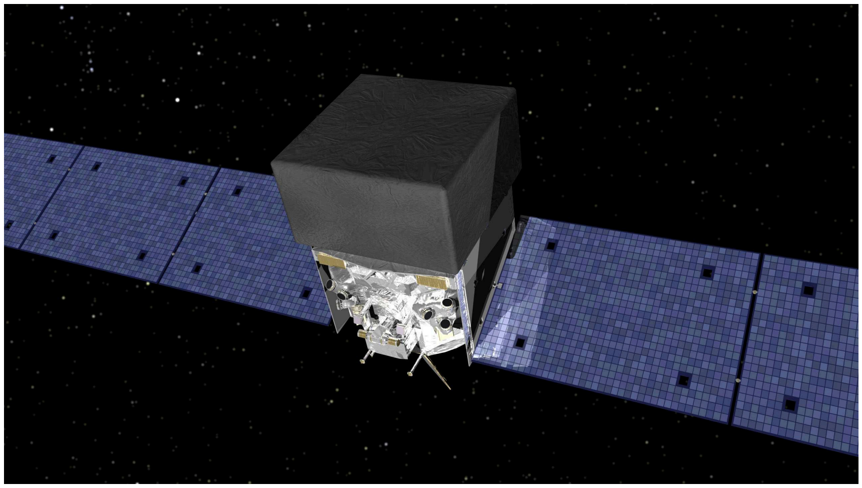

Figure 2–7 Artist's conception of the *Fermi* γ-ray telescope. Image thanks to NASA.

reprocess data with the most up-to-date calibration data and tools. In addition to calibration issues, users should be aware that the *XMM* mirrors can focus soft protons, so that soft proton flares frequently contaminate *XMM* data. These flares can be detected and should be removed before analysis.

2.6 γ-ray observations of pulsars

In addition to their X-ray emission, pulsars can also emit γ-rays. These γ-rays may originate from the pulsar magnetosphere and show pulsations (e.g. Kniffen et al., 1974; Guillemot et al., 2012), or they may originate from shocks, either intrabinary (e.g. Khangulyan et al., 2012) or nebular (e.g. Abdo et al., 2011). One substantial catalog of γ-ray emitting pulsars is presented Abdo et al. (2010), but many others have been discovered since then (The Fermi-LAT collaboration, 2013), and in particular, a number of black widows and redbacks (Roberts, 2011).

The primary telescope for studying γ-rays in the range 100 MeV to 300 GeV is the *Fermi* γ-ray space telescope, launched by NASA in 2008 (and pictured in Figure 2–7). This satellite has two instruments, the Gamma-ray Burst Monitor (which I will not discuss further) and the Large Area Telescope (LAT). The LAT has a roughly 8000 square degree field of view and is operated in a survey mode that rocks from north to south or vice versa every hour-and-a-half orbit, scanning the entire sky every three hours. Observing a point source with this telescope is therefore simply a matter of downloading calibration data and photon lists from the *Fermi* data server.[6]

The *Fermi* LAT, described in detail in Atwood et al. (2009), is a pair-conversion telescope. Incoming γ-rays collide with tungsten strips, producing electron-positron pairs which are tracked through the instrument. The top layer of the LAT consists of alternating layers of tungsten and silicon strip detectors: the tungsten layers induce pair conversion, and the silicon strips track the positions of the produced pairs; the layers alternate because it is important to measure the track initial position with as little multiple scattering as possible. Below the tracker layers is a calorimeter consisting of CsI(Tl) crystals; the function of the calorimeter is to reconstruct the total energy of the particle shower, and also to provide constraints on the particle track for high-energy events. Surrounding the whole instrument is an anticoincidence detector consisting of clear plastic scintillator blocks connected to photomultiplier tubes. Strong background rejection

[6] http://fermi.gsfc.nasa.gov/cgi-bin/ssc/LAT/LATDataQuery.cgi

— the design goal was 99.97% — is necessary because of the high background of charged particles potentially producing γ-ray-like cascades within the instrument. Data acquisition electronics reject the majority of this background aboard the spacecraft to limit telemetry volume; substantial additional processing on the ground classifies the remaining events.

Since the track reconstruction process is more effective for higher-energy photons, the LAT PSF is highly energy-dependent, but its 68% containment width decreases as a power-law from about 4° at 100 MeV to 0.07° at 40 GeV, above which it remains roughly constant. The telescope is sensitive to γ-rays from about 50 MeV up to more than 300 GeV; in practice the high end of the usable energy range is set by the scarcity of high-energy photons.

On-orbit calibration of the *Fermi* satellite is described in detail in Abdo et al. (2009), but in particular, timestamps on events are derived from an onboard GPS-stabilized clock, and have 1-μs accuracy.

Downloaded *Fermi* data are analyzed with a suite of software tools and calibration data, obtainable from the *Fermi* Science Support Center. Instructions for particular analysis tasks are provided in the form of a "Cicerone", a series of worked examples. In particular, for point sources with known (or fitted) spectra, the *Fermi* tool `gtsrcprob` is able to assign to each photon a probability that it is from the source of interest; this takes into account backgrounds and nearby sources whose PSF overlaps with the source of interest. Using these weighted photon lists provides the highest sensitivity for pulsation searches.

Certain mathematical tools are necessary when working with probability-weighted photons. In particular, when forming a histogram, the values in each bin are naturally $\sum_i p_i$ of all the photons falling in that bin, and the uncertainties are approximately (Pletsch et al., 2012)

$$\sigma = \sqrt{\sum_i p_i^2}, \tag{2.2}$$

generalizing the usual rule for unweighted Poisson statistics. The background number of counts per bin b is given by (Guillemot et al., 2012)

$$b = \sum_i p_i(1 - p_i). \tag{2.3}$$

Testing for pulsations can be carried out using a weighted version of the H test developed by Kerr (2011). The H test, originally developed by de Jager et al. (1989), estimates the total power in the first n Fourier harmonics, returning both a false positive probability and the number n of harmonics with significant signal. A natural, data-driven, binning-independent representation of the light curve can be constructed from these first n Fourier coefficients.

CHAPTER 3
Discovery of an X-ray binary/millisecond pulsar link

This chapter is based on the manuscript "A Radio Pulsar/X-ray Binary Link" published in Science *in 2009 October (Archibald et al., 2009). The paper describes the discovery of J1023 as a radio pulsar, discusses the optical history that makes it so unusual, and comments on the implications of the discovery for millisecond pulsar evolution.*

3.1 Introduction

The radio source FIRST J102347.67+003841.2 (hereafter J1023) has long been known to be unusual. It appears optically as a $V \sim 17.5$ magnitude star with a mid-G (solar-type) spectrum and mild 0.4-magnitude 0.198-day orbital variability (Thorstensen and Armstrong, 2005). However, in observations from May 2000 to Dec. 2001, J1023 had a blue spectrum with prominent emission lines (Bond et al., 2002; Szkody et al., 2003) and exhibited rapid flickering by ~ 1 magnitude (Bond et al., 2002). This optical behaviour is typical of an accretion flow, and the double-peaked nature of the lines (Szkody et al., 2003) suggested an accretion disc specifically. This history led to the system's classification as an accreting white dwarf (Bond et al., 2002) or neutron star (Thorstensen and Armstrong, 2005) binary. Since 2002 the system has shown no further sign of accretion, and, as described in Section 3.4, our optical spectroscopic observations of 2008 Dec. 23

and 25 (Figure 3–1) confirm the continued absence of emission lines. Thus J1023 is currently in a quiescent state.

3.2 Discovery of radio pulsations

In mid-2007, we carried out a 350-MHz pulsar survey with the Robert C. Byrd Green Bank Telescope (GBT) in West Virginia. This survey covered $\sim 12,000$ square degrees of sky in the declination range $-21°$ to $+26°$, and used the Spigot pulsar autocorrelation spectrometer (Kaplan et al., 2005) to synthesize a filterbank sampling a 50-MHz band with 2048 frequency channels every 81.92 μs. The survey followed the approach described in Section 2.3; further details of the survey and ongoing data analysis procedures are presented by Boyles et al. (2013) and Lynch et al. (2013). In the data from 2007 June 28, we found PSR J1023+0038, a bright MSP (mean flux density $\sim$ 75 mJy in the 140 s discovery observation, which is shown in Figure 3–2) with a spin period of 1.69 ms at a dispersion measure (DM) of 14.33 pc cm^{-3}.

This candidate was very promising, so we obtained follow-up observations with the GBT. A literature search revealed that the position of this candidate was compatible with that of FIRST J102347.67+003841.2, suggested by Thorstensen and Armstrong (2005) to be a neutron-star binary. After initial confirmation of the candidate at 350 MHz, we switched to 2000 MHz and pointed the telescope at the known position of FIRST J102347.67+003841.2. When we detected strong radio pulsations, it was clear that we had found a binary MSP with a

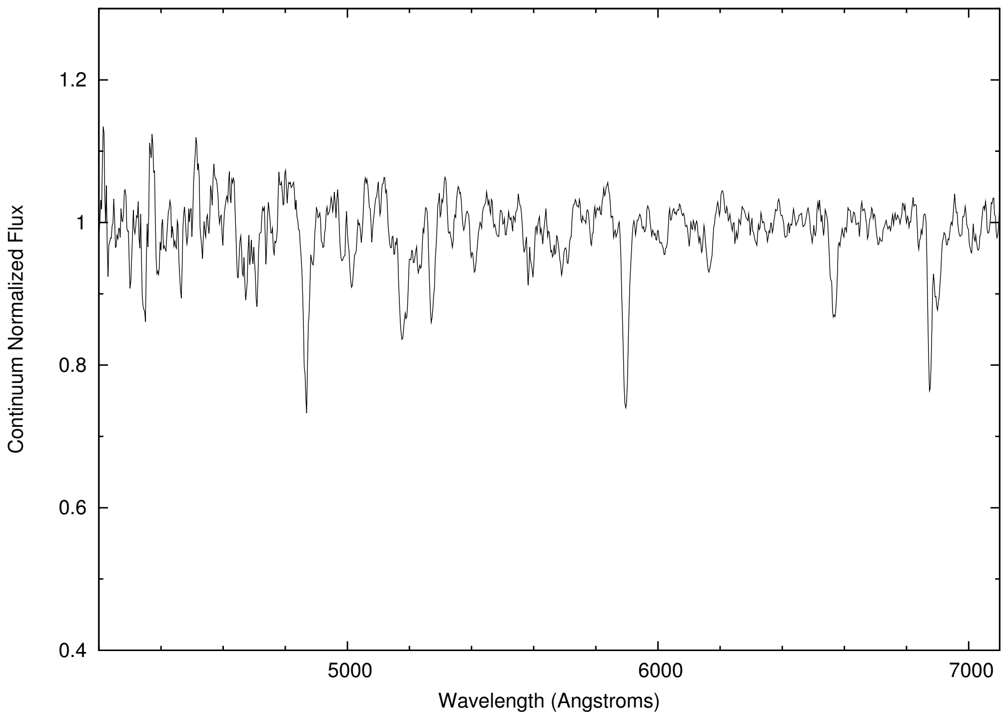

Figure 3–1 Optical spectrum for J1023, taken on 2008 Dec. 23 and 25. The flux was normalized by dividing by the best-fit continuum spectrum. Note the absence of emission lines; this spectrum is consistent with that observed in quiescence and not consistent with that observed during J1023's active phase (see Szkody et al., 2003, for comparison).

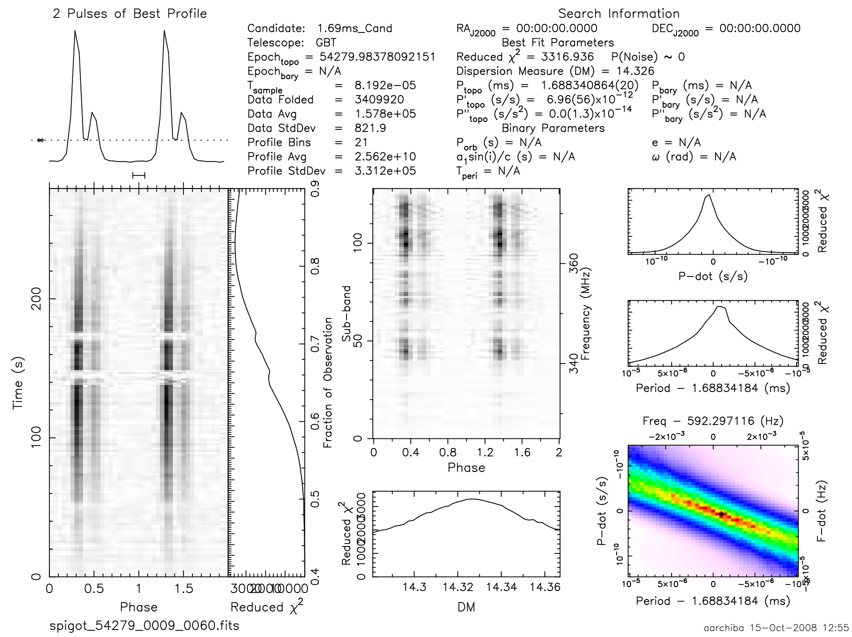

Figure 3–2 A plot of the discovery observation. This plot, output by PRESTO, is very similar to the plots presented for each periodicity candidate. For this particular plot we selected the best section of drift-scan data rather than the overlapping two-minute beams on which the search process operates. The top left panel shows two copies of the pulse profile. The leftmost panel shows flux density as a function of pulse phase and time, averaged over frequency (assuming the best-fit DM) and short time spans. The top middle panel shows flux density as a function of pulse phase and observing frequency, averaged over time (assuming the best-fit period and period derivative) and small frequency ranges. The middle bottom, top right, and middle right panels show root-mean-squared amplitude as a function of DM, acceleration, and pulse period. The bottom right panel shows root-mean-squared amplitude as a function of period and period derivative. Above the plots is a selection of numerical data, some of which describe the candidate and some of which are invalid or irrelevant in the context of a drift-scan search candidate. Note the excellent signal-to-noise of this candidate (quantified by having a reduced χ-squared of 3314 for a no-signal fit), its strong peak at a non-zero DM (evidence that it is less likely to be radio-frequency interference) and broad-band nature. Note also the presence of scintillation effects in the frequency-versus-phase plot and the presence of two short disappearances of the signal in the time-versus-phase plot; both effects are discussed in Chapter 5.

unique[1] history: it had an accretion disc in 2001. The observations described below, in particular the binary orbit derived from pulsar timing, quickly made it certain that the pulsar we had discovered was indeed the compact object in FIRST J102347.67+003841.2; in light of its current nature we generally call the system PSR J1023+0038.

3.3 Pulsar timing

We obtained follow-up radio observations in order to time the pulsar and follow the procedure described in Section 2.4 to obtain its spin and binary parameters (Table 3–1). These observations, summarized in Table 3–2, used several different instruments and backends, in several different radio-frequency bands (see the table notes for specifics of which band and which setup were used). Where possible, we observed simultaneously with both coherent dedispersion backends with online folding (for precise timing) and incoherent dedispersion backends with offline folding (for larger bandwidths and to retain single-pulse and transient information). For these observations we used the GBT, the Arecibo Observatory (AO) in Puerto Rico, the Westerbork Synthesis Radio Telescope (WSRT) in the Netherlands, and the Parkes telescope in Australia (see Section 2.2 for a description of

[1] The 2013 June discovery of M28I provides a second example of a system showing both acretion disc phases and radio pulsar emission; see Section 7.3.

these telescopes). The GBT Pulsar Spigot is an autocorrelation spectrometer[2] (Kaplan et al., 2005). The Green Bank Ultimate Pulsar Processing Instrument[3] (GUPPI) is a new digital filterbank for the GBT. The Wideband Arecibo Pulsar Processors (WAPPs) are a set of autocorrelation spectrometers usable on independent frequency bands (Dowd et al., 2000a). The Parkes 10/50 cm dual-band receiver feeds a set of 1-bit analog filterbanks. The Arecibo Signal Processor (ASP) and the Green Bank Astronomical Signal Processor (GASP) are systems that use a computer cluster to coherently dedisperse (Hankins and Rickett, 1975) and fold incoming data using the best available ephemeris. Further details of the operation of ASP and GASP and the reduction of their output data can be found in Demorest (2007a) and Ferdman (2008a). We used the Pulsar Machine II (PuMa-II) coherent dedispersion backend at the WSRT (an interferometer we operated in phased-array mode). The precise modes we used for each observation are described in the footnotes to Table 3–2.

[2] An autocorrelation spectrometer is an instrument designed to measure the power in each of many frequency channels by using the fact that the power spectrum is the Fourier transform of the autocorrelation. An autocorrelation spectrometer digitally computes the autocorrelation of the input signal up to some maximum lag, averaged over some time resolution. The first step of analysis is then to Fourier transform the data stream in the lag direction, producing a sequence of power spectra at the same time resolution. As computing technology has improved this technique is generally being replaced by the polyphase filterbank, a fast-Fourier-transform-based technique that provides better inter-channel isolation.

[3] https://safe.nrao.edu/wiki/bin/view/CICADA/NGNPP

Parameter	Value
Right ascension (α; J2000)	$10^{\mathrm{h}}23^{\mathrm{m}}47.687(3)^{\mathrm{s}}$
Declination (δ; J2000)	$00^\circ38'41.15(7)''$
Proper motion in α (μ_α)	$10(1)\,\mathrm{mas\,yr^{-1}}$
Proper motion in δ (μ_δ)	$-16(2)\,\mathrm{mas\,yr^{-1}}$
Epoch (MJD)	54802
Data span (MJD)	54757–54847
Dispersion Measure	$14.325(10)\,\mathrm{pc\,cm^{-3}}$
Pulsar period (P)	$1.6879874440059(4)\,\mathrm{ms}$
Pulsar period derivative ($\dot{P}$)	$1.2(8)\times10^{-20}$
Orbital period (P_b)	$0.1980962019(6)\,\mathrm{d}$
Orbital period derivative ($\dot{P}_b$)	$2.5(4)\times10^{-10}$
Orbital period second derivative ($\ddot{P}_b$)	$-5.21(14)\times10^{-11}\,\mathrm{s^{-1}}$
Time of ascending node (MJD)	54801.97065348(9)
$a\sin i$	$0.3433494(3)$ lt-s
Eccentricity	$\lesssim 2\times10^{-5}$
1600 MHz flux density	$\sim 14\,\mathrm{mJy}$
Spectral index (α; $S\propto\nu^{-\alpha}$)	~ -2.8
Surface magnetic field	$< 3\times10^{8}\,\mathrm{G}$
Spin-down luminosity	$< 3\times10^{35}\,\mathrm{erg\,s^{-1}}$

Table 3–1 Pulsar parameters. Parameters listed in the top section are held fixed in the timing fit that produces the parameters listed in the middle section. The position and proper motion values are from the USNO NOMAD optical catalog (Zacharias et al., 2004). With the exceptions of DM, $\dot{P}$, and eccentricity, uncertainties on timing quantities are twice the formal 1σ errors returned by `tempo`. The DM of the source varies substantially: around orbital phase 0.3 it occasionally increases by $\sim$0.15 pc cm^{-3}, and we observe apparent orbit-to-orbit variations as large as $\sim$0.01 pc cm^{-3}. Our estimate of the eccentricity is approximate because DM variations limit the orbital coverage of usable timing data. Our estimate for $\dot{P}$ is approximate because it is highly covariant with pulsar position; it was estimated by carrying out multiple fits for $\dot{P}$ while varying the position according to the optical uncertainties listed above. We have not subtracted Shklovskii or Galactic accelerations (see Section 7.2). The flux density and spectral index are estimated using standard system temperature and gain values for the GBT and the Parkes telescope. The flux density is subject to substantial variability intrinsic to the pulsar (we have observed brightening by a factor $\sim$4 within minutes) for this reason we estimated the spectral index using the simultaneous 600-MHz and 3000-MHz Parkes observations. Interstellar scintillation may nevertheless have affected this measurement. Spin-down luminosity and magnetic field upper limits are calculated from the 3σ upper limit on $\dot{P}$ using Equations 1.1 and 1.2.

Table 3–2. Our Observations of J1023.

Date	MJD range	Orbital phase	Instrument	Frequency (MHz)	Comment
2007 Jun 25	54276.00	0.85	GBT Spigot[a]	350	Nearby track
2007 Jun 28	54279.98	0.96	GBT Spigot[a]	350	Discovery
2008 Oct 18	54757.54	0.71	GBT Spigot[*,a]	350	
2008 Oct 18	54757.58–54757.66	0.91–1.32	GBT GUPPI/Spigot[b]	2000	
2008 Oct 18	54757.68	0.40	GBT GUPPI/Spigot[b]	2000	
2008 Oct 18	54757.74	0.69	GBT GUPPI/Spigot[b]	2000	
2008 Oct 21	54760.76–54760.79	0.95–1.09	GBT Spigot[*,a]	800	
2008 Oct 22	54761.47–54761.51	0.57–0.73	GBT Spigot[*,a]	800	
2008 Oct 22	54761.53–54761.55	0.83–0.96	AO ASP/WAPPs[c]	327	
2008 Oct 24	54763.82–54764.06	0.39–1.59	Parkes 10 cm/50 cm[d]	3000/700	
2008 Oct 27	54766.53–54766.54	0.09–0.14	AO ASP/WAPPs[e]	1400	
2008 Nov 1	54771.50–54771.68	0.20–1.10	GBT Spigot/GUPPI/GASP[*,f]	1400	
2008 Nov 6	54776.49–54776.50	0.37–0.43	AO ASP/WAPPs[e]	1400	
2008 Nov 11	54781.03–54781.25	0.30–1.39	WSRT PuMa-II[g]	150	
2008 Nov 12	54782.47–54782.50	0.55–0.70	AO ASP/WAPPs[e]	1400	
2008 Nov 18	54788.45–54788.48	0.77–0.91	AO ASP/WAPPs[e]	1400	
2008 Nov 24	54794.43–54794.46	0.92–1.08	AO ASP/WAPPs[e]	1400	
2008 Dec 2	54802.42–54802.44	0.28–0.40	AO ASP/WAPPs[e]	1400	
2008 Dec 12	54812.39–54812.42	0.62–0.74	AO ASP/WAPPs[e]	1400	
2008 Dec 22	54822.36–54822.39	0.95–1.09	AO WAPPs[h]	1400	
2008 Dec 31	54831.09–54831.33	0.03–1.22	WSRT PuMa-II[i]	350	

3.3.1 Obtaining a timing solution

From each of our radio observations of PSR J1023+0038, we produced folded profiles for time and frequency subintegrations. We then established a reference profile for each combination of band and backend, usually from the best observation available with that configuration, though we re-used the standard profile from the 1400 MHz ASP observations for the 1388 MHz and 1788 MHz GASP observations because the ASP data had a much higher signal-to-noise ratio. We produced a pulse time-of-arrival (TOA) for each subintegration (generally 30 s or 60 s) by cross-correlating the reference profile with the folded profile for that subintegration (Taylor, 1992). These arrival times are based on observatory time standards which were corrected to UTC using the GPS satellites. We fed these TOAs into the pulsar timing program `tempo`[4] to obtain a timing solution, allowing for an arbitrary time offset of each band/backend combination. The solution we obtain is referenced to the DE405 solar system ephemeris.

The initial timing solution we obtained showed the timing effects of DM variations, as well as possibly some other systematics. We excluded any data that showed signs of short-term DM variations, including all data between orbital phases 0.15 and 0.55, as well as most low-frequency data. Specifically, the only data below 1 GHz we used was one pair of 800-MHz GBT observations separated by a day and one 350-MHz GBT observation. These data sets did not show signs of significant short-term DM variations, and the inter-instrument jump we assumed

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

Table 3–2 (cont'd)

Date	MJD range	Orbital phase	Instrument	Frequency (MHz)	Comment
2009 Jan 2	54833.33–54833.36	0.31–0.48	AO ASP/WAPPs[e]	1400	
2009 Jan 5	54836.16–54836.33	0.60–1.48	GBT GASP*/GUPPI[j]	1700	
2009 Jan 16	54847.28–54847.30	0.76–0.84	AO ASP*/WAPPs[e]	1400	

*Used in the timing solution.

[a]We used the Spigot with a 50 MHz band with 2048 frequency channels and a time resolution of 81.92 μs.

[b]We used the Spigot with an 800 MHz band with 2048 frequency channels and a time resolution of 81.92 μs. GUPPI recorded full Stokes parameters for an 800 MHz band with 2048 frequency channels and a time resolution of 40.96 μs. In both cases only about 600 MHz of the bandwidth is usable.

[c]We used two WAPPs, each with a 12.5 MHz band with 512 frequency channels and a time resolution of 64 μs. ASP recorded a bandwidth of 16 MHz centered at 327 MHz.

[d]The dual-band receiver synthesized two filterbanks, one 64 MHz wide with 256 channels and centered at 590 MHz, and the other 768 MHz wide with 256 channels centered at 3032 MHz. Both had a time resolution of 80 μs.

[e]We used three WAPPs, with center frequencies 1170, 1370, and 1570 MHz, using bandwidths of 100 MHz with 256 channels and a time resolution of 64 μs. ASP recorded a 52 or 56 MHz bandpass centered on 1412 MHz.

[f]We used the Spigot with an 800 MHz band with 2048 frequency channels and a time resolution of 81.92 μs. GUPPI recorded full Stokes parameters for an 800 MHz band with 2048 frequency channels and a time resolution of 40.96 μs. In both cases only about 600 MHz of the bandwidth is usable. GASP recorded an 84 MHz bandpass centered at 1388 MHz.

[g]We used the PuMa-II backend to coherently dedisperse 8 2.5 MHz channels.

[h]We used the same three-WAPP configuration as in f above, but the ASP data was accidentally corrupted.

[i]We used the PuMa-II backend to coherently dedisperse 6 10 MHz channels.

[j]We used GUPPI in an online folding mode, recording full Stokes parameters for an 800 MHz band with 2048 frequency channels and a time resolution of 40.96 μs; only 600 MHz of the bandwidth was usable. GASP recorded a 76 MHz bandpass centered at 1788 MHz.

absorbed any effects of long-term DM variations. To mitigate any remaining systematics, we artificially increased the uncertainties on each individual data set until it was fit with a reduced χ^2 of approximately one.

3.3.2 Estimating DM variations

The relatively strong radio signal from J1023 makes it possible to do detailed analysis of the DM variations. In this section we detail the procedure we used.

The global DM value we report is obtained from the 2008 Oct. 18 GUPPI observation. We produced TOAs for four frequency subbands for each time subintegration, aligning the folded data with a common template, and used the pulsar timing package `tempo` to fit for the DM that minimized the weighted residuals. This process yields a very precise DM, but when we applied an analogous procedure to several other data sets, the results were inconsistent with the precise value and varied from 14.32 to 14.33 $\mathrm{pc\,cm^{-3}}$. We applied a similar process to the data taken at Arecibo with the WAPPs, obtaining a DM value for each time subintegration. These values showed the large DM excesses around orbital phase 0.3, but they also confirmed the variability of the average DM, giving values stable within an orbit but varying from orbit to orbit by as much as 0.01 $\mathrm{pc\,cm^{-3}}$. We therefore note that the DM of the signal from J1023 varies by 0.01 $\mathrm{pc\,cm^{-3}}$ on both short and long timescales.

3.4 Historical observations and current state

Since the unusual nature of J1023 depends on its historical active phase, and since this phase is indicated by a complex sequence of observations, in this section and in Table 3–3 we provide a detailed history of the observations that lead us to believe there was an accretion disc in the system circa 2001.

The source FIRST J102347.67+003841 was detected as a point source in 1998 August in the FIRST VLA 1.4 GHz continuum radio sky survey (Becker et al., 1995). Three observations over the course of a week showed that the flux density of J1023 varied from less than a few mJy to roughly 7 mJy (Bond et al., 2002). The source was also identified with a $V \sim 17.5$ magnitude star having a history of observations back to a 1952 Palomar Observatory Sky Survey observation (Zacharias et al., 2004). In 1999 March, a Sloan Digital Sky Survey filter-photometry observation showed that this star had colours consistent with those of a G star (Szkody et al., 2003).

However, observations at optical wavelengths showed a dramatic change: spectra taken in 2000 May (Bond et al., 2002) were substantially bluer, and had prominent emission lines, leading to classification of J1023 as a cataclysmic variable. Further spectroscopic observations (Szkody et al., 2003; Bond et al., 2002) confirmed this spectrum, and revealed emission lines with double peaks and orbital radial velocity modulation as late as 2001 December. Optical photometry (Bond et al., 2002) showed flickering, varying by $\sim$1 mag on timescales from ten seconds to one day.

Optical observations after 2001 December showed a return to the previous state: a 2002 May observation (Homer et al., 2006) showed a G-like spectrum without emission lines, and photometry taken in late 2002 and early 2003 (Woudt et al., 2004) shows a smooth light curve with a modulation period of 0.198 days, having amplitude $\sim$0.3 mag in V, though with some low-level flickering. Observations in 2004 (Thorstensen and Armstrong, 2005) found a smooth light curve consistent with that of 2003, though with no flickering, and a spectrum matching that of a mid-G star. Radial velocity measurements showed a regular Doppler shift with 0.198 day period and peak velocity $268 \pm 4\,\mathrm{km\,s^{-1}}$ (Thorstensen and Armstrong, 2005). Light curve and colour-variation modelling indicated that the primary was too massive and too luminous to be a white dwarf (given the observed red optical spectrum), so it was suggested that the source was probably an LMXB with a neutron star primary, now in quiescence (Thorstensen and Armstrong, 2005). This idea was supported by *XMM-Newton* observations in 2004, which found a hard power-law X-ray spectrum (Homer et al., 2006).

In order to verify that J1023 is currently in a quiescent state, we took an optical spectroscopic observation 2008 Dec. 23 and 25 with the Goodman high-throughput spectrograph on the Southern Astrophysical Research telescope. Since no flux calibration observation was taken, we produced a spectrum by averaging our seven 900-s exposures and dividing by the best-fit continuum. The resulting spectrum (Figure 3–1) confirms the continued absence of emission lines, and hence quiescent state of the system.

Table 3–3. Published observations of J1023.

Date	Type	Reference	Comment
1992 Nov 7	436 MHz pulsar survey[a]	(Lyne et al., 1998)	No detection ($\lesssim$ 3 mJy)
1998 Aug 3	1.4 GHz continuum[b]	(Bond et al., 2002)	$<$ 1.8 mJy
1998 Aug 8	1.4 GHz continuum[b]	(Bond et al., 2002)	$<$ 3.4 mJy
1998 Aug 10	1.4 GHz continuum[b]	(Bond et al., 2002)	6.56 mJy
1999 Mar 22	u', g', r', i', z' photometry[c]	(Adelman-McCarthy et al., 2008)	G star colour
2000 May 6	Spectroscopy[d]	(Bond et al., 2002)	Blue, emission lines
2000 May 9	Spectroscopy[e]	ESO archive	Blue, emission lines
2000 Nov 21	Fast photometry[f]	(Bond et al., 2002)	Flickering by $\sim$ 1 mag
2000 Nov 22	Fast photometry[f]	(Bond et al., 2002)	Flickering by $\sim$ 1 mag
2000 Dec 22	Fast photometry[f]	(Bond et al., 2002)	Flickering by $\sim$ 1 mag
2000 Dec 23	Fast photometry[f]	(Bond et al., 2002)	Flickering by $\sim$ 1 mag
2000 Dec 24	Fast photometry[f]	(Bond et al., 2002)	Flickering by $\sim$ 1 mag
2001 Feb 1	Spectroscopy[c]	(Adelman-McCarthy et al., 2008)	Blue, emission lines
2001 Dec 10	Spectroscopy[g]	(Szkody et al., 2003)	Blue, double-peaked emission lines
2002 May 11	Spectropolarimetry[h]	(Homer et al., 2006)	G star spectrum, no significant polarization
2002 Dec 30	Fast photometry[i]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Oct 16	8.4 GHz continuum[b]	(Mason and Gray, 2007)	No detection
2003 Jan 28	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Jan 29	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Jan 30	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Jan 31	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering

Table 3–3 (cont'd)

Date	Type	Reference	Comment
2003 Feb 03	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Feb 05	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Feb 21	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2003 Feb 24	Fast photometry[j]	(Woudt et al., 2004)	Smooth, low-amplitude flickering
2004 Feb 16	Spectropolarimetry[k]	(Homer et al., 2006)	G star spectrum, no significant polarization
2004 May 12	X-ray[l]	(Homer et al., 2006)	Hard power-law spectrum
2004 May 12	X-ray[m]	(Homer et al., 2006)	Irregular, 60% modulation
2004 May 12	B photometry[n]	(Homer et al., 2006)	Smooth periodic 40% modulation
2004 May 13	Photometry[o]	(Homer et al., 2006)	Smooth periodic 30% modulation
2004 May 23	Spectroscopy[p]	(Homer et al., 2006)	
2003 Jan 31	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Jan 18	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Jan 19	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Jan 20	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Mar 8	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Mar 9	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Nov 18	Spectroscopy[p]	(Thorstensen and Armstrong, 2005)	
2004 Feb 29	Filter photometry[p]	(Thorstensen and Armstrong, 2005)	
2004 Feb 29	Fast I photometry[p]	(Thorstensen and Armstrong, 2005)	
2004 Mar 1	Fast I photometry[p]	(Thorstensen and Armstrong, 2005)	
2004 May 12	Fast B, V, I photometry[q]	(Thorstensen and Armstrong, 2005)	

3.5 X-ray upper limit

In order to improve on the upper bound of Homer et al. (2006) on J1023's X-ray brightness during its active phase, we searched the archival *Rossi X-ray Timing Explorer* All Sky Monitor data for the known position of J1023. These data cover the energy range 2–10 keV for all of the span 1996–2006 with the exception of brief periods when the source passed close to the Sun. The average annual flux for each year 1996–2008 must be beneath the systematic upper limit of $2.4 \times 10^{-11}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$ ($4.8 \times 10^{33}\,\mathrm{erg\,s^{-1}}(d/1.3\,\mathrm{kpc})^2$). Moreover, since LMXBs are variable sources, we searched individual days as well, obtaining varying upper limits; in particular, the average X-ray flux during Feb 1, 2001, when an optical emission-line spectrum was observed (Szkody et al., 2003), did not exceed $1.2 \times 10^{-10}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$ (corresponding to a luminosity of $2.4 \times 10^{34}\,\mathrm{erg\,s^{-1}}(d/1.3\,\mathrm{kpc})^2$, using our distance estimate from Section 3.6.1).

3.6 Discussion

3.6.1 Inferred parameters

Combining the pulsar parameters we obtained from timing (Table 3–1) with optical radial velocity measurements (Thorstensen and Armstrong, 2005), we can compute a ratio of radial velocities in the orbit around the center of mass; the ratio of these velocities is the reciprocal of the ratio of masses. We find that the pulsar is a factor of 7.1 ± 0.1 more massive than its companion. The values of the individual masses depend only on the inclination angle i of the orbital plane with respect to the plane of the sky. The pulsar's Keplerian mass function $(m_c \sin i)^3/(m_p + m_c)^2$ is $1.1 \times 10^{-3}\,M_\odot$, where m_c and m_p are the masses of the

Table 3–3 (cont'd)

Date	Type	Reference	Comment
2004 May 16	Fast B, V, I photometry[q]	(Thorstensen and Armstrong, 2005)	
2004 May 20	Fast B, V, I photometry[q]	(Thorstensen and Armstrong, 2005)	
···	H_α, l', r' imaging[r]	(Witham et al., 2006)	No H_α excess

Note. — Horizontal lines indicate the beginning and end of the active phase.

