Binary and Millisecond Pulsars

Duncan R. Lorimer
Department of Physics
West Virginia University
P.O. Box 6315
Morgantown
WV 26506, U.S.A.
email: Duncan.Lorimer@mail.wvu.edu
http://astro.phys.wvu.edu

Abstract

We review the main properties, demographics and applications of binary and millisecond radio pulsars. Our knowledge of these exciting objects has greatly increased in recent years, mainly due to successful surveys which have brought the known pulsar population to over 1800. There are now 83 binary and millisecond pulsars associated with the disk of our Galaxy, and a further 137 pulsars in 25 of the Galactic globular clusters. Recent highlights include the discovery of a young relativistic binary system, a rejuvination in globular cluster pulsar research including growing numbers of pulsars with masses in excess of 1.5 $M_\odot$, a precise measurement of relativistic spin precession in the double pulsar system and a Galactic binary millisecond pulsar with an orbital eccentricity of 0.44.

This is a close-to-final draft of an online manuscript submitted to Living Reviews in Relativity. It is intended for beginning students and scientists new to the field of pulsar research, and should provide you with a lot of the background material that we will cover at the institute. The article should appear in the Journal later this summer and will be available as a free resource (no journal subscription necessary!). In the meantime, you may access the URL of the previous version at http://relativity.livingreviews.org/Articles/lrr-2005-7 I would be delighted to hear from you with comments or suggestions for clarification about the article.

1 Preamble

Pulsars – rapidly rotating highly magnetised neutron stars – have many applications in physics and astronomy. Striking examples include the confirmation of the existence of gravitational radiation [369, 370, 371], the first extra-solar planetary system [406, 297] and the first detection of gas in a globular cluster [118]. The diverse zoo of radio pulsars currently known is summarized in Figure 1.

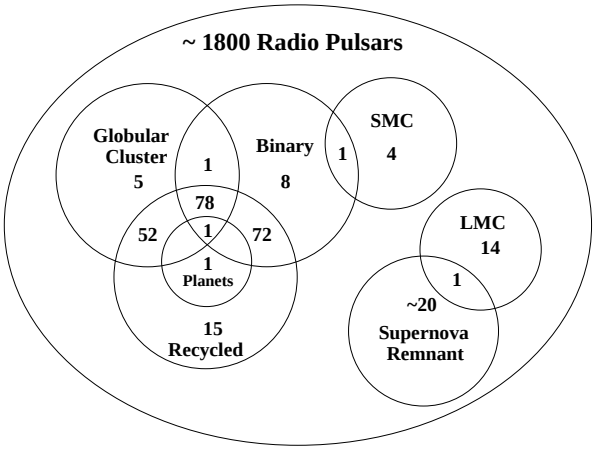

Figure 1: *Venn diagram showing the numbers and locations of the various types of radio pulsars known as of June 2008. The large and small Magellanic clouds are denoted by LMC and SMC.*

Pulsar research has proceeded at a rapid pace during the first three versions of this article [225, 226, 228]. Surveys with the Parkes radio telescope [311], at Green Bank [282], Arecibo [278] and the Giant Metre Wave Radio Telescope [281] have more than doubled the number of pulsars known a decade ago. With new instrumentation coming online, and new telescopes planned, these are exciting times for pulsar astronomy.

The aims of this review are to introduce the reader to the field and to focus on some of the many applications of pulsar research in relativistic astrophysics. We begin in Section 2 with an overview of the pulsar phenomenon, a review of the key observed population properties, the origin and evolution of pulsars and the main search strategies. In Section 3, we review present understanding in pulsar demography, discussing selection effects and their correction techniques. This leads to empirical estimates of the total number of normal and millisecond pulsars (see Section 3.3) and relativistic binaries (see Section 3.4) in the Galaxy and has implications for the detection of gravitational radiation by current and planned telescopes. Our review of pulsar timing in Section 4 covers the basic techniques (see Section 4.2), timing stability (see Section 4.3), binary pulsars (see Section 4.4), and using pulsars as sensitive detectors of long-period gravitational waves (see Section 4.7). We conclude with a brief outlook to the future in Section 5. Up-to-date tables of parameters of binary and millisecond pulsars are included in Appendix A.

2 Pulsar Phenomenology

Most of the basic observational facts about radio pulsars were established shortly after their discovery [146] in 1967. Although there are still many open questions, the basic model has long been established beyond all reasonable doubt, i.e. pulsars are rapidly rotating, highly magnetised neutron stars formed during the supernova explosions of massive stars.

2.1 The lighthouse model

Figure 2 shows an animation depicting the rotating neutron star, also known as the "lighthouse" model. As the neutron star spins, charged particles are accelerated along magnetic field lines in the magnetosphere (depicted by the light blue cones). The accelerating particles emit electromagnetic radiation, most readily detected at radio frequencies as a sequence of observed pulses. Each pulse is produced as the magnetic axis (and hence the radiation beam) crosses the observer's line of sight each rotation. The repetition period of the pulses is therefore simply the rotation period of the neutron star. The moving "tracker ball" on the pulse profile in the animation shows the relationship between observed intensity and rotational phase of the neutron star.

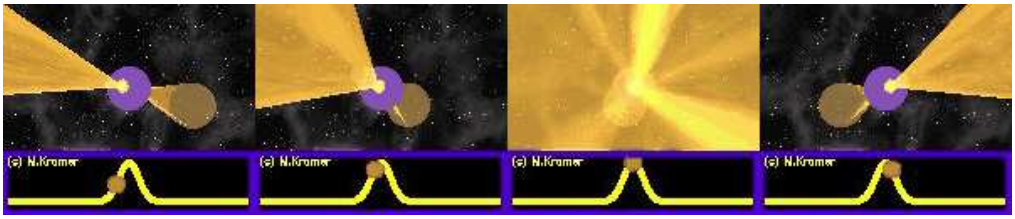

Figure 2: *Animated GIF showing the rotating neutron star (or "lighthouse") model for pulsar emission. Animation designed by Michael Kramer.*

Neutron stars are essentially large celestial flywheels with moments of inertia $\sim 10^{38}$ kg m^2. The rotating neutron star model [294, 126] predicts a gradual slowdown and hence an increase in the pulse period as the outgoing radiation carries away rotational kinetic energy. This idea gained strong support when a period increase of 36.5 ns per day was measured for the pulsar B0531+21[1] in the Crab nebula [324], which implied that a rotating neutron star with a large magnetic field must be the dominant energy supply for the nebula [127].

2.2 Pulse periods and slowdown rates

The pulsar catalogue [252, 15] contains up-to-date parameters for 1775 pulsars. Most of these are "normal" with pulse periods $P \sim 0.5$ s which increase secularly at rates $\dot{P} \sim 10^{-15}$ s/s. A growing fraction are "millisecond pulsars", with 1.4 ms $\lesssim P \lesssim$ 30 ms and $\dot{P} \lesssim 10^{-19}$ s/s. As shown in the "P–$\dot{P}$ diagram" in Figure 3, normal and millisecond pulsars are distinct populations.

The differences in P and $\dot{P}$ imply fundamentally different magnetic field strengths and ages. Treating the pulsar as a rotating magnetic dipole, one may show [232] that the surface magnetic field strength $B \propto (P\dot{P})^{1/2}$ and the characteristic age $\tau_c = P/(2\dot{P})$. Lines of constant B and τ_c are drawn on Figure 3, from which we infer typical values of 10^{12} G and 10^7 yr for the normal pulsars

[1]Like most astronomical sources, pulsars are named after their position in the sky. Those pulsars discovered prior to the mid 1990s are named based on the Besselian equatorial coordinate system (B1950) and have a B prefix followed by their right ascension and declination. More recently discovered pulsars follow the Julian (J2000) epoch. Pulsars in globular clusters, where the positional designation is not unique, have additional characters to distinguish them. For example PSR J1748−2446ad in the globular cluster Terzan 5 [145].

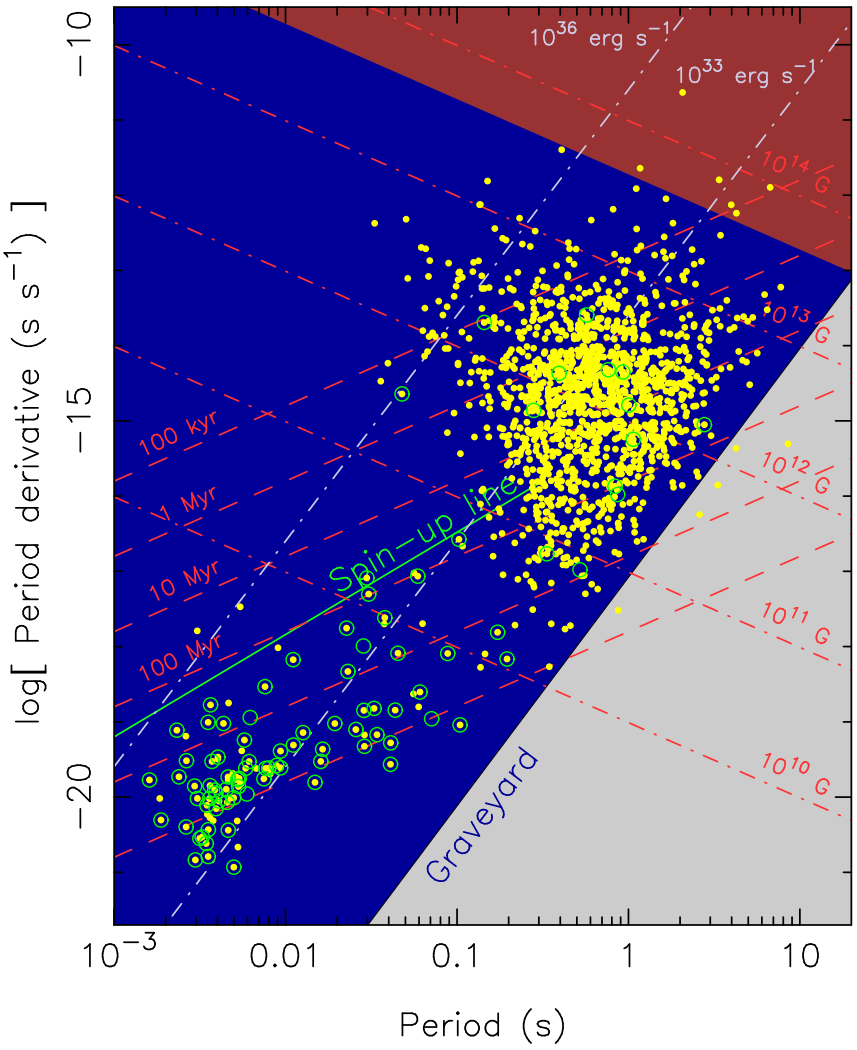

Figure 3: *The P–$\dot{P}$ diagram showing the current sample of radio pulsars. Binary pulsars are highlighted by open circles. Theoretical models [77] do not predict radio emission outside the dark blue region. Figure provided by Michael Kramer.*

and 10^8 G and 10^9 yr for the millisecond pulsars. For the rate of loss of kinetic energy, sometimes called the spin-down luminosity, we have $\dot{E} \propto \dot{P}/P^3$. The lines of constant $\dot{E}$ shown on Figure 3 indicate that the most energetic objects are the very young normal pulsars and the most rapidly spinning millisecond pulsars.

The most rapidly rotating neutron star currently known, J1748−2446ad,with a spin rate of 716 Hz, resides in the globular cluster Terzan 5 [145]. As discussed by Lattimer & Prakash [209], the limiting (non-rotating) radius of a 1.4 $M_\odot$ neutron star with this period is 14.3 km. If a precise measurement for the pulsar mass can be made through future timing measurements (Section 4), then this pulsar could be a very useful probe of the equation of state of super dense matter.

While the hunt for more rapidly rotating pulsars and even "sub-millisecond pulsars" continues, and most neutron star equations of state allow higher spin rates than 716 Hz, it has been suggested [41] that the dearth of pulsars with $P < 1.5$ ms is caused by gravitational wave emission from Rossby-mode instabilities [8]. The most rapidly rotating pulsars [145] are predominantly members of eclipsing binary systems which could hamper their detection in radio surveys. Independent constraints on the limiting spin frequencies of neutron stars come from studies of millisecond X-ray binaries [69] which are not thought to be selection-effect limited [72]. This analysis does not predict a significant population of neutron stars with spin rates in excess of 730 Hz.

2.3 Pulse profiles

Pulsars are weak radio sources. Measured intensities vary between 20μJy and 5 Jy (1 Jy $\equiv$ 10^{-26} W m^{-2} Hz^{-1}). As a result, even with a large radio telescope, the coherent addition of many hundreds or even thousands of pulses is usually required in order to produce a detectable "integrated profile". Remarkably, although the individual pulses vary dramatically, the integrated profile at a given observing frequency is very stable and can be thought of as a fingerprint of the neutron star's emission beam. Profile stability is of key importance in pulsar timing measurements discussed in Section 4.

The selection of integrated profiles in Figure 4 shows a rich diversity in morphology including two examples of "interpulses" – a secondary pulse separated by about 180 degrees from the main pulse. One interpretation for this phenomenon is that the two pulses originate from opposite magnetic poles of the neutron star (see however [253]). Since this is an unlikely viewing angle, we would expect interpulses to be a rare phenomenon. Indeed, the fraction of known pulsars in which interpulses are observed in their pulse profiles is only a few percent [198]. Recent work [394] on the statistics of interpulses among the known population show that the interpulse fraction depends inversely on pulse period. The origin of this effect could be due to alignment between the spin and magnetic axis of neutron stars on timescales of order 10^7 yr [394]. If we take this result at face value, and assume that all millisecond pulsars descend from normal pulsars (Section 2.6), then the implication is that millisecond pulsars should be preferentially aligned rotators. However, their appears to be no strong evidence in favour of this expectation based on the pulse profile morphology of millisecond pulsars [206].

Two contrasting phenomenological models have been proposed to explain the observed pulse shapes. The "core and cone" model [314] depicts the beam as a core surrounded by a series of nested cones. Alternatively, the "patchy beam" model [247, 132] has the beam populated by a series of randomly-distributed emitting regions. Recent work suggests that the observational data can be better explained by a hybrid empirical model depicted in Figure 5 which employs patchy beams in a core and cone structure [183].

A key feature of this new model is that the emission height of young pulsars is radically different from that of the older population. Monte Carlo simulations using this phenomenological model appear to be very successful at explaining the rich diversity of pulse shapes. Further work in this area is necessary to understand the origin of this model, and improve our understanding of the

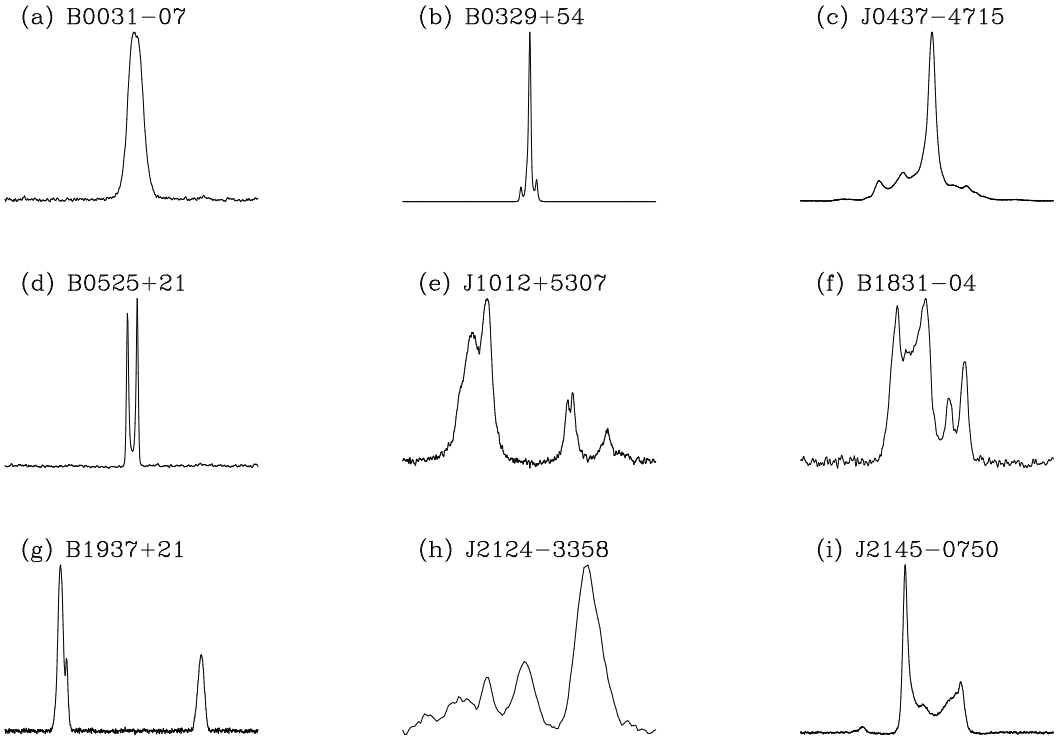

Figure 4: *A variety of integrated pulse profiles taken from the available literature. References: Panels a, b, d, f [128], Panel c [26], Panels e, g, i [206], Panel h [34]. Each profile represents 360 degrees of rotational phase. These profiles are freely available from an on-line database [381].*

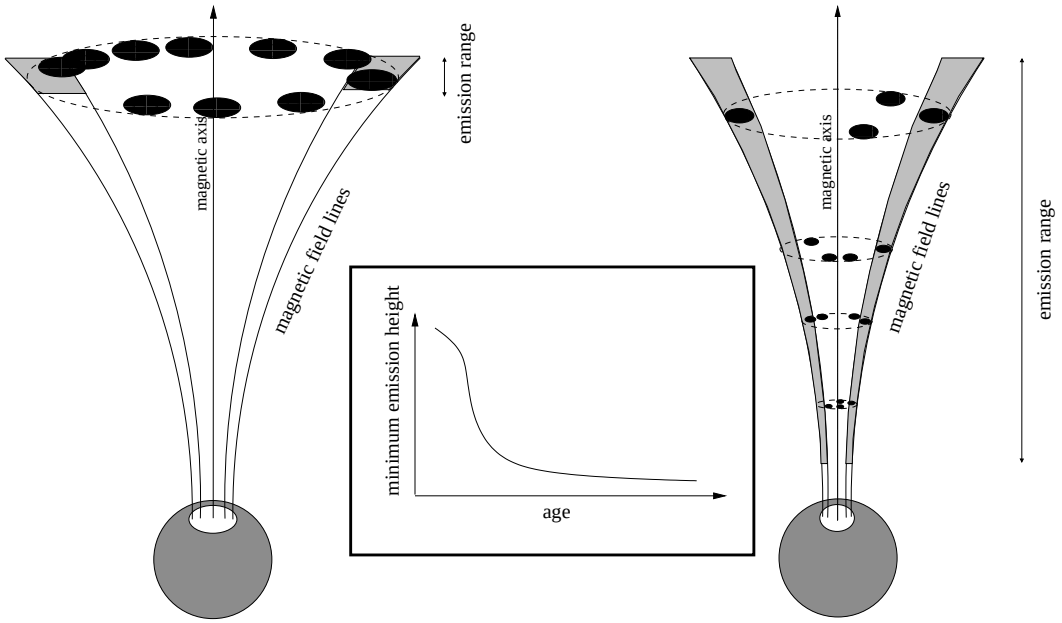

Figure 5: *A recent phenomenological model for pulse shape morphology. The neutron star is depicted by the grey sphere and only a single magnetic pole is shown for clarity. Left: a young pulsar with emission from a patchy conal ring at high altitudes from the surface of the neutron star. Right: an older pulsar in which emission emanates from a series of patchy rings over a range of altitudes. Centre: schematic representation of the change in emission height with pulsar age. Figure designed by Aris Karastergiou and Simon Johnston [183].*

shape and evolution of pulsar beams and fraction of sky they cover. This is of key importance to the results of population studies reviewed in Section 3.2.

2.4 The pulsar distance scale

Quantitative estimates of the distance to each pulsar can be made from the measurement of *pulse dispersion* – the delay in pulse arrival times across a finite bandwidth. Dispersion occurs because the group velocity of the pulsed radiation through the ionised component of the interstellar medium is frequency dependent. As shown in Figure 6, pulses emitted at lower radio frequencies travel slower through the interstellar medium, arriving later than those emitted at higher frequencies.

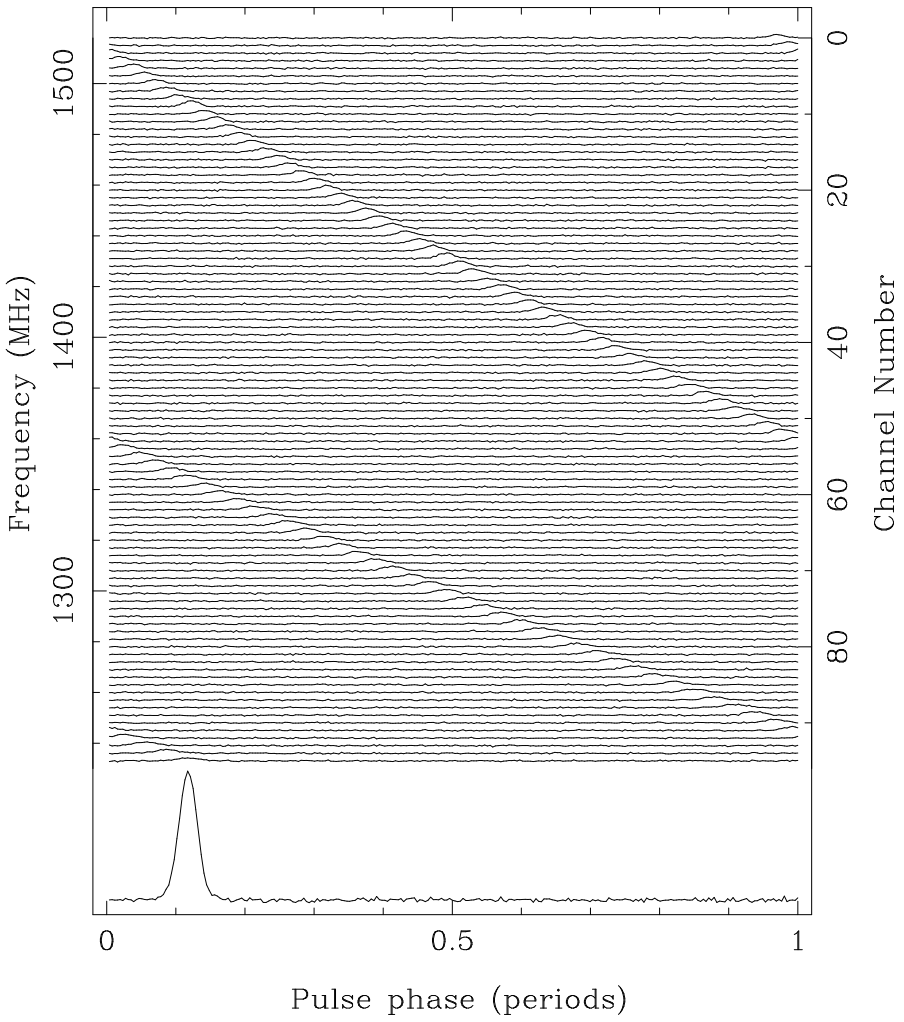

Figure 6: *Pulse dispersion shown in this Parkes observation of the 128 ms pulsar B1356–60. The dispersion measure is 295 cm^{-3} pc. The quadratic frequency dependence of the dispersion delay is clearly visible. Figure provided by Andrew Lyne.*

Quantitatively, the delay Δt in arrival times between a high frequency $\nu_{\rm hi}$ and a low frequency

$\nu_{\rm lo}$ pulse can be shown [232] to be

$$\Delta t = 4.15 \text{ ms} \times \left[\left(\frac{\nu_{\rm lo}}{\rm GHz} \right)^{-2} - \left(\frac{\nu_{\rm hi}}{\rm GHz} \right)^{-2} \right] \times \left(\frac{\rm DM}{\rm cm^{-3}\ pc} \right), \quad (1)$$

where the dispersion measure

$$\rm DM = \int_0^d n_e \, dl, \quad (2)$$

is the integrated column density of electrons, $n_{\rm e}$, out to the pulsar at a distance d. This equation may be solved for d given a measurement of DM and a model of the free electron distribution calibrated from the 100 or so pulsars with independent distance estimates and measurements of scattering for lines of sight towards various Galactic and extragalactic sources [367, 392]. A recent model of this kind, known as NE2001 [88, 89], provides distance estimates with an average uncertainty of $\sim 30\%$.

Because the electron density models are only as good as the scope of their input data allow, one should be mindful of systematic uncertainties. For example, studies of the Parkes multibeam pulsar distribution [203, 221] suggest that the NE2001 model underestimates the distances of pulsars close to the Galactic plane. This suspicion has been dramatically confirmed recently by an extensive analysis [124] of pulsar DMs and measurements of Galactic Hα emission. This work shows that the distribution of Galactic free electrons to be exponential in form with a scale height of 1830^{+120}_{-250} pc. This value is a factor of two higher than previously thought. A revised version of the NE2001 model which takes into account these and other developments is currently in preparation [84].

2.5 Pulsars in binary systems

As can be inferred from Figure 1, only a few percent of all known pulsars in the Galactic disk are members of binary systems. Timing measurements (see Section 4) place useful constraints on the masses of the companions which, often supplemented by observations at other wavelengths, tell us a great deal about their nature. The present sample of orbiting companions are either white dwarfs, main sequence stars or other neutron stars. Two notable hybrid systems are the "planet pulsars" B1257+12 and B1620−26. PSR B1257+12 is a 6.2-ms pulsar accompanied by at least three terrestrial-mass bodies [406, 297, 405] while B1620−26, an 11-ms pulsar in the globular cluster M4, is part of a triple system with a 1–2 $M_{\rm Jupiter}$ planet [373, 19, 372, 341] orbiting a neutron star–white dwarf stellar binary system. The current limits from pulsar timing favour a roughly 45-yr mildly-eccentric ($e \sim 0.16$) orbit with semi-major axis ~ 25 AU. Despite several tentative claims over the years, no other convincing cases for planetary companions to pulsars exist. Orbiting companions are much more common around millisecond pulsars ($\sim 80\%$ of the observed sample) than around the normal pulsars ($\lesssim 1\%$). In general, binary systems with low-mass companions ($\lesssim 0.7\, M_\odot$ – predominantly white dwarfs) have essentially circular orbits: $10^{-5} \lesssim e \lesssim 0.01$. Binary pulsars with high-mass companions ($\gtrsim 1\, M_\odot$ – massive white dwarfs, other neutron stars or main sequence stars) tend to have more eccentric orbits, $0.15 \gtrsim e \gtrsim 0.9$.

2.6 Evolution of normal and millisecond pulsars

A simplified version of the presently favoured model [42, 114, 342, 5] to explain the formation of the various types of systems observed is shown in Figure 7. Starting with a binary star system, a neutron star is formed during the supernova explosion of the initially more massive star. From the virial theorem, in the absence of any other factors, the binary will be disrupted if more than half the total pre-supernova mass is ejected from the system during the (assumed symmetric) explosion [147, 39]. In practice, the fraction of surviving binaries is also affected by the magnitude

and direction of any impulsive "kick" velocity the neutron star receives at birth from a slightly asymmetric explosion [147, 23]. Binaries that disrupt produce a high-velocity isolated neutron star and an OB runaway star [43]. The high probability of disruption explains qualitatively why so few normal pulsars have companions. Those that survive will likely have high orbital eccentricities due to the violent conditions in the supernova explosion. There are currently four known normal radio pulsars with massive main sequence companions in eccentric orbits which are examples of binary systems which survived the supernova explosion [173, 187, 351, 352, 240]. Over the next 10^{7-8} yr after the explosion, the neutron star may be observable as a normal radio pulsar spinning down to a period $\gtrsim$ several seconds. After this time, the energy output of the star diminishes to a point where it no longer produces significant amounts of radio emission.

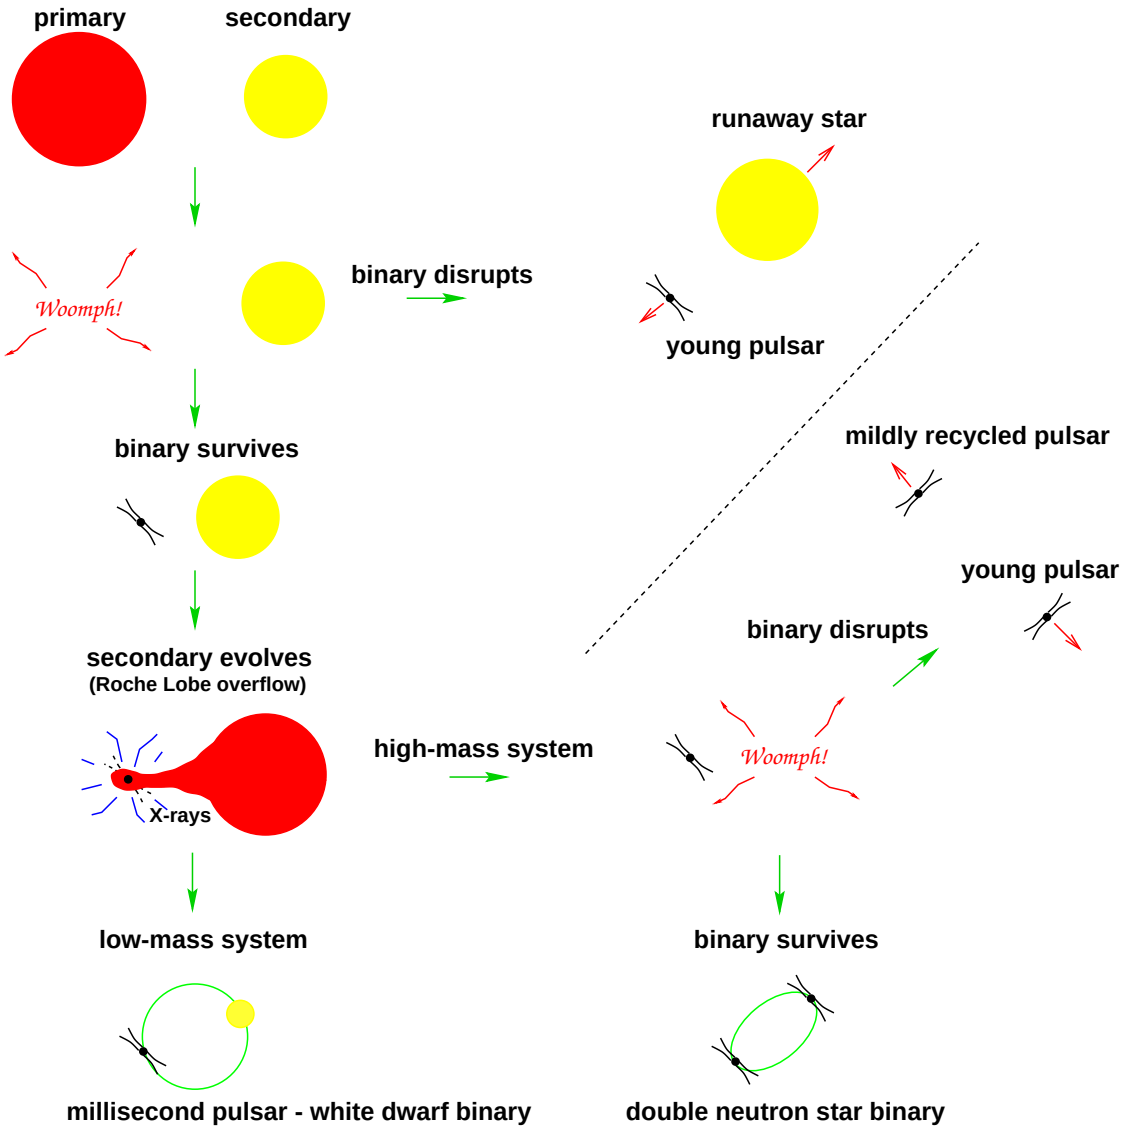

Figure 7: *Cartoon showing various evolutionary scenarios involving binary pulsars.*

For those few binaries that remain bound, and in which the companion is sufficiently massive

to evolve into a giant and overflow its Roche lobe, the old spun-down neutron star can gain a new lease of life as a pulsar by accreting matter and angular momentum at the expense of the orbital angular momentum of the binary system [5]. The term "recycled pulsar" is often used to describe such objects. During this accretion phase, the X-rays produced by the frictional heating of the infalling matter onto the neutron star mean that such a system is expected to be visible as an X-ray binary. Two classes of X-ray binaries relevant to binary and millisecond pulsars exist: neutron stars with high-mass or low-mass companions. Detailed reviews of the X-ray binary population, including systems likely to contain black holes, can be found elsewhere [39, 323].

2.6.1 High-mass systems

In a high-mass X-ray binary, the companion is massive enough that it also explodes as a supernova, producing a second neutron star. If the binary system is lucky enough to survive the explosion, the result is a double neutron star binary. Nine such systems are now known, the original example being PSR B1913+16 [155] – a 59-ms radio pulsar which orbits its companion every 7.75 hr [369, 370]. In this formation scenario, PSR B1913+16 is an example of the older, first-born, neutron star that has subsequently accreted matter from its companion.

For many years, no clear example was known where the second-born neutron star was observed as a pulsar. The discovery of the double pulsar J0737−3039 [51, 243], where a 22.7-ms recycled pulsar "A" orbits a 2.77-s normal pulsar "B" every 2.4 hr, has now provided a dramatic confirmation of this evolutionary model in which we identify A and B as the first and second-born neutron stars respectively. Just how many more observable double pulsar systems exist in our Galaxy is not clear. While the population of double neutron star systems in general is reasonably well understood (see Section 3.4.1), a number of effects conspire to reduce the detectability of double pulsar systems. First, the lifetime of the second born pulsar cannot be prolonged by accretion and spin-up and hence is likely to be less than one tenth that of the recycled pulsar. Second, the radio beam of the longer period second-born pulsar is likely to be much smaller than its spun-up partner making it harder to detect (see Section 3.2.3). Finally, as discussed in Section 4.4, the PSR J0737–3039 system is viewed almost perfectly edge-on to the line of sight and there is strong evidence [263] that the wind of the more rapidly spinning A pulsar is impinging on B's magnetosphere which is causing it to radiate [220]. In summary, the prospects of ever finding more than a few 0737-like systems appears to be rather low.

A notable recent addition to the sample of double neutron star binaries is PSR J1906+0746, a 144-ms pulsar in a highly relativistic 4-hr orbit with an eccentricity of 0.085 [223]. Timing observations show the pulsar to be young with a characteristic age of only 10^5 yr. Measurements of the relativistic periastron advance and gravitational redshift (see Section 4.4) constrain the masses of the pulsar and its companion to be 1.25 $M_\odot$ and 1.37 $M_\odot$ respectively [184]. We note, however, that the pulsar exhibits significant amounts of timing noise (see Section 4.3), making the analysis of the system far from trivial. Taken at face value, these parameters suggest that the companion is also a neutron star, and the system appears to be a younger version of the double pulsar. Despite intensive searches, no pulsations from the companion star (expected to be a recycled radio pulsar) have so far been observed [223]. This could be due to unfavourable beaming and/or intrinsic radio faintness.

2.6.2 Low-mass systems

The companion in a low-mass X-ray binary evolves and transfers matter onto the neutron star on a much longer time-scale, spinning it up to periods as short as a few ms [5]. Tidal forces during the accretion process serve to circularize the orbit. At the end of the spin-up phase, the secondary sheds its outer layers to become a white dwarf in orbit around a rapidly spinning millisecond pulsar. This model has gained strong support in recent years from the discoveries of quasi-periodic kHz

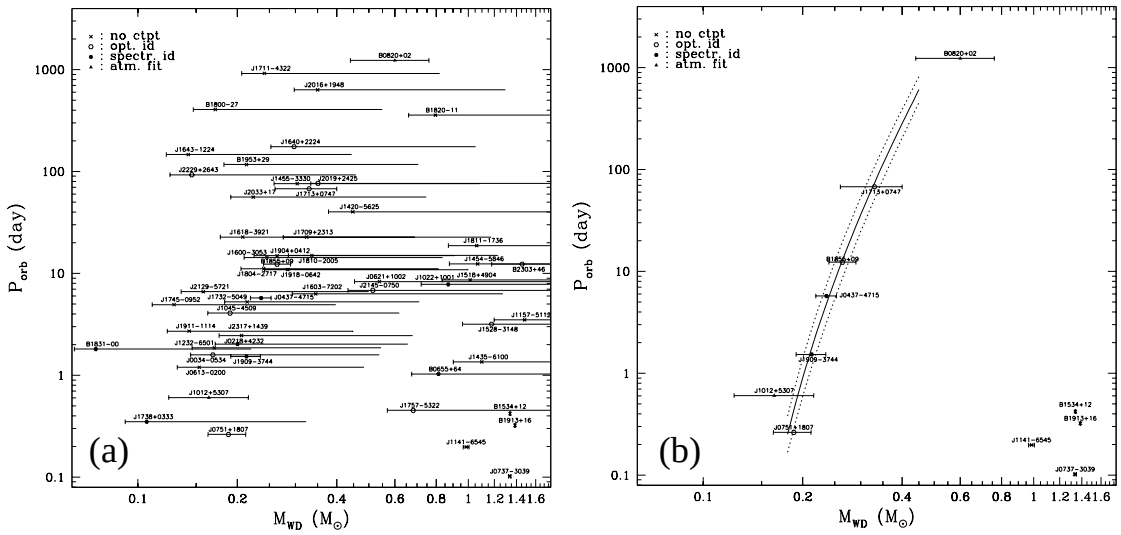

Figure 8: *Orbital period versus companion mass for binary pulsars showing the whole sample where, in the absence of mass determinations, statistical arguments based on a random distribution of orbital inclination angles (see Section 4.4) have been used to constrain the masses as shown (Panel a), and only those with well determined companion masses (Panel b). The dashed lines show the uncertainties in the predicted relation [364]. This relationship indicates that as these systems finished a period of stable mass transfer due to Roche-lobe overflow, the size and hence period of the orbit was determined by the mass of the evolved secondary star. Figure provided by Marten van Kerkwijk [386].*

oscillations in a number of low-mass X-ray binaries [398], as well as Doppler-shifted 2.49-ms X-ray pulsations from the transient X-ray burster SAX J1808.4−3658 [400, 68]. Seven other "X-ray millisecond pulsars" are now known with spin rates and orbital periods ranging between 185–600 Hz and 40 min–4.3 hr respectively [399, 267, 177, 207].

Numerous examples of these systems in their post X-ray phase are now seen as the millisecond pulsar–white dwarf binary systems. Presently, 20 of these systems have compelling optical identifications of the white dwarf companion, and upper limits or tentative detections have been found in about 30 others [386]. Comparisons between the cooling ages of the white dwarfs and the millisecond pulsars confirm the age of these systems and suggest that the accretion rate during the spin-up phase was well below the Eddington limit [137].