[a]The Australia Telescope National Facility Parkes radio telescope

[b]The Very Large Array, operated by the National Radio Astronomy Observatory

[c]The SDSS 2.5 m telescope at the Apache Point Observatory

[d]Lick Observatory 3 m telescope

[e]ESO Faint Object Spectrograph and Camera on the 2.2 m telescope at La Silla

[f]Kitt Peak National Observatory 2.1 m

[g]Double-imaging spectrograph on the Apache Point Observatory 3.5 m telescope

[h]Kitt Peak Bok Telescope, the spectropolarimeter SPOL

[i]The 1.9 m telescope at the Sutherland site of the South African Astronomical Observatory

[j]The 1.0 m telescope at the Sutherland site of the South African Astronomical Observatory

[k]The 6.5 m Multi-Mirror Telescope

[l]XMM-Newton EPIC-pn

[m]XMM-Newton EPIC-MOS1/2

[n]XMM-Newton OM

[o]US Naval Observatory Flagstaff Station 1 m telescope

[p]The Hiltner 2.4 m telescope

[q]The 1.3 m McGraw-Hill telescope

[r]The 2.5 m Isaac Newton Telescope

companion and the pulsar, respectively. For a neutron star mass in the range $\sim$1.0 to $\sim$3.0 $M_{\odot}$ (Lattimer and Prakash, 2004), i is restricted to the range $\sim 53^{\circ}$ to $\sim 34^{\circ}$, and the companion mass is restricted to $\sim$0.14 to $\sim 0.42 M_{\odot}$.

Stellar models, based on the assumption (reasonable in light of the evidence for a recent accretion flow) that the companion is filling its Roche lobe (see Section 1.2), show that it is possible for the companion star to exhibit the observed orbital modulation of colour, magnitude, and radial velocity, if the primary has an isotropic luminosity of $\sim$2 $L_{\odot}$ (Thorstensen and Armstrong, 2005). These models also imply $i < 55^{\circ}$, which is compatible with the range obtained from theoretical neutron star mass limits. Similar models account for the light curve modulation during quiescence of the low-mass companion to SAX J1808.4−3658 (Deloye et al., 2008), an accreting X-ray pulsar in a two-hour orbit (e.g. Wijnands and van der Klis, 1998; Chakrabarty and Morgan, 1998), and which has been suspected of harbouring an MSP when in quiescence (Burderi et al., 2003; Heinke et al., 2009; Hartman et al., 2008).

The brightness, radial velocities, and presumed Roche-lobe-filling companion of J1023 imply that the distance to the system is $0.9\ \mathrm{kpc}/\sin i$; for the allowed range of i, the distance is between $\sim$ 1.1 and $\sim$ 1.6 kpc. A pulsar mass of $1.4 M_{\odot}$ implies a distance of 1.3 kpc. For comparison, given the observed DM of $14.325\,\mathrm{pc\,cm^{-3}}$ the standard Galactic free-electron density model (Cordes and Lazio, 2002) predicts a distance of 0.6 kpc, consistent with our range given the known limitations of this model, in particular the problems with scale height, relevant for sources out of the Galactic plane.

Ordinarily one can estimate a pulsar's spin-down luminosity from its observed spin-down $\dot{P}$, but for MSPs this is subject to contamination by several factors, including accelerations of J1023 and the solar system barycenter in the Galactic gravitational potential. From the measured optical proper motion μ of J1023 (Table 3–1) we expect a significant positive contribution to $\dot{P} = Pd\mu^2/c \simeq 1.8 \times 10^{-21}(d/1.3\,\mathrm{kpc})$ due to the Shklovskii effect (Shklovskii, 1970). While in some binary systems the accelerations can be constrained independently using orbital period variations, in J1023 we observe very large orbital period variability (binary period derivative $\sim 3 \times 10^{-10}$), almost certainly due to classical tidal torquing from gas motion in the extended envelope of the companion star (see Section 7.6 for more discussion). The near-zero eccentricity of the pulsar's orbit (Table 3–1) also suggests that the system has undergone tidal circularization.

3.6.2 Complex radio behaviour

This picture of a Roche-lobe-filling companion is supported by evidence for the presence of ionized material in the system. Our broad-band radio observations reveal substantial DM variations at certain orbital phases (Figure 3–3 G) as well as smaller DM variations on time scales of minutes throughout the orbit. Moreover, we observe regular eclipses ranging from very brief at 3000 MHz to most of the orbit in our 150 MHz observations with the Westerbork Synthesis Radio Telescope (Figure 3–3 A–F). We also see very brief eclipses at all orbital phases (Figure 3–3 E in particular). Given the established range of i, the line-of-sight between the pulsar and the Earth will not intersect the Roche lobe of the companion at any point in the orbit. If we assume that during the DM variations near orbital

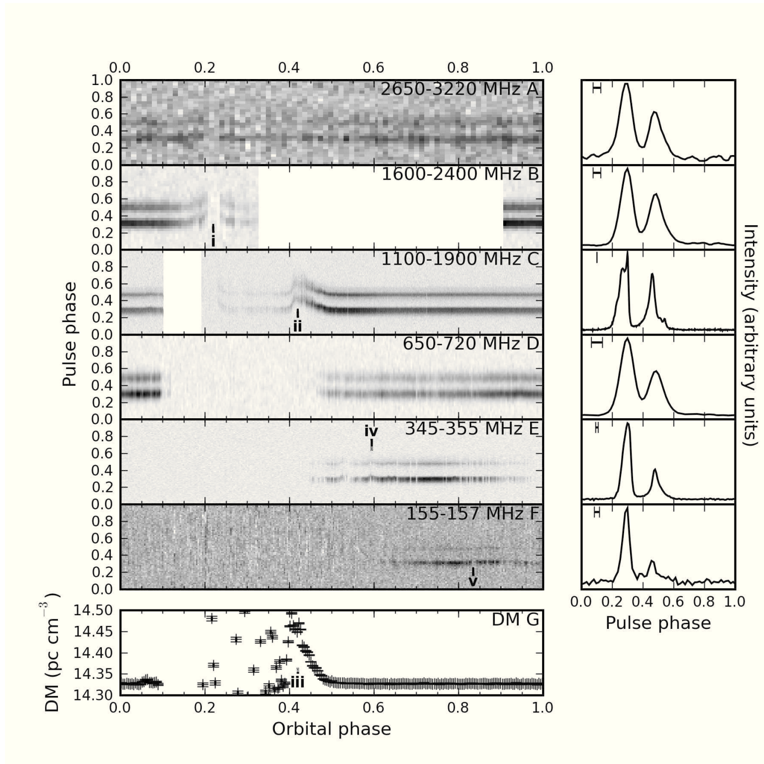

Figure 3–3 Frequency dependence of eclipses, DM variations, and pulse profiles. **(A–F)** show flux density as a function of pulse phase and orbital phase at 3000, 2000, 1600, 700, 350, and 156 MHz, respectively. **(G)** shows DM, as estimated from the 1600 MHz observations (see Section 3.3.2 for details). Orbital phase is defined to be zero at the pulsar's ascending node, so that the companion passes closest to our line of sight at orbital phase 0.25. To the right of **(A)** through **(F)** are pulse profiles at the respective bands; the instrumental smearing time is indicated with a horizontal bar in each panel. The exception is that the profile **(C)** is instead based on a 1410 MHz observation with higher time resolution. Note that the eclipse is nearly absent at 3000 MHz, but is longer at 700 MHz than at 1600 MHz. At low frequencies, random short eclipses, indicated by **(iv)** and **(v)**, are also visible. Pulse phase is as predicted by a phase-coherent timing solution, so that large short-term timing variations are visible as vertical motions of the pulse peak. Note in particular the pulse arrival time variations at eclipse ingress **(i)** and egress **(ii)**; the variation at egress appears to be due to the substantial DM variation **(iii)** visible in **(G)**.

phase 0.4 (Figure 3–3 ii and iii) the obstructing plasma is uniform, has a plasma frequency of 700 MHz (based on the loss of the 700 MHz signal at this phase, presumed to be due to free-free absorption), and contributes the $0.15\,\mathrm{pc\,cm^{-3}}$ excess DM, we predict a thickness of 3×10^4 km (thin compared to the companion size of $\sim 3 \times 10^5$ km). This suggests that a thin but dense layer of material, perhaps a shock due to the pulsar wind meeting either winds from the companion or material overflowing the companion's Roche lobe, crosses the line of sight at this orbital phase.

Similar eclipses and DM variations have been seen in the Galactic "black widow" systems, tight neutron star binary systems in the Galactic field with companions an order of magnitude less massive than that of J1023 (e.g. Fruchter et al., 1988a). Such eclipses have also been seen in several radio MSPs in globular clusters (e.g. D'Amico et al., 2001; Stappers et al., 1996). In both cases they have been attributed to the presence of gas flowing out from the irradiated companion, as is probably the case for J1023.

3.6.3 Accretion disc behaviour during the active phase

For an accreting neutron star the accretion luminosity is proportional to the accretion rate, which is unknown for J1023. But if the matter is to reach the neutron star surface, it must overcome magnetic pressure at the Keplerian corotation radius, so there is a minimum accretion rate and hence a minimum luminosity for a given neutron-star magnetic field (Campana et al., 1998). For a magnetic field equal to our timing-derived upper limit (Table 3–1) this minimum luminosity for J1023 is 1.1×10^{37} erg s^{-1} (using formulae from Campana et al., 1998). Such a high luminosity can be clearly ruled out by our analysis of archival *Rossi X-ray Timing Explorer* all sky monitor data (see Section 3.5 for details): the average annual X-ray luminosity for each year 1996–2008 was less than $4.8 \times 10^{33}\,\mathrm{erg\,s^{-1}}(d/1.3\,\mathrm{kpc})^2$. The variable nature of LMXB X-ray emission cannot explain this low luminosity: on Feb 1, 2001, when an optical emission line spectrum was observed (Szkody et al., 2003), the average flux was less than $2.4 \times 10^{34}\,\mathrm{erg\,s^{-1}}(d/1.3\,\mathrm{kpc})^2$. If accretion occurred during the active phase, the annual upper limit would imply that the magnetic field of the neutron star in J1023 was less than 6×10^6 G, smaller than that of any known MSP.

It therefore seems more likely that infalling matter did not reach the neutron star surface, but instead underwent what is known as "propeller-mode accretion" (Illarionov and Sunyaev, 1975): infalling matter entered the light cylinder but was stopped by magnetic pressure outside the corotation radius, which prevented it from falling further inward. This process also has a minimum accretion rate and minimum luminosity: if the infalling matter does not reach the light cylinder, the radio pulsar mechanism will presumably become active. If this occurs no stable balance exists between outward radiation and wind pressure and ram pressure of infalling material, and the disc will be cleared from the system (Campana et al., 1998). The minimum luminosity for propeller-mode accretion is substantially lower than that for standard accretion: our upper limit on the magnetic field (Table 3–1), derived from timing in quiescence, implies a minimum luminosity of just 2.7×10^{32} erg s^{-1} (using formulae from Campana et al., 1998). Thus it seems likely that during J1023's active phase, the mass transfer rate was high enough for propeller-mode accretion, but not high enough for material to reach the neutron-star surface.

3.6.4 Relation to similar systems

It seems likely that during its active phase, J1023 did not actually accrete material onto its surface, but instead the accretion disc operated in propeller mode. In this scenario, a small drop in the mass transfer rate would clear the disc from the system and return J1023 to its current quiescent state. Overflowing matter from the companion would then encounter a shock near the inner Lagrange point and be carried out of the system (Campana et al., 1998), possibly explaining the observed variable hard power-law X-ray spectrum of J1023 (Homer et al., 2006, and Chapter 4) and the presence of gas in the system. Such shocks have been suggested to explain the variable power-law X-ray emission in 47 Tuc W (Bogdanov et al., 2005a) and PSR J1740−5340 in NGC 6397 (Grindlay et al., 2002), which also have low-mass companions, and have previously been compared with SAX J1808.4−3658, though with no evidence for the recent presence of an accretion disc.

Despite the apparent resemblance of J1023 to 47 Tuc W and PSR J1740−5340, the latter two systems reside in globular clusters, and are likely to have acquired their current companions in exchange interactions after the pulsars were "recycled" (Camilo and Rasio, 2005). J1023 was the first highly recycled MSP found in the field of the Galaxy with both a non-degenerate companion and an orbit that has been circularized through tidal interactions. A globular cluster origin for J1023 is extremely unlikely due to its large distance from the nearest globular cluster as well as from the Galactic bulge. The evidence points to J1023's having been recycled by its current companion, which has not yet completed the transformation to a white dwarf.

The observed transition of J1023 suggests that it is in a bistable state: for certain rates of Roche lobe overflow, if the radio pulsar mechanism is quenched, propeller-mode accretion can occur, but if the radio pulsar mechanism is active, that same mass accretion rate cannot overcome the radiation pressure and no accretion occurs. Should the mass transfer rate of J1023 rise sufficiently, then, it may enter another LMXB phase: a disc will form, the radio emission may be quenched, and the X-ray luminosity may increase dramatically, due to either propeller-mode accretion, with a net spin-down of the pulsar, or even brighter accretion onto the surface, with a net spin-up.[5]

3.7 Summary

J1023 is currently a radio MSP with a nondegenerate companion. It seems that in 2001, it went through an active phase. During this active phase an accretion disc appears to have formed and been hot enough to show bright emission lines; nevertheless it does not appear that matter actually accreted onto the neutron-star surface. Instead it seems likely that the accretion disc was in "propeller mode," in which inflowing material was stopped by the pulsar's magnetic field far enough from the neutron star that the magnetic field at the

[5] In mid-2013 the system M28I was discovered to have gone from radio pulsar to accreting binary back to radio pulsar; see Section 7.3 for further discussion.

inner edge of the disc was rotating more rapidly than the Keplerian motion of the disc material. Coupling between the two would have transferred energy to the inner edge of the disc, expelling material from the system and slowing the neutron star. Since the emission lines have disappeared and the optical behaviour of the system is now consistent with that of a mildly heated companion, we infer that the disc has been cleared from the system, possibly by a particle wind from the pulsar. Nevertheless an increase in the amount of Roche lobe overflow from the companion might overcome this pressure and trigger another active episode, possibly even in the next few years. In the meantime J1023 exhibits complex radio behaviour presumably driven by the presence of ionized material in and near the system, which may shed some light on the physics driving this transient object.

CHAPTER 4
X-ray observations of PSR J1023+0038

This chapter is based on the manuscript "X-ray Variability and Evidence for Pulsations from the Unique Radio Pulsar/X-ray Binary Transition Object FIRST J102347.6+003841," published in The Astrophysical Journal *in 2010 August (Archibald et al., 2010). This paper describes our X-ray observations of J1023 and what they let us infer about the system.*

4.1 Introduction

Prior to the discovery that J1023 contained a millisecond pulsar, and after it had returned to quiescence, Homer et al. (2006) observed J1023 with *XMM-Newton*. They concluded that it had a hard power-law spectrum, variability not necessarily correlated with the orbital phase, and a 0.5–10 keV X-ray luminosity of $2.5 \times 10^{32}\,\mathrm{erg\,s^{-1}}$ ($0.065 L_{\odot}$) at a distance, based on estimates available at the time, of 2 kpc. This supported the suggestion of Thorstensen and Armstrong (2005) that J1023 had a neutron-star primary rather than a white dwarf as had previously been thought.

Not only does the detection of J1023 as a radio millisecond pulsar confirm that it is a neutron-star binary, radio timing provides an ephemeris predicting the rotation of the pulsar. For this reason, we observed J1023 for a longer time and with high time resolution using *XMM-Newton* (see Section 2.5 for a description of this telescope). Our observations confirm the spectral observations of Homer et al.

(2006), show evidence for orbital variability, and show for the first time probable X-ray pulsations at the pulsar period.

4.2 Observations

We obtained an *XMM-Newton* observation (OBSID 0560180801) of J1023 of 33 ks, corresponding to 1.94 binary orbits, on 2008 Nov 26. We operated the PN camera in fast timing mode and the MOS cameras in "prime full window" (imaging) mode, and we used the thin filter on all the EPIC cameras. The data were free of soft proton flares, so we were able to use the entire exposure time. We also retrieved the data set (OBSID 0203050201) used by Homer et al. (2006) from the *XMM-Newton* archive; it consists of a single 16-ks exposure (0.92 binary orbits) obtained on 2004 May 12, with all EPIC instruments using "prime full window" mode and the thin filter. This data set is also free of soft proton flares, but since all three EPIC instruments were operated in imaging mode, in this case high-resolution timing data are not available.

The Optical Monitor was operated in both observations with the B filter and a mode with 10-s time resolution. Pipeline processing showed optical modulation consistent with that reported by Thorstensen and Armstrong (2005), so we did no further analysis of the optical data.

We analyzed the X-ray data with the *XMM-Newton* Science Analysis Software[1] version 9.0.0[2] and `xspec` version 12.5.0ac[3] .

We reprocessed the MOS and PN data with the `emchain` and `epchain` pipelines, respectively, then extracted photons meeting the recommended pattern, pulse invariant, and flag criteria. In particular, this restricted MOS photons to the 0.2–12 keV energy range, PN imaging photons to the 0.13–15 keV energy range, and PN timing photons to the 0.25–15 keV energy range. For all analyses except the spectral fitting, we restricted the imaging photons from either type of camera to 0.2–10 keV for consistency, since this involved discarding only a handful of photons.

The default pipeline processing showed no evidence for extended emission from J1023, so we used the SAS tool `eregionanalyse` to select extraction regions to optimize the ratio of source photons to the square root of source plus background photons. These regions were circles of radius 35″ (enclosing about 90% and 85% of the 1.5 keV flux on the MOS and PN cameras, respectively) in imaging observations. For the timing-mode PN observation the one-dimensional averaging of the energy-dependent point-spread function makes region selection

[1] The *XMM-Newton* SAS is developed and maintained by the Science Operations Centre at the European Space Astronomy Centre and the Survey Science Centre at the University of Leicester.

[2] http://xmm.esac.esa.int/sas/

[3] http://heasarc.nasa.gov/docs/xanadu/xspec/

more complicated, so we adjusted the extraction region by hand to optimize detection significance (see below), selecting a band $12''$ wide (enclosing about 60% of the 1.5 keV flux). We selected background regions from nearby large circles in the imaging-mode observations, and from a band $50''$ wide, to one side of the source, in the timing-mode PN data. All together, this yielded approximately 5700 background-subtracted source photons in imaging mode and approximately 2400 background-subtracted source photons in timing mode.

We computed photon arrival times using the SAS barycentering tool, `barycen`, the precise optical position $10^{\mathrm{h}}23^{\mathrm{m}}47.67^{\mathrm{s}}$ $+0°38'41.2''$ (J2000) given in the NOMAD catalog (Zacharias et al., 2004), and the DE405 solar system ephemeris. To produce pulsar rotational phases during the 2008 observation, we used the tool `tempo`[4] and the ephemeris given in Table 3–1, which is based on radio pulsar timing data bracketing the X-ray observation epoch. The extremely sensitive phase-coherent timing procedure used to generate this table revealed minuscule variations in the orbital period of J1023. These are modelled in the published ephemeris by expressing the orbital period as a quadratic polynomial obtained by a fit to several months of 2008 timing data. Extrapolating this polynomial back to the 2004 *XMM-Newton* observation results in an unreasonable and poorly constrained orbital period at the time of ~ -6 days, a clear indication that the orbital period variations are not actually given by such a quadratic polynomial. While it is necessary to model the orbital period derivatives to obtain good quality

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

Table 4–1. Spectral fits to all EPIC image data

	power-law	neutron star atmosphere plus power law
N_{H} (cm^{-2})	$< 5 \times 10^{19}$	$< 1 \times 10^{20}$
photon index	1.26(4)	0.99(11)
kT (keV)	$\cdots$	0.12(2)
Thermal emission radius (km)[a]	$\cdots$	$0.7^{+0.5}_{-0.1}$
0.5-10 keV unabsorbed flux ($\mathrm{erg\,cm^{-2}\,s^{-1}}$)	$4.66(17) \times 10^{-13}$	$4.9(3) \times 10^{-13}$
0.5-10 keV luminosity[a]($\mathrm{erg\,s^{-1}}$)	$9.4(4) \times 10^{31}$	$9.9(5) \times 10^{31}$
Thermal fraction (0.5–10 keV luminosity)	$\cdots$	0.06(2)
χ^2/degrees of freedom	230/214	208/212
Null hypothesis probability	0.21	0.56

[a]Assuming a distance of 1.3 kpc; see Section 3.6.1.

pulsar phase predictions, it is possible to obtain adequate *orbital* phase predictions by using a simple model with constant orbital period of 0.19809620 days and (pulsar) ascending node of MJD 54801.970652993, based on the model from Table 3–1. This model gives adequate orbital phase predictions back to 2004, matching the orbital phases observed in Thorstensen and Armstrong (2005) to within 1% of a period as well as matching the ephemeris given in Table 3–1 exactly on 2008 Dec 1. We therefore used this model to evaluate the orbital phase of each barycentered photon in both our observations.

4.3 Results

4.3.1 Spectral fitting

As described in Section 2.5, all the modes of the X-ray cameras can provide spectral data, but the timing mode does not have good spectral calibration. We therefore used `xspec` to fit each of several models to all EPIC imaging data sets

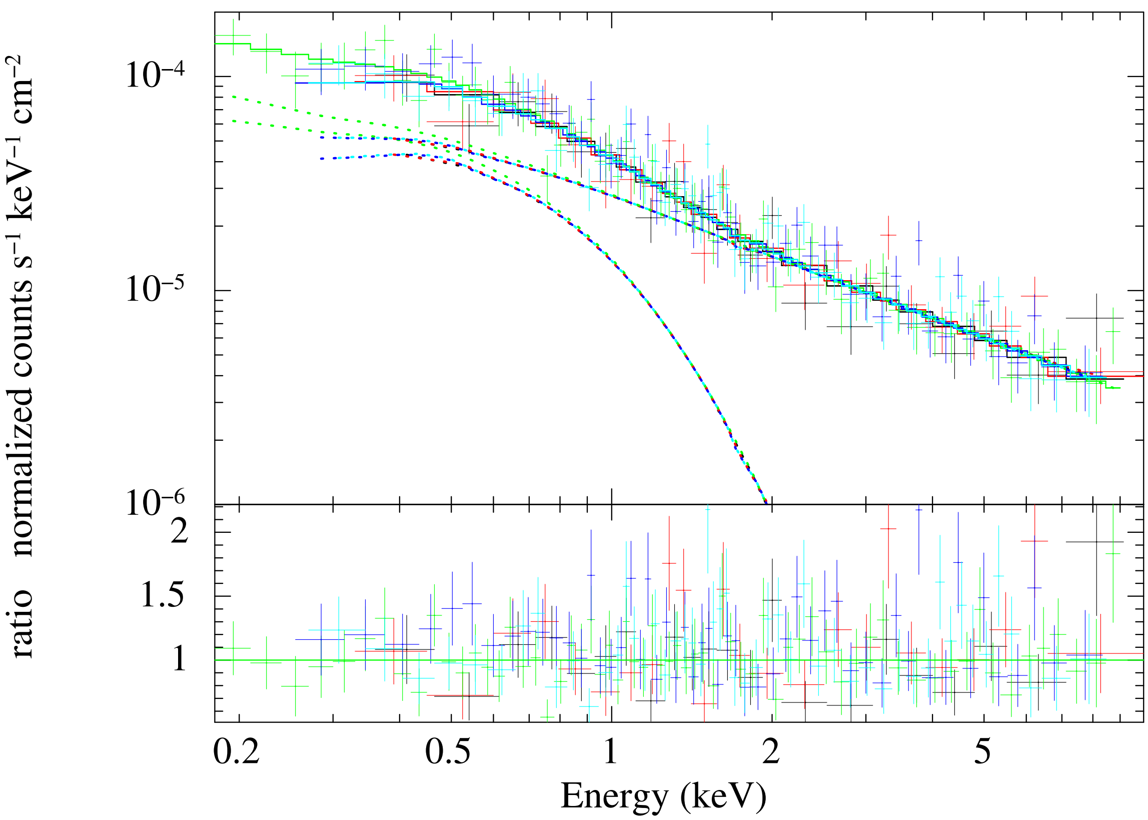

Figure 4–1 Absorbed neutron-star atmosphere plus power-law spectral fit to all EPIC imaging-mode photons from both observations. Each data set is plotted in a different colour: black is the 2004 MOS1, red the 2004 MOS2, green the 2004 PN, blue the 2008 MOS1, and cyan the 2008 MOS2 data set. The vertical axis is normalized by the effective area, and the expected responses from the two additive components are overplotted separately (the harder component is the power law). The lower panel is the ratio of model predicted flux to observed flux in each group of detector channels.

from both epochs simultaneously. To verify that the data sets were compatible, we also fit different absorbed power laws to the two observations, obtaining compatible spectral indices and normalizations. For the MOS cameras we restricted photons to the recommended energy range, namely 0.2–10 keV, and for the PN camera we used the recommended lower limit of 0.13 keV but used the upper limit of 10 keV as there were not enough photons above this to provide any constraint. We grouped the photons to obtain at least 20 counts per spectral bin so as to have approximately Gaussian statistics, allowing least-squares fitting to be used and simplifying error calculations (the loss of spectral resolution was unimportant for the broad thermal and power-law models we used).

We obtained an adequate fit to all imaging data sets with a power-law model (coming naturally from many non-thermal electron distributions and emission mechanisms) including photoelectric absorption by the interstellar medium. We did not obtain adequate fits with an absorbed single black body ($\chi^2 = 1203.7$ with 213 degrees of freedom, for a null hypothesis probability of $\sim 10^{-138}$) or neutron-star atmosphere model ($\chi^2 = 1007.1$ with 213 degrees of freedom, for a null hypothesis probability of $\sim 10^{-103}$).

The simple power-law model, with a photon index of 1.26(4), already has a null hypothesis probability of 0.21, which is satisfactory from a statistical point of view. Although the statistics do not call for a thermal component, we may expect one on physical grounds (for example from the surface of the neutron star). We therefore also consider an absorbed neutron-star atmosphere plus power-law model to fit these data, as shown in Figure 4–1. This model, specified in `xspec` as `phabs(nsa+pow)`, is based on Zavlin et al. (1996). During fitting we froze the neutron-star mass at $1.4 M_\odot$ and the radius at 10 km for the purposes of redshift and light bending; the smaller effective thermal emission radius suggests that the emission comes from a small "hot spot" on the surface. In light of the fact that the estimated $B < 10^8$ G (Chapter 3), we also selected a model in which the magnetic field has negligible effects on the atmosphere, i.e. $\lesssim 10^9$ G.

These two models are summarized in Table 4–1. Uncertainties given are 90% intervals returned by `xspec`'s `error` command; luminosities are obtained by using the `cflux` model component to estimate unabsorbed and thermal fluxes in the 0.5–10 keV range and then using an assumed distance of 1.3 kpc (see Section 3.6.1) to determine a luminosity. "Thermal fraction" is the fraction of this luminosity due to the thermal component of the spectrum, if any.

In both cases, fitting for photoelectric absorption in the model gave a very low upper bound on the neutral hydrogen column density, as noted in Homer et al. (2006). Such a low value is to be expected since the entire Galactic column density in this direction is estimated[5] to be only $1.9 \times 10^{20}\,\mathrm{cm}^{-2}$ (based on Kalberla et al., 2005).

Although both models are statistically adequate, the Akaike information criterion (Akaike, 1974) suggests one should prefer the model containing a thermal component. The F test gives a null probability of 2×10^{-5} that the more complex

[5] Using the HEASARC online calculator, `http://heasarc.nasa.gov/cgi-bin/Tools/w3nh/w3nh.pl`

model would produce such a large improvement in fit due to chance if the simple power-law model were correct. While there are concerns with using the F test in such a situation (Protassov et al. 2002, but see also Stewart 2009) it appears that the statistics do somewhat favour a model with a thermal component, although a purely non-thermal model cannot be excluded.

4.3.2 Orbital variability

To test for orbital variability, we selected 0.2–12 keV source photons from the imaging-mode data sets in both our observations and those of Homer et al. (2006) and reduced them to the solar-system barycenter. When combining data from multiple instruments, to take into account the different orbital coverages and sensitivities, we weighted bin values so that all average count rates match that observed in the MOS1 camera in the same observation. For the hardness ratio analysis, we omitted the PN data so that their different energy response and incomplete orbital coverage would not skew the results.

A weighted histogram of these counts is shown in Fig. 4–2, along with a best-fit sum of sinusoids based on a finely binned profile. We selected four sinusoids as this appears to give a good representation of the curve and a time resolution similar to the histogram. A Kuiper test (Paltani, 2004) gives a probability of 1.3×10^{-19} (9.0σ) that photons drawn from a uniform distribution in orbital phase would be this non-uniform; the reduced χ^2 for a constant fit to the histogram is 11.3, with 11 degrees of freedom, also strongly indicating variability. While the evidence for variability is strong, with only $\sim$3 orbits covered by the two observations it is not certain that this variability is linked to orbital phase,

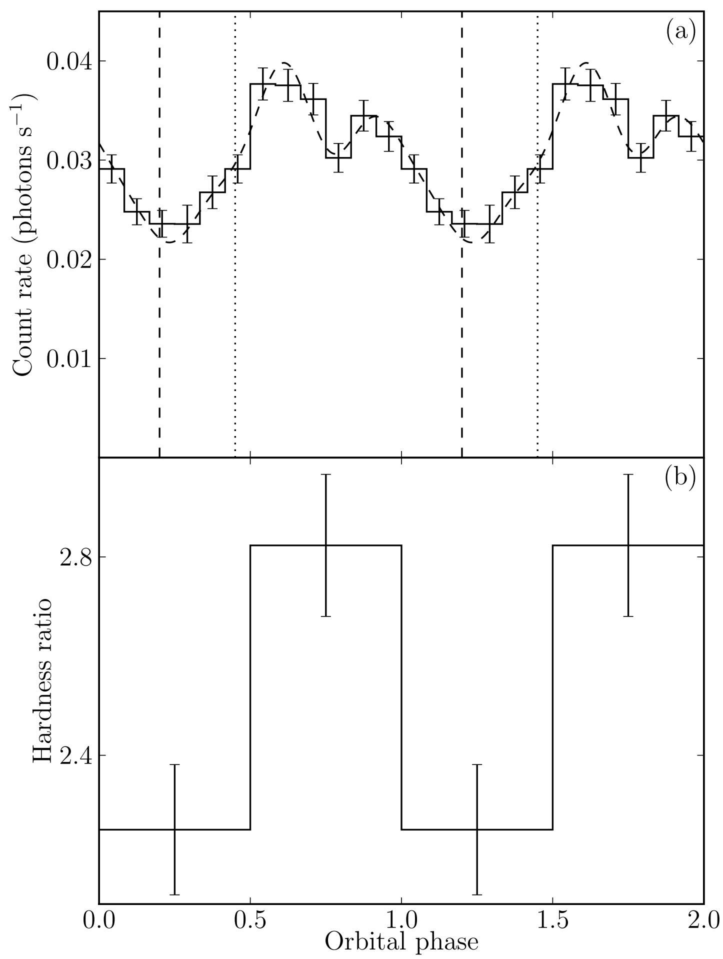

Figure 4–2 Panel (a): Average MOS-equivalent count rate in the 0.2–12 keV energy range as a function of orbital phase. These data are a weighted average of all the imaging-mode data. Dashed and dotted vertical lines indicate the beginning and end of the 1.4 GHz radio eclipse, respectively. The dashed curve is a best-fit combination of four sinusoids. Panel (b): number of MOS photons in the 1–12 keV range divided by number of MOS photons in the 0.2–1 keV range, also as a function of orbital phase. Both panels show two cycles for clarity. The companion's closest approach to our line of sight to the pulsar happens at orbital phase 0.25; that the radio eclipse is off-center may be due to eclipsing material being far out of the orbital plane, as shown in Figure 4–5.

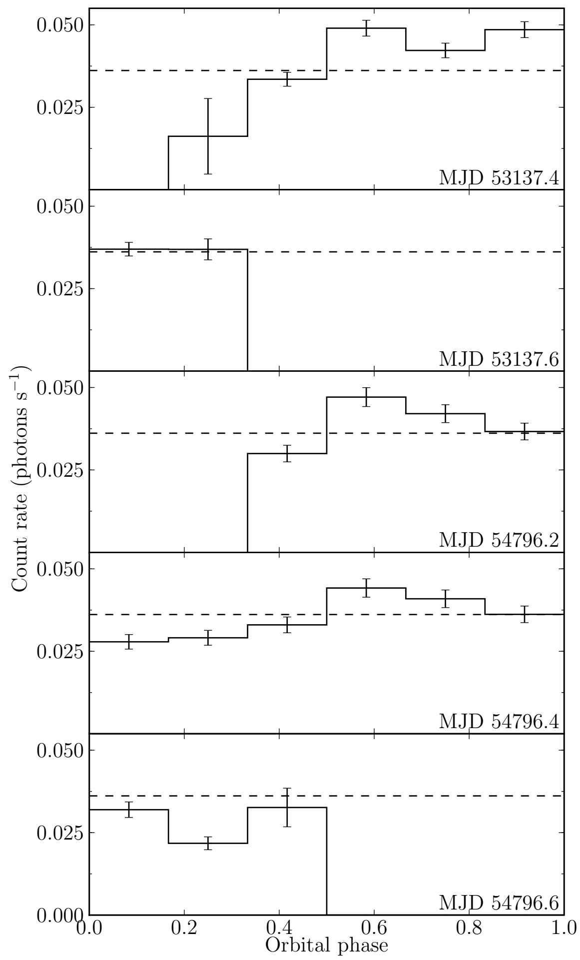

Figure 4–3 Average count rate in the 0.2–12 keV energy range as a function of orbital phase, for each orbit independently; the starting MJD for each orbit is indicated. The vertical axis is MOS-equivalent count rate (based on MOS1/2 and imaging-mode PN data, scaled to match the MOS1 count rate), and the dashed horizontal line indicates the average count rate over all orbital phases and data sets. Where no bar is plotted, no data are available.

although the fact that the minimum in the X-ray light curve occurs near orbital phase 0.25, when the companion passes closest to our line of sight to the pulsar, suggests a link. To test this we plotted the photon arrival rates for each orbital period separately. Figure 4–3 shows that the flux during each individual orbit appears to be lower during phase 0–0.5 than during phase 0.5–1, which suggests that the variability is indeed orbital. Note that Homer et al. (2006) detected variability, but having only limited orbital coverage, could not determine whether it was orbital. To test for spectral changes, given the paucity of photons, we computed a hardness ratio (Fig. 4–2b), dividing the number of photons harder than 1 keV by the number of photons softer than 1 keV and comparing the eclipse versus non-eclipse regions (for this purpose we defined the "eclipse" region to be phases 0–0.5). The reduced χ^2 for a fit of these values to a constant hardness ratio is 8.76 for 1 degree of freedom, and the probability of such a reduced χ^2 arising if the hardness ratio were constant is 3.1×10^{-3} (2.7σ). Thus we see marginally significant softening in the eclipse region.[6]

4.3.3 Pulsations at the pulsar period

To test for pulsations, we extracted source and background photons from the PN camera using the energy range 0.25–2.5 keV, selected to give the most significant detection. We barycentered the photon arrival times and used the

[6] Bogdanov et al. (2011) use a longer observation with the *Chandra X-ray Observatory* to observe J1023; they find evidence for orbital varibility in brightness but not in hardness, and they fit a model of an eclipsing intrabinary shock. See Section 7.3 for more discussion.