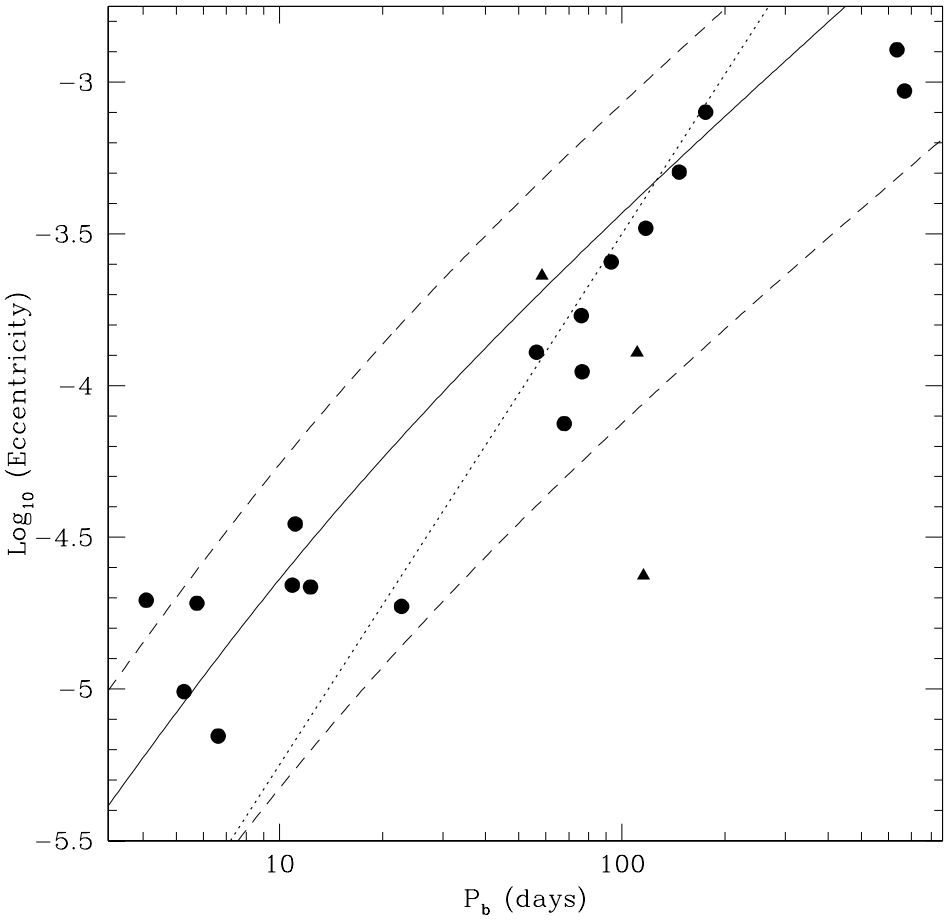

Figure 9: *Eccentricity versus orbital period for a sample of 21 low-mass binary pulsars which are not in globular clusters, with the triangles denoting three recently discovered systems [350]. The solid line shows the median of the predicted relationship between orbital period and eccentricity [302]. Dashed lines show 95% the confidence limit about this relationship. The dotted line shows* $P_{\rm b} \propto e^2$. *Figure provided by Ingrid Stairs [350] using an adaptation of the orbital period-eccentricity relationship tabulated by Fernando Camilo.*

Further support for the above evolutionary scenarios comes from two correlations in the observed sample of low-mass binary pulsars. Firstly, as seen in Figure 9, there is a strong correlation between orbital period and eccentricity. The data are in very good agreement with a theoretical

relationship which predicts a relic orbital eccentricity due to convective eddy currents in the accretion process [302]. Secondly, as shown in Panel b of Figure 8, where companion masses have been measured accurately, through radio timing (see Section 4.4) and/or through optical observations [386], they are in good agreement with a relation between companion mass and orbital period predicted by binary evolution theory [364]. A word of caution is required in using these models to make predictions, however. When confronted with a larger ensemble of binary pulsars using statistical arguments to constrain the companion masses (see Panel a of Figure 8), current models have problems in explaining the full range of orbital periods on this diagram [350].

2.7 Intermediate mass binary pulsars

The range of white dwarf masses observed is becoming broader. Since this article originally appeared [225], the number of "intermediate-mass binary pulsars" [56] has grown significantly [60]. These systems are distinct to the millisecond pulsar–white dwarf binaries in several ways: (i) the spin period of the radio pulsar is generally longer (9–200 ms); (ii) the mass of the white dwarf is larger (typically $\gtrsim 0.5\,M_{\odot}$); (iii) the orbit, while still essentially circular, is often significantly more eccentric ($e \gtrsim 10^{-3}$); (iv) the binary parameters do not necessarily follow the mass–period or eccentricity–period relationships. It is not presently clear whether these systems originated from low- or high-mass X-ray binaries. It was suggested by van den Heuvel [385] that they have more in common with high-mass systems. Subsequently, Li proposed [213] that a thermal-viscous instability in the accretion disk of a low-mass X-ray binary could truncate the accretion phase and produce a more slowly spinning neutron star.

2.8 Isolated recycled pulsars

The scenarios outlined qualitatively above represent a reasonable understanding of binary evolution. There are, however, a number of pulsars with spin properties that suggest a phase of recycling took place but have no orbiting companions. While the existence of such systems in globular clusters are more readily explained by the high probability of stellar interactions compared to the disk [340], it is somewhat surprising to find them in the Galactic disk. Out of a total of 72 Galactic millisecond pulsars, 16 are isolated (see Table 2). Although it has been proposed that these millisecond pulsars have ablated their companion via their strong relativistic winds [196] as may be happening in the PSR B1957+20 system [123], it is not clear whether the energetics or time-scales for this process are feasible [211]. There is some observational evidence that suggests that solitary millisecond pulsars are less luminous than binary millisecond pulsars [26, 206]. If confirmed by future discoveries, this would need to be explained by any viable evolutionary model.

There are four further "anomalous" isolated pulsars with periods in the range 28–60 ms [61, 235]. When placed on the P–$\dot{P}$ diagram, these objects populate the region occupied by the double neutron star binaries. The most natural explanation for their existence, therefore, is that they are "failed double neutron star binaries" which disrupted during the supernova explosion of the secondary [61]. A simple calculation [235], suggested that for every double neutron star we should see of order ten such isolated objects. Recent work [32] has investigated why so few are observed. Using the most recent population synthesis models to follow the evolution of binary systems [31], it appears that the discrepancy may not be as significant as previously supposed. In particular, the space velocity distribution of surviving binary systems is narrower than for the isolated objects that were during the second supernova explosion. The isolated systems occupy a larger volume of the Galaxy than the surviving binaries and are harder to detect. When this selection effect is accounted for [32], the relative sample sizes appear to be consistent with the disruption hypothesis.

2.9 A new class of millisecond pulsars?

One of the most remarkable recent discoveries is the binary pulsar J1903+0327 [71]. Found in an on-going multibeam survey with the Arecibo telescope [279, 83], this 2.15-ms pulsar is distinct from all other millisecond pulsars in that its 95-day orbit has an eccentricity of 0.43! In addition, timing measurements of the relativistic periastron advance and Shapiro delay in this system (see Section 4.4), show the mass of the pulsar to be 1.74 ± 0.04 $M_\odot$ and the companion star to be 1.051 ± 0.015 $M_\odot$. Optical observations show a possible counterpart which is consistent with a 1 $M_\odot$ star. While similar systems have been observed in globular clusters (e.g. PSR J0514−4002A in NGC 1851[119]), presumably a result of exchange interactions, the standard recycling hypothesis outlined in Section 2.6 cannot account for pulsars like J1903+0327 in the Galactic disk.

How could such an eccentric binary millisecond pulsar system form? One possibility is that the binary system was produced in an exchange interaction in a globular cluster and subsequently ejected, or the cluster has since disrupted. Statistical estimates [71] of the likelihood of both these channels are in the range 1–10%, implying that a globular cluster origin cannot be ruled out.

Another possibility is that the pulsar is a member of an hierarchical triple system with a one solar mass white dwarf in the 95-day orbit, and a main sequence star in a much wider and highly inclined orbit which has so far not been revealed by timing. The origin of the high eccentricity is through perturbations from the outer star, the so-called Kozai mechanism [201]. Formation estimates based on observational data on stellar multiplicity [308] find that around 4% of all binary millisecond pulsars are expected to be triple systems [71]. The existence of a single triple system among the current sample of millisecond pulsars appears to be consistent with this hypothesis.

If future observations of the proposed optical counterpart confirm it as the binary companion through spectral line measurements of orbital Doppler shifts, the above triple-system scenario will be ruled out. Such an observation would favour a hybrid scenario suggested by van den Heuvel [384] in which the white dwarf and pulsar merge due to gravitational radiation losses. Tidal disruption of the white dwarf in the inspiral would produce an accretion disk and induce an eccentricity in the orbit of the outer star leaving behind an eccentric binary system. This is the only currently proposed model that can naturally account for the high pulsar mass observed in this system which could arise from accretion of a white-dwarf debris disk following coalescence.

2.10 Pulsar velocities

Pulsars have long been known to have space velocities at least an order of magnitude larger than those of their main sequence progenitors, which have typical values between 10 and 50 km s^{-1}. Proper motions for over 250 pulsars have now been measured largely by radio timing and interferometric techniques [241, 27, 115, 138, 149, 408]. These data imply a broad velocity spectrum ranging from 0 to over 1000 km s^{-1} [246], with the current record holder being PSR B1508+55 [75], with a proper motion and parallax measurement implying a transverse velocity of 1083^{+103}_{-90} km s^{-1}. As Figure 10 illustrates, high-velocity pulsars born close to the Galactic plane quickly migrate to higher Galactic latitudes. Given such a broad velocity spectrum, as many as half of all pulsars will eventually escape the Galactic gravitational potential [246, 86].

Such large velocities are perhaps not surprising, given the violent conditions under which neutron stars are formed. If the explosion is only slightly asymmetric, an impulsive "kick" velocity of up to 1000 km s^{-1} can be imparted to the neutron star [339]. In addition, if the neutron star progenitor was a member of a binary system prior to the explosion, the pre-supernova orbital velocity will also contribute to the resulting speed of the newly-formed pulsar. The relative contributions of these two factors to the overall pulsar birth velocity distribution is currently not well understood.

The distribution of pulsar velocities has a high velocity component due to the normal pulsars [246, 149, 110], and a lower velocity component from binary and millisecond pulsars [224, 85, 248, 149]. The main reason for this dichotomy appears to be that, in order to survive and subsequently

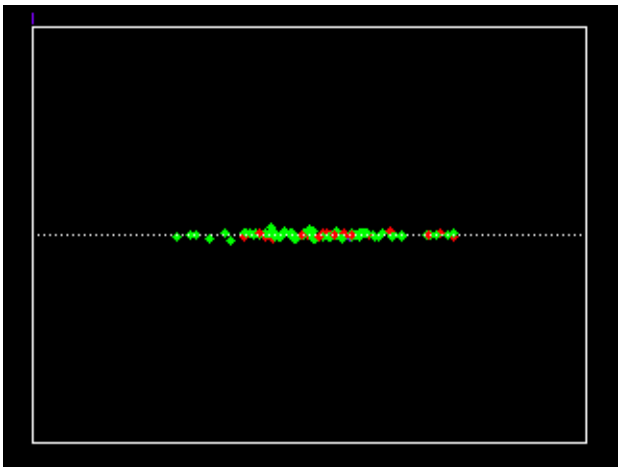

Figure 10: *GIF movie showing a simulation following the motion of 100 pulsars in a model gravitational potential of our Galaxy for* 200 Myr. *The view is edge-on, i.e. the horizontal axis represents the Galactic plane (*30 kpc *across) while the vertical axis represents* ±10 kpc *from the plane. This snapshot shows the initial configuration of young neutron stars.*

form recycled pulsars through the accretion process outlined above, the binary systems contain only those neutron stars with lower birth velocities.

Further insights into pulsar kicks from analyses of proper motion and polarization data [169, 170, 313] find strong evidence is found for an alignment between the spin axis and the velocity vector at birth. These data have recently been combined with modeling of pulsar-wind nebulae [285], where strong evidence is found for a model in which the natal impulse is provided by an anisotropic flux of neutrinos from the proto-neutron star on timescales of a few seconds.

2.11 Pulsar searches

The radio sky is being repeatedly searched for new pulsars in a variety of ways. In the following, we outline the major search strategies that are optimized for binary and millisecond pulsars.

2.11.1 All-sky searches

The oldest radio pulsars form a relaxed population of stars oscillating in the Galactic gravitational potential [136]. The scale height for such a population is at least 500 pc [228], about 10 times that of the massive stars which populate the Galactic plane. Since the typical ages of millisecond pulsars are several Gyr or more, we expect, from our vantage point in the Galaxy, to be in the middle of an essentially isotropic population of nearby sources. All-sky searches for millisecond pulsars at high Galactic latitudes have been very effective in probing this population.

Motivated by the discovery of two recycled pulsars at high latitudes with the Arecibo telescope [403, 406], surveys carried out at Arecibo, Parkes, Jodrell Bank and Green Bank by others in the 1990s [61, 62, 63, 255, 248] discovered around 30 further objects. Although further searching of this kind has been carried out at Arecibo in the past decade [237], much of the recent efforts have been concentrated along the plane of our Galaxy and in globular clusters discussed below. Very recently, however, 12,000 square degrees of sky was surveyed using a new 350-MHz receiver on the Green Bank telescope [395, 46]. Processing of these data is currently underway, with two millisecond pulsars found with around 10% of the dataset analysed.

2.11.2 Searches close to the plane of our Galaxy

Young pulsars are most likely to be found near to their places of birth close to the Galactic plane. This was the target region of the main Parkes multibeam survey and has so far resulted in the discovery of 783 pulsars [254, 268, 203, 150, 112, 221, 191, 52], almost half the number currently known! Such a large haul inevitably results in a number of interesting individual objects such as the relativistic binary pulsar J1141−6545 [188, 292, 28, 153, 38], a young pulsar orbiting an $\sim 11\,M_\odot$ star (probably a main sequence B-star [351, 352]), a young pulsar in a ~ 5 yr-eccentric orbit ($e = 0.955$; the most eccentric found so far) around a 10–20 $M_\odot$ companion [240, 221], several intermediate-mass binary pulsars [60], and two double neutron star binaries [244, 111]. Further analyses of this rich data set are now in progress and will ensure yet more discoveries in the near future.

Motivated by the successes at Parkes, a multibeam survey is now in progress with the Arecibo telescope [279] and the Effelsberg radio telescope. The Arecibo survey has so far discovered 46 pulsars [279, 83] with notable finds including a highly relativistic binary [223] and an eccentric millisecond pulsar binary [71]. Hundreds more pulsars could be found in this survey over the next five years. A significant fraction of this yield are expected to be distant millisecond pulsars in the disk of our Galaxy. With the advent of sensitive low-noise receivers at lower observing frequencies, surveys of the Galactic plane are being carried out with the GMRT [175], Green Bank [144] and Westerbork [329]. At the time of writing, no new millisecond pulsars have been found in these searches, though significant amounts of data remain to be fully processed.

2.11.3 Searches at intermediate and high Galactic latitudes

To probe more deeply into the population of millisecond and recycled pulsars than possible at high Galactic latitudes, the Parkes multibeam system was also used to survey intermediate latitudes [106, 104]. Among the 69 new pulsars found in the survey, 8 are relatively distant recycled objects. Two of the new recycled pulsars from this survey [104] are mildly relativistic neutron star-white dwarf binaries. An analysis of the full results from this survey should significantly improve our knowledge on the Galaxy-wide population and birth-rate of millisecond pulsars. Arecibo surveys at intermediate latitudes also continue to find new pulsars, such as the long-period binaries J2016+1948 and J0407+1607 [283, 237], and the likely double neutron star system J1829+2456 [73].

Although the density of pulsars decreases with increasing Galactic latitude, discoveries away from the plane provide strong constraints on the scale height of the millisecond pulsar population. Two recent surveys with the Parkes multibeam system [52, 160] have resulted in a number of interesting discoveries. Pulsars at high latitudes are especially important for the millisecond pulsar timing array (Section 4.7.3) which benefits from widely separated pulsars on the sky to search for correlations in the cosmic gravitational wave background on large angular scales.

2.11.4 Targeted searches of globular clusters

Globular clusters have long been known to be breeding grounds for millisecond and binary pulsars [66]. The main reason for this is the high stellar density and consequently high rate of stellar interaction in globular clusters relative to most of the rest of the Galaxy. As a result, low-mass X-ray binaries are almost 10 times more abundant in clusters than in the Galactic disk. In addition, exchange interactions between binary and multiple systems in the cluster can result in the formation of exotic binary systems [340]. To date, searches have revealed 137 pulsars in 25 globular clusters [280]. Early highlights include the double neutron star binary in M15 [309] and a low-mass binary system with a 95-min orbital period in 47 Tucanae [59], one of 23 millisecond pulsars currently known in this cluster alone [59, 230].

On-going surveys of clusters continue to yield new surprises [319, 93], with no less than 70 discoveries in the past five years [316]. Among these is the most eccentric binary pulsar in a globular cluster so far – J0514−4002 is a 4.99 ms pulsar in a highly eccentric ($e = 0.89$) binary system in the globular cluster NGC 1851 [121]. The cluster with the most pulsars is now Terzan 5 which boasts 33 [320, 280], 31 of which were found with the Green Bank Telescope [282]. The spin periods and orbital parameters of the new pulsars reveal that, as a population, they are significantly different to the pulsars of 47 Tucanae which have periods in the range 2–8 ms [230]. The spin periods of the new pulsars span a much broader range (1.4–80 ms) including the first, third and fourth shortest spin periods of all pulsars currently known. The binary pulsars include two systems with eccentric orbits and likely white dwarf companions. No such systems are known in 47 Tucanae. The difference between the two pulsar populations may reflect the different evolutionary states and physical conditions of the two clusters. In particular, the central stellar density of Terzan 5 is about twice that of 47 Tucanae, suggesting that the increased rate of stellar interactions might disrupt the recycling process for the neutron stars in some binary systems and induce larger eccentricities in others.

2.11.5 Targeted searches of other regions

While globular clusters are the richest targets for finding millisecond pulsars, other regions of interest have been searched. Recently, a search of error boxes from unidentified sources from the Energetic Gamma-Ray Experiment Telescope (EGRET) revealed three new binary pulsars J1614−2318, J1614−2230 and J1744−3922 [143, 92, 315]. None of these pulsars is likely to be energetic enough to be associated with their target EGRET sources [143]. While convincing EGRET associations with several young pulsars are now known [203], it is not clear whether millisecond pulsars are relevant to the energetics of these enigmatic sources [70]. Despite this lack of success, it is quite possible that the recent launches of the AGILE [3] and GLAST [274] gamma-ray observatories will provide further opportunities for follow-up.

Other targets of interest are X-ray point sources found with the Chandra [273] and XMM-Newton [107] observatories and TeV sources found with HESS [259]. The X-ray sources have been particularly fruitful targets for young pulsars, with a number of discoveries of extremely faint objects [57]. Although not directly relevant to the topic of this review, these searches are revolutionizing our picture of the young neutron star population and should provide valuable insights into the beaming fraction and birthrate of these pulsars.

2.11.6 Extragalactic searches

The only radio pulsars known outside of the Galactic field and its globular cluster systems are the 19 currently known in the Large and Small Magellanic Clouds [260, 91, 251]. The lack of millisecond pulsars in the sample so far is most likely due to the limited sensitivity of the searches and large distance to the clouds. Further surveys in the Magellanic clouds are warranted. Surveys of more distant galaxies have so far been fruitless. Current instrumentation is only sensitive to giant isolated pulsars of the kind observed from the Crab [134] and the millisecond pulsars [197]. While surveys for such events are on-going [199], detections of weaker periodic sources are likely to require the enhanced sensitivity of the next generation radio telescopes.

2.11.7 Surveys with new telescopes

All surveys that have so far been conducted, or will be carried out in the next few years, will ultimately be surpassed by the next generation of radio telescopes. In the near term, the low-frequency array [215, 108] is set to discover hundreds of faint nearby pulsars [389] in the next five years. While the Square Kilometre Array [158] is not expected to be completed until 2020, a

number of pathfinder instruments are now under development. In China, the Five hundred meter Aperture Spherical Telescope [109] is scheduled for completion in 2013 and will provide significant advances for pulsar research [269]. The Australian Square Kilometre Array Pathfinder, currently under construction, will provide large-area coverage allowing rapid surveys of the radio sky for pulsars and transient radio sources [168]. Similar capabilities are offered by the Allen Telescope Array in California [337], which is now in operation. In South Africa, an array of 50 dishes is currently under construction and should be operational by 2012 [190].

2.12 Going further

Two books, *Pulsar Astronomy* [245] and *Handbook of Pulsar Astronomy* [232], cover the literature and techniques and provide excellent further reading. The morphological properties of pulsars have recently been comprehensively discussed in recent reviews [343, 336]. Those seeking a more theoretical viewpoint are advised to read *The Theory of Neutron Star Magnetospheres* [265] and *The Physics of the Pulsar Magnetosphere* [37]. Our summary of evolutionary aspects serves merely as a primer to the vast body of literature available. Further insights can be found from more detailed reviews [39, 304, 348]. Pulsar resources available on the *Internet* are continually becoming more extensive and useful. A good starting point for pulsar-surfers is the *Handbook of Pulsar Astronomy* website [233], as well as the *Pulsar Astronomy* wiki [2] and the Cool Pulsars site [1].

3 Pulsar Statistics and Demography

The observed pulsar sample is heavily biased towards the brighter objects that are the easiest to detect. What we observe represents only the tip of the iceberg of a much larger underlying population [131]. The bias is well demonstrated by the projection of pulsars onto the Galactic plane shown in Figure 11. The clustering of sources around the Sun seen in the left panel of this figure is clearly at variance with the distribution of other stellar populations which show a radial distribution symmetric about the Galactic centre. Also shown in Figure 11 is the cumulative number of pulsars as a function of the projected distance from the Sun compared to the expected distribution for a simple model population with no selection effects. The observed number distribution becomes strongly deficient beyond a few kpc.

3.1 Selection effects in pulsar searches

3.1.1 The inverse square law and survey thresholds

The most prominent selection effect is the inverse square law, i.e. for a given intrinsic luminosity[2], the observed flux density varies inversely with the distance squared. This results in the observed sample being dominated by nearby and/or high luminosity objects. Beyond distances of a few kpc from the Sun, the apparent flux density falls below the detection thresholds $S_{\rm min}$ of most surveys. Following [102], we express this threshold as follows:

$$S_{\rm min} = \frac{\rm S/N_{min}}{\eta\sqrt{n_{\rm pol}}} \left(\frac{T_{\rm rec}+T_{\rm sky}}{\rm K}\right) \left(\frac{G}{\rm K\ Jy^{-1}}\right)^{-1} \left(\frac{\Delta\nu}{\rm MHz}\right)^{-1/2} \left(\frac{t_{\rm int}}{\rm s}\right)^{-1/2} \left(\frac{W}{P-W}\right)^{1/2} {\rm mJy}, \quad (3)$$

where $\rm S/N_{min}$ is the threshold signal-to-noise ratio, η is a generic fudge factor ($\lesssim 1$) which accounts for losses in sensitivity (e.g., due to sampling and digitization noise), $n_{\rm pol}$ is the number of

[2]Pulsar astronomers usually define the luminosity $L \equiv Sd^2$, where S is the mean flux density at 400 MHz (a standard observing frequency) and d is the distance derived from the DM (see Section 2.4). Since this ignores any assumptions about beaming or geometrical factors, it is sometimes referred to as a "pseudoluminosity" [9].

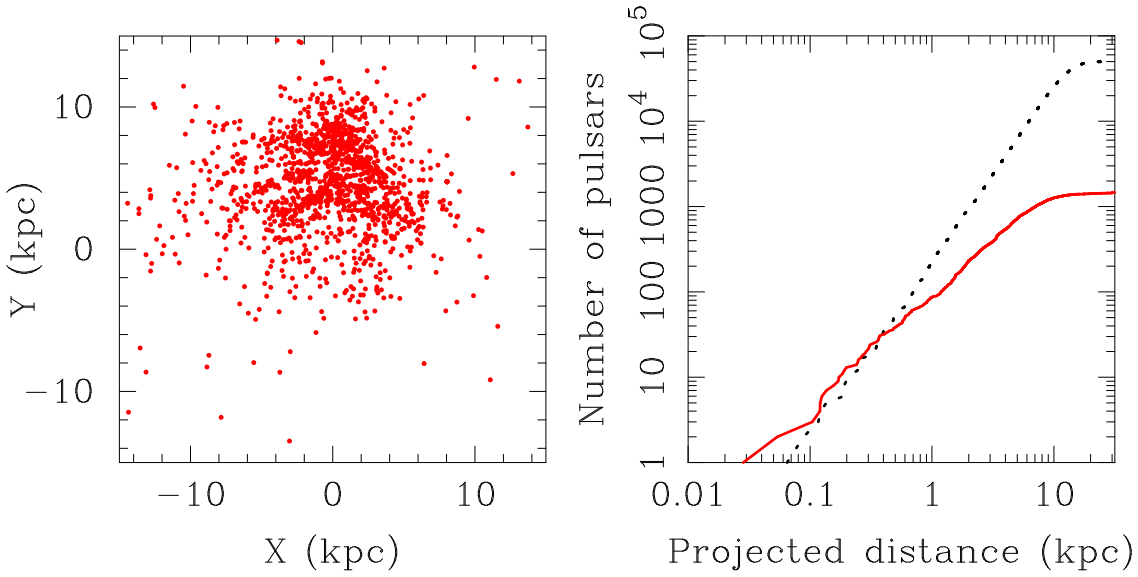

Figure 11: *Left panel: The current sample of all known radio pulsars projected onto the Galactic plane. The Galactic centre is at the origin and the Sun is at* $(0, 8.5)$ kpc. *Note the spiral-arm structure seen in the distribution which is now required by the most recent Galactic electron density model [88, 89]. Right panel: Cumulative number of pulsars as a function of projected distance from the Sun. The solid line shows the observed sample while the dotted line shows a model population free from selection effects.*

polarizations recorded (either 1 or 2), $T_{\rm rec}$ and $T_{\rm sky}$ are the receiver and sky noise temperatures, G is the gain of the antenna, $\Delta\nu$ is the observing bandwidth, $t_{\rm int}$ is the integration time, W is the detected pulse width and P is the pulse period.

3.1.2 Interstellar pulse dispersion and multipath scattering

It follows from Equation (3) that the sensitivity decreases as $W/(P-W)$ and hence W increases. Also note that if $W \gtrsim P$, the pulsed signal is smeared into the background emission and is no longer detectable, regardless of how luminous the source may be. The detected pulse width W may be broader than the intrinsic value largely as a result of pulse dispersion and multipath scattering by free electrons in the interstellar medium. The dispersive smearing scales as $\Delta\nu/\nu^3$, where ν is the observing frequency. This can largely be removed by dividing the pass-band into a number of channels and applying successively longer time delays to higher frequency channels *before* summing over all channels to produce a sharp profile. This process is known as incoherent dedispersion.

The smearing across the individual frequency channels, however, still remains and becomes significant at high dispersions when searching for short-period pulsars. Multipath scattering from electron density irregularities results in a one-sided broadening of the pulse profile due to the delay in arrival times. A simple scattering model is shown in Figure 12 in which the scattering electrons are assumed to lie in a thin screen between the pulsar and the observer [334]. The timescale of this effect varies roughly as ν^{-4}, which can not be removed by instrumental means.

Dispersion and scattering are most severe for distant pulsars in the inner Galaxy where the number of free electrons along the line of sight becomes large. The strong frequency dependence of both effects means that they are considerably less of a problem for surveys at observing frequencies $\gtrsim 1.4$ GHz [80, 172] compared to the 400-MHz search frequency used in early surveys. An added

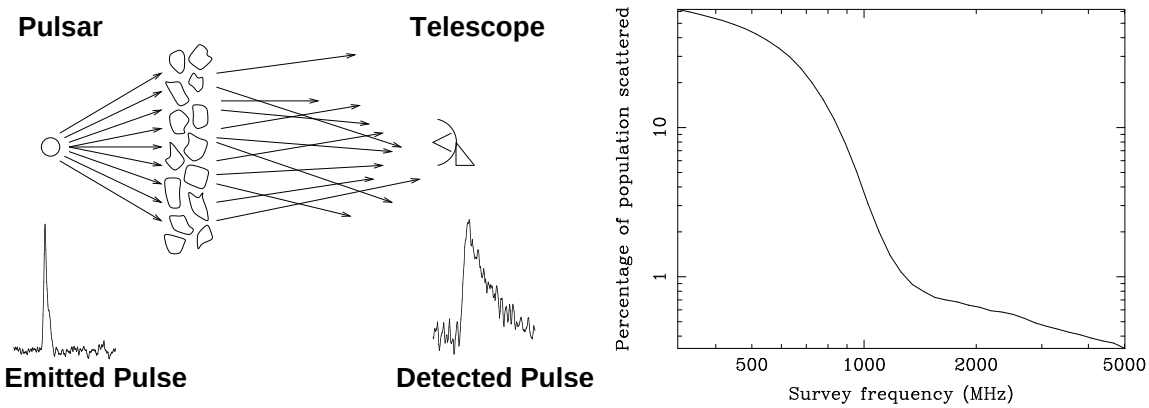

Figure 12: *Left panel: Pulse scattering caused by irregularities in the interstellar medium. The different path lengths and travel times of the scattered rays result in a "scattering tail" in the observed pulse profile which lowers its signal-to-noise ratio. Right panel: A simulation showing the percentage of Galactic pulsars that are likely to be undetectable due to scattering as a function of observing frequency. Low-frequency ($\lesssim$ 1 GHz) surveys clearly miss a large percentage of the population due to this effect.*

bonus for such observations is the reduction in $T_{\rm sky}$ which scales with frequency as approximately $\nu^{-2.8}$ [210]. Pulsar intensities also have an inverse frequency dependence, with the average scaling being $\nu^{-1.6}$ [238], so that flux densities are roughly an order of magnitude lower at 1.4 GHz compared to 400 MHz. Fortunately, this can be at least partially compensated for by the use of larger receiver bandwidths at higher radio frequencies. For example, the 1.4-GHz system at Parkes has a bandwidth of 288 MHz [244] compared to the 430-MHz system, where nominally 32 MHz is available [255].

3.1.3 Orbital acceleration

Standard pulsar searches use Fourier techniques [232] to search for *a-priori* unknown periodic signals and usually assume that the apparent pulse period remains constant throughout the observation. For searches with integration times much greater than a few minutes, this assumption is only valid for solitary pulsars or binary systems with orbital periods longer than about a day. For shorter-period binary systems, the Doppler-shifting of the period results in a spreading of the signal power over a number of frequency bins in the Fourier domain, leading to a reduction in S/N [166]. An observer will perceive the frequency of a pulsar to shift by an amount $aT/(Pc)$, where a is the (assumed constant) line-of-sight acceleration during the observation of length T, P is the (constant) pulsar period in its rest frame and c is the speed of light. Given that the width of a frequency bin in the Fourier domain is $1/T$, we see that the signal will drift into more than one spectral bin if $aT^2/(Pc) > 1$. Survey sensitivities to rapidly-spinning pulsars in tight orbits are therefore significantly compromised when the integration times are large.

As an example of this effect, as seen in the time domain, Figure 13 shows a 22.5-min search mode observation of the binary pulsar B1913+16 [155, 369, 370]. Although this observation covers only about 5% of the orbit (7.75 hr), the severe effects of the Doppler smearing on the pulse signal are very apparent. While the standard search code nominally detects the pulsar with S/N = 9.5 for this observation, it is clear that this value is significantly reduced due to the Doppler shifting of the pulse period seen in the individual sub-integrations.

It is clearly desirable to employ a technique to recover the loss in sensitivity due to Doppler

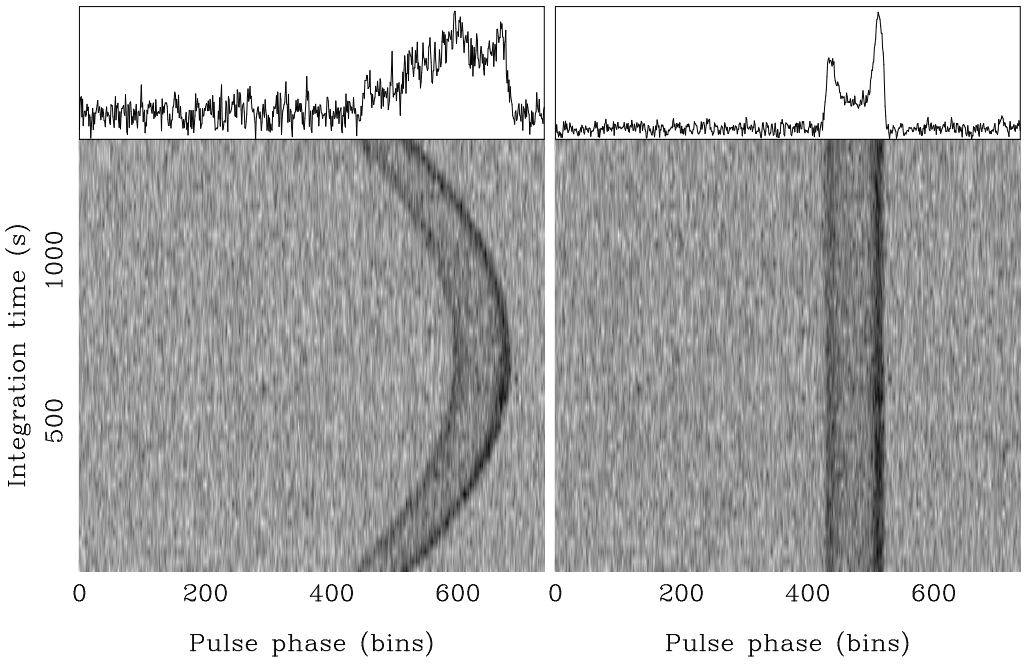

Figure 13: *Left panel: A* 22.5-min *Arecibo observation of the binary pulsar B1913+16. The assumption that the pulsar has a constant period during this time is clearly inappropriate given the quadratic drifting in phase of the pulse during the observation (linear grey scale plot). Right panel: The same observation after applying an acceleration search. This shows the effective recovery of the pulse shape and a significant improvement in the signal-to-noise ratio.*

smearing. One such technique, the so-called "acceleration search" [266], assumes the pulsar has a constant acceleration during the observation. Each time series can then be re-sampled to refer it to the frame of an inertial observer using the Doppler formula to relate a time interval τ in the pulsar frame to that in the observed frame at time t, as $\tau(t) \propto (1 + at/c)$. Searching over a range of accelerations is desirable to find the time series for which the trial acceleration most closely matches the true value. In the ideal case, a time series is produced with a signal of constant period for which full sensitivity is recovered (see right panel of Figure 13). This technique was first used to find PSR B2127+11C [7], a double neutron star binary in M15 which has parameters similar to B1913+16. Its application to 47 Tucanae [59] resulted in the discovery of nine binary millisecond pulsars, including one in a 96-min orbit around a low-mass ($0.15\,M_\odot$) companion. This is currently the shortest binary period for any known radio pulsar.

For intermediate orbital periods, in the range 30 min–several hours, another promising technique is the dynamic power spectrum search shown in Figure 14. Here the time series is split into a number of smaller contiguous segments which are Fourier-transformed separately. The individual spectra are displayed as a two-dimensional (frequency versus time) image. Orbitally modulated pulsar signals appear as sinusoidal signals in this plane as shown in Figure 14.

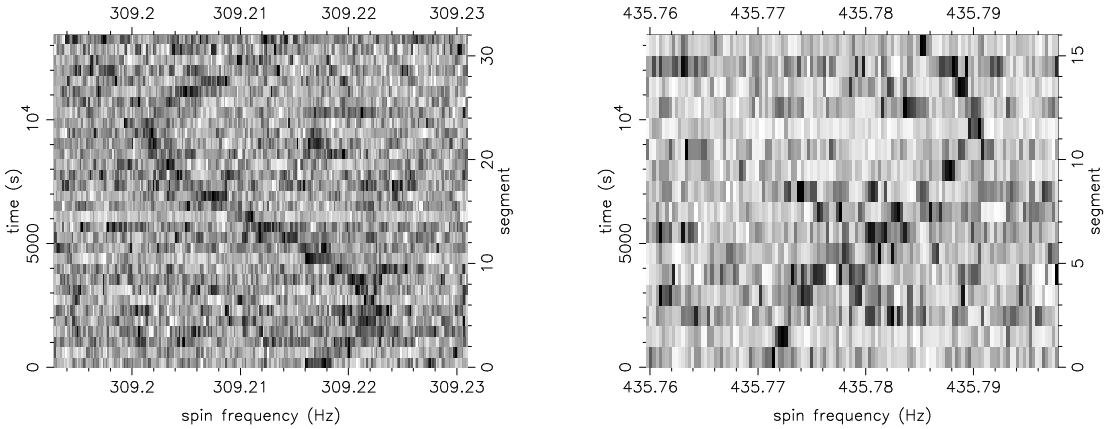

Figure 14: *Dynamic power spectra showing two recent pulsar discoveries in the globular cluster M62 showing fluctuation frequency as a function of time. Figure provided by Adam Chandler.*

This technique has been used by various groups where spectra are inspected visually [249]. Much of the human intervention can be removed using an hierarchical scheme for selecting significant events [74]. This approach was recently applied to a search of the globular cluster M62 resulting in the discovery of three new pulsars. One of the new discoveries – M62F, a faint 2.3-ms pulsar in a 4.8-hr orbit – was detectable only using the dynamic power spectrum technique.

For the shortest orbital periods, the assumption of a constant acceleration during the observation clearly breaks down. In this case, a particularly efficient algorithm has been developed [82, 317, 176, 318] which is optimised to finding binaries with periods so short that many orbits can take place during an observation. This "phase modulation" technique exploits the fact that the Fourier components are modulated by the orbit to create a family of periodic sidebands around the nominal spin frequency of the pulsar. While this technique has so far not resulted in any new discoveries, the existence of short period binaries in 47 Tucanae [59], Terzan 5 [320] and the 11-min X-ray binary X1820−303 in NGC 6624 [357], suggests that there are more ultra-compact radio binary pulsars that await discovery.

3.2 Correcting the observed pulsar sample

In the following, we review common techniques to account for many of the aforementioned selection effects and form a less biased picture of the true pulsar population.

3.2.1 Scale factor determination

A very useful approach [303, 391], is to define a scaling factor ξ as the ratio of the total Galactic volume weighted by pulsar density to the volume in which a pulsar is detectable:

$$\xi(P, L) = \frac{\int\int_{\rm G} \Sigma(R, z) R\, dR\, dz}{\int\int_{P,L} \Sigma(R, z) R\, dR\, dz}. \tag{4}$$

Here, $\Sigma(R, z)$ is the assumed pulsar space density distribution in terms of galactocentric radius R and height above the Galactic plane z. Note that ξ is primarily a function of period P and luminosity L such that short-period/low-luminosity pulsars have smaller detectable volumes and therefore higher ξ values than their long-period/high-luminosity counterparts.

In practice, ξ is calculated for each pulsar separately using a Monte Carlo simulation to model the volume of the Galaxy probed by the major surveys [271]. For a sample of $N_{\rm obs}$ observed pulsars above a minimum luminosity $L_{\rm min}$, the total number of pulsars in the Galaxy

$$N_{\rm G} = \sum_{i=1}^{N_{\rm obs}} \frac{\xi_i}{f_i}, \tag{5}$$

where f is the model-dependent "beaming fraction" discussed below in Section 3.2.3. Note that this estimate applies to those pulsars with luminosities $\gtrsim L_{\rm min}$. Monte Carlo simulations have shown this method to be reliable, as long as $N_{\rm obs}$ is reasonably large [229].