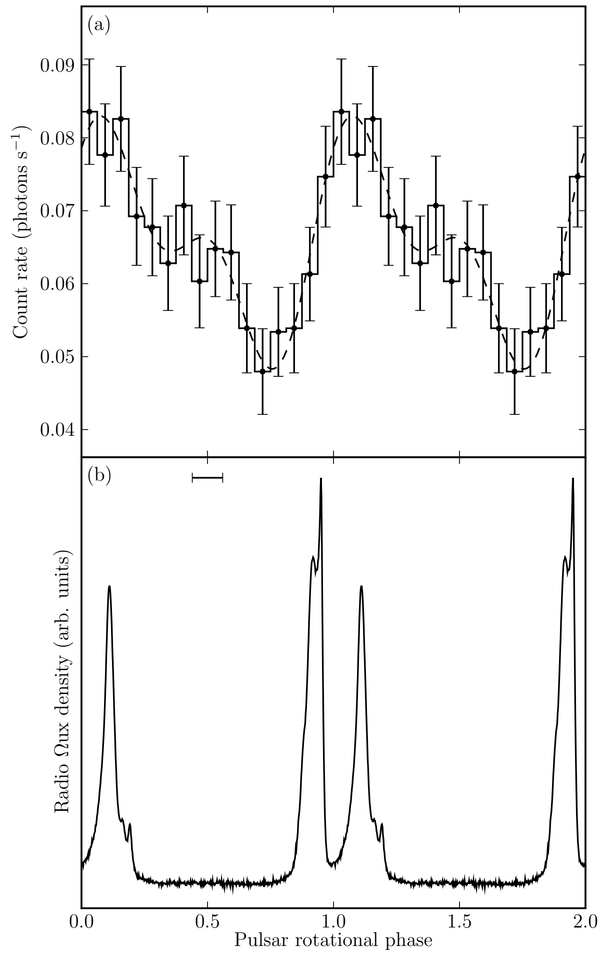

Figure 4–4 Panel (a) shows a background-subtracted light curve based on 0.25–2.5 keV PN photons folded according to the radio ephemeris; the sinusoid (drawn with a dotted line) is a two-component sinusoid fit to the unbinned photon arrival times. Panel (b) shows a 1400 MHz profile obtained with Arecibo (and shown in Chapter 3). Both panels show two cycles for clarity. The uncertainty in relative alignment, due primarily to the absolute timing uncertainty in the *XMM-Newton* data, is indicated by the horizontal error bar. The 20-μs uncertainty relative to the radio data due to dispersion measure uncertainty is relatively unimportant.

program `tempo`[7] and the contemporaneous radio ephemeris given in Table 3–1 to assign each photon a rotational phase. We then tested these photons for uniform distribution in phase. The Kuiper test (Paltani, 2004) gave a (single-trial) null hypothesis probability of 3.7×10^{-6} (4.5σ) and the H test (de Jager et al., 1989) gave a (single-trial) null hypothesis probability of 2.4×10^{-6} (4.6σ), using an optimal number of sinusoids (two). We confirmed that no significant pulsations were detected with the background photons or with an incorrect ephemeris. We also verified that the period predicted by the ephemeris is very close to the period at which the significance peaks (holding all other ephemeris parameters fixed).

To estimate the degree to which the X-ray flux is modulated at the pulse period,[8] one could simply take the lowest bin in the histogram as the background level and compute the fraction of photons above it; this yields a pulsed fraction of 0.27(9). However, this method is subject to large uncertainties, dependence on binning, and a large statistical upward bias due to the fact that we selected the bin in which signal plus noise is lowest, rather than that in which the (unknown) signal is lowest. To avoid these problems, we define a root-mean-squared pulsed flux by fitting a model $F(x)$ with two sinusoids:

$$f_{\mathrm{RMS}} = \sqrt{\int_0^1 (F(x) - \bar{F})^2 \mathrm{d}x},$$

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

[8] Our very limited orbital coverage limits the utility of such an analysis for orbital variability; Bogdanov et al. (2011) are able to analyze the orbital variability of J1023 in more detail.

where $\bar{F}$ is the model mean flux. This is in some sense a degree of modulation; if the signal consists of a constant background plus a variable emission process that drops to zero, this is not directly measuring the fraction of emission due to the variable process. (There is a conversion factor that depends on the exact shape of the pulse profile.) However, this quantity can be computed with substantially less uncertainty and bias. Computationally, we estimate the root-mean-squared amplitude in the Fourier domain, based on the amplitudes of the two complex Fourier coefficients (since higher-order Fourier coefficients are dominated by noise); since noise always contributes a positive power to Fourier coefficients, to reduce the bias we subtract the expected contribution of noise from the squared amplitude of each coefficient before taking the square root. We then convert this pulsed flux value to a pulsed fraction by dividing by the total background-subtracted flux from J1023.

For the energy range 0.25–2.5 keV we estimate a root-mean-squared pulsed fraction of 0.11(2). In the two subbands 0.25–0.6 keV and 0.6–2.5 keV, we find root mean squared pulsed fractions of 0.17(5) and 0.14(3), respectively, with profiles that are broadly similar and in phase. Above 2.5 keV we find no evidence for pulsations, though a 3σ upper limit on the pulsed fraction is only 0.20.

4.4 Discussion

To summarize our results for J1023, we observed an X-ray spectrum dominated by a hard power-law component, possibly with a small, soft, thermal contribution. The emission appears to be modulated at the 0.198-day orbital period, with substantial dips as the companion passes near the line of sight; these dips are accompanied by a possible softening in the spectrum. We also found that the X-ray emission is very likely modulated at the 1.69-ms rotational period of the radio pulsar.

4.4.1 Nature of the X-ray emission

In a review of X-ray emission from millisecond pulsars, Zavlin (2007) describes three primary sources of X-ray emission: emission from an intrabinary shock, emission from the neutron star itself, which is some combination of thermal and magnetospheric emission, and emission from a pulsar wind nebula outside the binary system. As all these mechanisms may be operating in J1023, we next consider their contributions, if any, to the observed X-rays from this system.

Emission From an Intrabinary Shock

J1023 likely has a strong pulsar wind, since the optical data suggest that the companion is being heated by a luminosity of $\sim 2L_{\odot}$ from the pulsar (Thorstensen and Armstrong, 2005). This is much greater than the X-ray luminosity of the system ($\sim 0.03L_{\odot}$), but well below the upper limit ($\sim 80L_{\odot}$) on the spin-down luminosity of the pulsar (Chapter 3). If material is leaving the companion, either through Roche-lobe overflow or a stellar wind, we should expect an intrabinary shock, where this flow meets the pulsar wind. Such a shock could readily produce power-law X-ray emission (Arons and Tavani, 1993). If localized, it could easily account for the orbital modulation we observe in the X-ray emission. If the emission were due to Roche-lobe-overflowing material meeting the pulsar wind we might expect it to be localized at or near L1. As discussed in Chapter 3, with only the assumption that the pulsar mass is $1.4M_{\odot}$ and that the companion fills

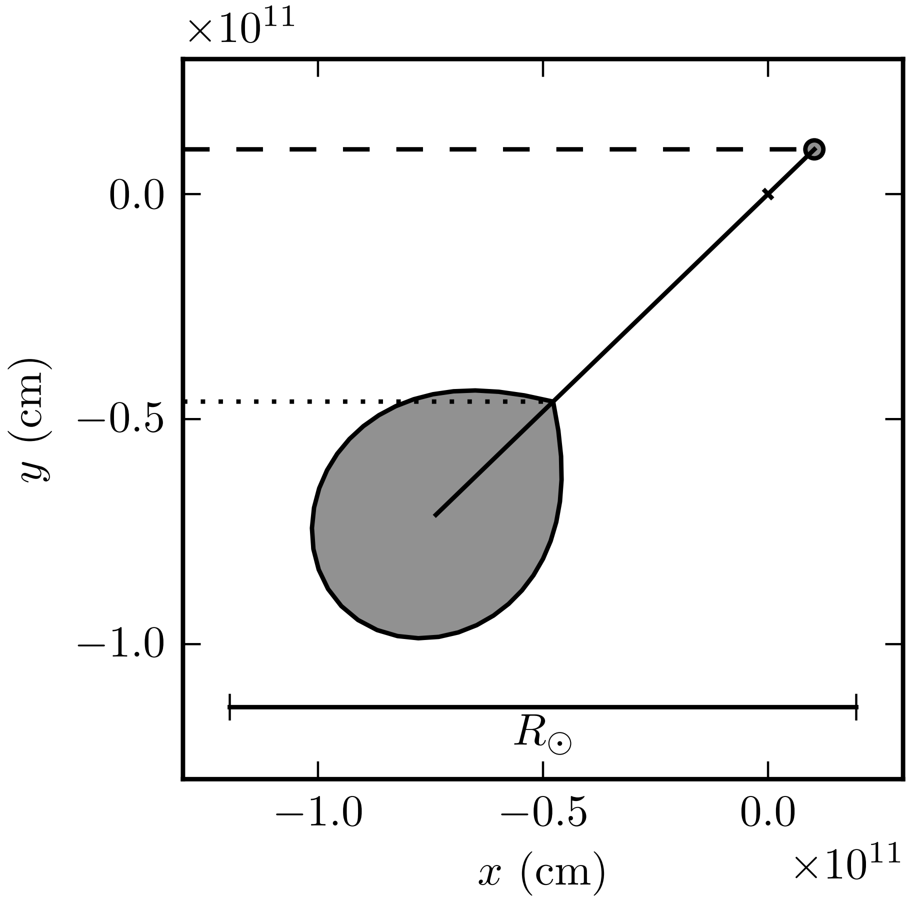

Figure 4–5 J1023 system geometry at phase 0.25 as seen from within the plane of the sky. The system orbit is in and out of the page, the dashed line is our line of sight to the pulsar, and the dotted line is our line of sight to L1, which is eclipsed by the Roche-lobe-filling companion. The pulsar is indicated by a not-to-scale circle, but the companion is drawn to scale. The tick mark on the line connecting pulsar and companion marks the center of mass of the system. This system geometry assumes a pulsar mass of $1.4M_{\odot}$; a higher mass would mean a larger system seen more nearly face-on, but the range of plausible angles is constrained to ∼53°–34° by the narrow range of plausible pulsar masses 1.0–3.0$M_{\odot}$ (Chapter 3). Note that at this orbital phase the X-rays are near their minimum, and low radio frequencies from the pulsar are eclipsed, indicating the presence of ionized material well out of the orbital plane.

its Roche lobe, we can infer the system geometry is that illustrated in Figure 4–5. Then the L1 point itself is eclipsed by the companion for 0.32 of the orbit, centered on orbital phase 0.25. We would expect spectral softening during this period, since the relatively hard power-law emission is blocked but the, presumably softer, emission from the neutron star is not. A larger X-ray-emitting region, due either to material streaming away from the L1 point or to a stellar wind shock, would result in a broader, shallower, and potentially asymmetric dip in the X-ray light curve.

The possibility of a wind from the companion is interesting, since it is not clear how Roche-lobe overflow would provide the ionized material causing the radio eclipses, which occur well above the orbital plane, roughly centered on the companion's closest approach to our line of sight. In Chapter 3 we reported DM changes of $\sim 0.15\,\mathrm{pc\,cm^{-3}}$ above a cutoff frequency of $\sim 1\,\mathrm{GHz}$ at radio eclipse ingress and egress. Assuming that this is the plasma frequency implies an electron density of $\sim 10^{10}\,\mathrm{cm^{-3}}$; combined with the excess DM measurement, this implies a layer of thickness $\sim 4 \times 10^{7}\,\mathrm{cm}$, relatively thin compared to the orbital separation of $1.2 \times 10^{11}\,\mathrm{cm}$. This thin sheet of material could in principle arise as the shock front where the pulsar wind meets a wind from the companion, but it is difficult to explain the companion's wind being strong enough to support such a shock: given the line of sight geometry of the system, the material causing the radio eclipses must be about as far from the companion as it is from the pulsar. If we assume that the highly relativistic pulsar wind provides the $\sim 2L_{\odot}$ required by the companion heating models of Thorstensen and Armstrong (2005), it seems difficult to explain how the companion could have a wind of comparable pressure anywhere along the line of sight. Thus the radio eclipses remain a mystery, and Roche-lobe overflow seems a more likely explanation for the origin of the intrabinary shock.

Emission From the Neutron Star

Some MSPs, binary or isolated, have X-ray emission modulated at the rotational period. This emission is some combination of magnetospheric emission, from high-energy particles in the magnetosphere, and thermal emission, from polar caps heated by bombardment by high-energy particles moving along the open field lines. Observationally, Zavlin (2007) draws the distinction that magnetospherically dominated emission has high pulsed fractions, narrow peaks, and power-law spectra, while polar-cap-dominated emission has lower pulsed fractions, broader peaks, and thermal spectra. Zavlin also suggests that pulsars with spin-down luminosities $\gtrsim 10^{35}\,\mathrm{erg\,s^{-1}}$ tend to be magnetospherically dominated while those with spin-down luminosities $\lesssim 10^{34}\,\mathrm{erg\,s^{-1}}$ tend to be thermally dominated, but since few examples are known, this classification is tentative, and the spin-down luminosity from J1023 is only known to be $< 3 \times 10^{35}\,\mathrm{erg\,s^{-1}}$ (Chapter 3). In any case, classification of the pulsations we observe in J1023 is difficult.

In systems in which thermal emission from the neutron star provides all the detectable X-rays, detailed modelling indicates that the pulsed fraction should normally be $\lesssim 50\%$ (that is, pulsed emission must be accompanied by roughly equal or greater unpulsed emission), since light bending makes the large polar caps visible from almost all angles, though for certain combinations of parameters higher pulsed fractions can occur (Bogdanov et al., 2008). If the pulsations are purely thermal, then our spectral upper limit on the thermal fraction of 0.06(2) is difficult to reconcile with the fact that we observe a fractional modulation of 0.11(2) at the pulsar period. Thus it seems likely that at least some magnetospheric emission is present, as it has a nonthermal spectrum and can more readily have higher pulsed fractions. The poor signal-to-noise in our observations, due to the high background (from the system itself as well as instrumental) and the scarce photons, makes it impractical to determine the sharpness of the pulse profile or the hardness of the pulsations. In any case the pulsations, if real, are a clear sign of X-ray emission from the pulsar itself.

Emission From a Pulsar Wind Nebula

Nebular emission from pulsars arises when the particle wind driven by the pulsar's magnetospheric activity flows out of the pulsar's immediate neighbourhood and meets the surrounding medium; for reviews see Gaensler and Slane (2006) or Kaspi et al. (2006). Pulsars like J1023 whose wind is confined by the ambient interstellar medium exhibit nebular emission that generally takes a cometary form, with an arc-like bow shock preceding the pulsar and a "trail" of ejected material streaming back along the pulsar's track.

The angular size of bow shock emission depends on the pulsar's spin-down luminosity, its proper motion, and the local interstellar medium density. Kargaltsev and Pavlov (2008) give a formula for predicting the "stand-off" angle of the X-ray bow shock from the pulsar:

$$\theta = 5.4'' \, n_{0.1}^{-1/2} \mu_{19}^{-1} \dot{E}_{35}^{1/2} D_{1.3}^{-2},$$

where $n_{0.1}$ is the mean density of the interstellar medium divided by $0.1\,\mathrm{cm}^{-3}$, μ_{19} is the proper motion divided by $19\,\mathrm{mas\,yr}^{-1}$, $\dot{E}_{35}$ is the spin-down luminosity divided by $10^{35}\,\mathrm{erg\,s}^{-1}$, and $D_{1.3}$ is the distance to J1023 divided by 1.3 kpc. We have supplied best-guess parameters for J1023, so that absent further information, we might expect a stand-off distance of $\sim 5''$ for the X-ray bow shock. Thus although nebular emission is not resolved in our *XMM-Newton* images, higher-resolution images with the *Chandra X-ray Observatory*, or in the $H\alpha$ or radio bands, might yet resolve a bow shock.[9]

Kargaltsev and Pavlov (2008) find that the ratio of neutron-star emission to nebular emission generally lies between $\sim$0.1 and 10, which suggests that nebular emission may produce some of the X-rays we observe. Moreover, Kargaltsev and Pavlov find that nebular emission from bow shocks has a power-law spectrum with photon indices $1 \lesssim \Gamma \lesssim 2$ (Kargaltsev and Pavlov, 2008), consistent with what we observe from J1023, though the latter is somewhat harder than most pulsar wind nebulae. The orbital variability, on the other hand, argues against that component of the emission arising from an extrabinary pulsar wind nebula.

4.4.2 Comparison to similar systems

In quiescence, most LMXBs have spectra that are dominated by thermal emission from the neutron star. This emission is typically consistent with a hydrogen atmosphere model with effective radius $\sim$10 km (Brown et al., 1998).

[9] In fact, Bogdanov et al. (2011) did observe J1023 with the *Chandra X-ray Observatory* and found no evidence for an X-ray bow shock.

Some LMXBs in quiescence also have a small power-law component of photon index $1 \lesssim \Gamma \lesssim 2$ in their spectrum (Campana et al., 1998), while a few quiescent LMXBs have spectra completely dominated by this non-thermal component. For example, the quiescent LMXB SAX J1808.4−3658 has a spectrum that is a fairly hard power law with no detectable thermal emission (Heinke et al., 2009). J1023 resembles this latter category, although if a thermal component is present its effective radius appears to be substantially less than 10 km and the power-law spectral index is fairly hard. On the other hand, J1023's radio and possible X-ray pulsations distinguish it from all known quiescent LMXBs.

SAX J1808.4−3658 is of particular interest because it is the prototype and best-studied example of a class of LMXBs from which millisecond X-ray pulsations have been detected during active phases. Many authors have suggested that SAX J1808.4−3658 should turn on as a radio pulsar during quiescence (e.g. Wijnands and van der Klis, 1998; Chakrabarty and Morgan, 1998; Campana et al., 2004), but no radio pulsations have been detected in spite of thorough searches (e.g. Burgay et al., 2003; Iacolina et al., 2009). The non-detection of radio pulsations could be because SAX J1808.4−3658 is shrouded by previously ejected ionized material, or it could simply be an issue of unfortunate beaming. In any case, it is natural to compare SAX J1808.4−3658 to J1023.

During its active phases, SAX J1808.4−3658 reaches a sustained luminosity of $\sim 4 \times 10^{36}\,\mathrm{erg\,s^{-1}}$ (Galloway and Cumming, 2006), in sharp contrast to the upper limit of $2 \times 10^{34}\,\mathrm{erg\,s^{-1}}$, available for J1023 during its disc phase (Chapter 3). In its quiescent state, SAX J1808.4−3658 shows X-ray emission dominated by a hard power law ($\Gamma = 1.74(11)$), with a small ($\lesssim 0.1$) fraction of thermal emission, and a total X-ray luminosity of $7.9(7) \times 10^{31}\,\mathrm{erg\,s^{-1}}$ (Heinke et al., 2009). This is somewhat similar to what we observe in J1023, which has a total luminosity of $9.9(5) \times 10^{31}\,\mathrm{erg\,s^{-1}}$, although the spectral index of 1.26(4) is somewhat harder than that seen from SAX J1808.4−3658. If J1023 has a thermal component, its power law is substantially harder, with a spectral index of 0.99(11). The thermal component would have an effective radius of $0.7^{+0.5}_{-0.1}\,\mathrm{km}$ and supply 0.06(2) of the total luminosity. While the spectral index seen in J1023 is harder than that seen from SAX J1808.4−3658, it is in line with the hard spectral indices seen from other X-ray-detected radio MSPs (for example 47 Tuc W, though PSR J1740-530 is softer; see below). Since SAX J1808.4−3658 is known to cool quickly, even thermal X-ray pulsations in quiescence would be evidence of magnetospheric activity heating its polar caps. Campana et al. (2002) searched for such X-ray pulsations from SAX J1808.4−3658 in quiescence, but were unable to conduct meaningful searches due to its faintness; SAX J1808.4−3658 is roughly 3.5 kpc away (Galloway and Cumming, 2006) compared to the roughly 1.3 kpc we assume for J1023 (Chapter 3).

The similarity of the X-ray properties of SAX J1808.4−3658 in quiescence to those of J1023 suggests that SAX J1808.4−3658 may indeed harbour a radio MSP in quiescence; conversely, the same similarity suggests that J1023 may undergo episodes of active accretion.

Among radio pulsars, the two best-studied examples which closely resemble J1023 are 47 Tuc W (Bogdanov et al., 2005b) and PSR J1740−5430 (D'Amico et al., 2001). Both pulsars exhibit large, variable radio-frequency eclipses, and both are thought to have somewhat massive unevolved companions ($\gtrsim 0.13 M_\odot$ and 0.19–$0.8 M_\odot$, respectively). For comparison, J1023's companion is thought to be $\sim 0.2 M_\odot$, and shows a spectrum characteristic of a main-sequence star (Thorstensen and Armstrong, 2005). The radio eclipses are in all three cases evidence for material leaving the companion and presumably being expelled from the system, and in all three cases the companion appears to fill, and perhaps overflow, its Roche lobe (Bogdanov et al., 2005b; Ferraro et al., 2001; Thorstensen and Armstrong, 2005).

Bogdanov et al. (2005b) studied 47 Tuc W and found an X-ray spectrum consisting of a dominant hard power-law ($\Gamma = 1.14(35)$, contributing $\sim$75% of the total flux) plus a thermal component. They also observed orbital variability, and in particular a substantial X-ray eclipse spanning phases[10] 0.15–0.45 and centred slightly ahead of the optical minimum, at phase 0.2. This is accompanied by substantial softening of the X-ray spectrum during eclipse. Given these observations, Bogdanov et al. suggest that the power-law X-ray emission originates from a shocked region where material overflowing from the companion of 47 Tuc W meets the pulsar wind and is blown out of the system. Bogdanov et al. (2010) studied *Chandra X-ray Observatory* observations of PSR J1740−5430. Although

[10] Note that these authors use a different convention for orbital phase, setting phase 0.5 to the optical minimum, while we set phase 0.25 to the optical minimum. All phases given here have been converted to our convention.

limited by the scarcity of photons, they found a spectrum consistent with a power law of index $\Gamma = 1.73(8)$, or with a somewhat harder power-law plus a black-body component. Their estimates of variability were limited by restricted orbital coverage as well as by a paucity of photons, but they found marginal evidence for orbital modulation, in particular a possible decrease in luminosity during the radio eclipses, possibly accompanied by a softening. Our observations of J1023 are more photon-starved than those of Bogdanov et al. (2005b) of 47 Tuc W, but we do see evidence for eclipses, and possibly for softening during eclipses. In the case of J1023, the possibility of overflowing gas is supported by the recent disc phase.

The similarity of the observational properties of these three MSPs suggests that their companions are all overflowing their Roche lobes, so that mass flows through the L1 point, where it encounters the pulsar wind and is swept away. On the other hand, the similarity also suggests the possibility that like J1023, 47 Tuc W or PSR J1740−5430 may undergo episodes in which the mass transfer rate from the companion increases and an accretion disc forms. Such episodes would presumably be signalled by optical changes, possibly by extinction of the radio pulsations, and presumably by modest X-ray brightening. All together, these observations suggest that as a system reaches the end of its life as an LMXB, its accretion rate drops and the MSP becomes active. After this point, the wind from the pulsar generally prevents mass leaving the companion from forming an accretion disc or entering the pulsar magnetosphere. Temporary increases in the mass accretion rate may allow the occasional suppression of the wind and formation of an accretion disc, as was observed in J1023 in 2001.[11]

4.5 Summary

J1023 shows X-ray emission probably consisting primarily of emission from an intrabinary shock, plus (we think likely) a smaller amount of emission from the pulsar itself and possibly some unresolved nebular emission. The shock emission shows substantial variability that appears to be linked to orbital phase, as well as possible spectral softening during eclipses. Similarities to 47 Tuc W and PSR J1740−5340 suggest that this may be due to gas overflowing from the companion meeting the pulsar wind in a shock close to the companion. The emission from the pulsar itself, if real, is probably due to some combination of radiation in the magnetosphere and thermal radiation from polar caps heated by high-energy particles streaming downward, though which effect dominates is unclear. Further observations to provide improved spectra should help clarify the origin of both pulsed and unpulsed emission, as would further high-time-resolution observations to allow better measurement of pulse profiles, pulsed hardness ratios, and orbital variations of pulsed fraction. X-ray observations with higher spatial resolution might also resolve extended nebular emission around J1023.

[11] The transition from MSP to LMXB and back was observed, on quite short timescales, in the system M28I. See Section 7.3 for more about this system, and Chapter 7 for more discussion of these transitions.

CHAPTER 5
Radio observations of PSR J1023+0038

This chapter is based on the manuscript "Long-Term Radio Timing Observations of the Transition Millisecond Pulsar PSR J1023+0038," to be submitted to The Astrophysical Journal*. This chapter includes the parts of the paper that discuss our long-term program of radio observations of J1023, including pulsar timing measurements and observations of the complex radio phenomena that occur in this system.*

5.1 Introduction

In order to better understand J1023, we carried out a long-term campaign of radio observations. The goal of this campaign was to produce and maintain a long-term phase-connected timing solution for J1023, and to observe and understand the complex radio phenomenology of this system. The long-term solution played a key role in a very long baseline interferometry[1] campaign, described in Deller et al. (2012),[2] and permitted a search for pulsations in *Fermi* data, described in

[1] Very long baseline interferometry is a technique where careful calibration allows signals from telescopes separated by thousands of kilometers to be combined in phase. The resulting images can have the highest spatial resolutions available in astronomy.

[2] As second author on Deller et al. (2012), I maintained the timing ephemeris and made pulse phase predictions for use in pulsar gating of the VLBI data. I also

Chapter 6. The complex phenomenology we describe in Section 5.4 sheds light on the processes occurring within the binary system; we discuss the implications in Chapter 7. Here we describe in detail our observations and the phenomena they reveal.

5.2 Observations

We have carried out radio timing observations of J1023 spanning four years with four different observing systems, described below. The observations are summarized in Figure 5–1.

Observations at the Arecibo Observatory in Puerto Rico were carried out with the 300-m William E. Gordon radio telescope. We used the L-band wide receiver, which makes available four independent 100-MHz intermediate-frequency channels (IFs). We fed one of these IFs into the ASP coherent dedispersion backend (Demorest, 2007b; Ferdman, 2008b), producing online-folded data in PSRFITS format spanning 1384 to 1430 MHz. We fed the other three IFs into three WAPPs (Dowd et al., 2000b), which produced search-mode[3] data in a custom format covering 1120 to 1220 MHz, 1320 to 1420 MHz, and 1520 to 1620 MHz. When possible, we validated observations of J1023 by accompanying them with observations of PSR J1022+1001, a nearby pulsar with a 10-ms period

worked out the implications of the distance measurement for the system geometry and wrote the corresponding section of the paper.

[3] Where possible we obtained data in search mode in order to retain information about short-timescale events like the brief outages described in Section 5.4.3.

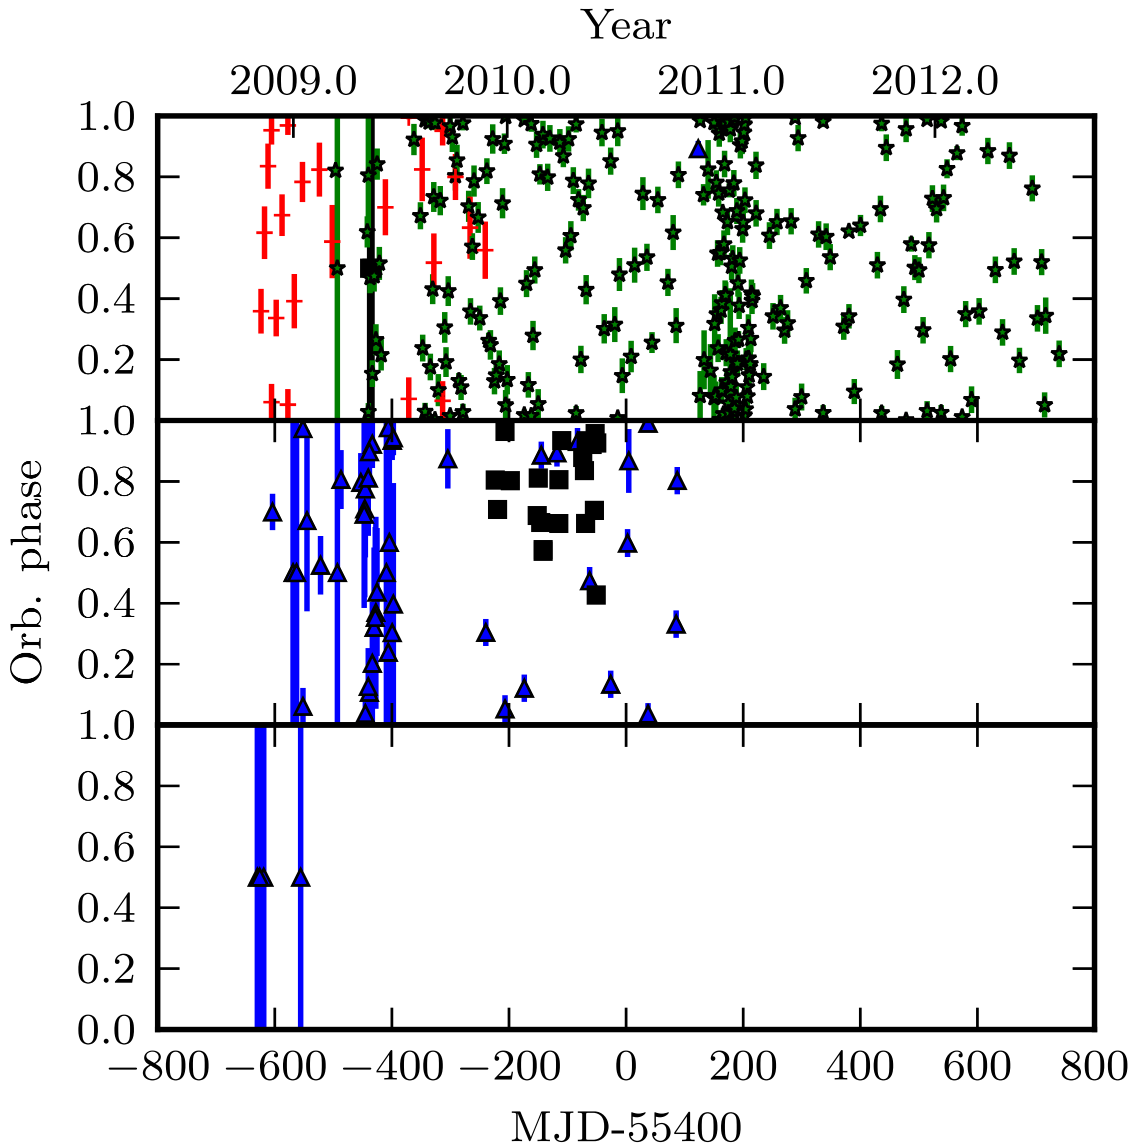

Figure 5–1 Our radio observations of J1023, showing the range of orbital phases covered and the date. Top panel shows 1400 MHz observations, middle panel shows 350 MHz observations, and bottom panel shows 150 MHz observations. Colours and markers indicate observatory, with red crosses denoting Arecibo, green stars denoting Jodrell Bank, black squares denoting Green Bank, and blue triangles denoting Westerbork. Observations longer than a full orbit are marked here as top-to-bottom bars.

that rises before and sets after J1023.[4] For the ASP data, we also took calibration observations with a gated noise diode. Our program consisted of an observation approximately every three weeks, lasting about an hour. This program stopped in early 2010, when the telescope was shut down for repairs, and did not resume thereafter.

Observations with the Westerbork Synthesis Radio Telescope (WSRT) in the Netherlands used the telescope's tied-array mode, which combines (up to) all 14 dishes of the array in phase to provide the equivalent collecting area of a 94-m single dish (see Section 2.2 for more on this mode). The observations were acquired in one of three bands centred at 150, 350, and 1380 MHz, with respective total recorded bandwidths of 20, 80, and 160 MHz. The WSRT observations were typically half an hour to a few hours in length. Baseband data were recorded and subsequently coherently dedispersed and folded offline using the PuMaII pulsar backend (Karuppusamy et al., 2008).

Observations at the Jodrell Bank Observatory in the United Kingdom used a dual-polarization cryogenic receiver on the 76-m Lovell telescope, having a system equivalent flux density of 25 Jy on cold sky. Data in the range 1350 MHz to 1700 MHz were processed by a digital filterbank for a bandwidth of 112 MHz producing 0.5 MHz channels up until August 2009; subsequently a bandwidth of

[4] J1023 is at the extreme southern limit of the declination range accessible from Arecibo, so it can be observed for no more than about an hour per day.

about 300 MHz with 0.25 MHz channels was produced. The data were incoherently dedispersed and folded online. Most observations were half an hour long but a few were longer, and in particular one long observation on MJD 54906 was simultaneous with a long WSRT 350 MHz observation.

Observations with the 110-m Robert C. Byrd telescope at Green Bank were carried out with the GUPPI pulsar instrument (Ransom et al., 2009) in two different modes. For our long observations, we recorded 1100 to 1900 MHz in incoherent dedispersion mode with online folding, producing files in PSRFITS format. Each observation began and ended with a calibration scan using a gated noise diode, though three of the four calibration scans proved unusable. In this configuration we carried out two-full-orbit observations on 2009 May 10 and 2009 May 14. We also carried out a number of short-term low-frequency observations, recording data covering 300 to 400 MHz in incoherently dedispersed search mode (that is, without online folding) and producing data in PSRFITS format (Hotan et al., 2004).

5.3 Long-term timing solutions

A key aspect of the current work was maintaining a long-term timing solution for J1023 (for example in support of Deller et al., 2012). In light of the complexities we will discuss in Section 5.4 and below, frequent observations were essential. Early on, the regular 1400-MHz monitoring with the Arecibo telescope provided the key timing observations, while later the Jodrell Bank 1400-MHz observations, at some times almost daily, provided essential timing information. Even with this large collection of observations, obtaining an adequate model for the rotation of J1023 was difficult.

Many observations were obtained in online folding mode, based on our then-best estimate of the pulsar's ephemeris. Others were taken in search mode (power measurements in each frequency channel every few tens of microseconds). We folded the search-mode data with preliminary versions of our long-term ephemeris, then refolded when the ephemeris was improved significantly. The result of this latter process was a folded pulsar observation consisting of one profile per subintegration (30 s to several minutes) per frequency channel.

For each combination of instrument and frequency band, we selected an observation in which the folded pulse had a good signal-to-noise ratio and negligible phase drift or DM variation. We then averaged these observations down to a single polarization, a single frequency channel, and a single time integration to produce a high signal-to-noise ratio template (like those in Figure 5–2). We then used this template and either the PSRCHIVE (Hotan et al., 2004) tool `pat` (for PSRCHIVE data) or the PRESTO (Ransom, 2001a) tool `get_toas.py` (for WAPP data) to produce pulse time-of-arrival measurements (TOAs) by cross-correlation. For some observations we produced two sets of TOAs, one set based on averaging across the whole observing band and the second based on averaging across a modest number of subbands. The former process produced TOAs with lower uncertainties and in a more modest number, while the latter process permitted testing for DM variations (see Section 5.4.2) but produced too many TOAs for use in a global timing solution.

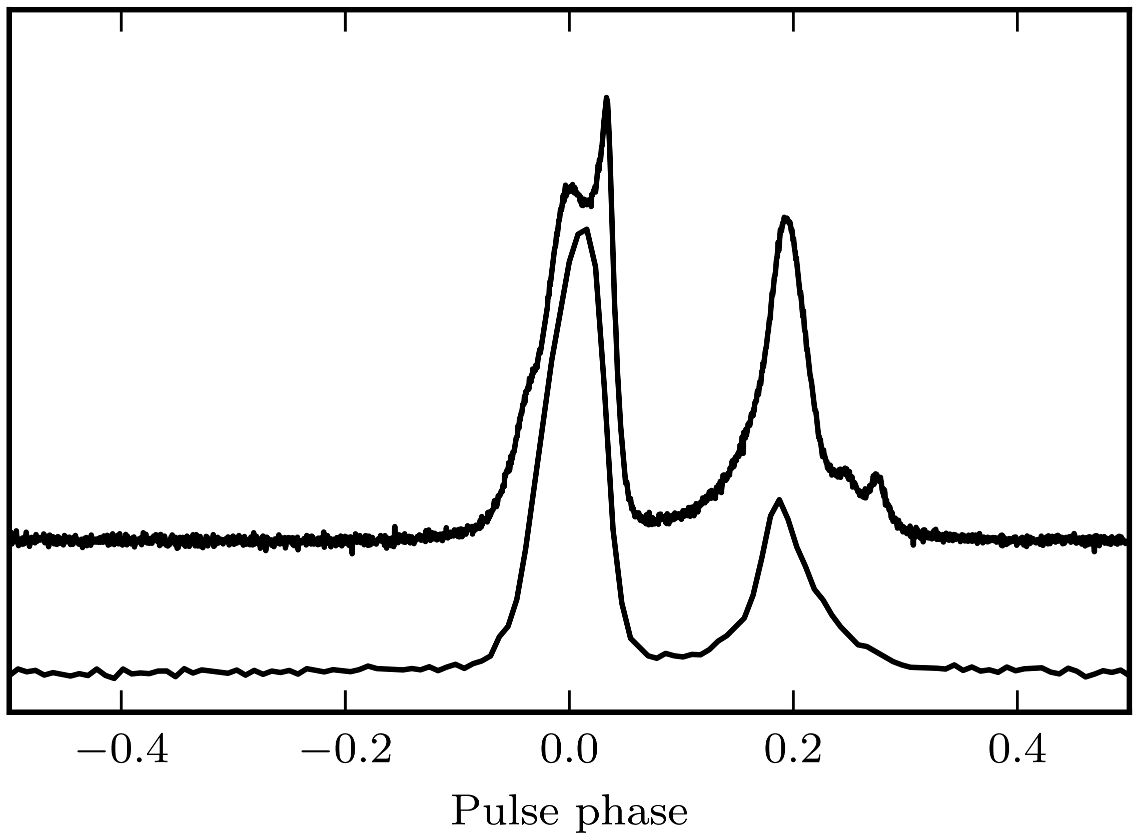

Figure 5–2 Radio pulse profiles of J1023. Top: 1400-MHz profile obtained with coherent dedispersion using the ASP at Arecibo. Bottom: 350-MHz profile obtained with coherent dedispersion using the WSRT. Vertical scale is flux density, scaled arbitrarily for each profile. Profiles are arranged horizontally to approximately align both major peaks.

When observing J1023, a number of phenomena (described in Section 5.4) may corrupt individual TOAs. Most obvious are eclipses and flux variations, which may make pulse phase measurements impossible. More subtle are DM variations, which introduce additional frequency-dependent delays. These variations can contribute substantially to the systematic errors in TOA measurements. In the 350-MHz data, these effects can be substantial at all orbital phases. Fortunately, in the 1400-MHz data these effects are small except at orbital phases close to inferior conjunction. We used only the 1400-MHz data for timing purposes, and we handled these various effects by rejecting all TOAs with phases too near inferior conjunction or with measured uncertainties larger than $10\,\mu$s. We also removed certain TOAs that showed evidence of being severely affected by excess DM (as described in Section 5.4.2).

Given a collection of good TOAs, the task of producing an ephemeris consists of two parts, in our case carried out using the pulsar timing software `tempo2` (Hobbs et al., 2006). The first is to determine the exact number of turns between each pair of TOAs, while the second is to adjust the ephemeris to model the rotation more closely. The first process is called "phase connection" and is normally a manual, iterative process of extending a partial ephemeris to cover new observations. Errors in the predicted number of turns generally reveal themselves by producing residuals with errors greater than a few tenths in phase. With the kind of dense sampling possible with J1023, they may also be visible as large trends within single observations. We carried out the process of phase connection on our observations of J1023, producing the ephemeris given in Table 5–1. In

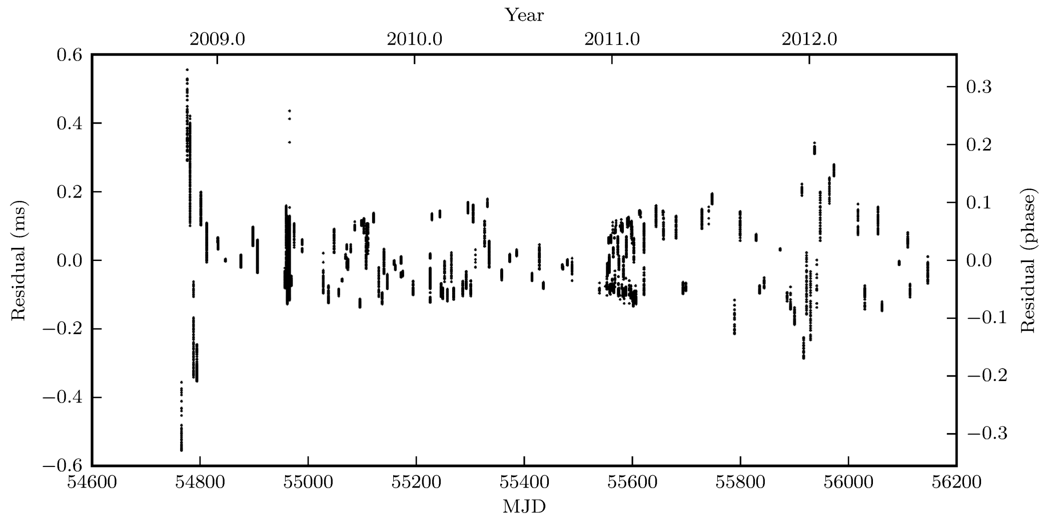

Figure 5–3 Residuals as a function of time for the simple long-term solution in Table 5–1. Vertical bars indicate uncertainties. In some cases they are smaller than the circle used to mark the value. In this plot a sequence of TOAs from a longer observation may have quite different residuals due to the orbital dependence shown in Figure 5–4; this can have the appearance of a long vertical bar.

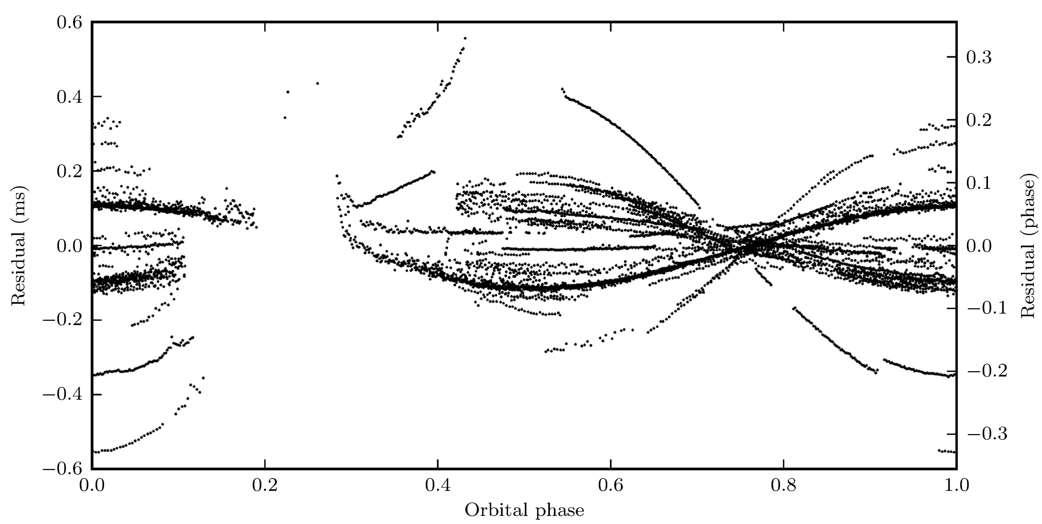

Figure 5–4 Residuals as a function of orbital phase for the simple long-term solution in Table 5–1. Vertical bars indicate uncertainties. In some cases they are smaller than the circle used to mark the value.

Table 5–1 Long-term ephemeris for J1023

Fit and data-set	
Pulsar name	J1023+0038
MJD range	54766.5—56146.6
Number of TOAs	7478
Rms timing residual (μs)	114.0
Weighted fit	No
Measured Quantities	
Pulse frequency, ν (s^{-1})	592.42145906986(10)
First derivative of pulse frequency, $\dot{\nu}$ (s^{-2})	$-2.432(3)\times10^{-15}$
Orbital period, P_b (d)	0.1980963569(3)
Epoch of periastron, T_0 (MJD)	54905.9713992(3)
Projected semi-major axis of orbit, x (lt-s)	0.343343(3)
First derivative of orbital period, $\dot{P}_b$	$-7.32(6)\times10^{-11}$
Set Quantities	
Right ascension, α	10:23:47.687198
Declination, δ	+00:38:40.84551
Epoch of frequency determination (MJD)	54906
Epoch of position determination (MJD)	54995
Dispersion measure, DM ($\mathrm{cm}^{-3}\mathrm{pc}$)	14.3308
Proper motion in right ascension, μ_α ($\mathrm{mas\,yr}^{-1}$)	4.76
Proper motion in declination, μ_δ ($\mathrm{mas\,yr}^{-1}$)	−17.34
Parallax, π (mas)	0.731
Orbital eccentricity, e	0

Note: Figures in parentheses are the nominal 1σ `tempo2` uncertainties in the least-significant digits quoted. Position, proper motion, and parallax values are held fixed at the values from Deller et al. (2012). Orbital eccentricity was held fixed at zero.

building this ephemeris, we fixed J1023's position, proper motion, and parallax at the values obtained with VLBI (Deller et al., 2012). We also fixed the DM at $14.3308\,\mathrm{pc\,cm^{-3}}$, a value obtained from fitting a WSRT 350-MHz observation that is far from eclipse and that has a good signal-to-noise ratio. The DM selected is relatively unimportant in the long-term fitting, since all the TOAs from a given telescope/backend/band combination are at the same frequency, and we fit for arbitrary offsets between telescope/backend/band combinations.

The ephemeris in Table 5–1 produced the phase residuals shown in Figures 5–3 and 5–4. While these residuals are in some cases as large as 0.2 turns, their smooth orbital dependence makes it clear that the ephemeris does correctly phase-connect all our observations — that is, we can unambiguously account for the exact number of pulsar rotations between each observing epoch. While the orbital dependence of the residuals indicates unmodelled orbital variations, the scale of these variations is rather small. As a compromise between model complexity and residual size, this ephemeris uses a very simple orbital model, allowing the orbital period to vary linearly over the span of our observation. This ephemeris is appropriate for orbital phase predictions and for online-folding observations of J1023 that are not too far in the future.

In the interest of better understanding the orbital variations and other phenomenology observed in J1023, it is valuable to use `tempo2` to fit many shorter-term observations. However, such automated use of `tempo2` risks introducing a spurious phase turn, rendering useless the results of the fitting procedure. In order to prevent this, we implemented a modification to `tempo2`. By default, when tempo2 computes a residual for a pulse time-of-arrival measurement, it simply chooses the turn number that produces the smallest residual, that is, it attempts to make the residual as close to zero as possible. If the proposed solution is too far from the true solution, so that the residual should be larger than 0.5 in phase, this results in an extra phase turn. To mitigate this problem we wrote code to allow us to control the turn number directly.

The tempo2 output plugin general2 allows the output of turn numbers for each TOA. We implemented a mode, activated with TRACK -3, which allows turn numbers to be specified as part of the TOA data. Thus we were able to use the phase-connected solution from Table 5–1 and the general2 plugin to annotate our TOAs with turn numbers. We were then able to activate TRACK -3 and be certain that no automated fit introduced spurious phase turns. We point out that this allows, for example, easy changing from one timing model to a differently parametrized timing model; wrong parameters will not result in phase wraps and can therefore be corrected by tempo2's fitting procedure. We aim to incorporate the code implementing TRACK -3 into the distributed version of the open-source tempo2.[5]

This new mode for tempo2 was useful at several stages in our analysis. During initial phase connection, it was useful for stitching together solutions each covering part of the data set: if the differences between turn numbers agreed for those

[5] Those interested in these modifications to tempo2 may also contact the author at aarchiba@physics.mcgill.ca.

TOAs for which the solutions overlapped, and the overlap was sufficient, then the solutions were compatible, and we could merge the lists of turn numbers to annotate all the TOAs in either data set. Then a timing solution could be fit to the data without the complication of phase wrapping. Second, once we had a phase-connected timing solution, we could "bake in" the turn numbers so that the automated short-term fitting we describe in Section 5.4.5 and Chapter 6 could not introduce a phase wrap and therefore converge on an incorrect solution.

5.4 Radio phenomenology

J1023 exhibits a number of phenomena which give some insight into the nature of this complex binary system, several of which are visible in Figure 5–5. The most obvious is the main eclipse, which occurs near inferior conjunction (orbital phase 0.25), when the companion passes nearest to our line of sight to the pulsar. This eclipse is accompanied by excess DM at ingress and egress. Excess DM is also seen at other orbital phases, varying from orbit to orbit, as is seen in PSR J1748−2446A (Terzan 5 A; Nice et al., 1990). In 350-MHz observations, we observe brief (seconds to minutes) periods during which the pulsar signal disappears. We also observe the effects of interstellar scintillation on the signal from J1023, which can complicate observing by introducing large intensity variations at 1400 MHz. Finally, the orbital parameters of J1023 are variable, severely complicating long-term timing.

5.4.1 Main eclipse

The most obvious feature of J1023's radio emission is the main eclipse. The system's viewing geometry is well known: based on the mass measurement of

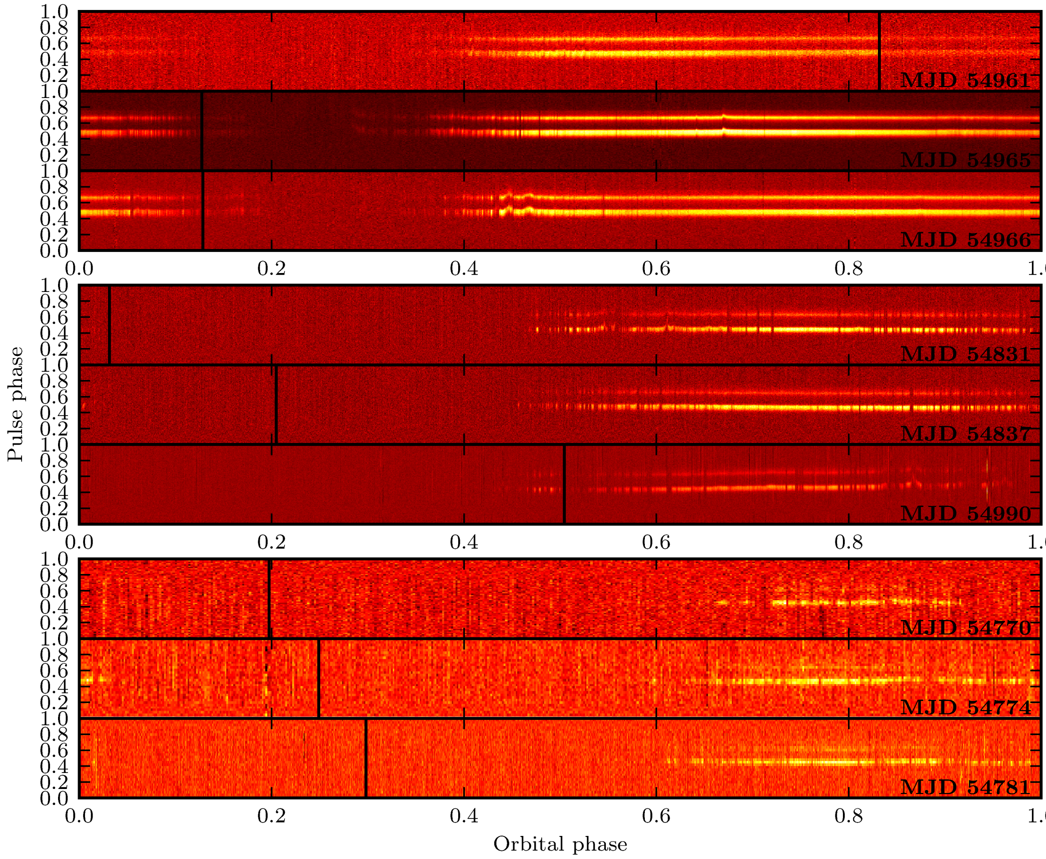

Figure 5–5 Several full-orbit observations of J1023. Colour indicates flux density as a function of pulse and orbital phase. The top group of three panels shows 1400-MHz observations with the GBT, the middle shows 350-MHz observations with the WSRT, and the bottom shows 150-MHz observations with the WSRT. Within each panel, the observation begins with the black vertical line, wraps around, and is cut off at the black line. The date of the observation is indicated within each panel; the second and third 1400 MHz panels are in fact taken from a single scan, so that the third panel takes up where the second leaves off. For each of the observations below 1400 MHz, we used the DM that gave the best signal-to-noise; values varied by $\sim 10^{-3}\,\mathrm{pc\,cm^{-3}}$. Differences between panels at the same frequency show the orbit-to-orbit variability of J1023; differences between observations at different frequencies show the same combined with the frequency dependence of the various effects described in Section 5.4. For a truly simultaneous dual-frequency comparison see Figure 5–7.

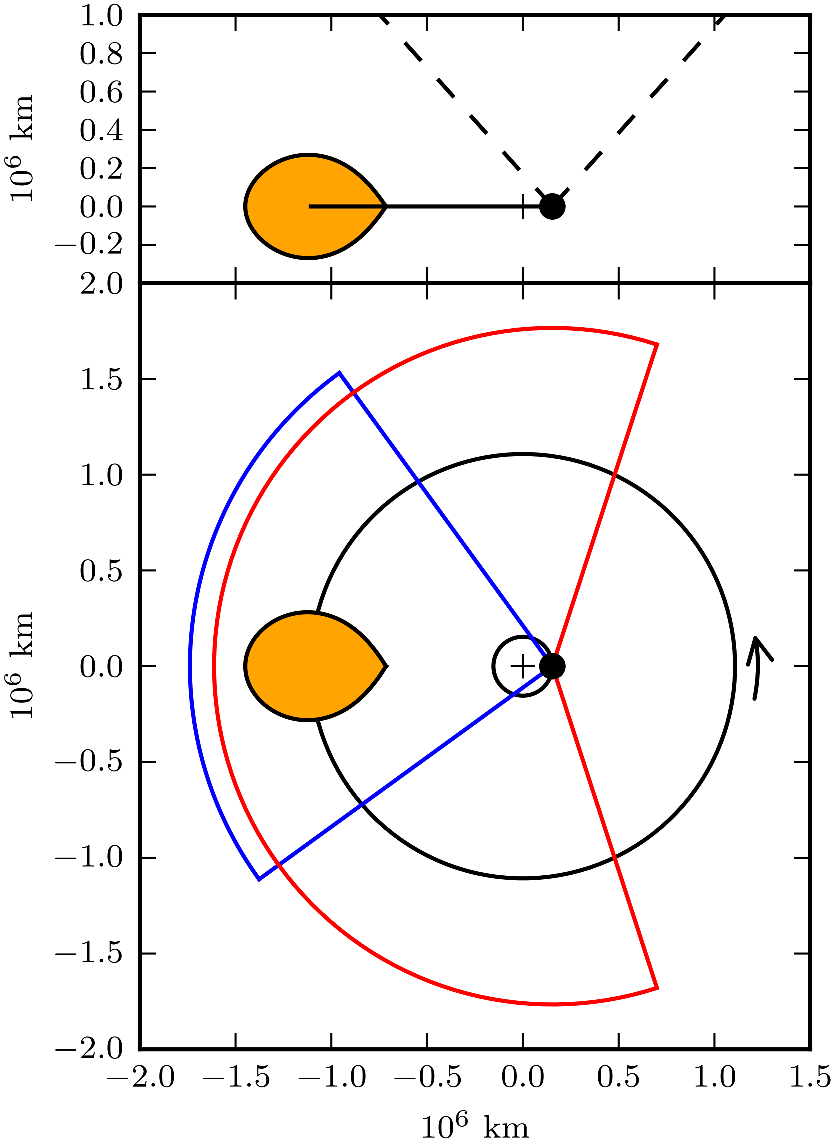

Figure 5–6 Geometry of the J1023 system and its eclipse. The orange teardrop shows the shape and size of the companion (assumed to fill its Roche lobe), while the black dot indicates the pulsar (not to scale). Top: Edge-on view of the system; dashed lines indicate the line-of-sight to Earth at orbital phases 0.25 and 0.75. Bottom: Face-on view of the system. Arcs indicate the main eclipse phase ranges from the dual-frequency observation, for a system rotating counter-clockwise. Red is for the 350-MHz eclipse while blue is for the 1400-MHz eclipse.

$1.71 \pm 0.16 M_{\odot}$ in Deller et al. (2012) and the Keplerian mass function of the system, we can infer that the inclination angle is $42 \pm 2°$. With this geometry, pictured in Figure 5–6, the companion's Roche lobe does not cross the line of sight to the pulsar. In fact, Figure 3–3 shows a 3-GHz full-orbit observation of J1023 in which the main eclipse is absent: approximately constant emission is visible throughout the orbit. The eclipse is therefore presumably due to relatively tenuous ionized material outside the companion's Roche lobe. On the other hand, at 1400 MHz the phase-averaged flux during eclipse is lower by a factor of several (in fact undetectable) than the flux outside eclipse (Deller 2012, personal communication).

A dual-frequency observation we acquired allows us to examine this effect more closely. Figure 5–7 shows both 350-MHz and 1400-MHz observations taken simultaneously with the WSRT and the Lovell telescope respectively. At 1400 MHz the eclipse length is about 0.25 in orbital phase, while at 350 MHz it is closer to 0.6 in orbital phase. It can also be seen that the eclipse ingress moves more than the eclipse egress as observations move to lower frequency, confirming what is seen in the non-simultaneous observations shown in Figure 5–5 and Figure 3–3.

5.4.2 Excess DM

A second feature of J1023's radio emission is excess DM. This is visible in Figure 5–5 and in timing observations as excess delays. Wide-band observations like those taken as part of our monitoring program at Arecibo allow us to confirm that these delays are in fact due to excess DM by computing TOAs from several subbands in the same subintegration, then by fitting for the DM within that

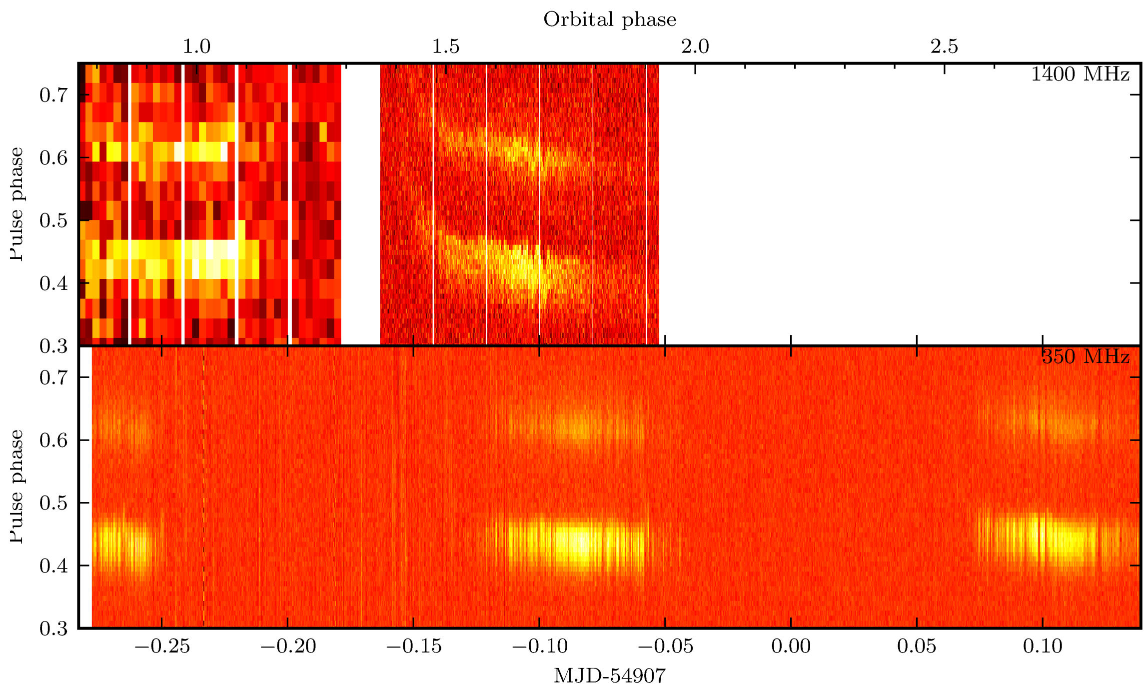

Figure 5–7 Simultaneous dual-frequency observation of J1023. Colour indicates intensity as a function of time and pulse phase. Top panel is observed with Jodrell Bank at 1400 MHz, bottom panel is observed with the WSRT at 350 MHz. Note that scintillation resulted in a drastically reduced signal-to-noise during the first part of the Jodrell Bank observation, so we have reduced the time and phase resolution.

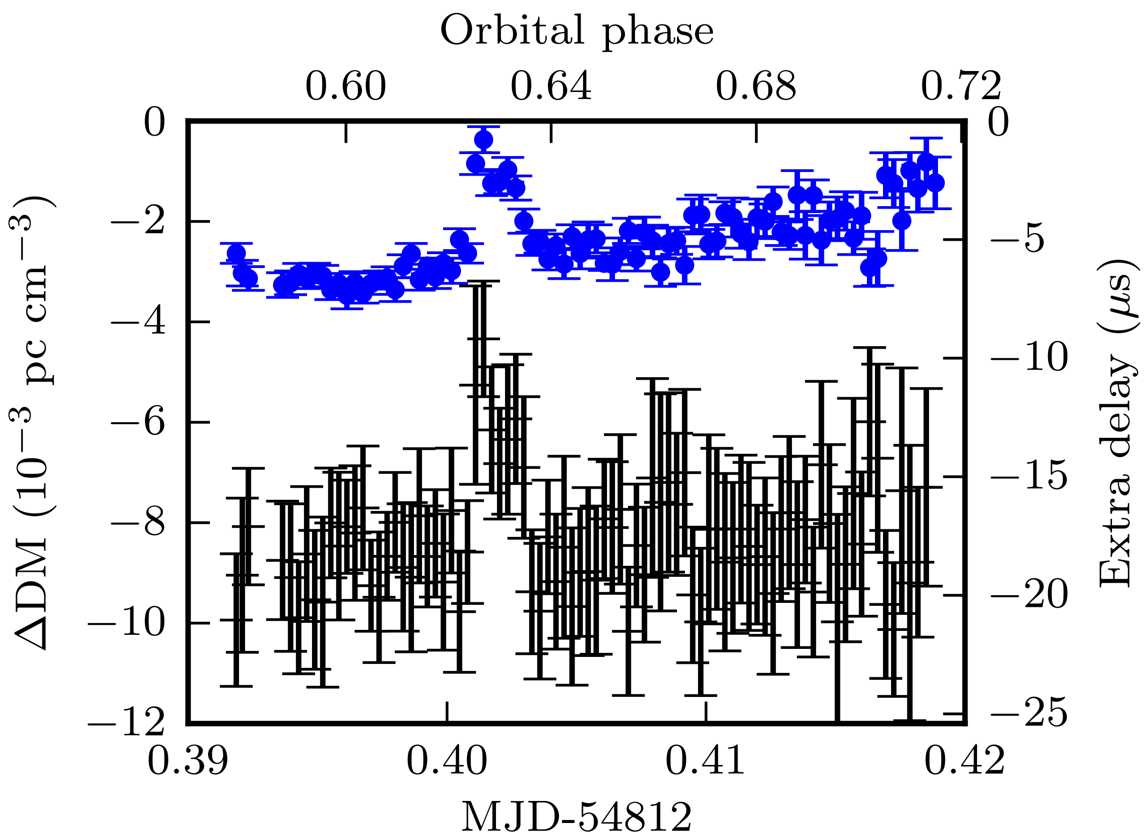

Figure 5–8 Short-term DM fits to an Arecibo 1400-MHz observation of J1023 on MJD 54812. Vertical error bars show the uncertainty returned by `tempo2`. The top sequence of points with blue circular markers are simple timing residuals from single-frequency observations assuming a single DM; their amplitude can be read off the scale to the right. The lower sequence of points with black crosses for markers are fits for dispersion measure within single subintegrations; their scale can be read off the left side of the plot. The two scales have been arranged so that the excess dispersion measure (relative to our standard assumed dispersion measure; the "excesses" are all negative in this plot) on the left corresponds to the excess delay (relative to a particular timing model; again negative in this plot) on the right for a 1400-MHz observation.

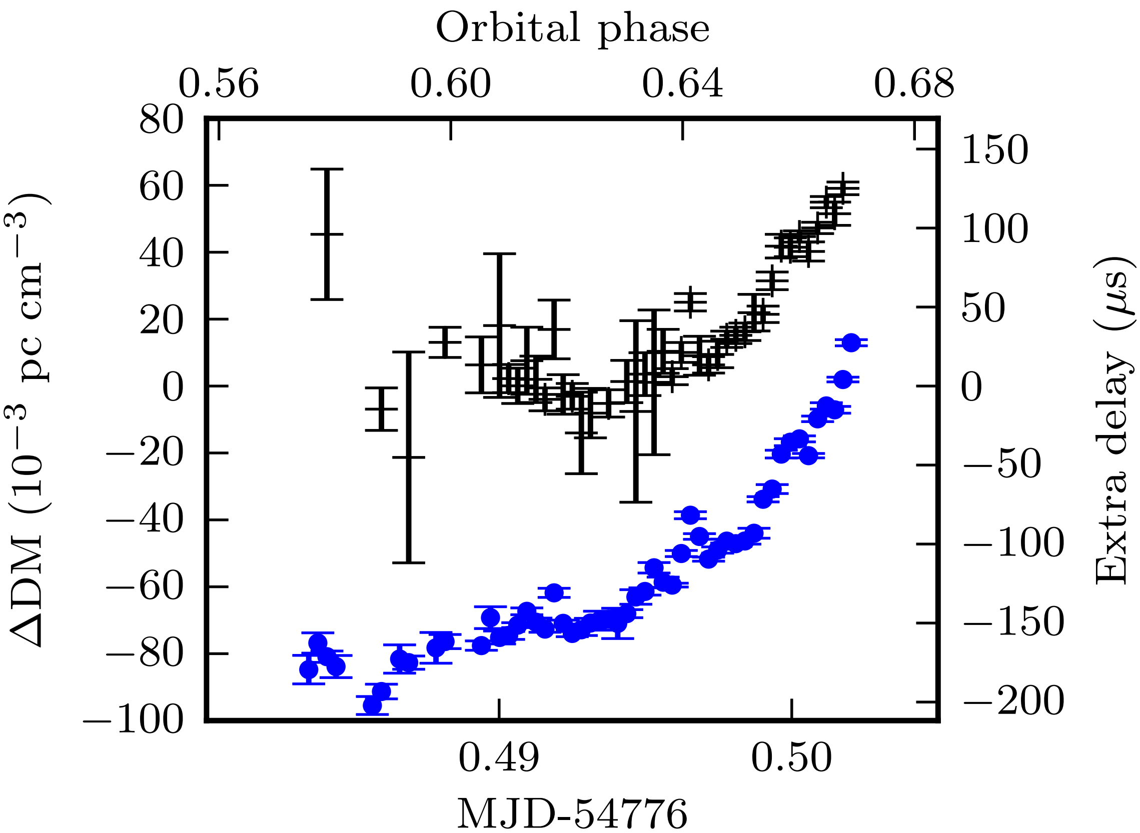

Figure 5–9 Short-term DM fits to an Arecibo 1400-MHz observation of J1023 on MJD 54776. Scales and traces are defined as in Figure 5–8, but note that because this observation was very near the beginning of our program and cannot be used for timing because of the strong excess DM, there may be a trend superimposed on the timing residuals.

single subintegration. The excess DM we observe takes several forms. First, we consistently observe excess DM surrounding the main eclipse. Specifically, we see the DM increase rapidly up to the moment the signal disappears, then fall rapidly down to the normal value after the signal reappears. The amount of excess DM is difficult to estimate, since the points at which we lose and reacquire signal depend on the signal-to-noise ratio, which varies substantially due to scintillation. Nevertheless it is clear that the amount of excess DM at a phase near the main eclipse also varies from orbit to orbit. Typical excess DM at loss of signal due to eclipse ingress is roughly $0.01\,\mathrm{pc\,cm^{-3}}$, while at signal reacquisition due to eclipse egress it is more usually around $0.15\,\mathrm{pc\,cm^{-3}}$. Second, in our hour-long monitoring observations, the average DM varies by a few times $10^{-3}\,\mathrm{pc\,cm^{-3}}$ from what we have adopted as the baseline DM for the system. Finally, in several cases, one of which is pictured in Figure 5–8, we see short-term (varying lengths but typically $\sim$300 s) excesses of DM (by varying amounts but typically $\sim 4 \times 10^{-3}\,\mathrm{pc\,cm^{-3}}$). Importantly, these DM excesses occur far from the main eclipse, and at different orbital phases, as do the signal disappearances discussed in Section 5.4.3. We also see one example of a substantially larger and longer-lasting increase in DM, shown in Figure 5–9; on MJD 54776, we observed the DM rise steadily by $\sim 5 \times 10^{-2}\,\mathrm{pc\,cm^{-3}}$ over the course of twenty minutes; our observation ended while the rise continued. This occurred around orbital phase 0.65, somewhat after the eclipse.

It should be noted that the delays introduced by this excess DM at 1400 MHz are on the order of tens of microseconds, while the residuals pictured in Figures 5–3 and 5–4 are on the order hundreds of microseconds, so excess DM is not sufficient to explain them (see Section 5.4.5 for their cause). On the other hand, we should expect the effect of DM variations to be some sixteen times greater for 350-MHz observations, producing delays on the order of hundreds of microseconds and phase shifts on the order of tenths of a turn. For this reason, it would be difficult to maintain a good timing solution for J1023 using only 350-MHz data, although given the high signal-to-noise ratio, it is possible to measure the DM and possibly correct for the delays within each observation. In fact, in order to obtain a good signal-to-noise ratio at 350 MHz and below, it is necessary to measure the DM for use in each observation.