3.2.2 The small-number bias

For small samples of observationally-selected objects, the detected sources are likely to be those with larger-than-average luminosities. The sum of the scale factors (5), therefore, will tend to underestimate the true size of the population. This "small-number bias" was first pointed out [178, 182] for the sample of double neutron star binaries where we know of only five systems relevant for calculations of the merging rate (see Section 3.4.1). Only when $N_{\rm obs} \gtrsim 10$ does the sum of the scale factors become a good indicator of the true population size.

Despite a limited sample size, it has been demonstrated [194] that rigorous confidence intervals of $N_{\rm G}$ can be derived using Bayesian techniques. Monte Carlo simulations verify that the simulated number of detected objects $N_{\rm detected}$ closely follows a Poisson distribution and that $N_{\rm detected} = \alpha N_{\rm G}$, where α is a constant. By varying the value of $N_{\rm G}$ in the simulations, the mean of this Poisson distribution can be measured. The Bayesian analysis [194] finds, for a single object, the probability density function of the total population is

$$P(N_{\rm G}) = \alpha^2 N_{\rm G} \exp(-\alpha N_{\rm G}). \tag{6}$$

Adopting the necessary assumptions required in the Monte Carlo population about the underlying pulsar distribution, this technique can be used to place interesting constraints on the size and, as we shall see later, birth rate of the underlying population.

3.2.3 The beaming correction

The "beaming fraction" f in Equation (5) is the fraction of 4π steradians swept out by a pulsar's radio beam during one rotation. Thus f is the probability that the beam cuts the line-of-sight

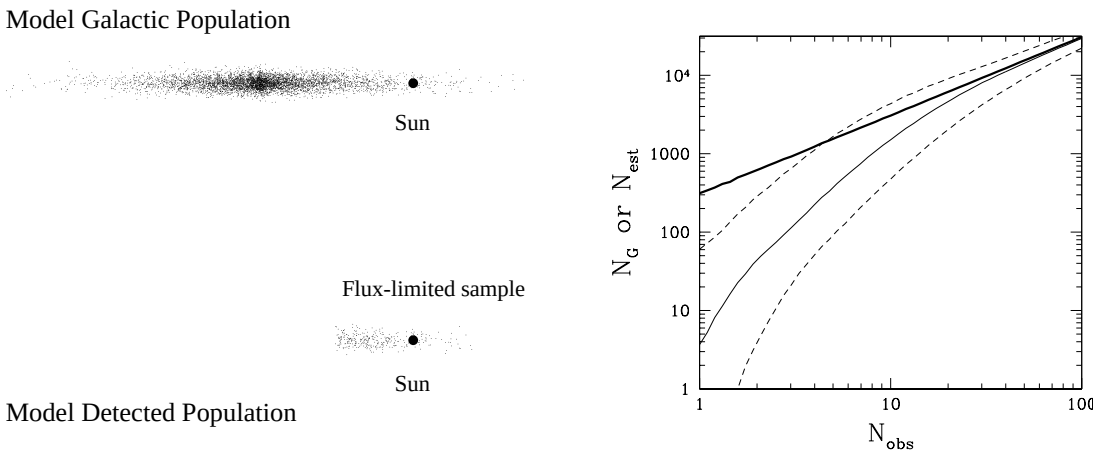

Figure 15: *Small-number bias of the scale factor estimates derived from a synthetic population of sources where the true number of sources is known. Left panel: An edge-on view of a model Galactic source population. Right panel: The thick line shows* $N_{\rm G}$*, the true number of objects in the model Galaxy, plotted against* $N_{\rm observed}$*, the number detected by a flux-limited survey. The thin solid line shows* $N_{\rm est}$*, the median sum of the scale factors, as a function of* $N_{\rm obs}$ *from a large number of Monte Carlo trials. Dashed lines show 25 and 75% percentiles of the* $N_{\rm est}$ *distribution.*

of an arbitrarily positioned observer. A naïve estimate for f of roughly 20% assumes a circular beam of width $\sim 10^\circ$ and a randomly distributed inclination angle between the spin and magnetic axes [368]. Observational evidence summarised in Figure 16 suggests that shorter period pulsars have wider beams and therefore larger beaming fractions than their long-period counterparts [272, 247, 40, 363].

When most of these beaming models were originally proposed, the sample of millisecond pulsars was $\lesssim 5$ and hence their predictions about the beaming fractions of short-period pulsars relied largely on extrapolations from the normal pulsars. An analysis of a large sample of millisecond pulsar profiles [206] suggests that their beaming fraction lies between 50 and 100%. Independent constraints on f for millisecond pulsars come from deep *Chandra* observations of the globular cluster 47 Tucanae [130] and radio pulsar surveys [59] which suggest that $f > 0.4$ and likely close to unity [140]. The large beaming fraction and narrow pulses often observed strongly suggests a fan beam model for millisecond pulsars [265].

3.3 The population of normal and millisecond pulsars

3.3.1 Luminosity distributions and local number estimates

Based on a number of all-sky surveys carried out in the 1990s, the scale factor approach has been used to derive the characteristics of the true normal and millisecond pulsar populations and is based on the sample of pulsars within 1.5 kpc of the Sun [248]. Within this region, the selection effects are well understood and easier to quantify than in the rest of the Galaxy. These calculations should therefore give a reliable *local pulsar population* estimate.

The 430-MHz luminosity distributions obtained from this analysis are shown in Figure 17. For the normal pulsars, integrating the corrected distribution above 1 mJy kpc^2 and dividing by $\pi \times (1.5)^2$ kpc^2 yields a local surface density, assuming a beaming model [40], of 156 ± 31pulsars kpc^{-2} for 430-MHz luminosities above 1 mJy kpc^2. The same analysis for the millisecond pulsars, assuming a mean beaming fraction of 75% [206], leads to a local surface density of 38 ± 16 pulsars kpc^{-2}

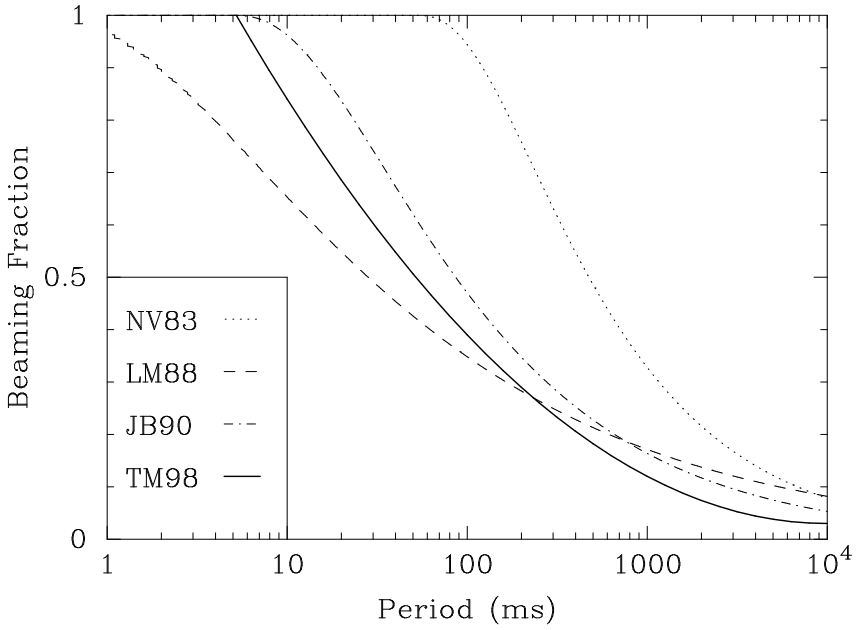

Figure 16: *Beaming fraction plotted against pulse period for four different beaming models: Narayan & Vivekanand 1983 (NV83) [272], Lyne & Manchester 1988 (LM88) [247], Biggs 1990 (JDB90) [40] and Tauris & Manchester 1998 (TM98) [363].*

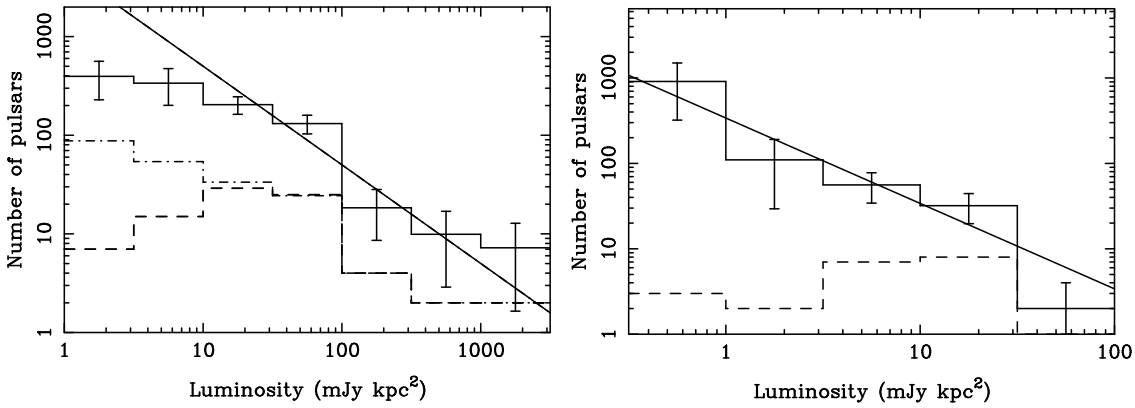

Figure 17: *Left panel: The corrected luminosity distribution (solid histogram with error bars) for normal pulsars. The corrected distribution* before *the beaming model has been applied is shown by the dot-dashed line. Right panel: The corresponding distribution for millisecond pulsars. In both cases, the observed distribution is shown by the dashed line and the thick solid line is a power law with a slope of* -1*. The difference between the observed and corrected distributions highlights the severe under-sampling of low-luminosity pulsars.*

also for 430-MHz luminosities above 1 mJy kpc^2.

3.3.2 Galactic population and birth-rates

Integrating the local surface densities of pulsars over the whole Galaxy requires a knowledge of the presently rather uncertain Galactocentric radial distribution [167, 227]. One approach is to assume that pulsars have a radial distribution similar to that of other stellar populations and to scale the local number density with this distribution in order to estimate the total Galactic population. The corresponding local-to-Galactic scaling is 1000 ± 250 kpc^2 [321]. This implies a population of $\sim 160,000$ active normal pulsars and $\sim 40,000$ millisecond pulsars in the Galaxy.

Based on these estimates, we are in a position to deduce the corresponding rate of formation or birth-rate required to sustain the observed population. From the P–$\dot{P}$ diagram in Figure 3, we infer a typical lifetime for normal pulsars of $\sim 10^7$ yr, corresponding to a Galactic birth rate of $\sim$ 1 per 60 yr – consistent with the rate of supernovae [383]. As noted in Section 2.2, the millisecond pulsars are much older, with ages close to that of the Universe $\tau_{\rm u}$ (we assume here $\tau_{\rm u} = 13.8$ Gyr [407]). Taking the maximum age of the millisecond pulsars to be $\tau_{\rm u}$, we infer a mean birth rate of at least 1 per 345,000 yr. This is consistent, within the uncertainties, of other studies of the millisecond pulsar population [113, 360] and with the birth-rate of low-mass X-ray binaries [236, 192].

3.4 The population of relativistic binaries

Although no radio pulsar has so far been observed in orbit around a black hole companion, we now know of several double neutron star and neutron star–white dwarf binaries which will merge due to gravitational wave emission within a reasonable timescale. The current sample of objects is shown as a function of orbital period and eccentricity in Figure 18. Isochrones showing various coalescence times $\tau_{\rm g}$ are calculated using the expression

$$\tau_{\rm g} \simeq 9.83 \times 10^6 \ {\rm yr} \left(\frac{P_{\rm b}}{\rm hr}\right)^{8/3} \left(\frac{m_1 + m_2}{M_\odot}\right)^{-2/3} \left(\frac{\mu}{M_\odot}\right)^{-1} \left(1 - e^2\right)^{7/2}, \qquad (7)$$

where m_1 and m_2 are the masses of the two stars, $\mu = m_1 m_2/(m_1 + m_2)$ is the "reduced mass", $P_{\rm b}$ is the binary period and e is the eccentricity. This formula is a good analytic approximation (within a few percent) to the numerical solution of the exact equations for $\tau_{\rm g}$ [298, 299].

In addition to tests of strong-field gravity through observations of relativistic binary systems (see Section 4.4), estimates of their Galactic population and merger rate are of great interest as one of the prime sources for current gravitational wave detectors such as GEO600 [135], LIGO [54], VIRGO [157] and TAMA [276]. In the following, we review empirical determinations of the population sizes and merging rates of binaries where at least one component is visible as a radio pulsar.

3.4.1 Double neutron star binaries

As discussed in Section 2.2, double neutron star (DNS) binaries are expected to be rare. This is certainly the case; as summarized in Table 1, only around ten DNS binaries are currently known. Although we only see both neutron stars as pulsars in J0737−3039 [243], we are "certain" of the identification in five other systems from precise component mass measurements from pulsar timing observations (see Section 4.4). The other systems listed in Table 1 have eccentric orbits, mass functions and periastron advance measurements that are *consistent* with a DNS identification, but for which there is presently not sufficient component mass information to confirm their nature.

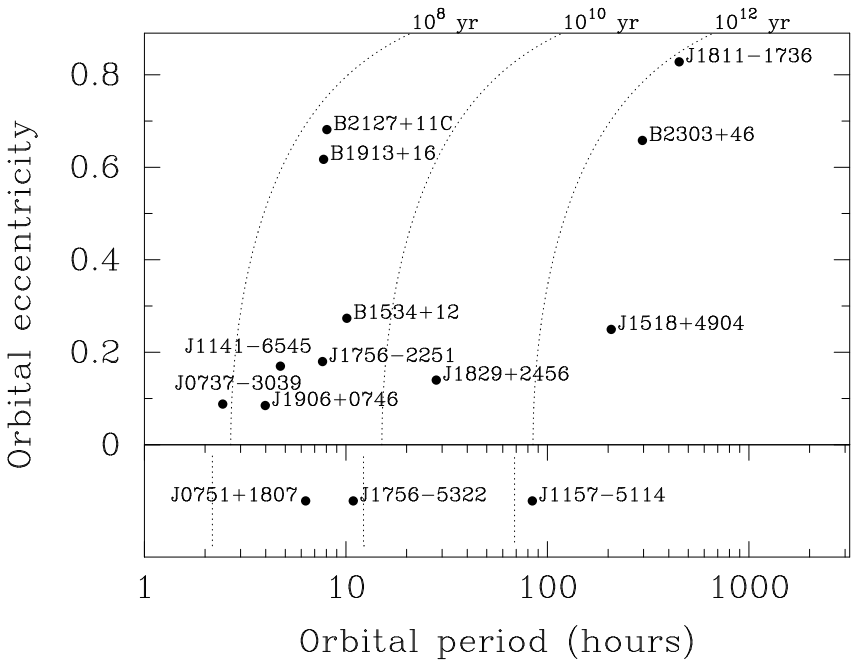

Figure 18: *The relativistic binary merging plane. Top: Orbital eccentricity versus period for eccentric binary systems involving neutron stars. Bottom: Orbital period distribution for the massive white dwarf–pulsar binaries.*

One further DNS candidate, the 95-ms pulsar J1753−2243 (see Table 3), has recently been discovered [191]. Although the mass function for this pulsar is lower than the DNS systems listed in Table 1, a neutron star companion cannot be ruled out in this case. Further observations should soon clarify the nature of this system. We note, however, that the 13.6-day orbital period of this system means that it will not contribute significantly to gravitational wave inspiral rate calculations discussed below.

Despite the uncertainties in identifying DNS binaries, for the purposes of determining the Galactic merger rate, the systems for which $\tau_{\rm g}$ is less than $\tau_{\rm u}$ (i.e. PSRs J0737−3039, B1534+12, J1756−2251, J1906+0746, B1913+16 and B2127+11C) are primarily of interest. Of these PSR B2127+11C is in the process of being ejected from the globular cluster M15 [309, 305] and is thought to make only a negligible contribution to the merger rate [301]. The general approach with the remaining systems is to derive scale factors for each object, construct the probability density function of their total population (as outlined in Section 3.2.1) and then divide these by a reasonable estimate for the lifetime. Getting such estimates is, however, difficult. It has been proposed [182] that the *observable lifetimes* for these systems are determined by the timescale on which the current orbital period is reduced by a factor of two [10]. Below this point, the orbital smearing selection effect discussed in Section 3.1.3 will render the binary undetectable by current surveys. More recent work [76] has suggested that a significant population of highly eccentric binary systems could easily evade detection due to their short lifetimes before gravitational wave inspiral. If this selection effect is significant, then the merger rate estimates quoted below could easily be underestimated by a factor of a few.

The results of the most recent DNS merger rate estimates kind [193], are summarised in the left panel of Figure 19. The combined Galactic merger rate, dominated by the double pulsar and J1906+0746 is found to be 118^{+174}_{-79} Myr^{-1}, where the uncertainties reflect the 95% confidence

	J0737−3039	J1518+4904	B1534+12	J1756−2251	J1811−1736
P [ms]	22.7/2770	40.9	37.9	28.5	104.2
$P_{\rm b}$ [d]	0.102	8.6	0.4	0.32	18.8
e	0.088	0.25	0.27	0.18	0.83
$\log_{10}(\tau_{\rm c}/[{\rm yr}])$	8.3/7.7	10.3	8.4	8.6	9.0
$\log_{10}(\tau_{\rm g}/[{\rm yr}])$	7.9	12.4	9.4	10.2	13.0
Masses measured?	Yes	No	Yes	Yes	Yes
	B1820−11	J1829+2456	J1906+0746	B1913+16	B2127+11C
P [ms]	279.8	41.0	144.1	59.0	30.5
$P_{\rm b}$ [d]	357.8	1.18	0.17	0.3	0.3
e	0.79	0.14	0.085	0.62	0.68
$\log_{10}(\tau_{\rm c}/[{\rm yr}])$	6.5	10.1	5.1	8.0	8.0
$\log_{10}(\tau_{\rm g}/[{\rm yr}])$	15.8	10.8	8.5	8.5	8.3
Masses measured?	No	No	Yes	Yes	Yes

Table 1: *Known and likely DNS binaries. Listed are the pulse period P, orbital period $P_{\rm b}$, orbital eccentricity e, characteristic age $\tau_{\rm c}$, and expected binary coalescence time-scale $\tau_{\rm g}$ due to gravitational wave emission calculated from Equation (7). To distinguish between definite and candidate DNS systems, we also list whether the masses of both components have been determined via the measurement of two or more post Keplerian parameters as described in Section 4.4.*

level using the techniques summarised in Section 3.2.2. Extrapolating this number to include DNS binaries detectable by LIGO in other galaxies [301], the expected event rate is $49^{+73}_{-33} \times 10^{-3}$ yr^{-1} for initial LIGO and 265^{+390}_{-178} yr^{-1} for advanced LIGO. Future prospects for detecting gravitational wave emission from binary neutron star inspirals are therefore very encouraging, especially if the population of highly eccentric systems is significant [76]. Since much of the uncertainty in the rate estimates is due to our ignorance of the underlying distribution of double neutron star systems, future gravitational wave detection could ultimately constrain the properties of this exciting binary species.

Although the double pulsar system J0737−3039 will not be important for ground-based detectors until its final coalescence in another 85 Myr, it may be a useful calibration source for the future space-based detector LISA [275]. It is calculated [179] that a 1-yr observation with LISA would detect (albeit with S/N $\sim$ 2) the continuous emission at a frequency of 0.2 mHz based on the current orbital parameters. Although there is the prospect of using LISA to detect similar systems systems through their continuous emission, current calculations [179] suggest that significant (S/N $>$ 5) detections are not likely. Despite these limitations, it is likely that LISA observations will be able to place independent constraints on the Galactic DNS binary population after several years of operation.