5.4.3 Short-term signal disappearances

At 350 MHz, we observe brief periods during which the pulsar signal disappears entirely. In some of our search-mode observations it is possible to see disappearances as short as a few seconds, but in most of our data (folded in 10-s subintegrations) we see disappearances varying in length from one subintegration to as long as minutes. These disappearances occur at all orbital phases (apart from during the main eclipse, when there is no signal to disappear), though they may be more frequent nearer to eclipse. In a few cases the signal disappearances begin or end with excess delays (presumably due to excess DM). The dynamic spectrum (see Figure 5–10) makes it clear that the disappearances span the entire bandpass ($\sim$80 MHz) and show no sign of spectral dependence. As is visible in Figure 5–5,

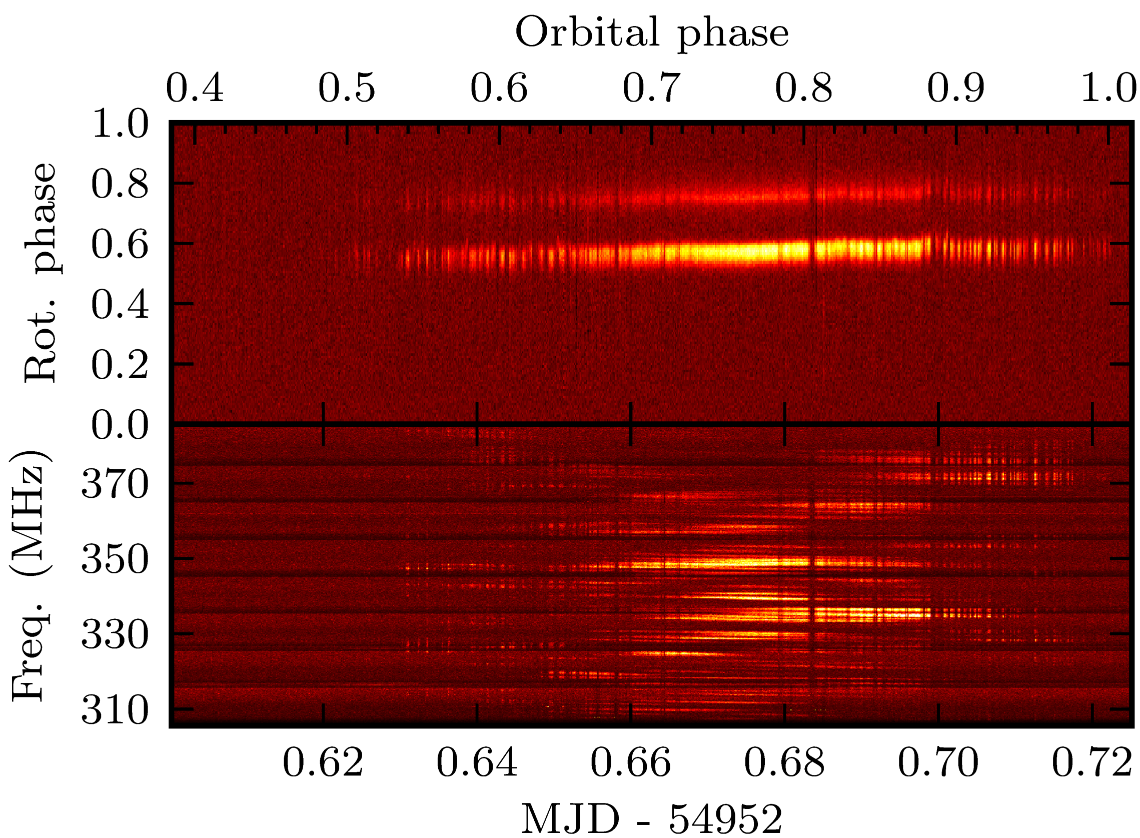

Figure 5–10 Top: Brightness as a function of time and pulse phase for a WSRT 350 MHz observation taken on MJD 54952. Note both short-term signal disappearances and episodes of delayed pulses, presumably due to excess DM. Bottom: Dynamic spectrum of the same observation. Note that the disappearances of signal are broadband and clearly distinct from the scintillation. This observation is assembled from eight subbands, so certain band-edge artifacts are visible, as is some radio-frequency interference.

these signal disappearances also occur at 1400 MHz and 150 MHz. That said, in our 350 MHz/1400 MHz dual-frequency observation, disappearances that occur at the lower frequency are not accompanied by disappearances at 1400 MHz; we are not able to test the converse.

5.4.4 Interstellar Scintillation

Interstellar scintillation results when the turbulent interstellar medium modulates the radio flux from a pulsar. The variation appears as a collection of *scintles* with a characteristic duration and bandwidth and with varying intensity. A few of these scintles, detected in our 350-MHz data, are visible in Figure 5–10. Autocorrelation analysis (Lorimer and Kramer, 2004) shows a bandwidth Δf_{DISS} of 1.4 MHz and a timescale Δt_{DISS} of 1100 s. The uncertainties on these quantities are dominated by the small number of bright scintles. Following Cordes (1986), we estimate the number of scintles in the observation we used as the total number of possible scintles times a small filling factor to account for the dominance of exceptionally bright scintles in our autocorrelation analysis: $N_s = 10^{-2}(T/\Delta t_{DISS})(B/\delta f_{DISS})$. In this case we are using a $T = 7000\,\mathrm{s}$ observation segment (aside from eclipsed times) with $B = 100\,\mathrm{MHz}$, so we estimate $N_s = 4.5$; we should therefore expect fractional uncertainties on scintillation parameters of $N_s^{-1/2} = 0.5$. Nevertheless, Lorimer and Kramer (2004), citing Cordes and Rickett (1998), give the following formula for inferring the pulsar's proper motion from scintillation parameters:

$$V_{\mathrm{ISS}} = 2.53 \times 10^4 \ \mathrm{km\,s^{-1}} \left(\frac{d}{1\mathrm{kpc}}\right)^{1/2} \times \tag{5.1}$$

$$\left(\frac{\Delta f_{\mathrm{DISS}}}{1\mathrm{MHz}}\right)^{1/2} \left(\frac{f}{1\mathrm{GHz}}\right)^{-1} \left(\frac{\Delta t_{\mathrm{DISS}}}{1\mathrm{s}}\right)^{-1}. \tag{5.2}$$

Applying this to the known distance (1.37 kpc; Deller et al., 2012) and a frequency of 350 MHz gives a velocity of $91\,\mathrm{km\,s^{-1}}$. Assuming a fractional uncertainty on the order of 50% makes this roughly consistent with the observed proper motion of $130\,\mathrm{km\,s^{-1}}$ (Deller et al., 2012), particularly as during this partial orbit the speed at which the sight line sweeps through the interstellar medium is also affected by the pulsar orbital motion, which varies with orbital phase from 40 to $60\,\mathrm{km\,s^{-1}}$ and is in an unknown direction. At 1400 MHz the scintles are too large to obtain good statistics, but visual inspection suggests a scintillation bandwidth of ∼50 MHz and timescale of ∼1 hour; this implies a velocity of $\sim 40\,\mathrm{km\,s^{-1}}$, with an even larger fractional uncertainty given the presence of only a handful of scintles within an observation. In practical terms, interstellar scintillation introduces a wide variation in observed radio flux at 1400 MHz, particularly for narrow-band observations, for example those described in Deller et al. (2012), but also visible in the broadband observations shown in Figure 5–5.

5.4.5 Timing Variations

In addition to pulse delays due to excess DM, timing of J1023 is made difficult by its erratic orbital behaviour. Figure 5–3 shows the poor quality of the fit obtainable by a simple model of the orbit including only a secular change in orbital period. Figure 5–4 also shows that the residuals have a substantial dependence on orbital phase, suggesting that varying the orbital parameters might improve the fit. While `tempo` offers a binary model allowing many orbital period derivatives, porting it to `tempo2` and fitting numerous orbital period derivatives failed to produce satisfactorily small residuals. Instead, we chose to fit short-term piecewise

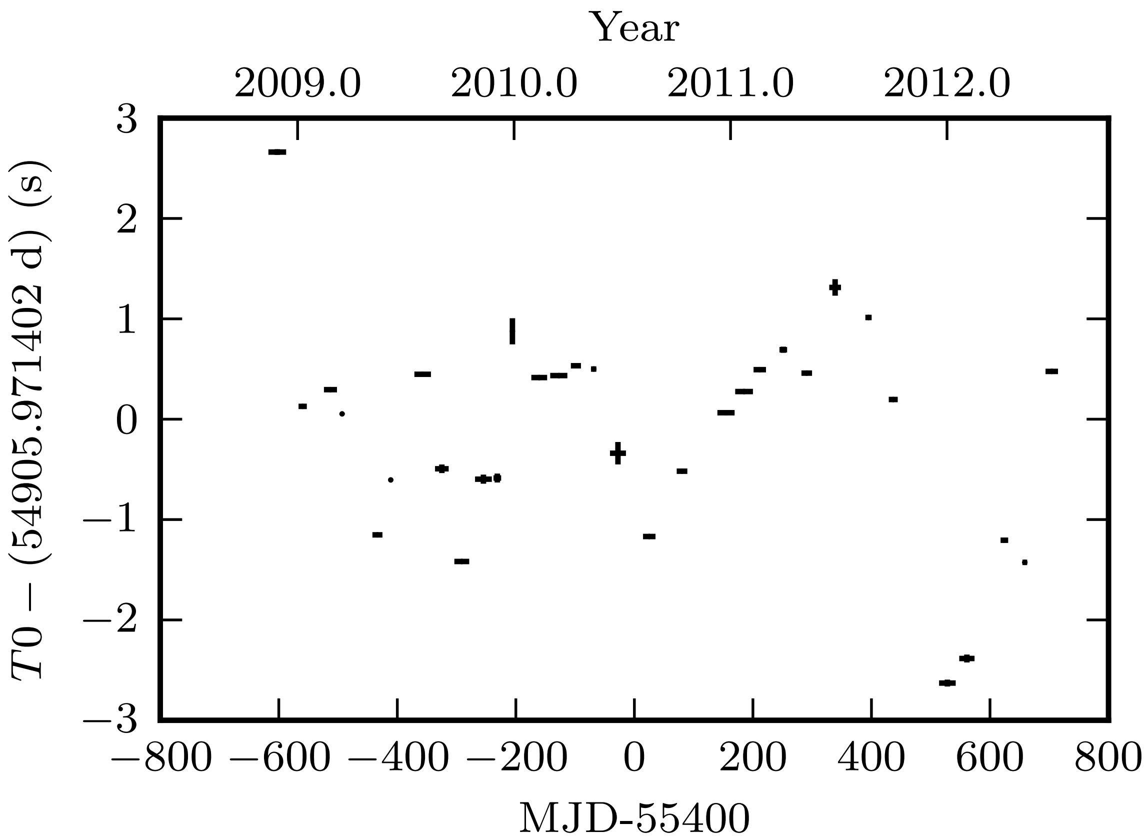

Figure 5–11 O-C (observed minus calculated) diagram of the time of orbital phase 0. A version of the long-term ephemeris was fit to each approximately 30-day section of data, allowing only orbital phase (and pulse phase) to vary, and locking the number of pulsar rotations to that computed from the original long-term ephemeris. Vertical bars indicate uncertainties reported by `tempo2` except where they are smaller than the circle used to indicate the value; horizontal bars indicate the range of data used in the fit. Since the original ephemeris was obtained by fitting an orbital period and its derivative to the whole data span, we have effectively removed the best-fit quadratic from these data points. Residual scatter appears to come from intrinsic orbital period variations.

solutions, in which we allowed orbital phase to vary. When using non-overlapping 30-day intervals, we obtained root-mean-squared residuals typically less than 1% of a turn. We also experimented with fitting binary period and/or amplitude within each segment, which reduced the residuals by a small additional amount. The fitted orbital parameters showed no discernable structure, either by eye or using a Lomb-Scargle periodogram, even when using intervals as short as 14 days. The scale of the orbital phase variations required for our 30-day fits is shown in Figure 5–11; they are of order 1 s.

In summary, the orbital variations we observe are substantial but difficult to quantify, let alone understand. In Section 7.6 we will discuss possible explanations for these orbital variations.

5.5 Summary

Our long-term radio monitoring campaign has provided much information on J1023. We have maintained a phase-coherent timing solution, which has permitted the parallax measurement of Deller et al. (2012) (including estimates of the pulsar mass and system geometry) as well as the analysis of γ-ray data presented in Chapter 6. Our monitoring campaign has also provided us with observations of J1023's complex radio behaviour. We have observed a primary eclipse roughly centred on the time the companion passes nearest to our line of sight; this eclipse becomes much longer, exceeding half the orbital period, at observing frequencies of 350 MHz and below. We have also observed excess DM, both near the edges of this eclipse and at other orbital phases; these episodes of excess DM vary in amount, duration, and orbital phase. We have observed short disappearances of signal at all orbital phases, sometimes but not always associated with episodes of excess DM. All these phenomena are clearly distinct from interstellar scintillation, which has sufficiently wide bandwidth and long timescale at 1400 MHz that it can substantially vary the observed flux density from orbit to orbit. The timing of J1023 is also complicated by apparently random variations in the orbital period. I will discuss the possible causes of this complex behaviour, and its relation to the similar behaviour seen in other redback and black widow systems, in Chapter 7.

CHAPTER 6
γ-ray emission from PSR J1023+0038

This chapter is also based on the manuscript "Long-Term Radio Timing Observations of the Transition Millisecond Pulsar PSR J1023+0038," to be submitted to The Astrophysical Journal*. This chapter includes the parts of the paper that discuss our analysis of the γ-ray emission from J1023, including searches for modulation at the orbital period and at the pulsar period.*

6.1 Introduction

Tam et al. (2010) noted that J1023 appears as a point source in γ-ray data from the *Fermi* space telescope (which was described in Section 2.6), and indeed in the 2FGL catalog J1023 is listed as having a 1–100 GeV flux of $5.4(9) \times 10^{-12}$ erg cm^{-2} s^{-1} and a power-law index of 2.5(3) (Nolan et al., 2012a). Tam et al. (2010) also report a flux of $5.5(9) \times 10^{-12}$ erg cm^{-2} s^{-1} for > 200 MeV. Either value corresponds to a γ-ray luminosity of 1.2×10^{33} erg s^{-1}. To better understand this γ-ray emission, we have obtained and folded the *Fermi* photons coming from J1023.

6.2 Data

We requested photon information and accompanying calibration data from the *Fermi* data server for the energy range 100 MeV to 300 GeV and a circular region of radius 15° centered on J1023's position. Following the *Fermi* Cicerone[1] , we selected photons with event class 2, zenith angle less than 100°, and a region of interest cut with the expression `(DATA_QUAL==1) && (LAT_CONFIG==1) && ABS(ROCK_ANGLE)<52`. Using the *Fermi* tool `gtsrcprob` and the spectral model from the 2FGL survey, we assigned to each photon a probability that it came from J1023 then culled all photons with probability less than 10^{-3}. This culling reduced the number of photons needing processing from 622882 to 47032 and the number of expected photons from the source from 623.6 to 521.6. Our Monte Carlo simulations make it clear that the low-probability photons we discarded make a negligible difference to detection significances.

6.3 Orbital modulation

To look for orbital modulation, we used the simple orbital ephemeris given in Table 5–1 and the `fermi` plugin that is part of `tempo2` to compute the orbital phase at which each photon was emitted. We used a weighted form of the H test (Kerr, 2011) and found a false positive probability of 0.78. The highest degree of orbital modulation consistent with this detection significance depends on the hypothetical pulse shape. As it takes into account multiple harmonics, the H test is least sensitive to a simple sinusoid, so to estimate an upper limit on the fraction of the photons that could be pulsed, we ran simulations of sinusoidally modulated signals. We found that if 35% of the source photons were pulsed, there would be

[1] The Fermi Cicerone is a set of data analysis guides available from `http://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/`.

at least a 95% chance that we would have obtained a result more significant than what we actually observed.

6.4 Rotational modulation

To look for rotational modulation of the γ-rays, we needed a reliable phase prediction for each pulse. To model the complex rotational behaviour of J1023, we assembled a collection of overlapping ephemerides. Each ephemeris spanned 90 days, and a new ephemeris was started approximately every 14 days. We used the `fermi` plug-in for `tempo2` to assign an arrival phase to each photon for each such ephemeris. For each *Fermi* photon and each ephemeris, we assigned a weight between zero and one based on its proximity to the middle of the interval used to fit the ephemeris. Photons outside the interval were assigned zero weight, and photons whose total weight was zero were discarded, limiting the photons used to the times covered by our radio timing observations, MJDs 54766–56114. We computed a weighted mean phase for each photon, effectively producing a piecewise ephemeris predicting rotational phase. The weighting procedure ensured that the predicted ephemeris was continuous and phase predictions were dominated by piecewise ephemerides using observations surrounding the prediction time. We also computed a weighted standard deviation as an estimate of phase uncertainties, plotted in Figure 6–1. All such standard deviations are less than about 10% in phase, and most are below 1%, indicating that ephemeris uncertainties probably do not smear the pulse profile significantly below approximately the tenth harmonic. This collection of piecewise ephemerides, or the spliced combination, is appropriate for short-term phase prediction throughout the span of time it covers.

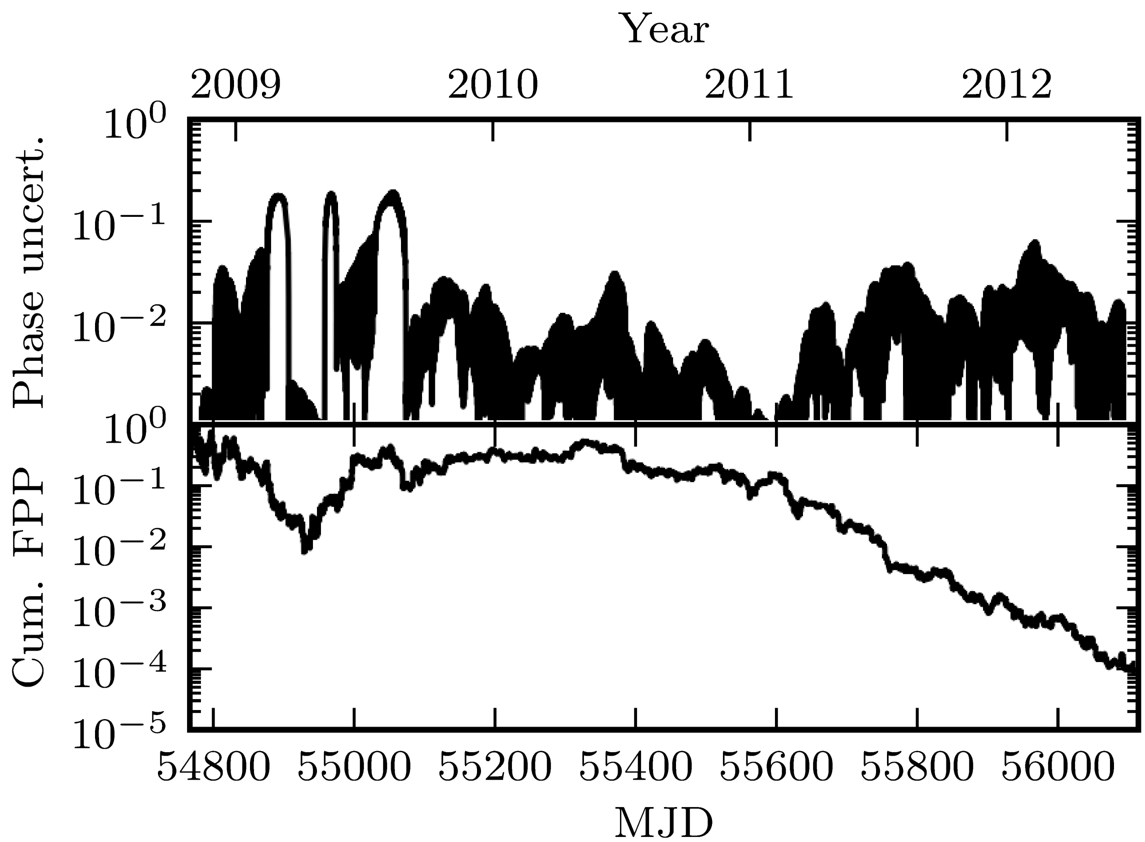

Figure 6–1 Top panel: Estimated phase uncertainty. We computed and plotted the weighted standard deviation of the phase predictions at the time of each *Fermi* photon. Bottom panel: Cumulative false positive probability for pulse phase. Values are based on the weighted H test applied to all usable *Fermi* photons before a given MJD.

The resulting profile is plotted in Figure 6–2, again in the form of a weighted histogram. The H test produces a false positive probability of 9.2×10^{-5}, equivalent to a significance of 3.7σ. The best-fit truncated Fourier series has 2 harmonics and is plotted over the histogram. These significances are single-trial significances, appropriate because we carried out only a single trial, since both the ephemeris and the spectral model were given. We did examine the dependence of the significance on energy and time cuts. A very modest improvement in significance (to 4.2σ) is available by discarding all photons below 400 MeV, possibly because these are more poorly localized (see Section 2.6) and therefore more contaminated by nearby sources. When we compute a cumulative detection significance, see Figure 6–1, it wanders for the first year, during which our ephemeris has the largest uncertainty, and increases steadily after that. We estimated the degree of pulse modulation that would lead to such a false positive probability by running Monte Carlo simulations taking into account the uncertainty in which photons come from the source and assuming that photon arrival phases are scattered independently by the phase uncertainty we calculated for each photon. Assuming a sinusoidal pulse profile affecting 90% of the source photons, the observed significance is greater than that observed in about 50% of the simulations.

6.5 Summary

In summary, using 100 MeV–300 GeV data from the *Fermi* LAT, we detect no orbital modulation of the γ-ray emission, with an upper limit of about 35% on the fraction of photons that can be pulsed. We detect ambiguous evidence (statistically 3.7σ) for γ-ray pulsations at the pulsar's rotational period. Even such

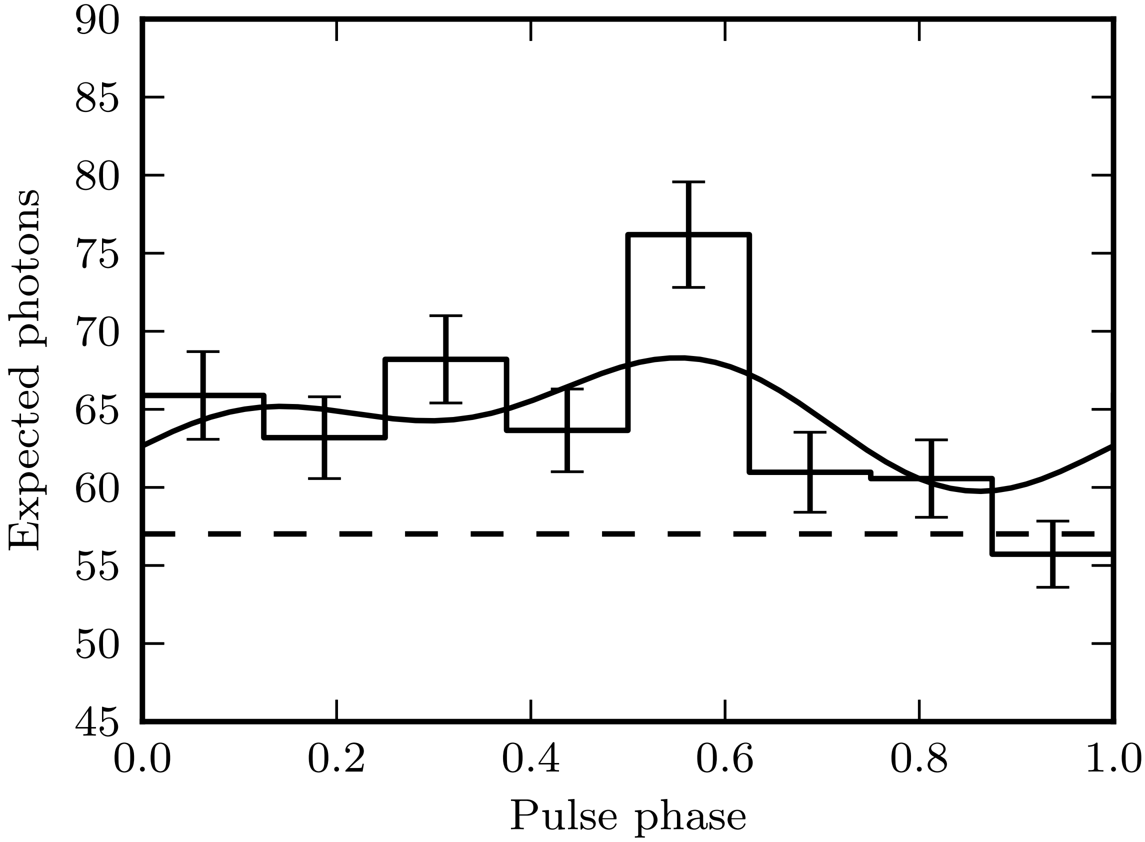

Figure 6–2 γ-ray light curve for J1023 as a function of rotational phase, based on weighted 100 MeV to 300 GeV *Fermi* photons. The smooth curve is a best-fit sum of sinusoids using the number of harmonics (two) suggested by the H test. The uncertainties and background level are computed using Equations 2.2 and 2.3, respectively.

a marginal detection, if real, would suggest that most of the photons participate in the modulation. A continued radio timing campaign would provide the timing data needed to analyze more recent *Fermi* data in search of pulsations, but the improvement in sensitivity will be slow.

CHAPTER 7
Conclusions

This chapter is partly based on the manuscript "Long-Term Radio Timing Observations of the Transition Millisecond Pulsar PSR J1023+0038," to be submitted to The Astrophysical Journal. *This chapter includes the parts of the manuscript that tie together all our observations of J1023, in this paper and in previous papers. This chapter also draws on the discussion in the manuscript to try to understand what they tell us about this and other related systems.*

7.1 Introduction

In light of the 2001 active phase described in Chapter 3, J1023 serves as an example of a system in transition from being an LMXB to being a radio MSP. As described in Section 1.5, it is important to improve our current limited understanding of this transition if we are to explain how the current population of MSPs arose from something like the current population of LMXBs. Our observations of the J1023 system allow us to clarify many aspects of its physics, and at least make evident the puzzling nature of others. Using these results to compare J1023 to other similar systems, including both the classical black widows and the newly named redbacks, should shed light on the transition from LMXBs to MSPs.

7.2 Basic parameters of J1023

Chapter 3 described the J1023 system based on only a few months of radio observations (and its optical history). Since then, its X-ray and γ-ray properties have been reported (Chapter 4, Bogdanov et al., 2011; Tam et al., 2010, Chapter 6), and a precise distance measurement has been carried out (Deller et al., 2012). It seems valuable, then, to summarize the basic system parameters in light of all these data plus our own long-term timing.

Although the orbital parameters vary, as discussed in Section 5.4.5, these variations are tiny in absolute terms. Combined with the mass measurement obtained in Deller et al. (2012), we can describe the system geometry more or less completely. In particular, the system inclination is 42 ± 2 degrees, and the separation between the companion's center of mass and the pulsar is 1.3×10^6 km ($1.9 R_\odot$). The system geometry is illustrated in Figure 5–6, along with the ranges of eclipse phase we observed in our dual-frequency observation (see Section 5.4.1).

This orbital geometry was inferred from the mass measurement in Deller et al. (2012). This mass measurement depended on a key assumption: that the companion filled its Roche lobe. This allowed a measurement of the companion radius, inferred from its spectrum, brightness, and distance, to be converted into a measurement of the Roche lobe size and therefore the system masses. Modelling of the optical light curves of another redback, PSR J2215+5135, constrains its companion's Roche lobe filling factor to be 0.99 ± 0.03 (Breton et al., 2013). While no such fitting process has yet been carried out for J1023, we can make some inferences about the system nonetheless. If the companion underfills its

Roche lobe, the masses and system geometry will be different from those assumed elsewhere in this thesis. Specifically, the masses will be larger than those reported in Deller et al. (2012), and the system will be closer to face-on (the inclination angle i will be less than 42 degrees). That said, it seems unlikely that any pulsar can have a mass larger than $3M_\odot$ (Lattimer and Prakash, 2004); this limits the Roche lobe filling factor (companion radius over Roche lobe radius) to be larger than 0.82, and i to be at least 36 degrees.

The measured spin period derivative, $6.930(8) \times 10^{-21}$ (see Table 5–1), must be corrected for the Shklovskii effect (Shklovskii, 1970) and the effect of acceleration due to the Galactic gravitational field. Both these corrections are computed in Deller et al. (2012); applying them to our period derivative measurement gives an intrinsic period derivative of $5.39(5) \times 10^{-21}$. Using the standard nominal moment of inertia for pulsar calculations ($10^{45}\,\mathrm{g\,cm^2}$) we can calculate a spin-down luminosity of $4.43(4) \times 10^{34}\,\mathrm{erg\,s^{-1}}$. The magnetic field we infer, using the standard formula $B = 3.2 \times 10^{19}\,\mathrm{G}\sqrt{(P/1\,\mathrm{s})\dot{P}}$, is $9.7 \times 10^{7}\,\mathrm{G}$, and the characteristic age $\tau = P/2\dot{P}$ is $4.96(5) \times 10^{9}\,\mathrm{y}$.

7.3 Recent developments relevant to J1023

In addition to the research described in this thesis, there has been substantial other study of J1023. In this section I will summarize the most relevant research for easy reference in later parts of this chapter.

J1023 has recently been the target of intensive multi-wavelength studies. Following the work in Chapter 4, Bogdanov et al. (2011)[1] confirmed the variability and orbital modulation of J1023, modelling the X-ray emission as coming from a shock near L1. Tam et al. (2010) noted that J1023 is a γ-ray source, and indeed J1023 appears in the second *Fermi* catalog (2FGL; Nolan et al., 2012b). Wang et al. (2012) searched for evidence in the infrared of a circumstellar dust cloud due to ejected disc material, but found none. Deller et al. (2012) carried out a very-long-baseline radio interferometry (VLBI) campaign (supported by the timing program desribed in Chapter 5) to measure the position, proper motion, and parallax of J1023, measuring its distance to be 1368^{+42}_{-39} pc and estimating its mass as $1.71 \pm 0.16\ M_{\odot}$.

As of 2013 May, the link between radio MSPs and LMXBs been strengthened by observations of the radio MSP PSR J1824$-$2452I in the globular cluster M28. Radio timing observations of this MSP are consistent with it having a main-sequence companion, perhaps qualifying it as a redback (Papitto et al., 2013a). The LMXB IGR J18245$-$2452, in the midst of an accretion episode including type I X-ray bursts, was recently observed Papitto et al. (2013a) to exhibit coherent X-ray pulsations matching the ephemeris of the radio MSP. The identification of these two sources shows that an object can swing back and forth between a radio MSP and a LMXB. In fact, it appears that the transition from LMXB to MSP can

[1] As second author on Bogdanov et al. (2011), I helped propose for time, model the system, and discuss the implications of the X-ray results.

happen in no more than a few weeks (Papitto et al., 2013b). This object promises to provide invaluable insights into the transition process, but it is much more distant (5.5 kpc) and consequently fainter (50 μJy in radio) than J1023, so our own observations should complement observations of this new system. The rapid transitions of PSR J1824−2452I also support the possibility that J1023 itself will transition back into a LMXB phase in the near future.

7.4 Energetics

Given the distance measurement of Deller et al. (2012), we are able to estimate an energy budget for the pulsar. The total spin-down power, as computed in Section 7.2, is $4.43 \times 10^{34}\,\mathrm{erg\,s^{-1}}$ ($11.5L_\odot$). For comparison, the companion's bolometric luminosity is approximately $6 \times 10^{32}\,\mathrm{erg\,s^{-1}}$ ($0.2L_\odot$), based on the temperature from Thorstensen and Armstrong (2005) and the radius of the Roche lobe.

Thorstensen and Armstrong (2005) estimated that the amount of irradiation received by the companion was consistent with the compact object having an isotropic luminosity of $2L_\odot$ ($6 \times 10^{33}\,\mathrm{erg\,s^{-1}}$); we therefore infer an irradiation efficiency of 10%. That said, the companion only intercepts about 1.3% of any isotropic emission, so that the actual heating of the companion is closer to $10^{32}\,\mathrm{erg\,s^{-1}}$ ($0.03L_\odot$; about 20% of the companion's bolometric luminosity). The average X-ray flux reported in Bogdanov et al. (2011) (choosing a power-law plus hydrogen atmosphere spectral model from among their possible assumptions) corresponds to an X-ray luminosity of $9.3 \times 10^{31}\,\mathrm{erg\,s^{-1}}$ ($0.024L_\odot$), and an efficiency of 0.21%. The γ-ray luminosity deduced from the γ-ray flux reported in the 2FGL catalog (Nolan et al., 2012b), $1.2 \times 10^{32}\,\mathrm{erg\,s^{-1}}$ ($0.3L_\odot$), is 3% of the nominal spin-down power.

Comparing these efficiencies to those of the general population of rotation-powered pulsars reported in Vink et al. (2011), we find the X-ray efficiency of J1023 is three times the typical value, but within the scatter (though their data include no MSPs; they report on point-source plus nebular emission, and since Archibald et al. (2010) and Bogdanov et al. (2011) detect no extended emission from J1023, if there is nebular emission it is included in the point source values given above). The γ-ray efficiency of J1023 is low but not unreasonably so compared to those reported for MSPs in Abdo et al. (2010), including MSPs with higher and lower spin-down powers; J1023's γ-ray efficiency similarly fits in with those reported by Espinoza et al. (2013).

7.4.1 Irradiation

We can also investigate the process by which J1023 irradiates its companion. Takata et al. (2012) suggest that γ-ray emission from an active radio pulsar is responsible for the heating of the near side of the companion in many systems, and in J1023 in particular. We find that in J1023 the γ-ray luminosity, if isotropic, is less than that needed to explain the heating of the companion (Thorstensen and Armstrong, 2005). One normally expects a strong pulsar wind consisting of a relativisic pair plasma to carry most of the pulsar's spin-down power. This is therefore a natural candidate to produce irradiation, either directly if it (or heavy ions entrained in it) can reach the companion's surface or indirectly through heating by emission from a shock powered by the pulsar wind.

In fact, Bogdanov et al. (2011) argue that in J1023 there is an X-ray-emitting shock near the first Lagrange point (L1). If this shock provides the X- or γ-rays that heat the companion, then from the geometry we should expect the companion to intercept roughly half of the total X- or γ-ray luminosity of the shock. If so, then either the X- or γ-ray luminosities would be sufficient to explain the companion's heating. While the observed X-ray orbital variability strongly suggests that a substantial fraction of the X-rays come from this shock, the lack of detectable γ-ray variability suggests that only a small fraction of the γ-rays can come from close to the companion (or we would expect eclipses). Nevertheless the γ-ray luminosity is sufficient that even if only a small fraction come from the shock, γ-rays could be responsible for heating the companion. For the original black widow pulsar, PSR B1957+20, such an intrabinary shock was predicted by Arons and Tavani (1993), and Stappers et al. (2003) examined *Chandra X-ray Observatory* data looking for evidence of one. Specifically, in line with the predictions, they observe a peak in the X-ray brightness (due to Doppler boosting) immediately following the radio eclipse; the same effect was also observed for J1023 by Bogdanov et al. (2011). It therefore seems quite plausible that the companion is being heated by X-rays generated in an intrabinary shock near L1 powered by the particle wind from the pulsar.

In addition to the orbital modulation of the X-ray emission, we note in Chapter 4 (and Bogdanov et al. (2011) confirm) that J1023 shows substantial orbit-to-orbit variability. We do not expect much variability from the MSP in J1023, either in X-ray emission or in particle wind flux. It therefore seems likely that the origin of the orbit-to-orbit X-ray variability comes from the other side of the shock: from the companion. If the amount or magnetization of the material leaving the companion is varying, then an unusually large variation might push material into the pulsar's Roche lobe, form a disc, and tip the system over into another active phase. In light of the fact that the irradiation of the companion — equal in amplitude to 20% of the companion's bolometric luminosity — seems likely to be coming from this shock, it seems that a feed-forward effect might be possible: if an increase in material brightens the shock, the shock will heat the face of the companion more, driving more material off the companion and into the shock. It is equally possible that an increase in material might dim the shock, but nevertheless the shock variability and the corresponding variations in irradiation may well be linked to the system's active phases.

7.4.2 Ejection

Energetics can also serve to place an upper bound on how much material can be leaving the system, through evaporation or through Roche lobe overflow followed by expulsion from the system. If the pulsar's entire spin-down luminosity were used accelerating material to the system escape velocity, the companion could be losing $4 \times 10^{-7} M_{\odot}\,\mathrm{yr}^{-1}$. Adding the companion's bolometric luminosity — driving a stellar wind, say — provides only a small increase. On the other hand, the companion is not much affected by heating, so we should not expect evaporation to dominate; we would expect Roche lobe overflow to be concentrated in the plane of the system, so it should intercept only a modest fraction of the pulsar's spin-down luminosity. Nevertheless, this upper bound is probably not appropriate for understanding the average mass loss rate of the companion, which we expect to be set by whatever mechanism drives Roche lobe overflow. In the short term, though, this upper limit is of particular interest: the system PSR J1824−2452I appears to transition rapidly (within three weeks) from an LMXB to a radio MSP. If this involves clearing an accretion disc from the system, then there cannot have been more mass in the accretion disc than can be ejected at this maximum rate over the course of three weeks. Different systems can obviously not be compared so directly, but if in the future J1023 undergoes an LMXB phase, the timescale of its transition back to a radio MSP may set a limit on the total mass of the accretion disc it can contain.

7.5 Ionized material in the system

It is clear from the eclipses, shorter signal disappearances, and DM excesses that there is much ionized material in or around the J1023 system. The nature and location of this material, however, are unclear.

The main eclipse is the most consistent feature in our radio observations. It is clear from Figure 5–6 that the material lies far outside the Roche lobe, and in fact the eclipsing material cannot be much closer to the companion than it is to the pulsar. It may be bound in the companion's magnetosphere, or it may consist of a wind flowing out from the companion. It should also be noted that since the eclipse region is larger at lower frequencies, the density of the eclipsing material must fall off gradually, so that if the eclipse region is bounded by a shock it must encompass the regions eclipsed at all frequencies.

7.5.1 Magnetically supported material

If the eclipsing material is trapped in the companion's magnetosphere, the latter must contain a strong enough magnetic field to resist the pulsar wind. The pulsar wind pressure at the companion's distance (assuming that the pulsar wind is isotropic and its power equals the spin-down power) is $7\,\mathrm{dyne\,cm^{-2}}$; the magnetic field required to produce an equal pressure is $13\,\mathrm{G}$. Whether such a field can be produced by the companion is unclear.

The companion appears to be a G star (Thorstensen and Armstrong, 2005). Since the irradiation produces only $\sim$400 K difference in temperature between the sides of the companion (the companion's mean temperature is $\sim$5700 K; Thorstensen and Armstrong, 2005), we expect it to play a fairly minor role in the companion's internal processes. Nevertheless, we have several reasons to think the companion is not a typical G star: it has presumably become helium-enriched through mass transfer, it has about 1.8 times the radius of a typical $0.24M_\odot$ main-sequence star (Feiden and Chaboyer, 2012), and it is not even spherical (since it fills or nearly fills its Roche lobe). If we assume regardless that the companion has a convective layer similar to that of solar-type stars, we should expect a dynamo process to generate a magnetic field. Since the companion is presumably tidally locked (given the viscous dissipation implied by rotating material filling the non-spherical Roche lobe), its rotational period is only 0.2 days (equatorial velocity $110\,\mathrm{km\,s^{-1}}$), and such rapidly rotating stars tend to have stronger magnetic fields than ordinary solar-type stars (Berdyugina, 2005). Near the Sun we find a magnetic field on the order of $10\,\mathrm{G}\,(R/R_\odot)^{-2}$ (Banaszkiewicz et al., 1998). The

J1023 system is small enough that the eclipsing regions are within a few solar radii, so it seems plausible that magnetic pressure could support a bubble against the pulsar wind. Relatively tenuous material leaving the surface of the companion would then stream along field lines and might produce the eclipses we observe. Filamentary structures in the companion's magnetosphere might also account for the short-term disappearances of signal; given their tens-of-seconds duration and the velocity of the pulsar on the plane of the sky, they should be hundreds to thousands of kilometers across. In short, it seems possible that a stellar wind from the companion could be supported against the pulsar wind by the magnetic fields that thread it.

7.5.2 Ram-pressure-supported material

On the other hand, we also know that solar-type stars often have substantial winds. This star in particular fills or nearly fills its Roche lobe, so material is readily removed from its surface. The ram pressure of such a wind might sustain a region in which material from the companion could produce eclipses. (If the stellar wind is supported against the pulsar wind by the magnetic fields that thread it, then we return to the model in the previous section.) Bednarek and Sitarek (2013) discuss pulsar and companion winds in redback and other systems, in the interest of understanding bow shocks these systems may produce. They do provide a description of the wind ram pressure balance within the binary system, parameterized by η:

$$\eta = \frac{L_p/c}{\dot{M} v_w},$$

where L_p is the pulsar wind luminosity, $\dot{M}$ is the mass loss rate of the companion, and v_w is the velocity with which the companion expels its wind. If we assume that the $\sim$0.15 pc cm^{-3} of post-eclipse excess DM we noted in Section 5.4.2 comes from a region of length roughly the orbital separation, we obtain an electron density of 4×10^6 cm^{-3}. If we assume also that the wind velocity is roughly the escape velocity the companion would have if it were not in a binary, we obtain $\dot{M} = 7 \times 10^{-13} M_\odot\,\mathrm{y}^{-1}$ and $\eta = 700$. This implies that the pulsar's wind strongly dominates the companion's wind; in fact the companion's wind only reaches a hyperboloidal region with half-opening angle $\psi = 2.1(1 - \eta^{-2/5})\eta^{-1/3}$ (radians; Bednarek and Sitarek, 2013), which in this case is 13°. If this is the case it is hard to see how material from the companion could reach far enough out of the plane to produce any eclipses at all.

If we reverse the process and ask what $\dot{M}$ would produce $\eta = 1$ for this v_w, we find an $\dot{M}$ of $5 \times 10^{-10} M_\odot\,\mathrm{y}^{-1}$. This would require a much higher electron density, 2×10^9 cm^{-3}, which limits the eclipse region to an implausible thickness of only 2000 km. Thus if the eclipse region is sustained by the companion's wind pressure, it would appear that the wind velocity is $\sim$30 times higher than the companion's escape velocity (which is not very different from the escape velocity from the binary system as a whole, or the companion's orbital velocity). It is unclear why the companion should have such a strong wind, since irradiation is mild and by assumption the wind is occurring in a region separated by a shock from the pulsar wind. Measurements of the DM constrain only the density of free electrons; a much higher wind density might be possible if the wind contains a large fraction of neutral material. Sufficient amounts of neutral material would absorb soft X-rays in a way inconsistent with the X-ray observations of Chapter 4 or Bogdanov et al. (2011); these observations limit the total column density to $\lesssim 10^{21}\,\mathrm{cm}^{-2}$, compared to the DM-inferred in-system electron column density of $\sim 5 \times 10^{17}\,\mathrm{cm}^{-2}$. It is therefore consistent with these observations to have an ionization fraction as little as 0.05% and therefore a mostly neutral wind. It is not clear, though, that it is plausible that the wind could be so weakly ionized while in the vicinity of an X-ray source emitting $0.024 L_\odot$, not to mention that the winds of solar-type stars are almost completely ionized (Cranmer, 2008). It therefore seems difficult to explain how ram pressure could support the companion's wind against the pulsar's wind.

7.5.3 Roche-lobe-overflowing material

There is a third possible source for eclipsing material. In PSR J1740−5340 (NGC 6397A), a globular-cluster redback, the Hα line profile suggests material streaming out from L1, trailing behind the companion (Sabbi et al., 2003). Since J1023 likely had an accretion disc in 2001, it seems very plausible that some small amount of Roche lobe overflow may be continuing in this system as well. It is, however, difficult to explain the observed eclipse behaviour using only such material. While the 1400-MHz eclipse is somewhat asymmetrical, it definitely begins before L1 crosses the line of sight at orbital phase 0.25. The 350-MHz eclipse region is even larger and symmetrical. Finally, it is unclear how such escaping material could come to be so far above the plane of the system as to affect our line of sight.

7.5.4 Puzzles

There remain problems with all these eclipse models. In particular, Bogdanov et al. (2011) modelled the X-ray variability as the eclipse of an emitting region by the companion and concluded that most of the X-rays originate from a shocked region near the L1 point or on the pulsar-facing surface of the companion. It is difficult to explain how such a shocked region could be so close to the companion if the eclipsing material, necessarily on the companion side of the shock, reaches as far out as we see in the low-frequency eclipses. Bogdanov et al. (2011) argue that the pulsar wind is concentrated in the orbital plane, but it would require very strong beaming to push the shock front down close to the companion's surface within the plane but allow the shock to reach so far around the pulsar above the plane. If the eclipsing material is bound in the star's magnetosphere, this magnetosphere may be anisotropic, and in particular, less extensive around the star's equator. Or, if the eclipsing material is in the form of a wind from the companion, equatorial focusing of the pulsar wind may produce a large η in this plane and a much smaller value out of the plane.

7.6 Orbital variations

Orbital period variations have been observed in many of the MSP systems that show the effects of ionized material around the system. In the classic black widow systems, this seems to take the form of slow, quasi-periodic orbital period variations (Arzoumanian et al., 1994; Lazaridis et al., 2011). The redbacks and some of the black widows seem to show more complicated orbital period variations (e.g. Hessels et al. 2013, in prep), and it has been suggested that the systems whose companions have higher Roche-lobe filling factors also show stronger orbital period variations (Breton et al., 2013). The accreting millisecond X-ray pulsar SAX J1808.4−3658, which is thought to contain a (not yet detected) radio pulsar in quiescence (Burderi et al., 2003; Iacolina et al., 2009), also shows evidence for a nonzero second derivative of orbital period, possibly indicating that it too undergoes orbital period variations (Patruno et al., 2012). J1023 fits into this picture, since it appears to have a high Roche lobe filling factor and also substantial orbital period variations.

If the companion in J1023 has a substantial magnetic field, it may undergo magnetic drag that slows it slightly below corotation — in fact, this is the "magnetic braking" mechanism thought to drive the long-term orbital evolution, and hence accretion, in such systems (Patruno and Watts, 2012). Beyond any possible differential rotation this might induce, given the companion's non-spherical shape, we would expect substantial and complex mass motions within the companion. If these are coupled to the orbital motion, orbital period variations might result.

7.6.1 Gravitational quadrupole coupling

One specific mechanism, gravitational quadrupole coupling (GQC), by which motions within the companion might be coupled to the orbital motion was set out by Applegate (1992) in order to understand close optical binaries such as Algol. The model was applied to the PSR B1957+20 system (the original black widow) by Applegate and Shaham (1994). In this model, the companion would have a surface layer rotating at a different speed from the underlying layers. The magnetic field within the companion would produce a torque coupling the layers, and variations in this magnetic field would then vary the rotation of the outer layer; if the outer layer spins more rapidly, the star becomes more oblate, while if it slows the star becomes less oblate. This change in the mass distribution within the companion, specifically in the quadrupole moment of the mass distribution, would then change the orbital period. Normally, such changes in the differential rotation require substantial energy transfer, and Applegate (1992) argues that this requires a corresponding fractional change in luminosity of the star of about 10%. Such variations are observed in a number of non-degenerate binary systems, and the luminosity variations are correlated in time with the orbital period variations as predicted by the model (Applegate, 1992).

In J1023, we see orbital phase variations on the order of 1 s. The apparent random wander in orbital period seems to produce this variation on the timescale of a few months, though we see no evidence for the quasi-periodicity described in Applegate (1992) and seen in many systems with orbital period variations (for example Applegate, 1992; Arzoumanian et al., 1994; Lazaridis et al., 2011). The lack of visible periodicity could come from a timescale shorter than we can resolve, so that we would see only the net result of more than one quasi-periodic cycle. However, applying the calculations below to such a short modulation period predicts a tremendous energy flow within the companion, implying luminosity variations over a hundred times the bolometric luminosity of the companion. It should be said that in Applegate (1992) the measured parameters of the source RS CVn imply luminosity variations almost ten times the luminosity of the star, so deviations from the model are known to occur. Nevertheless it seems implausible in this model that the timescale could be shorter than our observation cadence so as to hide the quasi-periodicity.

For the calculations that follow we will assume a time scale for the modulation of 200 days, plausible in light of Figure 5–11 but chosen to obtain reasonable luminosity variations. We will also assume, a necessity in this model, that the variations are quasi-periodic with period P_{mod}. The amplitude of the phase variations $\Delta t \sim 1\,\mathrm{s}$ implies fractional orbital period variations $\Delta P/P$ on the order of $2\pi\Delta t/P_{\mathrm{mod}} = 4 \times 10^{-7}$. Following Applegate (1992) and Applegate and Shaham (1994), we assume that the participating shell contains about a tenth of the companion's mass. They link the orbital period variations to changes in the quadrupole moment ΔQ (specifically increases and decreases in the oblateness of the companion) by

$$\frac{\Delta P}{P} = -9\left(\frac{R_c}{a}\right)^2 \frac{\Delta Q}{M_c R_c^2},$$

where M_c, R_c and a are the mass and radius of the companion and the semimajor axis of the orbit, respectively. Changes in the quadrupole moment are dominated by changes in the rotation rate of the outer shell, so they write

$$\Delta Q = \frac{2}{9}\Omega\Delta\Omega\frac{M_s R_c^5}{G M_c},$$

where M_s is the mass of the shell and Ω is its angular velocity. Using this we infer a fractional change in the shell's angular velocity on the order of 2×10^{-4}. Again following the above works, we assume that the companion is slowed below corotation by the same fractional amount, 2×10^{-4}. In this model, the angular momentum transfer between layers is produced by a torque coming from the subsurface magnetic field, acting quasi-periodically over the modulation period P_{mod}. They estimate the strength of this subsurface magnetic field B_s as

$$B_s^2 \sim 10 \frac{GM_c^2}{R_c^4} \left(\frac{a}{R_c} \right)^2 \frac{\Delta P}{P_{\mathrm{mod}}}.$$

In J1023 this calculation finds $B_s \sim 4 \times 10^4\,\mathrm{G}$, greater by an order of magnitude than those in the systems considered in Applegate (1992). While this subsurface field is difficult to constrain observationally, most of the systems considered by Applegate (1992) have values of this and other parameters that are surprisingly similar to each other, so such a large value for J1023 is unusual. On the other hand, the observations summarized in Thorstensen and Armstrong (2005) constrain luminosity variations to less than about 0.04 magnitudes, which is about 4%. The GQC model predicts that the transfer of kinetic energy between layers should be accompanied by luminosity variations on the same scale, in this case about $1.5 \times 10^{31}\,\mathrm{erg\,s^{-1}}$; by our selection of P_{mod} this is about 2% of the companion's luminosity, below the observational limits on such variations. Nevertheless, this low upper limit is a little puzzling in this model, since it implies that the companion is less efficient than usual at converting differential rotation into luminosity, in spite of a higher-than-normal subsurface magnetic field. Most systems discussed in Applegate (1992) undergo luminosity variations of more than about 10%, which would have been detected by Thorstensen and Armstrong (2005) unless they were on a much longer timescale than the orbital period variations.

7.6.2 Possible relation to active phases

If the orbital period variations are connected to mass motions within the companion, whether through the mechanism of Applegate (1992) or some other process, they may be connected to J1023's active phases. In particular, if the star very nearly fills its Roche lobe, shape distortions, whether global or local, might easily cause substantial amounts of material to spill out of the Roche lobe through L1. Going back to the GQC model, the fractional variations in the equatorial radius we expect are on the order of 10^{-4} (proportional to the fractional change in the shell's angular velocity). Substantial enough amounts of material spilling through L1 can overcome the pulsar wind pressure and penetrate the light cylinder; once this occurs, the process by which the wind is generated becomes ineffective, and an accretion disc can continue to exist even if the rate of mass transfer drops somewhat. This bistable mechanism, proposed by Burderi et al. (2003), could be made to switch states by random or ordered changes within the companion.

One particularly interesting possibility is connected to the companion's magnetic field. If this magnetic field is produced by a dynamo process, it may undergo periodic field reversals, as we observe in both the Earth and the Sun. The solar magnetic field reversals lead to the 11-year sunspot cycle, which affects the solar wind and phenomena like coronal mass ejections. If this occurred in J1023, it could have several possible effects. If the radio eclipses are the result of a stellar magnetosphere, a magnetic field reversal might allow the pulsar wind to penetrate much closer to the companion, possibly heating its surface and producing Roche lobe overflow. On the other hand, the magnetic field may drive the orbital period variations, both through drag producing differential rotation, and through the coupling it produces between layers of the companion. In this case, magnetic field reversals might be accompanied by unusually large distortions in the companion's shape, which could lead to Roche lobe overflow. Magnetic field reversals might also explain the timescale of J1023's active phase, which lasted a year and has not recurred in more than a decade. This timescale for active phases and recurrence times is within the range observed among accreting millisecond X-ray pulsars (Patruno and Watts, 2012), although these systems may have a persistent disc whose state may control accretion episodes.

If the orbital period variations are related to the active phase, it is possible that monitoring these orbital period variations might give some warning of the next active phase. Patruno et al. (2012) suggest that the orbital period variations in SAX J1808.4−3658 may in fact modulate the accretion episodes, producing more intense accretion when the orbit (and thus the companion's Roche lobe) is shrinking and less intense or no accretion when the orbit is growing. The situation is not exactly analogous, since in SAX J1808.4−3658 the orbital period variations, if quasi-periodic, would be on a timescale of decades, while recurrence times are on the order of two years (the orbital period is currently increasing); in J1023 the recurrence time is at least a decade, and orbital period variations (which are of similar magnitude to those in SAX J1808.4−3658) seem to happen on shorter timescales.

More generally, it seems difficult to produce orbital period variations without the companion being slowed below corotation. This slowing may be due to magnetic braking or some other process, but it would seem to imply loss of angular momentum from the binary over the long term. If this is occurring, then the Roche lobe must be shrinking. If a system has a main-sequence companion that is close to filling its Roche lobe, then it seems that orbital period variations should imply that Roche lobe overflow is continuing unless there is evaporation sufficient to outpace the Roche lobe shrinkage. In any case, if there is some relation between orbital period variations and active phases, and if it can be understood, then it might be possible to compare the orbital period variations seen in J1023 to those seen in other redbacks to gauge which of them may undergo active phases.

7.7 Summary and future work

J1023 is a transition object between LMXBs and radio MSPs: although it is now a radio MSP, it had an active phase in 2001. During this active phase an accretion disc formed, but it appears that matter was expelled from the system rather than accreted onto the surface of the neutron star. Another system, M28I, has now been discovered that shows this transition. These systems probe the poorly understood end of accretion in LMXBs.

J1023 is detected as both an X-ray and a γ-ray source. The origin of the X-rays appears to be partly from the neutron star itself, either the magnetosphere or hotspots on the surface heated by return currents, and partly from a shock close to the companion's surface. The origin of the γ-rays is less clear, but as they do not appear to show much orbital modulation and may be modulated at the pulsar period, we tentatively identify them as coming from the pulsar magnetosphere. In light of the energetics of the system, it therefore seems likely that the companion's irradiation is by X-rays coming from the shock where the pulsar wind meets the companion's wind or magnetosphere. Variations in this shock presumably originate from the companion, and may involve a feed-forward process that, if a runaway occurs, might tip the system into an active phase. That said, in order for an accretion disc to form in J1023, material from the companion would need to pass through the region where the shock currently occurs to form a Keplerian disc within the pulsar's Roche lobe.

J1023 exhibits a variety of radio phenomena that are evidence for binary interaction even in quiescence: eclipses, excess dispersion measure, short-term disappearances of signal, and orbital period variations. The origin of these phenomena is not clear, but a magnetically active companion may be able to explain several: the eclipses may be the result of ionized material trapped in a bubble supported against the pulsar wind by the companion's magnetic field. The excess dispersion measure and disappearances of signal may be due to blobs or tendrils of magnetically trapped material extending outside the primary bubble in the manner of solar prominences. Finally, the orbital period variations may be the result of the mechanism of Applegate (1992), in which this same magnetic field slows the companion slightly below corotation, and in which magnetic fields internal to the companion transfer angular momentum to and from the surface layers, modifying the companion's quadrupole moment and thus the system's orbital period. Since this mechanism operates by changing the companion's shape, variations might also show up in the amount of Roche lobe overflow, and particularly large shape variations might drive the system into an active phase. A magnetically active companion is plausible in light of the fact that the companion resembles a main-sequence star and, with a rotational period approximately equal to the system's 0.198-day orbit, the companion is also a rapid rotator. That said, we do not understand the internal processes of the companion, which is helium-enriched, substantially larger than expected for a main-sequence star of its mass, and not even spherical. Further understanding of the companion would clarify the behaviour of this complex system, might clarify the conditions that led to the system's 2001 active phase, and might hint at whether and when the system will enter another active phase.

Careful optical observations to check for luminosity variations in the companion would help clarify whether the GQC mechanism can plausibly be acting in the system, and if not, might help clarify other possible mechanisms. Multicolour light curve modeling and possibly orbital-phase-resolved optical spectroscopy might clarify how close the companion is to filling its Roche lobe, and asymmetries in the light curve might give some indication of the mechanism by which the companion is irradiated. Further X-ray observations, particularly if high time resolution is available, might make it possible to distinguish X-rays from the pulsar itself from X-rays originating in a shock region, possibly even mapping the shock region more clearly. Great progress could be made if the system were to enter another active phase; in such a case one would want to study J1023 with all the techniques developed to understand LMXBs. It would therefore also be valuable to monitor J1023 with pulsar timing and optical photometry or Hα measurements in order to catch another active phase as early as possible and to know the system's state leading up to its activity. We hope the description of J1023's behaviour we have provided in this thesis will help future observers in carrying out such studies.

APPENDIX A
A search for radio pulsations from the magnetar in Kes 75

This chapter is based on the manuscript "No detectable radio emission from the magnetar-like pulsar in Kes 75" published in The Astrophysical Journal *in 2008 November (Archibald et al., 2008). This paper illustrates pulsar searching techniques similar to those used to discover J1023, as applied to a rather different problem: searching for radio pulsations from a known magnetar. The original abstract appears below.*

The rotation-powered pulsar PSR J1846−0258 in the supernova remnant Kes 75 was recently shown to have exhibited magnetar-like X-ray bursts in mid-2006. Radio emission has not yet been observed from this source, but other magnetar-like sources have exhibited transient radio emission following X-ray bursts. We report on a deep 1.9 GHz radio observation of PSR J1846−0258 with the 100-m Green Bank Telescope in late 2007 designed to search for radio pulsations or bursts from this target. We have also analyzed three shorter serendipitous 1.4 GHz radio observations of the source taken with the 64-m Parkes telescope during the 2006 bursting period. We detected no radio emission from PSR J1846−0258 in either the Green Bank or Parkes datasets. We place an upper limit of 4.9 μJy on coherent pulsed emission from PSR J1846−0258 based on the 2007 November 2 observation, and an upper limit of 27 μJy around the time of the X-ray bursts. Serendipitously, we observed radio pulses from the

nearby RRAT J1846−02, and place a 3σ confidence level upper limit on its period derivative of 1.7×10^{-13}, implying its surface dipole magnetic field is less than 2.6×10^{13} G.

A.1 Introduction

Soft Gamma Repeaters (SGRs) and Anomalous X-ray Pulsars (AXPs) are now well accepted as being different though similar manifestations of "magnetars" – isolated, young ultra-highly magnetized neutron stars whose radiation is powered by their magnetic field (see Woods and Thompson, 2006, for a review). However, there remain some outstanding puzzles in the magnetar picture. One is the connection between magnetars and high-magnetic-field rotation-powered pulsars. There are now known seven otherwise ordinary rotation-powered pulsars having inferred surface dipole magnetic field $B > 4 \times 10^{13}$ G (computed using Equation 1.2; the quantum critical field $B_{\rm QED} = 4.4 \times 10^{13}$ G). Most of these high-magnetic-field pulsars, for example PSRs J1718−3718, J1847−0130 and J1814−1744, have B similar to or greater than those measured for *bona fide* magnetars, yet show only faint X-ray emission, if any (McLaughlin et al., 2003; Kaspi and McLaughlin, 2005; Pivovaroff et al., 2000). The Rotating Radio Transient (RRAT) J1819−1458 also has a field comparable to magnetars, but it exhibits modest X-ray emission (McLaughlin et al., 2007). Gonzalez et al. (2005) suggest that another high-B radio pulsar, PSR J1119−6127, shows evidence for possibly anomalous X-ray emission in the form of a high surface temperature and high pulsed fraction for thermal emission.

"Transient" magnetars, such as XTE J1810−197, are typically X-ray faint but occasionally have major AXP-like outbursts (e.g. Ibrahim et al., 2004). Similarly, the candidate transient AXP AX J1845−0258 is either extremely faint or undetectable in quiescence, but was seen to be at least several hundred times brighter than usual in a 1993 outburst (Vasisht et al., 2000; Tam et al., 2007). Kaspi and McLaughlin (2005) suggest such objects could be related to high-B radio pulsars, noting the similarity of the spectrum of the high-B radio pulsar PSR J1718−3718 to that of XTE J1810−197 in quiescence. This suggestion was supported by the discovery of radio pulsations from XTE J1810−197 after its major outburst (Camilo et al. 2006), albeit with an unusual radio spectrum and unusual radio variability properties. Camilo et al. (2007) report a second magnetar, 1E 1547.0−5408, in outburst with similar radio properties.

Very recently, the proposed connection between high-B radio pulsars and magnetars was given a major boost by the discovery of SGR-like X-ray bursts and a several-month-long flux enhancement from PSR J1846−0258, which was previously thought to be a purely rotation-powered pulsar (Gavriil et al., 2008). This source has a quiescent X-ray luminosity that could be rotation-powered, and has other properties of rotation-powered pulsars such as a pulsar wind nebula (PWN; Helfand et al., 2003; Sanjeev Kumar and Safi-Harb, 2008; Ng et al., 2008) and an unremarkable braking index of 2.65 ± 0.01 (Livingstone et al., 2006). PSR J1846−0258 has a period of 326 ms, an estimated dipole magnetic field of 4.9×10^{13} G, and an estimated spin-down age of 884 years (Livingstone et al., 2006). Its accurate position, obtained with the *Chandra* observatory, clearly associates it with the supernova remnant (SNR) Kes 75 (Helfand et al., 2003). However, no radio emission has yet been detected from the pulsar (Kaspi et al., 1996).

Here we report on radio observations of PSR J1846−0258, obtained fortuitously on the same day as the onset of its observed magnetar-like X-ray behaviour, as well as over a year after this episode. Using these data, we have searched for coherent radio pulsations. We also report on our search for single radio bursts from this source. We find neither radio pulsations nor bursts and set upper limits on both. However we do detect and describe radio bursts from a nearby, unrelated Rotating Radio Transient (RRAT) J1846−02 (McLaughlin et al., 2006).

A.2 Observations and Results

We obtained a deep observation of PSR J1846−0258 with the 100-m Robert C. Byrd telescope at Green Bank, WV, operated by the NRAO[1] . The RRAT J1846−02 (McLaughlin et al., 2006) is estimated to be within two arcminutes of PSR J1846−0258, though its position is uncertain by some seven arcminutes. Since the well-known periods of RRAT J1846−02 (4.4767 s) and PSR J1846−0258 (326.29 ms) are incommensurable (their ratio is 13.720), they are clearly distinct sources. We were also able to analyze several archival observations,

[1] The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

taken with the 64-m Parkes radio telescope in New South Wales, Australia, covering the key period during and just after the the X-ray bursts described by Gavriil et al. (2008). Together these observations constrain radio emission associated with the X-ray bursts on both long and short timescales.

There are two distance estimates in the literature for PSR J1846−0258, 21 kpc (Becker and Helfand, 1984) and $5 - 7.5$ kpc (Leahy and Tian, 2008). Using the NE2001 free electron density model (Cordes and Lazio, 2002), these distances predict a range of dispersion measures (DMs) from 210 pc cm^{-3} to 1441 pc cm^{-3}; the DM through the entire Galaxy in this direction is estimated to be 1464 pc cm^{-3}, although these figures are quite uncertain. The interstellar scattering times predicted by the NE2001 model range from 0.03 ms to 17 ms for a 1.9 GHz observing frequency, and from 0.1 ms to 65 ms for a 1.4 GHz observing frequency.

The intrinsic radio pulse width of PSR J1846−0258 is uncertain; the X-ray pulse is quite broad (Gotthelf et al., 2000), but X-ray pulse morphology can be very different from radio pulse morphology (e.g. as described in Gotthelf and Halpern, 2005; Camilo et al., 2007b, XTE J1810−197 had a broad X-ray profile but a small radio duty cycle). For the (mostly radio) pulsars listed in Manchester et al. (2005) the average duty cycle was about 5%, and the transient AXP XTE J1810−197 was observed with a variable duty cycle with typical value of 2% (Camilo et al., 2007b). Single pulses from the RRATs had lengths from 2 to 30 ms (McLaughlin et al., 2006). We elected to focus our search on a duty cycle of about 1%, that is, pulses of length 3 ms. We ran folding and single-pulse searches (see below) that were sensitive to somewhat shorter pulses, but pulses shorter than the time resolution of our searches would be detected (or not) according to their flux averaged over our time resolution; if the pulse length is $\alpha < 1$ times our search's time resolution, our sensitivity to the peak flux is reduced by a factor of $\sqrt{\alpha}$. Our tools are generally fairly sensitive to pulses longer than the time resolution, so we have tried to choose parameters that allow us to detect a broad range of duty cycles around 1%.

A.2.1 Deep single observation with the Green Bank Telescope

We observed PSR J1846−0258 while pointing at RA $18^{\mathrm{h}}46^{\mathrm{m}}24.96^{\mathrm{s}}$, DEC $-2°58'30.72''$ (J2000, Helfand et al., 2003) for 201 minutes, beginning at MJD 54406.969 (2007 November 2 23:14 UTC), using the Green Bank Telescope. We used the S-band receiver, observing a bandwidth of 600 MHz centered at 1950 MHz, feeding into the SPIGOT pulsar backend (Kaplan et al., 2005), which recorded 1024 channels (only 768 of which were within our bandpass) of 16-bit samples with a time resolution of 81.92 μs. The two linear polarization channels were summed. We performed RFI excision using the program `rfifind` from the software package PRESTO (Ransom, 2001b; Ransom et al., 2002). RFI conditions were mild, requiring only $\sim 4\%$ of the data to be discarded.

We analyzed the data in three ways: by trial folding, by using a Fourier-domain blind periodicity search, and by searching for bright dispersed single pulses.

Trial folding was carried out using the program `prepfold` from PRESTO. We folded the data into 64 pulse phase bins at 1281 evenly-spaced DMs ranging from zero to 4288 pc cm^{-3}. These DM spacings correspond to a shift of a single profile bin over the whole observation time per step. In case the pulsar has a very large or very small duty cycle we also repeated the search with 16 phase bins and 256 phase bins, with corresponding numbers of DM trials and the same minimum and maximum DM. The ephemeris we used for folding was taken from contemporaneous *Rossi X-ray Timing Explorer* (*RXTE*) observations, part of the program described in Livingstone et al. (2006). The timing model we used specifies the pulsar's rotational frequency as

$$\nu(t) = \nu_0 + \dot{\nu}_0(t - t_0) + \ddot{\nu}_0(t - t_0)^2/2, \tag{A.1}$$

where $\nu_0 = 3.064756108(4)$ Hz, $\dot{\nu}_0 = -6.6888(4) \times 10^{-11}$ Hz s^{-1}, $\ddot{\nu}_0 = 4.81(12) \times 10^{-20}$ Hz s^{-2}, and t_0 is MJD 54376.452350. Period and period derivative uncertainties obtained from this ephemeris are insignificant over the course of this observation.

We detected no radio pulsations at or near the trial ephemeris. The GBT observing guide describes a system temperature $T_{\rm sys}$ of approximately 19.5 K for this band and a gain G of approximately 1.75 K Jy^{-1}. Since the band at which we observed, 2 GHz, is more or less free of both ionospheric and tropospheric effects and the GBT receivers are quite stable, this number should be reasonably accurate. We modeled the background $T_{\rm BG}$ as a continuum of 400 K at 408 MHz (Haslam et al., 1982), which we assumed to have a spectral index of -2.6, plus a contribution $S_{\rm SNR}$ from the supernova remnant of 10 Jy at 1.4 GHz with a spectral index of -0.7 (Green, 2006). To compute the upper limit on the pulsar's mean flux, $S_{\rm min}$, we use the expression (e.g. Lorimer and Kramer, 2004):

$$S_{\rm min} = (S/N)_{\rm min}\beta\sqrt{\frac{W}{P-W}}\sigma, \tag{A.2}$$

where $(S/N)_{\rm min}$ is the signal-to-noise threshold at which we would certainly have detected the pulsations, β is a unitless factor describing quantization losses, P is the pulsar's period, W is the time per period during which the pulsar is on, and σ is the noise RMS amplitude, given by:

$$\sigma = \frac{(T_{\rm sys} + T_{\rm BG})/G + S_{\rm SNR}}{\sqrt{n_p t \Delta f}}, \tag{A.3}$$

where n_p is the number of polarizations added (two in our case), t is the observation length (201 minutes), and Δf is the bandwidth (600 MHz). Since the Spigot uses a three-level digitzer, we used $\beta = 1.16$ (Lorimer and Kramer, 2004). For this calculation, we assumed the duty cycle was 1/64, and that we would have detected any pulsar with a peak more than four times the RMS noise per bin. This gives an upper limit of 4.9 μJy. Longer duty cycles give higher upper limits, up to 43 μJy (assuming a square-wave profile).

We carried out the blind periodicity search using the program `accelsearch`, again from the PRESTO toolkit. We produced a dedispersed time series with time resolution 0.32768 ms for 2000 dispersion measures spaced by 2 pc cm^{-3} from zero to 4000 pc cm^{-3}. We searched only for unaccelerated pulsations, summing up to sixteen harmonics (the maximum supported by `accelsearch`). We detected no pulsations. We examined all candidates with a signal-to-noise ratio above 6. Applying Equations A.2 and A.3 as we did for our coherent pulsation search, and assuming a duty cycle of 1/32 (since the use of sixteen harmonics smears the folded profile by this much) we place an upper limit on coherent pulsations at any frequency of 10 μJy.

We carried out the single pulse search using `single_pulse_search.py`, again from the PRESTO toolkit. This software operates on the same dedispersed time series described above and computes running boxcar averages at a range of widths from our dedispersed sample time of 0.32768 ms up to 9.8304 ms. Statistically significant detections are recorded (after some automated sifting to reduce the effects of radio frequency interference and receiver gain changes), and the results are plotted in Figure A–1. The peak-flux threshold for detection of single pulses depends on the length of the pulses; see Figure A–2. The single-pulse searching code we used removes certain strong pulses which it interprets as receiver gain changes. As a result, the last and brightest pulse was erroneously removed from the list of single-pulse detections between a DM of 170 pc cm^{-3} and a DM of 300 pc cm^{-3}. It was nevertheless present in our dedispersed time series and we used it in our timing analysis.

The signature of a single bright astrophysical pulse should be a collection of single-pulse detections well above the DM = 0 axis. Indeed we find a number of such pulses, at a DM of approximately 237 pc cm^{-3}. Closer examination reveals twelve bright single pulses, all clustered around this DM. A number of other pulse detections were observed, but all appear to be either terrestrial RFI (groups of detections at the same time and strongest at zero DM) or noise.