3.4.2 White dwarf–neutron star binaries

Although the population of white dwarf–neutron star (WDNS) binaries in general is substantial, the fraction which will merge due to gravitational wave emission is small. Like the DNS binaries, the observed WDNS sample suffers from small-number statistics. From Figure 18, we note that only three WDNS systems are currently known that will merge within $\tau_{\rm u}$, PSRs J0751+1807 [239], J1757−5322 [104] and J1141−6545 [188]. Applying the same techniques as used for the DNS population, the merging rate contributions of the three systems can be calculated [195] and are shown in Figure 19. The combined Galactic coalescence rate is 4^{+5}_{-3} Myr^{-1} (at 68% confidence

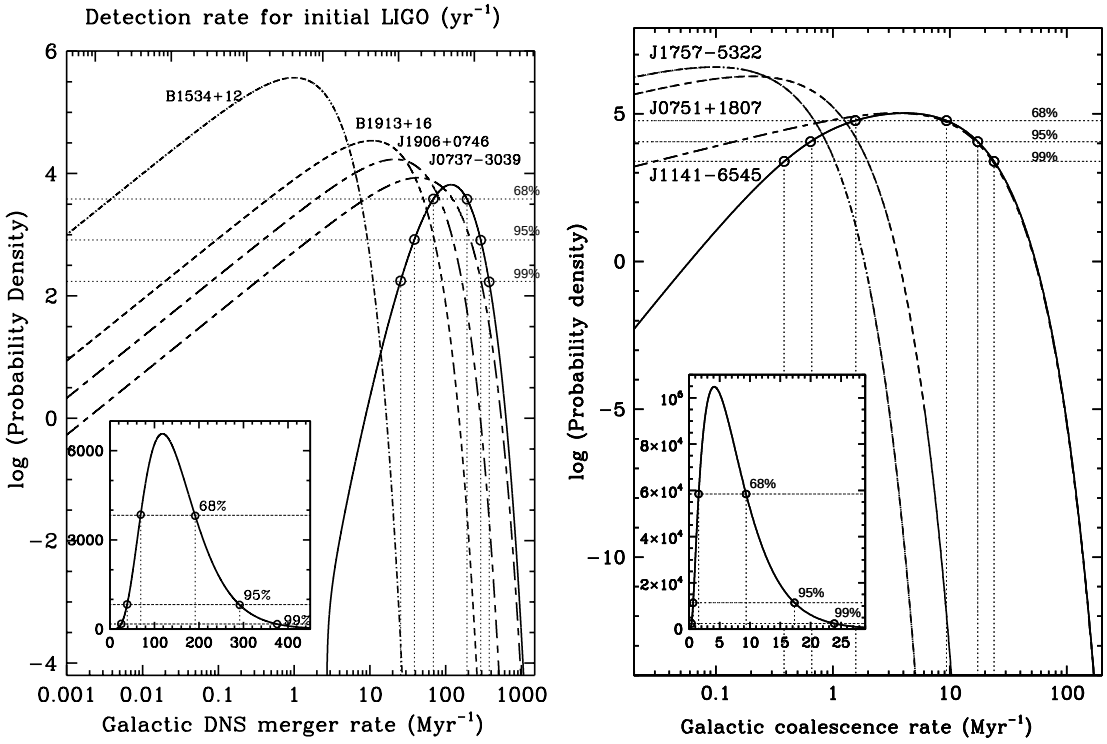

Figure 19: *The current best empirical estimates of the coalescence rates of relativistic binaries involving neutron stars. The individual contributions from each known binary system are shown as dashed lines, while the solid line shows the total probability density function on a logarithmic and (inset) linear scale. The left panel shows the most recent analysis for DNS binaries [193], while the right panel shows the equivalent results for NS–WD binaries [195]. Figures provided by Chunglee Kim.*

interval). This result is not corrected for beaming and therefore should be regarded as a lower limit on the total event rate. Although the orbital frequencies of these objects at coalescence are too low to be detected by LIGO, they do fall within the band planned for LISA [275]. Unfortunately, an extrapolation of the Galactic event rate out to distances at which such events would be detectable by LISA does not suggest that these systems will be a major source of detection [195]. Similar conclusions were reached by considering the statistics of low-mass X-ray binaries [81].

3.5 Going further

Good starting points for further reading can be found in other review articles [304, 193]. Our coverage of compact object coalescence rates has concentrated on empirical methods. Two software packages are freely available which allow the user to synthesize and search radio pulsar populations [291, 396, 325]. An alternative approach is to undertake a full-blown Monte Carlo simulation of the most likely evolutionary scenarios described in Section 2.6. In this scheme, a population of primordial binaries is synthesized with a number of underlying distribution functions: primary mass, binary mass ratio, orbital period distribution etc. The evolution of both stars is then followed to give a predicted sample of binary systems of all the various types. Although selection effects are not always taken into account in this approach, the final census is usually normalized to the star formation rate.

Numerous population syntheses (most often to populations of binaries where one or both members are NSs) can be found in the literature [100, 327, 377, 307, 214, 30, 31]. A group in Moscow has made a web interface to their code [358]. Although extremely instructive, the uncertain assumptions about initial conditions, the physics of mass transfer and the kicks applied to the compact object at birth result in a wide range of predicted event rates which are currently broader than the empirical methods [180, 195, 181]. Ultimately, the detection statistics from the gravitational wave detectors could provide far tighter constraints on the DNS merging rate than the pulsar surveys from which these predictions are made. Very recently [293] the results from empirical population constraints and full-blown binary population synthesis codes have been combined to constrain a variety of input parameters and physical conditions. The results of this work are promising, with stringent constraints being placed on the kick distributions, mass-loss fraction during mass transfer and common envelope assumptions.

4 Principles and Applications of Pulsar Timing

Pulsars are excellent celestial clocks. The period of the first pulsar [146] was found to be stable to one part in 10^7 over a few months. Following the discovery of the millisecond pulsar B1937+21 [21] it was demonstrated that its period could be measured to one part in 10^{13} or better [98]. This unrivaled stability leads to a host of applications including uses as time keepers, probes of relativistic gravity and natural gravitational wave detectors.

4.1 Observing basics

Each pulsar is typically observed at least once or twice per month over the course of a year to establish its basic properties. Figure 20 summarises the essential steps involved in a "time-of-arrival" (TOA) measurement. Pulses from the neutron star traverse the interstellar medium before being received at the radio telescope where they are de-dispersed and added to form a mean pulse profile.

During the observation, the data regularly receive a time stamp, usually based on a caesium time standard or hydrogen maser at the observatory plus a signal from the Global Positioning System of satellites (GPS; see [97]). The TOA is defined as the arrival time of some fiducial point

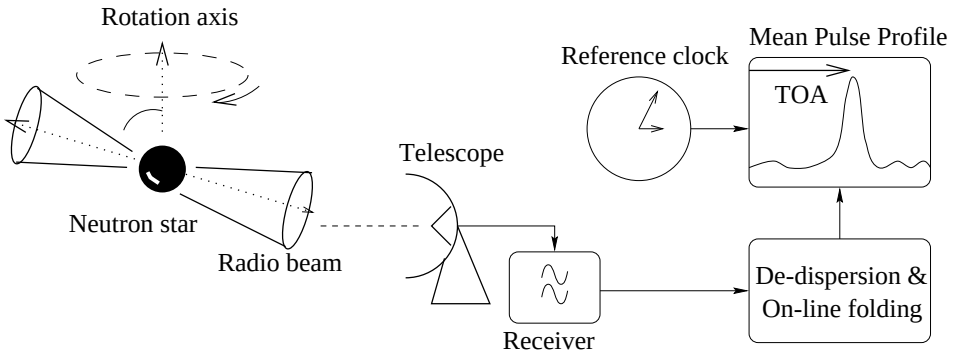

Figure 20: *Schematic showing the main stages involved in pulsar timing observations.*

on the integrated profile with respect to either the start or the midpoint of the observation. Since the profile has a stable form at any given observing frequency (see Section 2.3), the TOA can be accurately determined by cross-correlation of the observed profile with a high S/N "template" profile obtained from the addition of many observations at the particular observing frequency.

Successful pulsar timing requires optimal TOA precision which largely depends on the signal-to-noise ratio (S/N) of the pulse profile. Since the TOA uncertainty $\epsilon_{\rm TOA}$ is roughly the pulse width divided by the S/N, using Equation (3) we may write the fractional error as

$$\frac{\epsilon_{\rm TOA}}{P} \simeq \left(\frac{S_{\rm psr}}{\rm mJy}\right)^{-1} \left(\frac{T_{\rm rec}+T_{\rm sky}}{\rm K}\right) \left(\frac{G}{\rm K\ Jy^{-1}}\right)^{-1} \left(\frac{\Delta\nu}{\rm MHz}\right)^{-1/2} \left(\frac{t_{\rm int}}{\rm s}\right)^{-1/2} \left(\frac{W}{P}\right)^{3/2}. \quad (8)$$

Here, $S_{\rm psr}$ is the flux density of the pulsar, $T_{\rm rec}$ and $T_{\rm sky}$ are the receiver and sky noise temperatures, G is the antenna gain, $\Delta\nu$ is the observing bandwidth, $t_{\rm int}$ is the integration time, W is the pulse width and P is the pulse period (we assume $W \ll P$). Optimal results are thus obtained for observations of short period pulsars with large flux densities and small duty cycles (i.e. small W/P) using large telescopes with low-noise receivers and large observing bandwidths.

One of the main problems of employing large bandwidths is pulse dispersion. As discussed in Section 2.4, pulses emitted at lower radio frequencies travel slower and arrive later than those emitted at higher frequencies. This process has the effect of "stretching" the pulse across a finite receiver bandwidth, increasing W and therefore increasing $\epsilon_{\rm TOA}$. For normal pulsars, dispersion can largely be compensated for by the incoherent de-dispersion process outlined in Section 3.1.

The short periods of millisecond pulsars offer the ultimate in timing precision. In order to fully exploit this, a better method of dispersion removal is required. Technical difficulties in building devices with very narrow channel bandwidths require another dispersion removal technique. In the process of coherent de-dispersion [133, 232] the incoming signals are de-dispersed over the whole bandwidth using a filter which has the inverse transfer function to that of the interstellar medium. The signal processing can be done on-line either using high speed devices such as field programmable gate arrays [67, 295] or completely in software [346, 353]. Off-line data reduction, while disk-space limited, allows for more flexible analysis schemes [24].

The maximum time resolution obtainable via coherent dedispersion is the inverse of the total receiver bandwidth. The current state of the art is the detection [134] of features on nanosecond timescales in pulses from the 33-ms pulsar B0531+21 in the Crab nebula shown in Figure 21. Simple light travel-time arguments can be made to show that, in the absence of relativistic beaming effects [125], these incredibly bright bursts originate from regions less than 1 m in size.

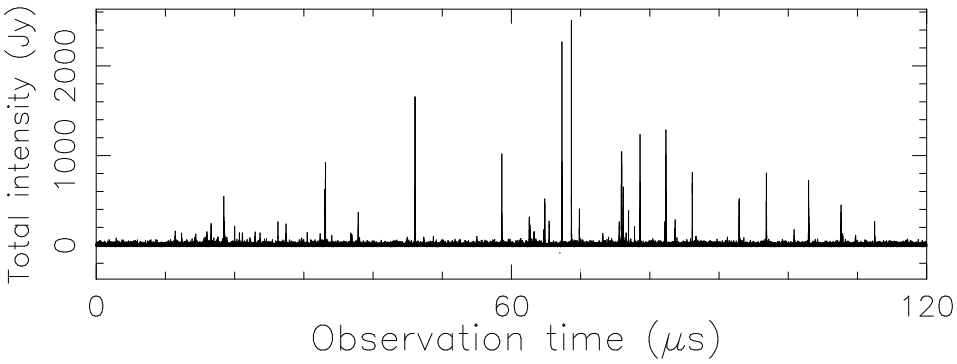

Figure 21: *A 120 μs window centred on a coherently-dedispersed giant pulse from the Crab pulsar showing high-intensity nanosecond bursts. Figure provided by Tim Hankins [134].*

4.2 The timing model

To model the rotational behaviour of the neutron star we ideally require TOAs measured by an inertial observer. An observatory located on Earth experiences accelerations with respect to the neutron star due to the Earth's rotation and orbital motion around the Sun and is therefore not in an inertial frame. To a very good approximation, the solar system centre-of-mass (barycentre) can be regarded as an inertial frame. It is now standard practice [156] to transform the observed TOAs to this frame using a planetary ephemeris such as the JPL DE405 [355]. The transformation between barycentric ($\mathcal{T}$) and observed (t) takes the form

$$\mathcal{T} - t = \frac{\underline{r}.\hat{\underline{s}}}{c} + \frac{(\underline{r}.\hat{\underline{s}})^2 - |\underline{r}|^2}{2cd} + \Delta t_{\rm rel} - \Delta t_{\rm DM}. \quad (9)$$

Here $\underline{r}$ is the position of the observatory with respect to the barycentre, $\hat{\underline{s}}$ is a unit vector in the direction of the pulsar at a distance d and c is the speed of light. The first term on the right hand side of Equation (9) is the light travel time from the observatory to the solar system barycentre. Incoming pulses from all but the nearest pulsars can be approximated by plane wavefronts. The second term, which represents the delay due to spherical wavefronts, yields the parallax and hence d. This has so far only been measured for five nearby millisecond pulsars [332, 376, 216, 345, 217]. The term $\Delta t_{\rm rel}$ represents the Einstein and Shapiro corrections due to general relativistic time delays in the solar system [20]. Since measurements can be carried out at different observing frequencies with different dispersive delays, TOAs are generally referred to the equivalent time that would be observed at infinite frequency. This transformation is the term $\Delta t_{\rm DM}$ (see also Equation (1)).

Following the accumulation of a number of TOAs, a surprisingly simple model is usually sufficient to account for the TOAs during the time span of the observations and to predict the arrival times of subsequent pulses. The model is a Taylor expansion of the rotational frequency $\Omega = 2\pi/P$ about a model value Ω_0 at some reference epoch $\mathcal{T}_0$. The model pulse phase

$$\phi(\mathcal{T}) = \phi_0 + (\mathcal{T} - \mathcal{T}_0)\Omega_0 + \frac{1}{2}(\mathcal{T} - \mathcal{T}_0)^2\dot{\Omega}_0 + \ldots, \quad (10)$$

where $\mathcal{T}$ is the barycentric time and ϕ_0 is the pulse phase at $\mathcal{T}_0$. Based on this simple model, and using initial estimates of the position, dispersion measure and pulse period, a "timing residual" is calculated for each TOA as the difference between the observed and predicted pulse phases.

A set of timing residuals for the nearby pulsar B1133+16 spanning almost 10 years is shown in Figure 22. Ideally, the residuals should have a zero mean and be free from any systematic

trends (see Panel a of Figure 22). To reach this point, however, the model needs to be refined in a bootstrap fashion. Early sets of residuals will exhibit a number of trends indicating a systematic error in one or more of the model parameters, or a parameter not incorporated into the model.

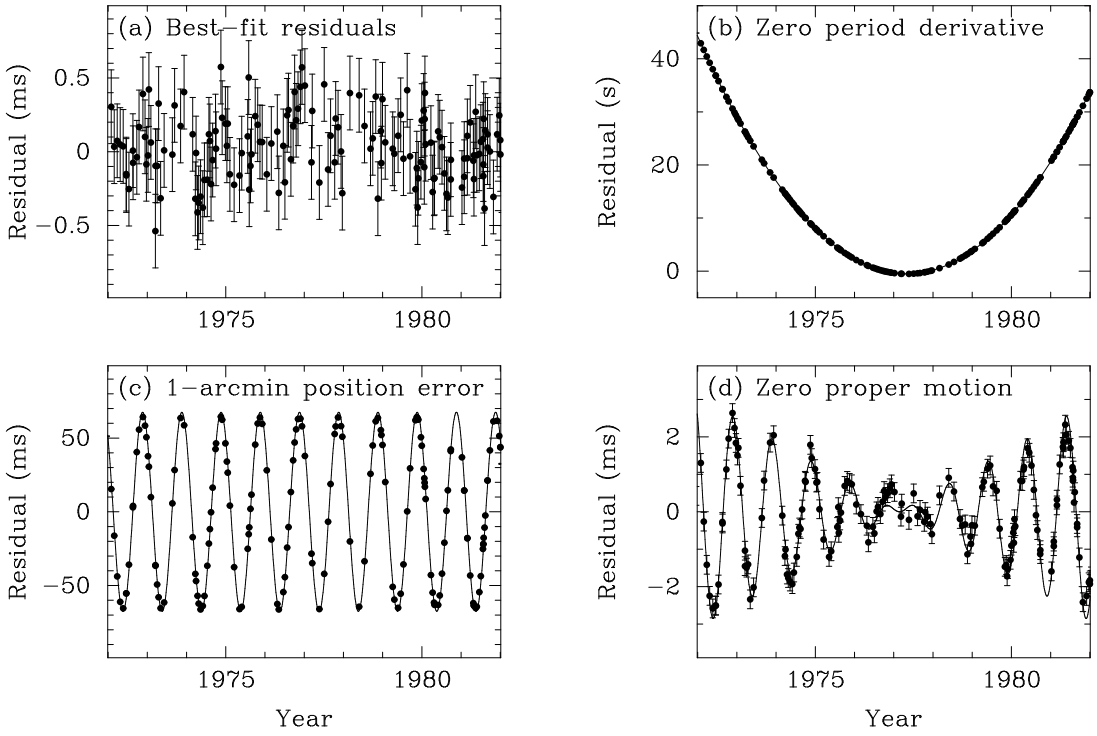

Figure 22: *Timing model residuals for PSR B1133+16. Panel a: Residuals obtained from the best-fitting model which includes period, period derivative, position and proper motion. Panel b: Residuals obtained when the period derivative term is set to zero. Panel c: Residuals showing the effect of a* 1-arcmin *position error. Panel d: Residuals obtained neglecting the proper motion. The lines in Panels b–d show the expected behaviour in the timing residuals for each effect. Data provided by Andrew Lyne.*

From Equation (10), an error in the assumed Ω_0 results in a linear slope with time. A parabolic trend results from an error in $\dot{\Omega}_0$ (see Panel b of Figure 22). Additional effects will arise if the assumed position of the pulsar (the unit vector $\hat{\underline{s}}$ in Equation (9)) used in the barycentric time calculation is incorrect. A position error results in an annual sinusoid (see Panel c of Figure 22). A proper motion produces an annual sinusoid of linearly increasing magnitude (see Panel d of Figure 22).

After a number of iterations, and with the benefit of a modicum of experience, it is possible to identify and account for each of these various effects to produce a "timing solution" which is phase coherent over the whole data span. The resulting model parameters provide spin and astrometric information with a precision which improves as the length of the data span increases. For example, timing observations of the original millisecond pulsar B1937+21, spanning almost 9 years (exactly 165,711,423,279 rotations!), measure a period of $1.5578064688197945 \pm 0.0000000000000004$ ms [189, 186] defined at midnight UT on December 5 1988! Measurements of other parameters are no less impressive, with astrometric errors of ~ 3 μarcsec being presently possible for the bright millisecond pulsar J0437−4715 [390].

4.3 Timing stability

Ideally, after correctly applying a timing model, we would expect a set of uncorrelated timing residuals with a zero mean and a Gaussian scatter with a standard deviation consistent with the measurement uncertainties. As can be seen in Figure 23, this is not always the case; the residuals of many pulsars exhibit a quasi-periodic wandering with time.