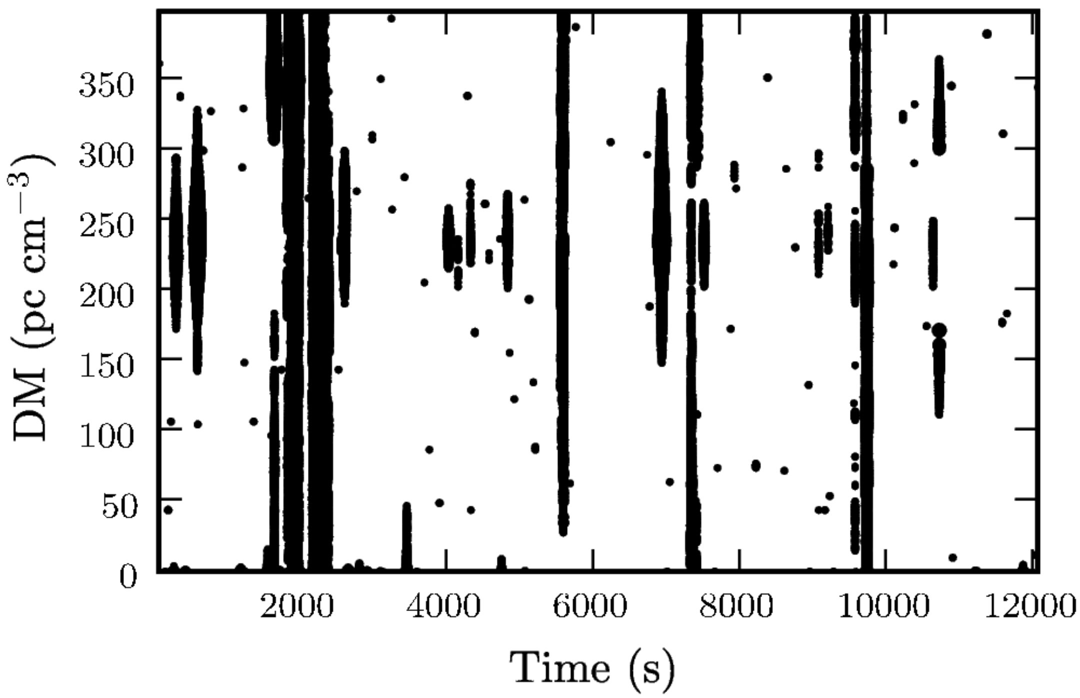

Figure A–1 Single-pulse search results plot from our three-hour Green Bank Telescope observation. Each single pulse detection above a threshold of 6σ is plotted as a circle whose diameter indicates the significance of the detection. Note the series of pulses from the RRAT J1846−02 around a DM of 237 pc cm^{-3}. The final, brightest, pulse from the RRAT, just before 11000 s, appears only below a DM of 170 pc cm^{-3} and above DM 300 pc cm^{-3} because it was so strong it was erroneously identified as a receiver gain change in dedispersed time series closer to the correct DM. Vertical groups of strong detections extending down to zero DM are RFI.

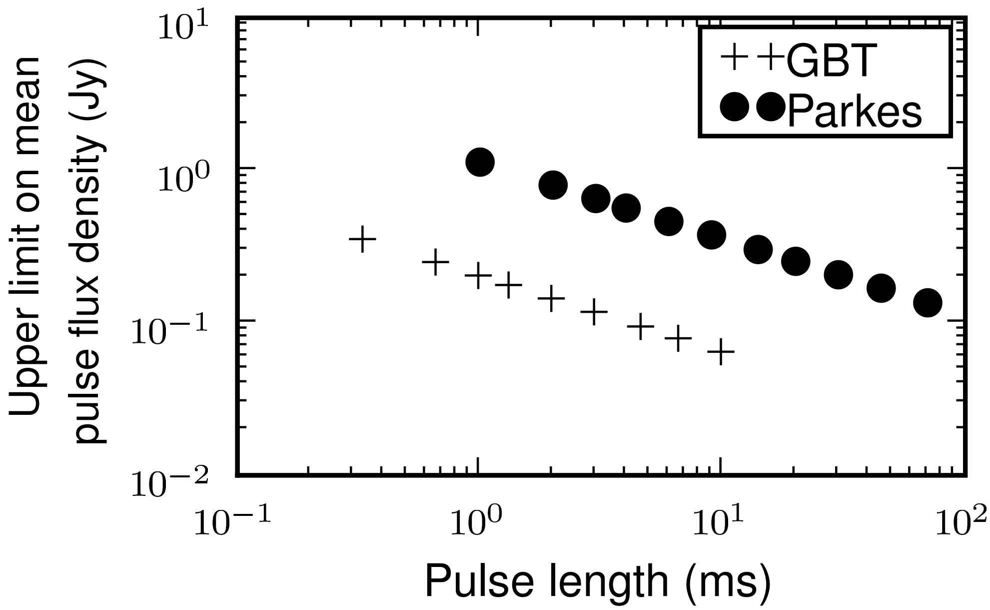

Figure A–2 Upper limits on the mean flux density for single pulses of different lengths for our GBT (Sec. A.2.1) and Parkes (Sec. A.2.2) observations.

When we fold the arrival times of all single-pulse candidates (apart from certain obvious RFI) according to our ephemeris for the X-ray pulsar in Kes 75, we find that they fall at random phases; in fact a Kuiper test gives a probability of 0.22 that a uniform distribution would give rise to arrival times more unevenly distributed in phase than this. However, the RRAT J1846−02, described in McLaughlin et al. (2006), falls within the GBT's 7′ beam when it is pointed at PSR J1846−0258. The DM reported in McLaughlin et al. (2006) is 239 pc cm^{-3}, which closely matches that of our single pulses. Moreover, if we fold the single-pulse arrival times at the reported period of the RRAT, 4.476739(6) s, we find that they fall within four milliperiods of the same phase. We therefore infer that the only significant cosmic single pulses in our observation are from RRAT J1846−02.

Since we detected twelve bright pulses from the RRAT J1846−02 over the course of approximately three hours, we can compute a period for the RRAT based on the timing of single pulses. We selected the single brightest pulse, smoothed it by convolution with a von Mises distribution (e.g. Mardia, 1975) of full width at half maximum 10 ms, and used it as a template. The barycentric arrival time of each pulse was estimated by Fourier-domain cross-correlation with this template. We then used the published period to compute the number of turns between each pair of pulses, and adjusted the period and starting phase to minimize the root-mean-squared residual phase.

While we would expect rather small formal errors on the arrival times for data of this quality, we do observe about 4 ms residual jitter, possibly due to pulse shape variations. We have used the RMS variation of the residuals as an estimate of the uncertainty on each arrival-time measurement. Taking this into account, we obtain a barycentric period estimate of 4.4767435(2) s at epoch MJD 54407. Subtracting this from the reported period of 4.476739(3) s (epoch MJD 53492, McLaughlin et al., 2006) and dividing by the elapsed time gives a period derivative estimate of $(5.5 \pm 3.8) \times 10^{-14}$, implying a 3σ upper limit on the period derivative of 1.7×10^{-13}. Using Equation 1.2 this implies an upper limit on the surface dipole magnetic field of 2.6×10^{13} G.

A.2.2 Observations with the Parkes telescope near the bursting epoch

As part of a program to monitor RRAT J1846−02, three one-hour observations including both it and PSR J1846−0258 were taken at MJDs 53886.64, 53923.54, and 53960.45. Bursts from PSR J1846−0258 were detected in X-ray observations at MJD 53886.92–53886.94 and 53943.46 (Gavriil et al., 2008); fortuitously, one radio observation was taken only six hours before the first of these bursts.

The radio observations were taken with a 256 MHz bandwidth centered at 1390 MHz. Each observation was 60 minutes long, and was recorded with the centre beam of the Parkes Multibeam receiver. They were one-bit digitized with a time resolution of 0.1 ms and 512 spectral channels. We analyzed them in the same three ways as for the GBT data — folding at the known period, blind periodicity searching, and single-pulse searching — using the same software tools. Unfortunately, because of the timing anomaly that the source underwent around this time (Gavriil et al., 2008), we were not able to produce a phase-coherent timing solution from contemporaneous *RXTE* data. We were able to obtain period estimates from periodograms of the *RXTE* data, and we searched a range of periods in the radio data to be certain of including the true period.

We folded each observation using 64 phase bins at 129 periods centered at the peridogram frequency of the nearest *RXTE* observation and spaced over 1.4 μs, and 1025 DMs spaced from 0 to 4008 pc cm^{-3}. As before we also folded with 16 and 256 bins. The center frequencies we used for this folding are 3.067684 Hz, 3.067457 Hz, and 3.067240 Hz, from RXTE observations at MJDs 53886.91, 53920.95, and 53955.57 for the Parkes observations on MJD 53886.64, 53923.54, and 53960.45 respectively. We saw no coherent pulsations. Using the same background figures and duty cycle assumptions as above, and telescope parameters obtained from the Parkes Radio Telescope Users' Guide ($T_{\rm sys} = 23.5$ K, $G = 0.67$ K Jy^{-1}, $\beta = 1.25$), we estimate a flux upper limit of 27 μJy for each observation.

Our blind searches were carried out as above, and we set an upper limit of ~ 58 μJy on all coherent pulsations from the source. For single-pulse search upper limits, see Figure A–2. We detected single pulses from RRAT J1846$-$02 in one of these data sets. Only one pulse in our three-hour GBT observation was brighter than our threshold for detection in the Parkes observations, so a single detection in the three one-hour Parkes observations is not unexpected.

A.3 Discussion

We did not detect radio emission from PSR J1846$-$0258, in spite of a deep GBT observation and a contemporaneous timing solution. Our upper limit of

4.9 μJy is a substantial improvement over that published in Kaspi et al. (1996), which quotes an upper limit at 1520 MHz of 100 μJy.

Two distance estimates for the pulsar/supernova remnant system are found in the literature: 21 kpc (Becker and Helfand, 1984), and more recently 5.1 kpc–7.5 kpc (Leahy and Tian, 2008). This newer result is based on more recent HI and ^{13}CO maps. The smaller distance yields a smaller diameter for the remnant, consistent with the indications that PSR J1846−0258 is very young (Livingstone et al., 2006). The smaller distance also yields a smaller X-ray luminosity for the pulsar, which had previously appeared to be unusually high (Helfand et al., 2003, 4.1×10^{35} erg s^{-1}, second only to the Crab). We will assume this more recent measurement is correct.

The transient AXP XTE J1810−197 was detected with 1.4 GHz radio flux densities peaking at 1500 mJy and fading to 10 mJy over the course of about 6 months (Camilo et al., 2007), with a flat enough spectrum that flux densities at 1.4 GHz and 1.9 GHz are comparable (Camilo et al., 2006). Its estimated distance is 3.5 kpc (Camilo et al., 2007). If the source we observed had the same luminosity at the far end of the Leahy and Tian (2008) estimated distance range, 7.5 kpc, we would expect a flux density ranging from 320 mJy down to 2 mJy. The transient AXP 1E 1547.0−5408 was observed to have a flux density $\sim$ 3 mJy at 1.4 GHz (Camilo et al., 2007a), rising with frequency; the source is at an estimated distance of 9 kpc (Camilo et al., 2007a). If PSR J1846−0258 had the same luminosity we would expect a flux of 4 mJy. All of these figures are significantly higher than our GBT detection threshold of 4.9 μJy, so we conclude that if PSR J1846−0258 was emitting radio pulsations during our observations, they must either be much weaker than those observed from transient magnetars, or they must be beamed elsewhere. A recent paper by Ng et al. (2008) computes an angle of 60° between the line of sight and the spin axis of the pulsar, and they combine this with a tentative suggestion that in order to obtain the observed braking index the magnetic inclination should be approximately 9° (Melatos, 1997; Livingstone et al., 2006). This would imply that our line of sight is, at its closest, 54° from the magnetic pole, making it unlikely that we would be in the pulsar beam. However, in light of the evidence given by Camilo et al. (2007) that the X-ray and radio emission beams of XTE J1810−197 are nearly parallel, the fact that we see X-ray emission from PSR J1846−0258 indicates that if it behaved like XTE J1810−197, any radio emission would likely be beamed in our direction as well.

Thus it appears likely that if PSR J1846−0258 is emitting radio pulsations, they are much weaker than those emitted by the known radio-emitting transient AXPs. However, since our GBT observation was taken some eighteen months after the X-ray bursting activity, it is possible that the radio emission had faded by the time of our observation. The upper limits obtained from our Parkes observations — 21 μJy — are also much less than we would have expected to observe from the known transient AXPs, so we can also constrain the brightness for the first two months after the X-ray bursts began. In particular, the X-ray bursts (four were observed in a one-hour observation, taking place only six hours after the first of our Parkes observations) are probably not accompanied by radio bursts.

Leaving aside the AXP-like behaviour of PSR J1846−0258, its radio emission appears to be very faint. The young pulsar PSR J0205+6449 in the supernova remnant 3C58 is detected in X-rays and very faintly at radio wavelengths. It is at an estimated distance of 3.2 kpc. The radio flux at 1.4 GHz is 45 μJy with a spectral index of -2.1 (Camilo et al., 2002). If PSR J1846−0258 had the same luminosity and spectral index we would have received a flux of 4.3 μJy at 1.9 GHz, just comparable to our upper limit. On the other hand, the X-ray luminosity of PSR J0205+6449 is estimated to be 2.84×10^{33}erg s^{-1} (Murray et al., 2002), while the X-ray luminosity of PSR J1846−0258 is estimated to be 7×10^{34}erg s^{-1} (Morton et al., 2007, adjusted to the Leahy et al. distance estimate). Thus in spite of having an X-ray luminosity more than 20 times that of PSR J0205+6449, PSR J1846−0258 appears to have a smaller radio luminosity than PSR J0205+6449. Beaming may account for this difference. More generally, the pseudoluminosity limit we set is 0.2 mJy kpc^2. Only eighteen pulsars have been detected at radio wavelengths that are fainter than this limit; their luminosities range down to 30 μJy kpc^2 (Manchester et al., 2005).

It is possible that PSR J1846−0258 produced radio emission that peaked several months after the X-ray event and faded by the time of our GBT observation. XTE J1810−197 was observed to brighten in radio about a year after its X-ray brightening, and it faded over the course of about a year (Camilo et al., 2007). We intend to search other observations from the monitoring program of RRAT J1846−02 for pulsations coming from PSR J1846−0258.

By comparing the period of RRAT J1846−02 measured from our observation with the published period, we were able to place a 3σ upper limit on the spin-down rate of 1.7×10^{-13}. This gives a 3σ upper limit on the surface dipole magnetic field of 2.6×10^{13} G, smaller than any known magnetar, and less than those of 24 more highly magnetized rotation-powered pulsars (Manchester et al., 2005), including PSR J1846−0258. We hope to combine this relatively long observation with a timing program being carried out by McLaughlin et al. to yield a phase-coherent timing solution for the RRAT J1846−02.

Facilities: GBT (S-band receiver), Parkes (multibeam receiver)

References

Abdo, A. A., Ackermann, M., Ajello, M., Allafort, A., Baldini, L., Ballet, J., Barbiellini, G., Bastieri, D., Bechtol, K., Bellazzini, R., Berenji, B., Blandford, R. D., Bloom, E. D., Bonamente, E., Borgland, A. W., Bouvier, A., Brandt, T. J., Bregeon, J., Brez, A., Brigida, M., Bruel, P., Buehler, R., Buson, S., Caliandro, G. A., Cameron, R. A., Cannon, A., Caraveo, P. A., Casandjian, J. M., Çelik, Ö., Charles, E., Chekhtman, A., Cheung, C. C., Chiang, J., Ciprini, S., Claus, R., Cohen-Tanugi, J., Costamante, L., Cutini, S., D'Ammando, F., Dermer, C. D., de Angelis, A., de Luca, A., de Palma, F., Digel, S. W., do Couto e Silva, E., Drell, P. S., Drlica-Wagner, A., Dubois, R., Dumora, D., Favuzzi, C., Fegan, S. J., Ferrara, E. C., Focke, W. B., Fortin, P., Frailis, M., Fukazawa, Y., Funk, S., Fusco, P., Gargano, F., Gasparrini, D., Gehrels, N., Germani, S., Giglietto, N., Giordano, F., Giroletti, M., Glanzman, T., Godfrey, G., Grenier, I. A., Grondin, M.-H., Grove, J. E., Guiriec, S., Hadasch, D., Hanabata, Y., Harding, A. K., Hayashi, K., Hayashida, M., Hays, E., Horan, D., Itoh, R., Jóhannesson, G., Johnson, A. S., Johnson, T. J., Khangulyan, D., Kamae, T., Katagiri, H., Kataoka, J., Kerr, M., Knödlseder, J., Kuss, M., Lande, J., Latronico, L., Lee, S.-H., Lemoine-Goumard, M., Longo, F., Loparco, F., Lubrano, P., Madejski, G. M., Makeev, A., Marelli, M., Mazziotta, M. N., McEnery, J. E., Michelson, P. F., Mitthumsiri, W., Mizuno, T., Moiseev,

A. A., Monte, C., Monzani, M. E., Morselli, A., Moskalenko, I. V., Murgia, S., Nakamori, T., Naumann-Godo, M., Nolan, P. L., Norris, J. P., Nuss, E., Ohsugi, T., Okumura, A., Omodei, N., Ormes, J. F., Ozaki, M., Paneque, D., Parent, D., Pelassa, V., Pepe, M., Pesce-Rollins, M., Pierbattista, M., Piron, F., Porter, T. A., Rainò, S., Rando, R., Ray, P. S., Razzano, M., Reimer, A., Reimer, O., Reposeur, T., Ritz, S., Romani, R. W., Sadrozinski, H. F.-W., Sanchez, D., Parkinson, P. M. S., Scargle, J. D., Schalk, T. L., Sgrò, C., Siskind, E. J., Smith, P. D., Spandre, G., Spinelli, P., Strickman, M. S., Suson, D. J., Takahashi, H., Takahashi, T., Tanaka, T., Thayer, J. B., Thompson, D. J., Tibaldo, L., Torres, D. F., Tosti, G., Tramacere, A., Troja, E., Uchiyama, Y., Vandenbroucke, J., Vasileiou, V., Vianello, G., Vitale, V., Wang, P., Wood, K. S., Yang, Z., and Ziegler, M. (2011). Gamma-Ray Flares from the Crab Nebula. *Science*, 331:739–. 2.6

Abdo, A. A., Ackermann, M., Ajello, M., Ampe, J., Anderson, B., Atwood, W. B., Axelsson, M., Bagagli, R., Baldini, L., Ballet, J., and et al. (2009). The on-orbit calibration of the Fermi Large Area Telescope. *Astroparticle Physics*, 32:193–219. 2.6

Abdo, A. A., Ackermann, M., Ajello, M., Atwood, W. B., Axelsson, M., Baldini, L., Ballet, J., Barbiellini, G., Baring, M. G., Bastieri, D., and et al. (2010). The First Fermi Large Area Telescope Catalog of Gamma-ray Pulsars. *ApJS*, 187:460–494. 2.6, 7.4

Adelman-McCarthy, J. K., Agüeros, M. A., Allam, S. S., Allende Prieto, C., Anderson, K. S. J., Anderson, S. F., Annis, J., Bahcall, N. A., Bailer-Jones,

C. A. L., Baldry, I. K., Barentine, J. C., Bassett, B. A., Becker, A. C., Beers, T. C., Bell, E. F., Berlind, A. A., Bernardi, M., Blanton, M. R., Bochanski, J. J., Boroski, W. N., Brinchmann, J., Brinkmann, J., Brunner, R. J., Budavári, T., Carliles, S., Carr, M. A., Castander, F. J., Cinabro, D., Cool, R. J., Covey, K. R., Csabai, I., Cunha, C. E., Davenport, J. R. A., Dilday, B., Doi, M., Eisenstein, D. J., Evans, M. L., Fan, X., Finkbeiner, D. P., Friedman, S. D., Frieman, J. A., Fukugita, M., Gänsicke, B. T., Gates, E., Gillespie, B., Glazebrook, K., Gray, J., Grebel, E. K., Gunn, J. E., Gurbani, V. K., Hall, P. B., Harding, P., Harvanek, M., Hawley, S. L., Hayes, J., Heckman, T. M., Hendry, J. S., Hindsley, R. B., Hirata, C. M., Hogan, C. J., Hogg, D. W., Hyde, J. B., Ichikawa, S.-i., Ivezić, Ž., Jester, S., Johnson, J. A., Jorgensen, A. M., Jurić, M., Kent, S. M., Kessler, R., Kleinman, S. J., Knapp, G. R., Kron, R. G., Krzesinski, J., Kuropatkin, N., Lamb, D. Q., Lampeitl, H., Lebedeva, S., Lee, Y. S., Leger, R. F., Lépine, S., Lima, M., Lin, H., Long, D. C., Loomis, C. P., Loveday, J., Lupton, R. H., Malanushenko, O., Malanushenko, V., Mandelbaum, R., Margon, B., Marriner, J. P., Martínez-Delgado, D., Matsubara, T., McGehee, P. M., McKay, T. A., Meiksin, A., Morrison, H. L., Munn, J. A., Nakajima, R., Neilsen, Jr., E. H., Newberg, H. J., Nichol, R. C., Nicinski, T., Nieto-Santisteban, M., Nitta, A., Okamura, S., Owen, R., Oyaizu, H., Padmanabhan, N., Pan, K., Park, C., Peoples, J. J., Pier, J. R., Pope, A. C., Purger, N., Raddick, M. J., Re Fiorentin, P., Richards, G. T., Richmond, M. W., Riess, A. G., Rix, H.-W., Rockosi, C. M., Sako, M., Schlegel, D. J., Schneider, D. P., Schreiber, M. R., Schwope, A. D., Seljak, U., Sesar, B., Sheldon, E., Shimasaku, K., Sivarani, T., Smith, J. A., Snedden, S. A., Steinmetz, M., Strauss, M. A., SubbaRao, M., Suto, Y., Szalay, A. S., Szapudi, I., Szkody, P., Tegmark, M., Thakar, A. R., Tremonti, C. A., Tucker, D. L., Uomoto, A., Vanden Berk, D. E., Vandenberg, J., Vidrih, S., Vogeley, M. S., Voges, W., Vogt, N. P., Wadadekar, Y., Weinberg, D. H., West, A. A., White, S. D. M., Wilhite, B. C., Yanny, B., Yocum, D. R., York, D. G., Zehavi, I., and Zucker, D. B. (2008). The Sixth Data Release of the Sloan Digital Sky Survey. *ApJS*, 175:297–313. 3–3

Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control*, 19:716–723. 4.3.1

Alpar, M. A., Cheng, A. F., Ruderman, M. A., and Shaham, J. (1982). A new class of radio pulsars. *Nature*, 300:728–730. 1.2

Andrade, M. C. B. (1996). Sexual Selection for Male Sacrifice in the Australian Redback Spider. *Science*, 271:70–72. 6

Antoniadis, J., Freire, P. C. C., Wex, N., Tauris, T. M., Lynch, R. S., van Kerkwijk, M. H., Kramer, M., Bassa, C., Dhillon, V. S., Driebe, T., Hessels, J. W. T., Kaspi, V. M., Kondratiev, V. I., Langer, N., Marsh, T. R., McLaughlin, M. A., Pennucci, T. T., Ransom, S. M., Stairs, I. H., van Leeuwen, J., Verbiest, J. P. W., and Whelan, D. G. (2013). A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340:448. 1.2

Applegate, J. H. (1992). A mechanism for orbital period modulation in close binaries. *ApJ*, 385:621–629. 5, 7.6.1, 7.6.2, 7.7

Applegate, J. H. and Shaham, J. (1994). Orbital period variability in the eclipsing pulsar binary PSR B1957+20: Evidence for a tidally powered star. *ApJ*, 436:312–318. 1.4, 7.6.1

Archibald, A. M., Kaspi, V. M., Bogdanov, S., Hessels, J. W. T., Stairs, I. H., Ransom, S. M., and McLaughlin, M. A. (2010). X-ray Variability and Evidence for Pulsations from the Unique Radio Pulsar/X-ray Binary Transition Object FIRST J102347.6+003841. *ApJ*, 722:88–95. (document), 4, 7.4

Archibald, A. M., Kaspi, V. M., Livingstone, M. A., and McLaughlin, M. A. (2008). No Detectable Radio Emission from the Magnetar-Like Pulsar in Kes 75. *ApJ*, 688:550–554. (document), A

Archibald, A. M., Stairs, I. H., Ransom, S. M., Kaspi, V. M., Kondratiev, V. I., Lorimer, D. R., McLaughlin, M. A., Boyles, J., Hessels, J. W. T., Lynch, R., van Leeuwen, J., Roberts, M. S. E., Jenet, F., Champion, D. J., Rosen, R., Barlow, B. N., Dunlap, B. H., and Remillard, R. A. (2009). A Radio Pulsar/X-ray Binary Link. *Science*, 324:1411–. (document), 3

Arons, J. and Tavani, M. (1993). High-energy emission from the eclipsing millisecond pulsar PSR 1957+20. *ApJ*, 403:249–255. 4.4.1, 7.4.1

Arzoumanian, Z., Fruchter, A. S., and Taylor, J. H. (1994). Orbital variability in the eclipsing pulsar binary PSR B1957+20. *ApJ*, 426:L85. 7.6, 7.6.1

Atwood, W. B., Abdo, A. A., Ackermann, M., Althouse, W., Anderson, B., Axelsson, M., Baldini, L., Ballet, J., Band, D. L., Barbiellini, G., and et al. (2009). The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission. *ApJ*, 697:1071–1102. 2.6

Baade, W. and Zwicky, F. (1934). Remarks on Super-Novae and Cosmic Rays. *Physical Review*, 46:76–77. 1.1

Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., and Goss, W. M. (1982). A millisecond pulsar. *Nature*, 300:615–618. 1.2

Backer, D. C., Kulkarni, S. R., and Taylor, J. H. (1983). Timing observations of the millisecond pulsar. *Nature*, 301:314. 1.2

Banaszkiewicz, M., Axford, W. I., and McKenzie, J. F. (1998). An analytic solar magnetic field model. *A&A*, 337:940–944. 7.5.1

Becker, R. H. and Helfand, D. J. (1984). New radio observations of the composite supernova remnant G29.7-0.3. *ApJ*, 283:154–157. A.2, A.3

Becker, R. H., White, R. L., and Helfand, D. J. (1995). The FIRST survey: faint images of the radio sky at twenty centimeters. *ApJ*, 450:559–577. 3.4

Bednarek, W. and Sitarek, J. (2013). Gamma-rays from nebulae around binary systems containing energetic rotation-powered pulsars. *MNRAS*, 430:2951–2959. 7.5.2

Berdyugina, S. V. (2005). Starspots: A Key to the Stellar Dynamo. *Living Reviews in Solar Physics*, 2:8. 7.5.1

Bhattacharya, D. and van den Heuvel, E. P. J. (1991). Formation and evolution of binary and millisecond radio pulsars. *Phys. Rep.*, 203:1–124. 1.3, 1.3, 1.3

Bogdanov, S., Archibald, A. M., Hessels, J. W. T., Kaspi, V. M., Lorimer, D., McLaughlin, M. A., Ransom, S. M., and Stairs, I. H. (2011). A Chandra X-Ray Observation of the Binary Millisecond Pulsar PSR J1023+0038. *ApJ*, 742:97. 6, 8, 9, 7.2, 7.3, 1, 7.4, 7.4.1, 7.5.2, 7.5.4

Bogdanov, S., Grindlay, J. E., and Rybicki, G. B. (2008). Thermal X-Rays from Millisecond Pulsars: Constraining the Fundamental Properties of Neutron Stars.

ApJ, 689:407–415. 4.4.1

Bogdanov, S., Grindlay, J. E., and van den Berg, M. (2005a). An X-Ray Variable Millisecond Pulsar in the Globular Cluster 47 Tucanae: Closing the Link to Low-Mass X-Ray Binaries. *ApJ*, 630:1029–1036. 3.6.4

Bogdanov, S., Grindlay, J. E., and van den Berg, M. (2005b). An X-Ray Variable Millisecond Pulsar in the Globular Cluster 47 Tucanae: Closing the Link to Low-Mass X-Ray Binaries. *ApJ*, 630:1029–1036. 4.4.2

Bogdanov, S., van den Berg, M., Heinke, C. O., Cohn, H. N., Lugger, P. M., and Grindlay, J. E. (2010). A Chandra X-ray Observatory Study of PSR J1740-5340 and Candidate Millisecond Pulsars in the Globular Cluster NGC 6397. *ApJ*, 709:241–250. 4.4.2

Bond, H. E., White, R. L., Becker, R. H., and O'Brien, M. S. (2002). FIRST J102347.6+003841: The First Radio-selected Cataclysmic Variable. *PASP*, 114:1359–1363. 3.1, 3.4, 3–3

Boyles, J., Lynch, R. S., Ransom, S. M., Stairs, I. H., Lorimer, D. R., McLaughlin, M. A., Hessels, J. W. T., Kaspi, V. M., Kondratiev, V. I., Archibald, A., Berndsen, A., Cardoso, R. F., Cherry, A., Epstein, C. R., Karako-Argaman, C., McPhee, C. A., Pennucci, T., Roberts, M. S. E., Stovall, K., and van Leeuwen, J. (2013). The Green Bank Telescope 350 MHz Drift-scan survey. I. Survey Observations and the Discovery of 13 Pulsars. *ApJ*, 763:80. (document), 2.3, 2–4, 3.2

Breton, R. P., van Kerkwijk, M. H., Roberts, M. S. E., Hessels, J. W. T., Camilo, F., McLaughlin, M. A., Ransom, S. M., Ray, P. S., and Stairs, I. H. (2013).

Discovery of the Optical Counterparts to Four Energetic Fermi Millisecond Pulsars. *ArXiv e-prints*. 7.2, 7.6

Brown, E. F., Bildsten, L., and Rutledge, R. E. (1998). Crustal Heating and Quiescent Emission from Transiently Accreting Neutron Stars. *ApJ*, 504:L95+. 4.4.2

Burderi, L., Di Salvo, T., D'Antona, F., Robba, N. R., and Testa, V. (2003). The optical counterpart to SAX J1808.4-3658 in quiescence: Evidence of an active radio pulsar? *A&A*, 404:L43–L46. 3.6.1, 7.6, 7.6.2

Burderi, L., Possenti, A., D'Antona, F., Di Salvo, T., Burgay, M., Stella, L., Menna, M. T., Iaria, R., Campana, S., and d'Amico, N. (2001). Where May Ultrafast Rotating Neutron Stars Be Hidden? *ApJ*, 560:L71–L74. 1.5

Burgay, M., Burderi, L., Possenti, A., D'Amico, N., Manchester, R. N., Lyne, A. G., Camilo, F., and Campana, S. (2003). A Search for Pulsars in Quiescent Soft X-Ray Transients. I. *ApJ*, 589:902–910. 4.4.2

Burgay, M., Joshi, B. C., D'Amico, N., Possenti, A., Lyne, A. G., Manchester, R. N., McLaughlin, M. A., Kramer, M., Camilo, F., and Freire, P. C. C. (2006). The Parkes High-Latitude pulsar survey. *MNRAS*, 368:283–292. 1.4

Burke, B. F. and Graham-Smith, F. (2009). *An Introduction to Radio Astronomy*. Cambridge University Press. 2.2, 2.3

Camilo, F., Cognard, I., Ransom, S. M., Halpern, J. P., Reynolds, J., Zimmerman, N., Gotthelf, E. V., Helfand, D. J., Demorest, P., Theureau, G., and Backer, D. C. (2007). The magnetar xte j1810-197: Variations in torque, radio flux density, and pulse profile morphology. *ApJ*, 663:497–504. A.3

Camilo, F., Ransom, S. M., Halpern, J. P., and Reynolds, J. (2007a). 1E 1547.0-5408: A Radio-emitting Magnetar with a Rotation Period of 2 Seconds. *ApJ*, 666:L93–L96. A.3

Camilo, F., Ransom, S. M., Halpern, J. P., Reynolds, J., Helfand, D. J., Zimmerman, N., and Sarkissian, J. (2006). Transient pulsed radio emission from a magnetar. *Nature*, 442:892–895. A.3

Camilo, F., Ransom, S. M., Peñalver, J., Karastergiou, A., van Kerkwijk, M. H., Durant, M., Halpern, J. P., Reynolds, J., Thum, C., Helfand, D. J., Zimmerman, N., and Cognard, I. (2007b). The Variable Radio-to-X-Ray Spectrum of the Magnetar XTE J1810-197. *ApJ*, 669:561–569. A.2

Camilo, F. and Rasio, F. A. (2005). Pulsars in Globular Clusters. In Rasio, F. A. and Stairs, I. H., editors, *Binary Radio Pulsars*, volume 328 of *Astronomical Society of the Pacific Conference Series*, page 147. 3.6.4

Camilo, F., Reynolds, J., Johnston, S., Halpern, J. P., Ransom, S. M., and van Straten, W. (2007). Polarized radio emission from the magnetar xte j1810-197. *Astrophys. J. Lett.*, 659:L37–L40. A.3

Camilo, F., Stairs, I. H., Lorimer, D. R., Backer, D. C., Ransom, S. M., Klein, B., Wielebinski, R., Kramer, M., McLaughlin, M. A., Arzoumanian, Z., and Müller, P. (2002). Discovery of Radio Pulsations from the X-Ray Pulsar J0205+6449 in Supernova Remnant 3C 58 with the Green Bank Telescope. *ApJ*, 571:L41–L44. A.3

Campana, S., Colpi, M., Mereghetti, S., Stella, L., and Tavani, M. (1998). The neutron stars of Soft X-ray Transients. *A&A Rev.*, 8:279–316. 3.6.3, 3.6.4, 4.4.2

Campana, S., D'Avanzo, P., Casares, J., Covino, S., Israel, G., Marconi, G., Hynes, R., Charles, P., and Stella, L. (2004). Indirect Evidence of an Active Radio Pulsar in the Quiescent State of the Transient Millisecond Pulsar SAX J1808.4-3658. *ApJ*, 614:L49–L52. 4.4.2

Campana, S., Stella, L., Gastaldello, F., Mereghetti, S., Colpi, M., Israel, G. L., Burderi, L., Di Salvo, T., and Robba, R. N. (2002). An XMM-Newton Study of the 401 Hz Accreting Pulsar SAX J1808.4-3658 in Quiescence. *ApJ*, 575:L15–L19. 4.4.2

Chakrabarty, D. and Morgan, E. H. (1998). The 2 hour orbit of a binary millisecond X-ray pulsar. *Nature*, 394:346–348. 3.6.1, 4.4.2

Chandrasekhar, S. (1931). The highly collapsed configurations of a stellar mass. *MNRAS*, 91:456–466. 1.1

Chandrasekhar, S. (1984). On stars, their evolution and their stability. *Reviews of Modern Physics*, 56:137–147. 1.1

Cocke, W. J., Disney, M. J., and Taylor, D. J. (1969). Discovery of Optical Signals from Pulsar NP 0532. *Nature*, 221:525–527. 1.1

Comella, J. M., Craft, H. D., Lovelace, R. V. E., and Sutton, J. M. (1969). Crab Nebula Pulsar NP 0532. *Nature*, 221:453–454. 1.1

Cordes, J. M. (1986). Space velocities of radio pulsars from interstellar scintillations. *ApJ*, 311:183–196. 5.4.4

Cordes, J. M. and Lazio, T. J. W. (2002). NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations. *ArXiv Astrophysics e-prints*. available at http://arxiv.org/abs/astro-ph/0207156. 2.1, 3.6.1, A.2

Cordes, J. M. and Rickett, B. J. (1998). Diffractive Interstellar Scintillation Timescales and Velocities. *ApJ*, 507:846–860. 5.4.4

Cranmer, S. R. (2008). Winds of Main-Sequence Stars: Observational Limits and a Path to Theoretical Prediction. In van Belle, G., editor, *14th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun*, volume 384 of *Astronomical Society of the Pacific Conference Series*, page 317. 7.5.2

D'Amico, N., Possenti, A., Manchester, R. N., Sarkissian, J., Lyne, A. G., and Camilo, F. (2001). An Eclipsing Millisecond Pulsar with a Possible Main-Sequence Companion in NGC 6397. *ApJ*, 561:L89–L92. 3.6.2, 4.4.2

de Jager, O. C., Raubenheimer, B. C., and Swanepoel, J. W. H. (1989). A poweful test for weak periodic signals with unknown light curve shape in sparse data. *A&A*, 221:180–190. 2.6, 4.3.3