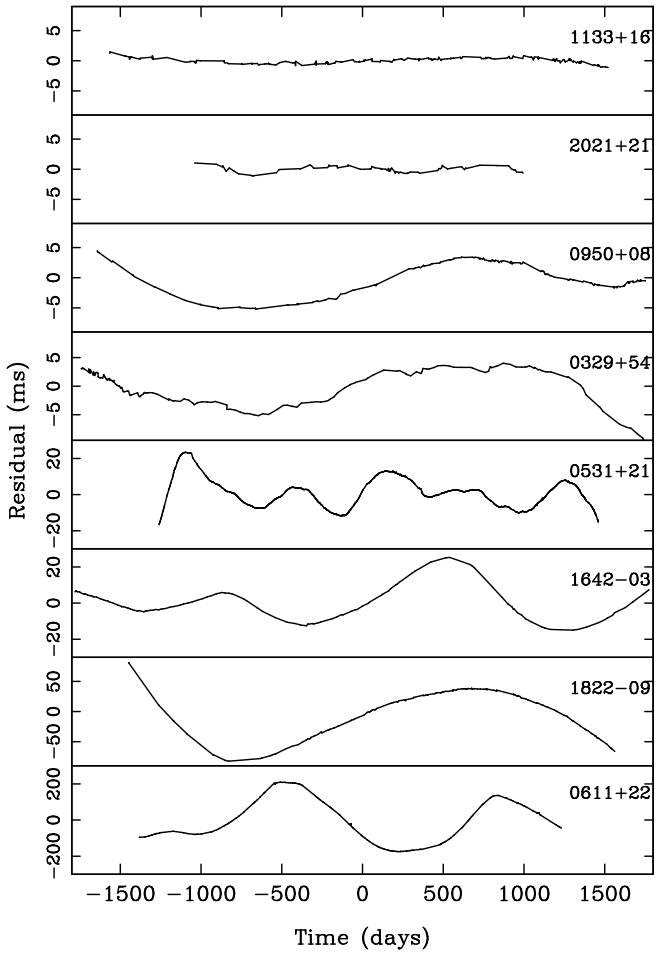

Figure 23: *Examples of timing residuals for a number of normal pulsars. Note the varying scale on the ordinate axis, the pulsars being ranked in increasing order of timing "activity". Data taken from the Jodrell Bank timing program [338, 151]. Figure provided by Andrew Lyne.*

Such "timing noise" is most prominent in the youngest of the normal pulsars [257, 87] and present at a lower level in the much older millisecond pulsars [189, 12]. While the physical processes of this phenomenon are not well understood, it seems likely that that may be connected to superfluid processes and temperature changes in the interior of the neutron star [6], or to processes in the magnetosphere [79, 78].

The relative dearth of timing noise for the older pulsars is a very important finding. It implies that the measurement precision presently depends primarily on the particular hardware constraints of the observing system. Consequently, a large effort in hardware development is presently being made to improve the precision of these observations using, in particular, coherent dedispersion

outlined in Section 4.1. Much progress in this area has been made by groups at Princeton [310], Berkeley [379], Jodrell Bank [380], UBC [378], Swinburne [362] and ATNF [16]. From high quality observations spanning over a decade [330, 331, 189], these groups have demonstrated that the timing stability of millisecond pulsars over such time-scales is comparable to terrestrial atomic clocks.

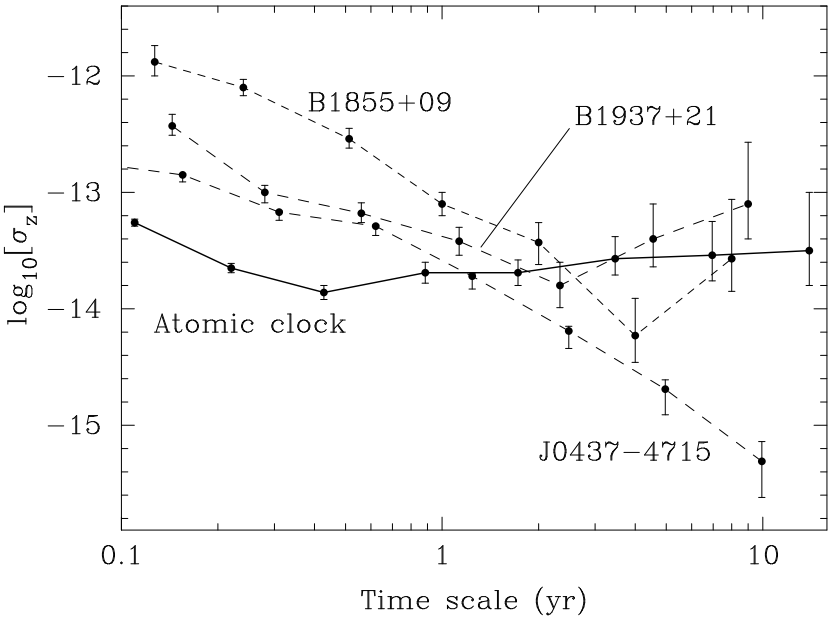

Figure 24: *The fractional stability of three millisecond pulsars compared to an atomic clock. Both PSRs B1855+09 and B1937+21 are comparable, or just slightly worse than, the atomic clock behaviour over timescales of a few years [258]. More recent timing of the millisecond pulsar J0437−4715 [390] indicates that it is inherently a very stable clock. Data for the latter pulsar provided by Joris Verbiest.*

This phenomenal stability is demonstrated in Figure 24 which shows σ_z [258, 382], a parameter closely resembling the Allan variance used by the clock community to estimate the stability of atomic clocks [365, 4]. Both PSRs B1937+21 and B1855+09 seem to be limited by a power law component which produces a minimum in σ_z after 2 yr and 5 yr respectively. This is most likely a result of a small amount of intrinsic timing noise [189]. The σ_z based on timing observations [390] of the bright millisecond pulsar J0437−4715 is now 1–2 orders of magnitude smaller than the other two pulsars or the atomic clock!

4.4 Timing binary pulsars

For binary pulsars, the timing model introduced in Section 4.2 needs to be extended to incorporate the additional motion of the pulsar as it orbits the common centre-of-mass of the binary system. Describing the binary orbit using Kepler's laws to refer the TOAs to the binary barycentre requires five additional model parameters: the orbital period $P_{\rm b}$, projected semi-major orbital axis x, orbital eccentricity e, longitude of periastron ω and the epoch of periastron passage T_0. This description, using five "Keplerian parameters", is identical to that used for spectroscopic binary stars. Analogous to the radial velocity curve in a spectroscopic binary, for binary pulsars the orbit is described by the apparent pulse period against time. An example of this is shown in Panel a

of Figure 25. Alternatively, when radial accelerations can be measured, the orbit can also be visualised in a plot of acceleration versus period as shown in Panel b of Figure 25. This method is particularly useful for determining binary pulsar orbits from sparsely sampled data [122].

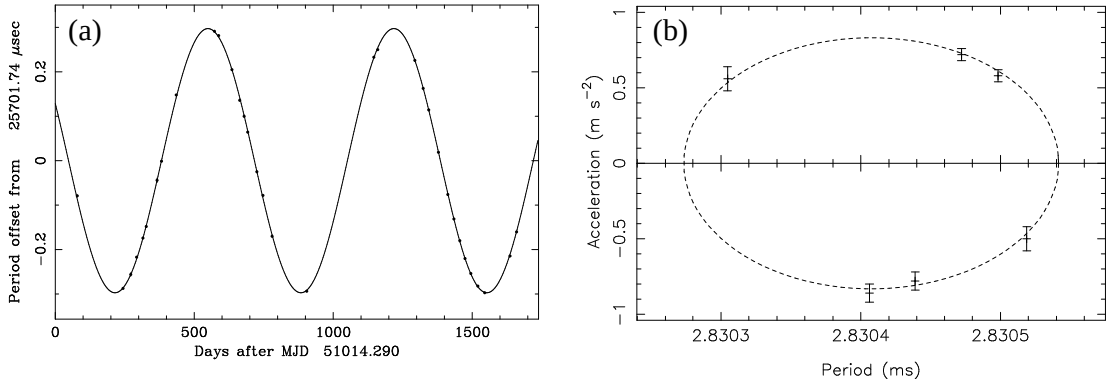

Figure 25: *Panel a: Keplerian orbital fit to the* 669-day *binary pulsar J0407+1607 [237]. Panel b: Orbital fit in the period-acceleration plane for the globular cluster pulsar 47 Tuc S [122].*

Constraints on the masses of the pulsar $m_{\rm p}$ and the orbiting companion $m_{\rm c}$ can be placed by combining x and $P_{\rm b}$ to obtain the mass function

$$f_{\rm mass} = \frac{4\pi^2}{G}\frac{x^3}{P_{\rm b}^2} = \frac{(m_{\rm c}\sin i)^3}{(m_{\rm p}+m_{\rm c})^2}, \tag{11}$$

where G is Newton's gravitational constant and i is the (initially unknown) angle between the orbital plane and the plane of the sky (i.e. an orbit viewed edge-on corresponds to $i = 90°$). In the absence of further information, the standard practice is to consider a random distribution of inclination angles. Since the probability that i is *less* than some value i_0 is $p(< i_0) = 1 - \cos(i_0)$, the 90% confidence interval for i is $26° < i < 90°$. For an assumed pulsar mass, the 90% confidence interval for $m_{\rm c}$ can be obtained by solving Equation (11) for $i = 26°$ and $90°$.

Although most of the presently known binary pulsar systems can be adequately timed using Kepler's laws, there are a number which require an additional set of "post-Keplerian" (PK) parameters which have a distinct functional form for a given relativistic theory of gravity [96]. In general relativity (GR) the PK formalism gives the relativistic advance of periastron

$$\dot{\omega} = 3\left(\frac{P_{\rm b}}{2\pi}\right)^{-5/3}(T_\odot M)^{2/3}(1-e^2)^{-1}, \tag{12}$$

the time dilation and gravitational redshift parameter

$$\gamma = e\left(\frac{P_{\rm b}}{2\pi}\right)^{1/3} T_\odot^{2/3} M^{-4/3} m_{\rm c}(m_{\rm p}+2m_{\rm c}), \tag{13}$$

the rate of orbital decay due to gravitational radiation

$$\dot{P}_{\rm b} = -\frac{192\pi}{5}\left(\frac{P_{\rm b}}{2\pi}\right)^{-5/3}\left(1+\frac{73}{24}e^2+\frac{37}{96}e^4\right)\left(1-e^2\right)^{-7/2} T_\odot^{5/3} m_{\rm p} m_{\rm c} M^{-1/3} \tag{14}$$

and the two Shapiro delay parameters

$$r = T_\odot m_{\rm c} \tag{15}$$

and

$$s = x \left(\frac{P_\mathrm{b}}{2\pi} \right)^{-2/3} T_\odot^{-1/3} M^{2/3} m_\mathrm{c}^{-1} \tag{16}$$

which describe the delay in the pulses around superior conjunction where the pulsar radiation traverses the gravitational well of its companion. In the above expressions, all masses are in solar units, $M \equiv m_\mathrm{p} + m_\mathrm{c}$, $x \equiv a_\mathrm{p} \sin i/c$, $s \equiv \sin i$ and $T_\odot \equiv GM_\odot/c^3 = 4.925490947\,\mu$s. Some combinations, or all, of the PK parameters have now been measured for a number of binary pulsar systems. Further PK parameters due to aberration and relativistic deformation [94] are not listed here but may soon be important for the double pulsar [204].

4.5 Testing general relativity

The key point in the PK definitions introduced in the previous section is that, given the precisely measured Keplerian parameters, the only two unknowns are the masses of the pulsar and its companion, m_p and m_c. Hence, from a measurement of just two PK parameters (e.g., $\dot{\omega}$ and γ) one can solve for the two masses and, using Equation (11), find the orbital inclination angle i. If three (or more) PK parameters are measured, the system is "overdetermined" and can be used to test GR (or, more generally, any other theory of gravity) by comparing the third PK parameter with the predicted value based on the masses determined from the other two.

The first binary pulsar used to test GR in this way was PSR B1913+16 discovered by Hulse & Taylor in 1974 [155]. Measurements of three PK parameters ($\dot{\omega}$, γ and $\dot{P}_\mathrm{b}$) were obtained from long-term timing observations at Arecibo [369, 370, 393]. The measurement of orbital decay, which corresponds to a shrinkage of about 3.2 mm per orbit, is seen most dramatically as the gradually increasing shift in orbital phase for periastron passages with respect to a non-decaying orbit shown in Figure 26. This figure includes recent Arecibo data taken since the upgrade of the telescope in the mid 1990s.

The measurement of orbital decay in the B1913+16 system, obtained from observations spanning a 30-yr baseline [393], is within 0.2% of the GR prediction and provided the first indirect evidence for the existence of gravitational waves. Hulse and Taylor were awarded the 1993 Nobel Physics prize [371, 154, 366] in recognition of their discovery of this remarkable laboratory for testing GR. A similar, though less precise, test of GR from this combination of PK parameters has recently been performed in the double neutron star binary PSR B2127+11C in the globular cluster M15 [161]. For this system, the measurement uncertainties permit a test of GR to 3% precision which is unlikely to improve due to various kinematic contaminations including the acceleration of the binary system in the cluster's gravitational potential?

Five PK parameters have been measured for the double pulsar discussed in detail below and for the double neutron star system PSR B1534+12 [354] where the test of GR comes from measurements of $\dot{\omega}$, γ and s. In this system, the agreement between theory and observation is within 0.7% [354]. This test will improve in the future as the timing baseline extends and a more significant measurement of r can be made.

Currently the best binary pulsar system for testing GR in the strong-field regime is the double pulsar J0737−3039. In this system, where two independent pulsar clocks can be timed, five PK parameters of the 22.7-ms pulsar "A" have been measured as well as two additional constraints from the measured mass function and projected semi-major axis of the 2.7-s pulsar "B". A useful means of summarising the limits so far is Figure 27 which shows the allowed regions of parameter space in terms of the masses of the two pulsars. The shaded regions are excluded by the requirement that $\sin i < 1$. Further constraints are shown as pairs of lines enclosing permitted regions as predicted by GR. The measurement [204] of $\dot{\omega} = 16.899 \pm 0.001$ deg yr^{-1} gives the total system mass $M = 2.5871 \pm 0.0002\,M_\odot$.

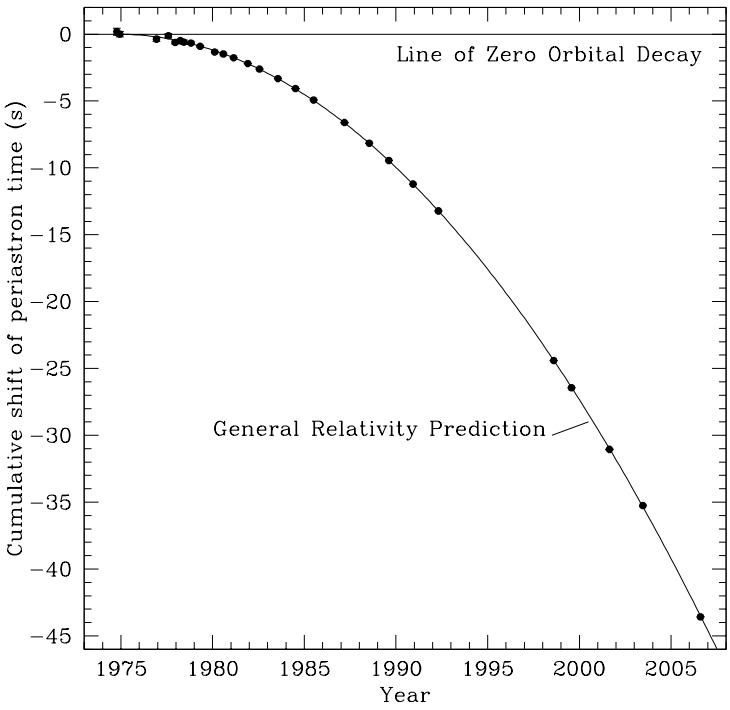

Figure 26: *Orbital decay in the binary pulsar B1913+16 system demonstrated as an increasing orbital phase shift for periastron passages with time. The GR prediction due entirely to the emission of gravitational radiation is shown by the parabola. Figure provided by Joel Weisberg.*

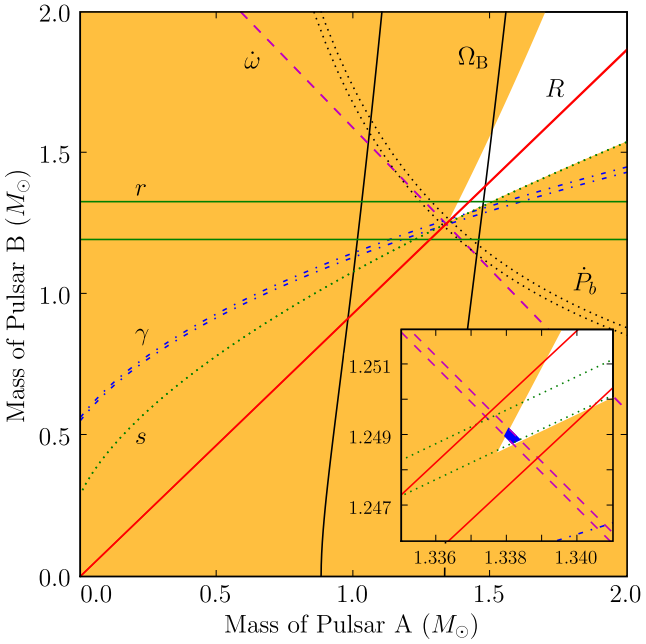

Figure 27: *'Mass–mass' diagram showing the observational constraints on the masses of the neutron stars in the double pulsar system J0737−3039. Inset is an enlarged view of the small square encompassing the intersection of the tightest constraints. Figure provided by Rene Breton [47].*

The measurement of the projected semi-major axes of both orbits gives the mass ratio $R = 1.071 \pm 0.001$. The mass ratio measurement is unique to the double pulsar system and rests on the basic assumption that momentum is conserved. This constraint should apply to any reasonable theory of gravity. The intersection between the lines for $\dot{\omega}$ and R yield the masses of A and B as $m_{\rm A} = 1.3381 \pm 0.007\,M_\odot$ and $m_{\rm B} = 1.2489 \pm 0.0007\,M_\odot$. From these values, using Equations (13–16) the expected values of γ, $\dot{P}_{\rm b}$, r and s may be calculated and compared with the observed values. These four tests of GR all agree with the theory to within the uncertainties. Currently the tightest constraint is the Shapiro delay parameter s where the observed value is in agreement with GR at the 0.05% level [204].

Another unique feature of the double pulsar system is the interactions between the two pulsars' radiation beams. Specifically, the signal from A is eclipsed for 30 s each orbit by the magnetosphere of B [243, 185, 262] and the radio pulses from B are modulated by the relativistic wind from A during two phases of the orbit [263]. These provide unique insights into plasma physics [129] and, as shown shown in Figure 27, a measurement of relativistic spin–orbit coupling in the binary system and, hence, another constraint on GR. By careful modeling of the change in eclipse profiles of A over a four-year baseline, Breton et al. [47] have been able to fit a remarkably simple model [3] and determine the precession of B's spin axis about the orbital angular momentum vector. This remarkable measurement agrees, within the 13% measurement uncertainty, to the GR prediction.

It took only two years for the double pulsar system to surpass the tests of GR possible from three decades of monitoring PSR B1913+16 and over a decade of timing PSR B1534+12. Ongoing precision timing measurements of the double pulsar system should soon provide even more stringent and new tests of GR. Crucial to these measurements will be the timing of the 2.7-s pulsar B, where the observed profile is significantly affected by A's relativistic wind [243, 263]. A careful decoupling of these profile variations is required to accurately measure TOAs for this pulsar and determine the extent to which the theory-independent mass ratio R can be measured. This task is compounded by the fact that pulsar B's signal is now [47] significantly weaker compared to the discovery observations five years ago [243]. This appears to be a direct result of its beam precessing out of our line of sight.