Deller, A. T., Archibald, A. M., Brisken, W. F., Chatterjee, S., Janssen, G. H., Kaspi, V. M., Lorimer, D., Lyne, A. G., McLaughlin, M. A., Ransom, S., Stairs, I. H., and Stappers, B. (2012). A Parallax Distance and Mass Estimate for the Transitional Millisecond Pulsar System J1023+0038. *ApJ*, 756:L25. 2–5, 5.1, 2, 5.3, 5–1, 5.3, 5.4.1, 5.4.4, 5.5, 7.2, 7.3, 7.4

Deloye, C. J. (2008). The Connection Between Low-Mass X-ray Binaries and (Millisecond) Pulsars: A Binary Evolution Perspective. In Bassa, C., Wang, Z., Cumming, A., and Kaspi, V. M., editors, *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, volume 983 of *American Institute of Physics Conference Series*, pages 501–509. 1.5

Deloye, C. J., Heinke, C. O., Taam, R. E., and Jonker, P. G. (2008). Optical observations of SAX J1808.4-3658 during quiescence. *MNRAS*, 391:1619–1628. 3.6.1

Demorest, P. B. (2007a). *Measuring the gravitational wave background using precision pulsar timing*. PhD thesis, University of California, Berkeley. 3.3

Demorest, P. B. (2007b). *Measuring the gravitational wave background using precision pulsar timing*. PhD thesis, University of California, Berkeley. 5.2

Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., and Hessels, J. W. T. (2010). A two-solar-mass neutron star measured using Shapiro delay. *Nature*, 467:1081–1083. 1.2

Dowd, A., Sisk, W., and Hagen, J. (2000a). WAPP — Wideband Arecibo Pulsar Processor. In Kramer, M., Wex, N., and Wielebinski, R., editors, *IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond*, volume 202 of *Astronomical Society of the Pacific Conference Series*, page 275. 3.3

Dowd, A., Sisk, W., and Hagen, J. (2000b). WAPP — Wideband Arecibo Pulsar Processor. In Kramer, M., Wex, N., and Wielebinski, R., editors, *IAU Colloq. 177: Pulsar Astronomy - 2000 and Beyond*, volume 202 of *Astronomical Society of the Pacific Conference Series*, page 275. 5.2

Edwards, R. T., Hobbs, G. B., and Manchester, R. N. (2006). TEMPO2, a new pulsar timing package - II. The timing model and precision estimates. *MNRAS*, 372:1549–1574. 2.4

Eggleton, P. P. (1983). Approximations to the radii of Roche lobes. *ApJ*, 268:368. 1.2

Espinoza, C. M., Guillemot, L., Çelik, Ö., Weltevrede, P., Stappers, B. W., Smith, D. A., Kerr, M., Zavlin, V. E., Cognard, I., Eatough, R. P., Freire, P. C. C., Janssen, G. H., Camilo, F., Desvignes, G., Hewitt, J. W., Hou, X., Johnston, S., Keith, M., Kramer, M., Lyne, A., Manchester, R. N., Ransom, S. M., Ray, P. S., Shannon, R., Theureau, G., and Webb, N. (2013). Six millisecond pulsars detected by the Fermi Large Area Telescope and the radio/gamma-ray connection of millisecond pulsars. *MNRAS*, 430:571–587. 7.4

Feiden, G. A. and Chaboyer, B. (2012). Reevaluating the Mass-Radius Relation for Low-mass, Main-sequence Stars. *ApJ*, 757:42. 7.5.1

Ferdman, R. D. (2008a). *Binary pulsars: Evolution and Fundamental Physics*. PhD thesis, University of British Columbia. 3.3

Ferdman, R. D. (2008b). *Binary pulsars: evolution and fundamental physics*. PhD thesis, University of British Columbia. 5.2

Ferraro, F. R., Possenti, A., D'Amico, N., and Sabbi, E. (2001). The Bright Optical Companion to the Eclipsing Millisecond Pulsar in NGC 6397. *ApJ*, 561:L93–L96. 4.4.2

Finzi, A. and Wolf, R. A. (1969). Possible Conversion of Rotational Energy of the Neutron Star in the Crab Nebula Into Energy of Relativistic Electrons. *ApJ*, 155:L107. 1.1

Freire, P. C. C., Bassa, C. G., Wex, N., Stairs, I. H., Champion, D. J., Ransom, S. M., Lazarus, P., Kaspi, V. M., Hessels, J. W. T., Kramer, M., Cordes, J. M., Verbiest, J. P. W., Podsiadlowski, P., Nice, D. J., Deneva, J. S., Lorimer, D. R., Stappers, B. W., McLaughlin, M. A., and Camilo, F. (2011). On the nature and evolution of the unique binary pulsar J1903+0327. *MNRAS*, 412:2763–2780. 2

Fritz, G., Henry, R. C., Meekins, J. F., Chubb, T. A., and Friedman, H. (1969). X-ray Pulsar in the Crab Nebula. *Science*, 164:709–712. 1.1

Fruchter, A. S., Gunn, J. E., Lauer, T. R., and Dressler, A. (1988a). Optical detection and characterization of the eclipsing pulsar's companion. *Nature*, 334:686–689. 3.6.2

Fruchter, A. S., Stinebring, D. R., and Taylor, J. H. (1988b). A millisecond pulsar in an eclipsing binary. *Nature*, 333:237–239. 1.4

Gaensler, B. M. and Slane, P. O. (2006). The Evolution and Structure of Pulsar Wind Nebulae. *ARA&A*, 44:17–47. 4.4.1

Galloway, D. K. and Cumming, A. (2006). Helium-rich Thermonuclear Bursts and the Distance to the Accretion-powered Millisecond Pulsar SAX J1808.4-3658. *ApJ*, 652:559–568. 4.4.2

Gavriil, F. P., Gonzalez, M. E., Gotthelf, E. V., Kaspi, V. M., Livingstone, M. A., and Woods, P. M. (2008). Magnetar-like Emission from the Young Pulsar in Kes 75. *ArXiv e-prints*. A.1, A.2, A.2.2

Gold, T. (1968). Rotating Neutron Stars as the Origin of the Pulsating Radio Sources. *Nature*, 218:731–732. 1.1

Gonzalez, M. E., Kaspi, V. M., Camilo, F., Gaensler, B. M., and Pivovaroff, M. J. (2005). Unusual Pulsed X-Ray Emission from the Young, High Magnetic Field Pulsar PSR J1119-6127. *ApJ*, 630:489–494. A.1

Gotthelf, E. V. and Halpern, J. P. (2005). The Spectral Evolution of Transient Anomalous X-Ray Pulsar XTE J1810-197. *ApJ*, 632:1075–1085. A.2

Gotthelf, E. V., Vasisht, G., Boylan-Kolchin, M., and Torii, K. (2000). A 700 Year-old Pulsar in the Supernova Remnant Kesteven 75. *ApJ*, 542:L37–L40. A.2

Green, D. A. (2006). A Catalogue of Galactic Supernova Remnants (2006 April version). *Astrophysics Group, Cavendish Laboratory, Cambridge, United Kingdom*. A.2.1

Grindlay, J. E., Camilo, F., Heinke, C. O., Edmonds, P. D., Cohn, H., and Lugger, P. (2002). Chandra Study of a Complete Sample of Millisecond Pulsars in 47 Tucanae and NGC 6397. *ApJ*, 581:470–484. 3.6.4

Guillemot, L., Johnson, T. J., Venter, C., Kerr, M., Pancrazi, B., Livingstone, M., Janssen, G. H., Jaroenjittichai, P., Kramer, M., Cognard, I., Stappers, B. W., Harding, A. K., Camilo, F., Espinoza, C. M., Freire, P. C. C., Gargano, F., Grove, J. E., Johnston, S., Michelson, P. F., Noutsos, A., Parent, D., Ransom, S. M., Ray, P. S., Shannon, R., Smith, D. A., Theureau, G., Thorsett, S. E., and Webb, N. (2012). Pulsed Gamma Rays from the Original Millisecond and Black Widow Pulsars: A Case for Caustic Radio Emission? *ApJ*, 744:33. 2.6, 2.6

Hankins, T. H. and Rickett, B. J. (1975). Pulsar signal processing. In *Methods in Computational Physics Volume 14 — Radio Astronomy*, pages 55–129, New York. Academic Press. 2.4, 3.3

Hartman, J. M., Patruno, A., Chakrabarty, D., Kaplan, D. L., Markwardt, C. B., Morgan, E. H., Ray, P. S., van der Klis, M., and Wijnands, R. (2008). The Long-Term Evolution of the Spin, Pulse Shape, and Orbit of the Accretion-powered Millisecond Pulsar SAX J1808.4-3658. *ApJ*, 675:1468–1486. 3.6.1

Haslam, C. G. T., Salter, C. J., Stoffel, H., and Wilson, W. E. (1982). A 408 MHz all-sky continuum survey. II - The atlas of contour maps. *A&AS*, 47:1–+. A.2.1

Heinke, C. O., Jonker, P. G., Wijnands, R., Deloye, C. J., and Taam, R. E. (2009). Further Constraints on Thermal Quiescent X-Ray Emission from SAX J1808.4-3658. *ApJ*, 691:1035–1041. 3.6.1, 4.4.2

Helfand, D. J., Collins, B. F., and Gotthelf, E. V. (2003). Chandra X-Ray Imaging Spectroscopy of the Young Supernova Remnant Kesteven 75. *ApJ*, 582:783–792. A.1, A.2.1, A.3

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., and Collins, R. A. (1968). Observation of a Rapidly Pulsating Radio Source. *Nature*, 217:709–713. 1.1

Hobbs, G., Archibald, A., Arzoumanian, Z., Backer, D., Bailes, M., Bhat, N. D. R., Burgay, M., Burke-Spolaor, S., Champion, D., Cognard, I., Coles, W., Cordes, J., Demorest, P., Desvignes, G., Ferdman, R. D., Finn, L., Freire, P., Gonzalez, M., Hessels, J., Hotan, A., Janssen, G., Jenet, F., Jessner, A., Jordan, C., Kaspi, V., Kramer, M., Kondratiev, V., Lazio, J., Lazaridis, K., Lee, K. J., Levin, Y., Lommen, A., Lorimer, D., Lynch, R., Lyne, A., Manchester, R., McLaughlin, M., Nice, D., Oslowski, S., Pilia, M., Possenti, A., Purver, M., Ransom, S., Reynolds, J., Sanidas, S., Sarkissian, J., Sesana, A., Shannon, R., Siemens, X., Stairs, I., Stappers, B., Stinebring, D., Theureau, G., van Haasteren, R., van Straten, W., Verbiest, J. P. W., Yardley, D. R. B., and You, X. P. (2010). The International Pulsar Timing Array project: using pulsars as a gravitational wave detector. *Classical and Quantum Gravity*, 27(8):084013. 1.2

Hobbs, G. B., Edwards, R. T., and Manchester, R. N. (2006). TEMPO2, a new pulsar-timing package - I. An overview. *MNRAS*, 369:655–672. 2.4, 5.3

Homer, L., Szkody, P., Chen, B., Henden, A., Schmidt, G., Anderson, S. F., Silvestri, N. M., and Brinkmann, J. (2006). XMM-Newton and Optical Follow-up Observations of SDSS J093249.57+472523.0 and SDSS J102347.67+003841.2. *AJ*, 131:562–570. 3.4, 3–3, 3.5, 3.6.4, 4.1, 4.2, 4.3.1, 4.3.2

Hotan, A. W., van Straten, W., and Manchester, R. N. (2004). PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis. *PASA*, 21:302–309. 2.4, 5.2, 5.3

Iacolina, M. N., Burgay, M., Burderi, L., Possenti, A., and di Salvo, T. (2009). Searching for pulsed emission from XTE J0929-314 at high radio frequencies. *A&A*, 497:445–450. 4.4.2, 7.6

Ibrahim, A. I., Markwardt, C. B., Swank, J. H., Ransom, S., Roberts, M., Kaspi, V., Woods, P. M., Safi-Harb, S., Balman, S., Parke, W. C., Kouveliotou, C., Hurley, K., and Cline, T. (2004). Discovery of a Transient Magnetar: XTE J1810-197. *ApJ*, 609:L21–L24. A.1

Illarionov, A. F. and Sunyaev, R. A. (1975). Why the Number of Galactic X-ray Stars Is so Small? *A&A*, 39:185. 3.6.3

Jansen, F., Lumb, D., Altieri, B., Clavel, J., Ehle, M., Erd, C., Gabriel, C., Guainazzi, M., Gondoin, P., Much, R., Munoz, R., Santos, M., Schartel, N., Texier, D., and Vacanti, G. (2001). XMM-Newton observatory. I. The spacecraft and operations. *A&A*, 365:L1–L6. 2.5

Kalberla, P. M. W., Burton, W. B., Hartmann, D., Arnal, E. M., Bajaja, E., Morras, R., and Pöppel, W. G. L. (2005). The Leiden/Argentine/Bonn (LAB) Survey of Galactic HI. Final data release of the combined LDS and IAR surveys with improved stray-radiation corrections. *A&A*, 440:775–782. 4.3.1

Kaplan, D. L., Escoffier, R. P., Lacasse, R. J., O'Neil, K., Ford, J. M., Ransom, S. M., Anderson, S. B., Cordes, J. M., Lazio, T. J. W., and Kulkarni, S. R. (2005). The Green Bank Telescope Pulsar Spigot. *PASP*, 117:643–653. 3.2, 3.3, A.2.1

Kargaltsev, O. and Pavlov, G. G. (2008). Pulsar Wind Nebulae in the Chandra Era. In C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi, editor, *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, volume 983 of *American Institute of Physics Conference Series*, pages 171–185. 4.4.1

Karuppusamy, R., Stappers, B., and van Straten, W. (2008). PuMa-II: A Wide Band Pulsar Machine for the Westerbork Synthesis Radio Telescope. *PASP*, 120:191–202. 5.2

Kaspi, V. M., Manchester, R. N., Johnston, S., Lyne, A. G., and D'Amico, N. (1996). A Search for Radio Pulsars in Southern Supernova Remnants. *AJ*, 111:2028–+. A.1, A.3

Kaspi, V. M. and McLaughlin, M. A. (2005). Chandra X-Ray Detection of the High Magnetic Field Radio Pulsar PSR J1718-3718. *ApJ*, 618:L41–L44. A.1

Kaspi, V. M., Roberts, M. S. E., and Harding, A. K. (2006). *Isolated neutron stars*, pages 279–339. Cambridge University Press. 2.5, 4.4.1

Kerr, M. (2011). Improving Sensitivity to Weak Pulsations with Photon Probability Weighting. *ApJ*, 732:38. 2.6, 6.3

Khangulyan, D., Aharonian, F. A., Bogovalov, S. V., and Ribó, M. (2012). Post-periastron Gamma-Ray Flare from PSR B1259-63/LS 2883 as a Result of Comptonization of the Cold Pulsar Wind. *ApJ*, 752:L17. 2.6

Kirsch, M. G. F., Becker, W., Benlloch-Garcia, S., Jansen, F. A., Kendziorra, E., Kuster, M., Lammers, U., Pollock, A. M. T., Possanzini, F., Serpell, E., and Talavera, A. (2004). Timing accuracy and capabilities of XMM-Newton. In K. A. Flanagan & O. H. W. Siegmund, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 5165 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pages 85–95. 2.5

Kniffen, D. A., Hartman, R. C., Thompson, D. J., Bignami, G. F., and Fichtel, C. E. (1974). Gamma radiation from the Crab Nebula above 35 MeV. *Nature*, 251:397–399. 2.6

Lattimer, J. M. and Prakash, M. (2004). The Physics of Neutron Stars. *Science*, 304:536–542. 3.6.1, 7.2

Lazaridis, K., Verbiest, J. P. W., Tauris, T. M., Stappers, B. W., Kramer, M., Wex, N., Jessner, A., Cognard, I., Desvignes, G., Janssen, G. H., Purver, M. B., Theureau, G., Bassa, C. G., and Smits, R. (2011). Evidence for gravitational quadrupole moment variations in the companion of PSR J2051-0827. *MNRAS*, 414:3134–3144. 7.6, 7.6.1

Leahy, D. A. and Tian, W. W. (2008). The distance of the SNR Kes 75 and PWN PSR J1846-0258 system. *A&A*, 480:L25–L28. A.2, A.3

Livingstone, M. A., Kaspi, V. M., Gotthelf, E. V., and Kuiper, L. (2006). A Braking Index for the Young, High Magnetic Field, Rotation-Powered Pulsar in Kesteven 75. *ApJ*, 647:1286–1292. A.1, A.2.1, A.3

Lorimer, D. R. and Kramer, M. (2004). *Handbook of Pulsar Astronomy*. 1.1, 1.3, 2, 2.1, 2.1, 2.3, 2.4, 5.4.4, A.2.1, A.2.1

Lynch, R. S., Boyles, J., Ransom, S. M., Stairs, I. H., Lorimer, D. R., McLaughlin, M. A., Hessels, J. W. T., Kaspi, V. M., Kondratiev, V. I., Archibald, A. M., Berndsen, A., Cardoso, R. F., Cherry, A., Epstein, C. R., Karako-Argaman, C., McPhee, C. A., Pennucci, T., Roberts, M. S. E., Stovall, K., and van Leeuwen, J. (2013). The Green Bank Telescope 350 MHz Drift-scan Survey II: Data Analysis and the Timing of 10 New Pulsars, Including a Relativistic Binary. *ApJ*, 763:81. 2.3, 3.2

Lyne, A. G., Manchester, R. N., Lorimer, D. R., Bailes, M., D'Amico, N., Tauris, T. M., Johnston, S., Bell, J. F., and Nicastro, L. (1998). The Parkes Southern Pulsar Survey – II. Final Results and Population Analysis. *MNRAS*, 295:743–755. 3–3

Manchester, R. N., Hobbs, G. B., Teoh, A., and Hobbs, M. (2005). The Australia Telescope National Facility Pulsar Catalogue. *AJ*, 129:1993–2006. 1–1, A.2, A.3

Manchester, R. N., Lyne, A. G., Camilo, F., Bell, J. F., Kaspi, V. M., D'Amico, N., McKay, N. P. F., Crawford, F., Stairs, I. H., Possenti, A., Kramer, M., and Sheppard, D. C. (2001). The Parkes multi-beam pulsar survey - I. Observing and data analysis systems, discovery and timing of 100 pulsars. *MNRAS*, 328:17–35. 2.3

Mardia, K. V. (1975). Statistics of directional data. *J. Roy. Statist. Soc. Ser. B*, 37(3):349–393. A.2.1

Mason, P. A. and Gray, C. L. (2007). AR Ursae Majoris Discovered to Be a Persistent Radio Polar: Results from a VLA Survey of Magnetic Cataclysmic Variables. *ApJ*, 660:662–668. 3–3

McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., Kramer, M., Faulkner, A. J., Manchester, R. N., Cordes, J. M., Camilo, F., Possenti, A., Stairs, I. H., Hobbs, G., D'Amico, N., Burgay, M., and O'Brien, J. T. (2006). Transient radio bursts from rotating neutron stars. *Nature*, 439:817–820. A.1, A.2, A.2.1

McLaughlin, M. A., Rea, N., Gaensler, B. M., Chatterjee, S., Camilo, F., Kramer, M., Lorimer, D. R., Lyne, A. G., Israel, G. L., and Possenti, A. (2007). Discovery of Pulsations and a Possible Spectral Feature in the X-Ray Emission from Rotating Radio Transient J1819-1458. *ApJ*, 670:1307–1313. A.1

McLaughlin, M. A., Stairs, I. H., Kaspi, V. M., Lorimer, D. R., Kramer, M., Lyne, A. G., Manchester, R. N., Camilo, F., Hobbs, G., Possenti, A., D'Amico, N., and Faulkner, A. J. (2003). PSR J1847-0130: A Radio Pulsar with Magnetar Spin Characteristics. *ApJ*, 591:L135–L138. A.1

Melatos, A. (1997). Spin-down of an oblique rotator with a current-starved outer magnetosphere. *MNRAS*, 288:1049–1059. A.3

Mereghetti, S. (2008). The strongest cosmic magnets: soft gamma-ray repeaters and anomalous X-ray pulsars. *A&A Rev.*, 15:225–287. 2.5

Morton, T. D., Slane, P., Borkowski, K. J., Reynolds, S. P., Helfand, D. J., Gaensler, B. M., and Hughes, J. P. (2007). Observations of X-Rays and Thermal Dust Emission from the Supernova Remnant Kes 75. *ApJ*, 667:219–225. A.3

Murray, S. S., Slane, P. O., Seward, F. D., Ransom, S. M., and Gaensler, B. M. (2002). Discovery of X-Ray Pulsations from the Compact Central Source in the Supernova Remnant 3C 58. *ApJ*, 568:226–231. A.3

Ng, C. ., Slane, P. O., Gaensler, B. M., and Hughes, J. P. (2008). Deep Chandra Observation of the PulsaR Wind Nebula Powered by the Pulsar J1846-0258 in the Supernova Remnant Kes 75. *ArXiv e-prints*. A.1, A.3

Nice, D. J., Thorsett, S. E., Taylor, J. H., and Fruchter, A. S. (1990). Observations of the eclipsing binary pulsar in Terzan 5. *ApJ*, 361:L61–L63. 5.4

Nolan, P. L., Abdo, A. A., Ackermann, M., Ajello, M., Allafort, A., Antolini, E., Atwood, W. B., Axelsson, M., Baldini, L., Ballet, J., and et al. (2012a). Fermi Large Area Telescope Second Source Catalog. *ApJS*, 199:31. 6.1

Nolan, P. L., Abdo, A. A., Ackermann, M., Ajello, M., Allafort, A., Antolini, E., Atwood, W. B., Axelsson, M., Baldini, L., Ballet, J., and et al. (2012b). Fermi Large Area Telescope Second Source Catalog. *ApJS*, 199:31. 7.3, 7.4

Pacini, F. (1967). Energy Emission from a Neutron Star. *Nature*, 216:567–568. 1.1

Paczyński, B. (1971). Evolutionary Processes in Close Binary Systems. *ARA&A*, 9:183. 1.2

Paltani, S. (2004). Searching for periods in X-ray observations using Kuiper's test. Application to the ROSAT PSPC archive. *A&A*, 420:789–797. 4.3.2, 4.3.3

Papitto, A., Ferrigno, C., Bozzo, E., Rea, N., Pavan, L., Campana, S., Romano, P., Burderi, L., Di Salvo, T., Riggio, A., Torres, D. F., Falanga, M., Hessels, J. W. T., Burgay, M., Sarkissian, J. M., Wieringa, M. H., Filipović, M. D., and Wong, G. F. (2013a). Swinging between rotation and accretion power in a millisecond binary pulsar. *ArXiv e-prints*. 7.3

Papitto, A., Hessels, J. W. T., Burgay, M., Ransom, S., Rea, N., Possenti, A., Stairs, I., Ferrigno, C., and Bozzo, A. (2013b). The transient low-mass X-ray binary IGR J18245-2452 is again active as a radio pulsar. *The Astronomer's Telegram*, 5069:1. 7.3

Patruno, A., Bult, P., Gopakumar, A., Hartman, J. M., Wijnands, R., van der Klis, M., and Chakrabarty, D. (2012). Accelerated Orbital Expansion and Secular Spin-down of the Accreting Millisecond Pulsar SAX J1808.4-3658. *ApJ*, 746:L27. 7.6, 7.6.2

Patruno, A. and Watts, A. L. (2012). Accreting Millisecond X-Ray Pulsars. *ArXiv e-prints*. 1.3, 1.3, 1.3, 7.6, 7.6.2

Pivovaroff, M. J., Kaspi, V. M., and Camilo, F. (2000). X-Ray Observations of the High Magnetic Field Radio Pulsar PSR J1814-1744. *ApJ*, 535:379–384. A.1

Pletsch, H. J., Guillemot, L., Allen, B., Kramer, M., Aulbert, C., Fehrmann, H., Ray, P. S., Barr, E. D., Belfiore, A., Camilo, F., Caraveo, P. A., Çelik, Ö., Champion, D. J., Dormody, M., Eatough, R. P., Ferrara, E. C., Freire, P. C. C., Hessels, J. W. T., Keith, M., Kerr, M., de Luca, A., Lyne, A. G., Marelli, M., McLaughlin, M. A., Parent, D., Ransom, S. M., Razzano, M., Reich, W., Saz Parkinson, P. M., Stappers, B. W., and Wolff, M. T. (2012). Discovery of Nine

Gamma-Ray Pulsars in Fermi Large Area Telescope Data Using a New Blind Search Method. *ApJ*, 744:105. 2.6

Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., and Siemiginowska, A. (2002). Statistics, Handle with Care: Detecting Multiple Model Components with the Likelihood Ratio Test. *ApJ*, 571:545–559. 4.3.1

Radhakrishnan, V. and Srinivasan, G. (1982). On the origin of the recently discovered ultra-rapid pulsar. *Current Science*, 51:1096–1099. 1.2

Ransom, S. M. (2001a). *New search techniques for binary pulsars*. PhD thesis, Harvard University. 2.3, 5.3

Ransom, S. M. (2001b). New Search Techniques for Binary Pulsars. In *Bulletin of the American Astronomical Society*, volume 33 of *Bulletin of the American Astronomical Society*, pages 1484–+. A.2.1

Ransom, S. M., Demorest, P., Ford, J., McCullough, R., Ray, J., DuPlain, R., and Brandt, P. (2009). GUPPI: Green Bank Ultimate Pulsar Processing Instrument. In *American Astronomical Society Meeting Abstracts #214*, volume 214 of *American Astronomical Society Meeting Abstracts*, page 605.08. 5.2

Ransom, S. M., Eikenberry, S. S., and Middleditch, J. (2002). Fourier Techniques for Very Long Astrophysical Time-Series Analysis. *AJ*, 124:1788–1809. 2, 2.3, A.2.1

Richards, D. W. and Comella, J. M. (1969). The Period of Pulsar NP 0532. *Nature*, 222:551–552. 1.1

Roberts, M. S. E. (2011). New Black Widows and Redbacks in the Galactic Field. In Burgay, M., D'Amico, N., Esposito, P., Pellizzoni, A., and Possenti, A., editors, *American Institute of Physics Conference Series*, volume 1357 of *American Institute of Physics Conference Series*, pages 127–130. 1.4, 2.6

Sabbi, E., Gratton, R., Ferraro, F. R., Bragaglia, A., Possenti, A., D'Amico, N., and Camilo, F. (2003). The Complex Hα Line Profile of the Bright Companion to PSR J1740-5340 in NGC 6397. *ApJ*, 589:L41–L44. 7.5.3

Sanjeev Kumar, H. and Safi-Harb, S. (2008). Variability of the High-Magnetic Field X-ray Pulsar PSR J1846-0258 Associated with the Supernova Remnant Kes 75 as Revealed by the Chandra X-ray Observatory. *ArXiv e-prints*. A.1

Shklovskii, I. S. (1970). Possible Causes of the Secular Increase in Pulsar Periods. *Soviet Ast.*, 13:562. 3.6.1, 7.2

Stappers, B. W., Bailes, M., Lyne, A. G., Manchester, R. N., D'Amico, N., Tauris, T. M., Lorimer, D. R., Johnston, S., and Sandhu, J. S. (1996). Probing the eclipse region of a binary millisecond pulsar. *ApJ*, 465:L119–L122. 1.4, 3.6.2

Stappers, B. W., Gaensler, B. M., Kaspi, V. M., van der Klis, M., and Lewin, W. H. G. (2003). An X-ray nebula associated with the millisecond pulsar B1957+20. *Science*, 299:1372–1374. 7.4.1

Stella, L., Campana, S., Colpi, M., Mereghetti, S., and Tavani, M. (1994). Do quiescent soft X-ray transients contain millisecond radio pulsars? *ApJ*, 423:L47–L50. 1.5

Stewart, I. M. (2009). Maximum-likelihood detection of sources among Poissonian noise. *A&A*, 495:989–1003. 4.3.1

Strüder, L., Briel, U., Dennerl, K., Hartmann, R., Kendziorra, E., Meidinger, N., Pfeffermann, E., Reppin, C., Aschenbach, B., Bornemann, W., Bräuninger, H., Burkert, W., Elender, M., Freyberg, M., Haberl, F., Hartner, G., Heuschmann, F., Hippmann, H., Kastelic, E., Kemmer, S., Kettenring, G., Kink, W., Krause, N., Müller, S., Oppitz, A., Pietsch, W., Popp, M., Predehl, P., Read, A., Stephan, K. H., Stötter, D., Trümper, J., Holl, P., Kemmer, J., Soltau, H., Stötter, R., Weber, U., Weichert, U., von Zanthier, C., Carathanassis, D., Lutz, G., Richter, R. H., Solc, P., Böttcher, H., Kuster, M., Staubert, R., Abbey, A., Holland, A., Turner, M., Balasini, M., Bignami, G. F., La Palombara, N., Villa, G., Buttler, W., Gianini, F., Lainé, R., Lumb, D., and Dhez, P. (2001). The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *A&A*, 365:L18–L26. 2.5

Szkody, P., Fraser, O., Silvestri, N., Henden, A., Anderson, S. F., Frith, J., Lawton, B., Owens, E., Raymond, S., Schmidt, G., Wolfe, M., Bochanski, J., Covey, K., Harris, H., Hawley, S., Knapp, G. R., Margon, B., Voges, W., Walkowicz, L., Brinkmann, J., and Lamb, D. Q. (2003). Cataclysmic Variables from the Sloan Digital Sky Survey. II. The Second Year. *AJ*, 126:1499–1514. 3.1, 3–1, 3.4, 3–3, 3.5, 3.6.3

Takata, J., Cheng, K. S., and Taam, R. E. (2012). X-Ray and Gamma-Ray Emissions from Rotation Powered Millisecond Pulsars. *ApJ*, 745:100. 7.4.1

Tam, C. R., Kaspi, V. M., Gaensler, B. M., and Gotthelf, E. V. (2007). Chandra monitoring of the candidate anomalous X-ray pulsar AX J1845.0-0258. *Ap&SS*, 308:519–523. A.1

Tam, P. H. T., Hui, C. Y., Huang, R. H. H., Kong, A. K. H., Takata, J., Lin, L. C. C., Yang, Y. J., Cheng, K. S., and Taam, R. E. (2010). Evidence for Gamma-ray Emission from the Low-mass X-ray Binary System First J102347.6+003841. *ApJ*, 724:L207–L211. 6.1, 7.2, 7.3

Tauris, T. M. and van den Heuvel, E. P. J. (2006). *Formation and evolution of compact stellar X-ray sources*, pages 623–665. Compact stellar X-ray sources. 1.3

Taylor, J. H. (1992). Pulsar Timing and Relativistic Gravity. *Royal Society of London Philosophical Transactions Series A*, 341:117–134. 3.3.1

The Fermi-LAT collaboration (2013). The Second Fermi Large Area Telescope Catalog of Gamma-ray Pulsars. *ArXiv e-prints*. 2.6

Thorstensen, J. R. and Armstrong, E. (2005). Is FIRST J102347.6+003841 Really a Cataclysmic Binary? *AJ*, 130:759–766. (document), 3.1, 3.2, 3.4, 3–3, 3.6.1, 4.1, 4.2, 4.4.1, 4.4.2, 7.4, 7.4.1, 7.5.1, 7.6.1

Turner, M. J. L., Abbey, A., Arnaud, M., Balasini, M., Barbera, M., Belsole, E., Bennie, P. J., Bernard, J. P., Bignami, G. F., Boer, M., Briel, U., Butler, I., Cara, C., Chabaud, C., Cole, R., Collura, A., Conte, M., Cros, A., Denby, M., Dhez, P., Di Coco, G., Dowson, J., Ferrando, P., Ghizzardi, S., Gianotti, F., Goodall, C. V., Gretton, L., Griffiths, R. G., Hainaut, O., Hochedez, J. F., Holland, A. D., Jourdain, E., Kendziorra, E., Lagostina, A., Laine, R., La Palombara, N., Lortholary, M., Lumb, D., Marty, P., Molendi, S., Pigot, C., Poindron, E., Pounds, K. A., Reeves, J. N., Reppin, C., Rothenflug, R., Salvetat, P., Sauvageot, J. L., Schmitt, D., Sembay, S., Short, A. D. T., Spragg, J., Stephen, J., Strüder, L., Tiengo, A., Trifoglio, M., Trümper, J., Vercellone, S., Vigroux, L., Villa, G., Ward, M. J., Whitehead, S., and Zonca, E. (2001). The European Photon Imaging Camera on XMM-Newton: The MOS cameras : The

MOS cameras. *A&A*, 365:L27–L35. 2.5

Vasisht, G., Gotthelf, E. V., Torii, K., and Gaensler, B. M. (2000). Detection of a Compact X-Ray Source in the Supernova Remnant G29.6+0.1: A Variable Anomalous X-Ray Pulsar? *ApJ*, 542:L49–L52. A.1

Venter, C., Harding, A. K., and Guillemot, L. (2009). Probing Millisecond Pulsar Emission Geometry Using Light Curves from the Fermi/Large Area Telescope. *ApJ*, 707:800–822. 1.2

Vink, J., Bamba, A., and Yamazaki, R. (2011). The Radiative X-ray and Gamma-ray Efficiencies of Rotation-powered Pulsars. *ApJ*, 727:131. 7.4

Wang, X., Wang, Z., and Morrell, N. (2012). Infrared Observations of The Millisecond Pulsar Binary J1023+0038: Evidence for Short-Term Nature of Its Interacting Phase in 2000–2001. *ArXiv e-prints*. 7.3

Wijnands, R. and van der Klis, M. (1998). A millisecond pulsar in an X-ray binary system. *Nature*, 394:344–346. 3.6.1, 4.4.2

Witham, A. R., Knigge, C., Gänsicke, B. T., Aungwerojwit, A., Corradi, R. L. M., Drew, J. E., Greimel, R., Groot, P. J., Morales-Rueda, L., Rodriguez-Flores, E. R., Rodriguez-Gil, P., and Steeghs, D. (2006). The properties of cataclysmic variables in photometric Hα surveys. *MNRAS*, 369:581–597. 3–3

Woods, P. M. and Thompson, C. (2006). Soft gamma repeaters and anomalous X-ray pulsars: magnetar candidates. *Compact stellar X-ray sources*, pages 547–586. A.1

Woudt, P. A., Warner, B., and Pretorius, M. L. (2004). High-speed photometry of faint cataclysmic variables - IV. V356 Aql, Aqr1, FIRST J1023+0038, Hα

0242-2802, GI Mon, AO Oct, V972 Oph, SDSS 0155+00, SDSS 0233+00, SDSS 1240-01, SDSS 1556-00, SDSS 2050-05, FH Ser. *MNRAS*, 351:1015–1025. 3.4, 3–3

Zacharias, N., Monet, D. G., Levine, S. E., Urban, S. E., Gaume, R., and Wycoff, G. L. (2004). The Naval Observatory Merged Astrometric Dataset (NOMAD). In *Bulletin of the American Astronomical Society*, volume 36 of *Bulletin of the American Astronomical Society*, page 1418. 3–1, 3.4, 4.2

Zavlin, V. E. (2007). Studying millisecond pulsars in X-rays. *Ap&SS*, 308:297–307. 4.4.1, 4.4.1

Zavlin, V. E., Pavlov, G. G., and Shibanov, Y. A. (1996). Model neutron star atmospheres with low magnetic fields. I. Atmospheres in radiative equilibrium. *A&A*, 315:141–152. 4.3.1