The relativistic effects observed in the double pulsar system are so large that corrections to higher post-Newtonian order may soon need to be considered. For example, $\dot{\omega}$ may be measured precisely enough to require terms of second post-Newtonian order to be included in the computations [95]. In addition, in contrast to Newtonian physics, GR predicts that the spins of the neutron stars affect their orbital motion via spin-orbit coupling. This effect would most clearly be visible as a contribution to the observed $\dot{\omega}$ in a secular [29] and periodic fashion [397]. For the J0737−3039 system, the expected contribution is about an order of magnitude larger than for PSR B1913+16 [243]. As the exact value depends on the pulsars' moment of inertia, a potential measurement of this effect allows the moment of inertia of a neutron star to be determined for the first time [95]. Such a measurement would be invaluable for studies of the neutron star equation of state and our understanding of matter at extreme pressure and densities [209].

Finally, in the neutron star–white dwarf binary J1141−6545, the measurement of $\dot{P}_b$, $\dot{\omega}$, γ and s permit a 6% test of GR [38]. Despite this being nominally only a relatively weak constraint for GR, the significantly different masses of the binary components (1.27 $M_\odot$ for the pulsar versus 1.02 $M_\odot$ for the companion) permits constraints on two coupling parameters between matter and the scalar gravitational field. These parameters are non-existent in GR, which describes gravity in terms of a tensor field, but are predicted by some alternative theories of gravity which consider tensor and scalar field components. The constraints provided by PSR J1141−6545 for the strongly non-linear coupling parameter are currently the most stringent to date and are set to improve dramatically over the next five years of timing observations.

[3] The reader is encouraged to view the beautiful movies explaining the eclipse model [261]

4.6 Pulsar timing and neutron star masses

Multiple PK parameters have now been measured for a number of binary pulsars which provide very precise measurements on the neutron star masses [374, 348]. Figure 28 shows the distribution of masses following updates to a recent compilation [348]. While the young pulsars and the double neutron star binaries are consistent with, or just below, the canonical $1.4\,M_\odot$, we note that the millisecond pulsars in binary systems have, on average, significantly larger masses. This provides strong support for their formation through an extended period of accretion in the past, as discussed in Section 2.6.

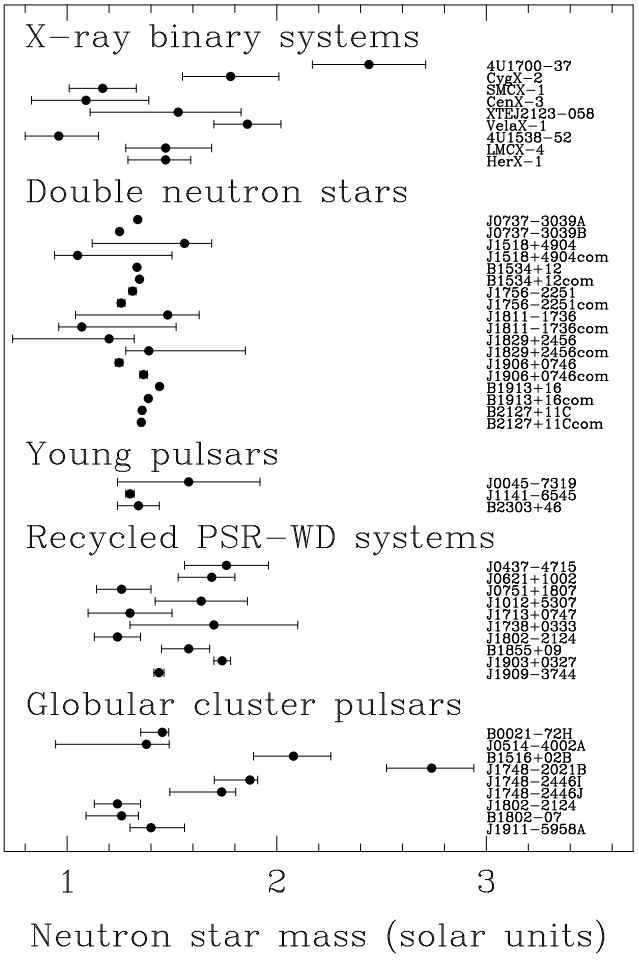

Figure 28: *Distribution of neutron star masses as inferred from timing observations of binary pulsars and X-ray binary systems. Figure adapted from an original version provided by Ingrid stairs. Additional information on globular cluster pulsar mass constraints provided by Paulo Freire.*

Also shown on this diagram are several eccentric binary systems in globular clusters which have their masses constrained from measurements of the relativistic advance of periastron and the Keplerian mass function. In these cases, the condition $\sin i < 1$ sets a lower limit on the companion mass $m_{\rm c} > (f_{\rm mass} M^2)^{1/3}$ and a corresponding upper limit on the pulsar mass. Probability density functions for both $m_{\rm p}$ and $m_{\rm c}$ can also be estimated in a statistical sense by *assuming* a

random distribution of orbital inclination angles. An example of this is shown in Figure 29 for the eccentric binary millisecond pulsar in M5 [120] where the nominal pulsar mass is 2.08 ± 0.19 $M_\odot$. If these estimations can be confirmed by the measurement of other relativistic parameters, these supermassive neutron stars will have important constraints on the equation of state of superdense matter.

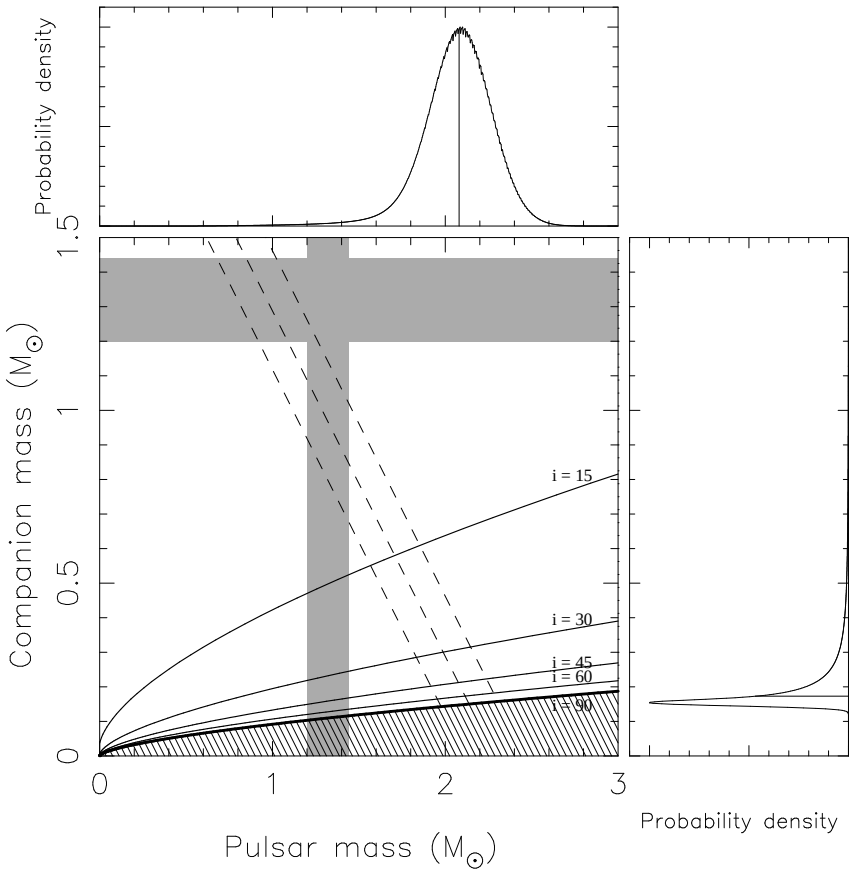

Figure 29: *Pulsar mass versus companion mass diagram showing mass constraints for the eccentric binary millisecond pulsar B1516+02B [120]. The allowed parameter space is bounded by the dashed lines which show the uncertainty on the total mass from the measurement of the relativistic advance of periastron. Masses within the hatched region are disallowed by the Keplerian mass function and the constraint that* $\sin i < 1$. *Assuming a random distribution of orbital inclination angles allows the probability density functions of the pulsar and companion mass to be derived, as shown by the top and right hand panels. Figure provided by Paulo Freire.*

Currently the largest measurement of a radio pulsar mass through multiple PK parameters is the eccentric millisecond pulsar binary system J1903+0327 [71] (see also Section 2.9). Thus, in addition to challenging models of millisecond pulsar formation, this new discovery has important implications for fundamental physics. When placed on the mass–radius diagram for neutron stars [209], this 1.74 $M_\odot$ pulsar appears to be incompatible with at least four equations of state for superdense matter. Future timing measurements are required to consolidate and verify this potentially very exciting result.

4.7 Pulsar timing and gravitational wave detection

We have seen in the Section 4.4 how pulsar timing can be used to provide indirect evidence for the existence of gravitational waves from coalescing stellar-mass binaries. In this final section, we look at how pulsar timing might soon be used to detect gravitational radiation directly. The idea to use pulsars as natural gravitational wave detectors was first explored in the late 1970s [333, 99]. The basic concept is to treat the solar system barycentre and a distant pulsar as opposite ends of an imaginary arm in space. The pulsar acts as the reference clock at one end of the arm sending out regular signals which are monitored by an observer on the Earth over some time-scale T. The effect of a passing gravitational wave would be to perturb the local spacetime metric and cause a change in the observed rotational frequency of the pulsar. For regular pulsar timing observations with typical TOA uncertainties of $\epsilon_{\rm TOA}$, this "detector" would be sensitive to waves with dimensionless amplitudes $h \gtrsim \epsilon_{\rm TOA}/T$ and frequencies $f \sim 1/T$ [36, 44].

4.7.1 Probing the gravitational wave background

Many cosmological models predict that the Universe is presently filled with a low-frequency stochastic gravitational wave background (GWB) produced during the big bang era [296]. A significant component [312, 163] is the gravitational radiation from massive black hole mergers. In the ideal case, the change in the observed frequency caused by the GWB should be detectable in the set of timing residuals after the application of an appropriate model for the rotational, astrometric and, where necessary, binary parameters of the pulsar. As discussed in Section 4.2, all other effects being negligible, the rms scatter of these residuals σ would be due to the measurement uncertainties and intrinsic timing noise from the neutron star.

For a GWB with a flat energy spectrum in the frequency band $f \pm f/2$ there is an additional contribution to the timing residuals $\sigma_{\rm g}$ [99]. When $fT \gg 1$, the corresponding wave energy density is

$$\rho_{\rm g} = \frac{243\pi^3 f^4 \sigma_{\rm g}^2}{208G}. \tag{17}$$

An upper limit to $\rho_{\rm g}$ can be obtained from a set of timing residuals by assuming the rms scatter is entirely due to this effect ($\sigma = \sigma_{\rm g}$). These limits are commonly expressed as a fraction of the energy density required to close the Universe

$$\rho_{\rm c} = \frac{3H_0^2}{8\pi G} \simeq 2 \times 10^{-29} h^2 \ {\rm g\ cm^{-3}}, \tag{18}$$

where the Hubble constant $H_0 = 100\,h$ km s^{-1} Mpc.

This technique was first applied [328] to a set of TOAs for PSR B1237+25 obtained from regular observations over a period of 11 years as part of the JPL pulsar timing programme [103]. This pulsar was chosen on the basis of its relatively low level of timing activity by comparison with the youngest pulsars, whose residuals are ultimately plagued by timing noise (see Section 4.3). By ascribing the rms scatter in the residuals ($\sigma = 240$ ms) to the GWB, the limit is $\rho_{\rm g}/\rho_{\rm c} \lesssim 4 \times 10^{-3} h^{-2}$ for a centre frequency $f = 7$ nHz.

This limit, already well below the energy density required to close the Universe, was further reduced following the long-term timing measurements of millisecond pulsars at Arecibo (see Section 4.3). In the intervening period, more elaborate techniques had been devised [36, 44, 359] to look for the likely signature of a GWB in the frequency spectrum of the timing residuals and to address the possibility of "fitting out" the signal in the TOAs. Following [36] it is convenient to define

$$\Omega_{\rm g} = \frac{1}{\rho_{\rm c}} \frac{{\rm d}\log\rho_{\rm g}}{{\rm d}\log f}, \tag{19}$$

i.e. the energy density of the GWB per logarithmic frequency interval relative to ρ_c. With this definition, the power spectrum of the GWB [152, 44] is

$$\mathcal{P}(f) = \frac{G\rho_g}{3\pi^3 f^4} = \frac{H_0^2 \Omega_g}{8\pi^4 f^5} = 1.34 \times 10^4 \Omega_g h^2 f_{\mathrm{yr}^{-1}}^{-5} \ \mu\mathrm{s}^2 \ \mathrm{yr}, \tag{20}$$

where $f_{\mathrm{yr}^{-1}}$ is frequency in cycles per year. The long-term timing stability of B1937+21, discussed in Section 4.3, limits its use for periods $\gtrsim 2$ yr. Using the more stable residuals for PSR B1855+09, Kaspi et al. [189] placed an upper limit of $\Omega_g h^2 < 1.1 \times 10^{-7}$. This limit is difficult to reconcile with most cosmic string models for galaxy formation [53, 375].

For binary pulsars, the orbital period provides an additional clock for measuring the effects of gravitational waves. In this case, the range of frequencies is not limited by the time span of the observations, allowing the detection of waves with periods as large as the light travel time to the binary system [36]. The most stringent results presently available are based on the B1855+09 limit $\Omega_g h^2 < 2.7 \times 10^{-4}$ in the frequency range $10^{-11} < f < 4.4 \times 10^{-9}$ Hz [200].

4.7.2 Constraints on massive black hole binaries

In addition to probing the GWB, pulsar timing is beginning to place interesting constraints on the existence of massive black hole binaries. Arecibo data for PSRs B1937+21 and J1713+0747 already make the existence of an equal-mass black hole binary in Sagittarius A* unlikely [218]. More recently, timing data from B1855+09 have been used to virtually rule out the existence of a proposed supermassive black hole as the explanation for the periodic motion seen at the centre of the radio galaxy 3C66B [361]. A simulation of the expected modulations of the timing residuals for the putative binary system, with a total mass of $5.4 \times 10^{10} M_\odot$, is shown along with the observed timing residuals in Figure 30. Although the exact signature depends on the orientation and eccentricity of the binary system, Monte Carlo simulations show that the existence of such a massive black hole binary is ruled out with at least 95% confidence [165].

4.7.3 A millisecond pulsar timing array

A natural extension of the single-arm detector concept discussed above is the idea of using timing data for a number of pulsars distributed over the whole sky to detect gravitational waves [142]. Such a "timing array" would have the advantage over a single arm in that, through a cross-correlation analysis of the residuals for pairs of pulsars distributed over the sky, it should be possible to separate the timing noise of each pulsar from the signature of the GWB common to all pulsars in the array. To illustrate this, consider the fractional frequency shift of the ith pulsar in an array

$$\frac{\delta\nu_i}{\nu_i} = \alpha_i \mathcal{A}(t) + \mathcal{N}_i(t), \tag{21}$$

where α_i is a geometric factor dependent on the line-of-sight direction to the pulsar and the propagation and polarisation vectors of the gravitational wave of dimensionless amplitude $\mathcal{A}$. The timing noise intrinsic to the pulsar is characterised by the function $\mathcal{N}_i$. The result of a cross-correlation between pulsars i and j is then

$$\alpha_i \alpha_j \langle \mathcal{A}^2 \rangle + \alpha_i \langle \mathcal{A}\mathcal{N}_j \rangle + \alpha_j \langle \mathcal{A}\mathcal{N}_i \rangle + \langle \mathcal{N}_i \mathcal{N}_j \rangle, \tag{22}$$

where the bracketed terms indicate cross-correlations. Since the wave function and the noise contributions from the two pulsars are independent, the cross correlation tends to $\alpha_i \alpha_j \langle \mathcal{A}^2 \rangle$ as the number of residuals becomes large. Summing the cross-correlation functions over a large number of pulsar pairs provides additional information on this term as a function of the angle on the sky [141] and allows the separation of the effects of clock and ephemeris errors from the GWB [116].

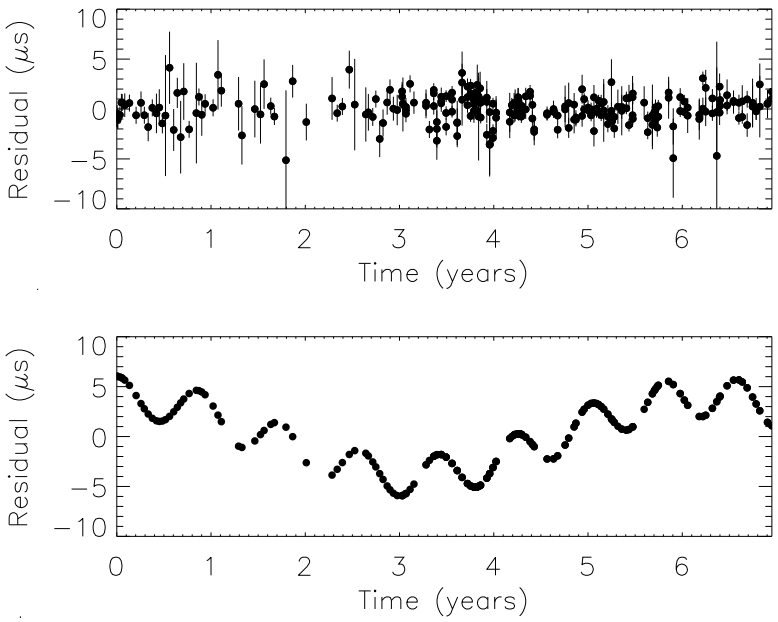

Figure 30: *Top panel: Observed timing residuals for PSR B1855+09. Bottom panel: Simulated timing residuals induced from a putative black hole binary in 3C66B. Figure provided by Rick Jenet [165].*

A recent analysis [164] applying the timing array concept to data for seven millisecond pulsars has reduced the energy density limit to $\Omega_g h^2 < 1.9 \times 10^{-8}$ for a background of supermassive black hole sources. The corresponding limits on the background of relic gravitational waves and cosmic strings are 2.0×10^{-8} and 1.9×10^{-8} respectively. These limits can be used to constrain the merger rate of supermassive black hole binaries at high redshift, investigate inflationary parameters and place limits on the tension of currently proposed cosmic string scenarios.

The region of the gravitational wave energy density spectrum probed by the current pulsar timing array is shown in Figure 31 where it can be seen that the pulsar timing regime is complementary to the higher frequency bands of LISA and LIGO.

A number of long-term timing projects are now underway to make a large-scale pulsar timing array a reality. The Parkes pulsar timing array [14, 148] observes twenty millisecond pulsars twice a month. The European Pulsar Timing Array [13] uses the Lovell, Westerbork, Effelsberg and Nancay radio telescopes to regularly observe a similar number. Finally, at Arecibo and Green Bank, regular timing of a dozen or more millisecond pulsars is carried out by a North American consortium [270]. It is expected that a combined analysis of all these efforts could reach a limits of $\Omega_g h^2 < 10^{-10}$ before 2010 [148]. Looking further ahead, the increase in sensitivity provided by the Square Kilometre Array [158, 202] should further improve the limits of the spectrum probed by pulsar timing. As Figure 31 shows, the SKA could provide up to two orders of magnitude improvement over current capabilities.

4.8 Going further

For further details on the techniques and prospects of pulsar timing, a number of excellent reviews are available [20, 365, 33, 347, 232]. Two audio files and slides from lectures at the centennial meeting of the American Physical Society are also available on-line [18, 401]. The tests of general

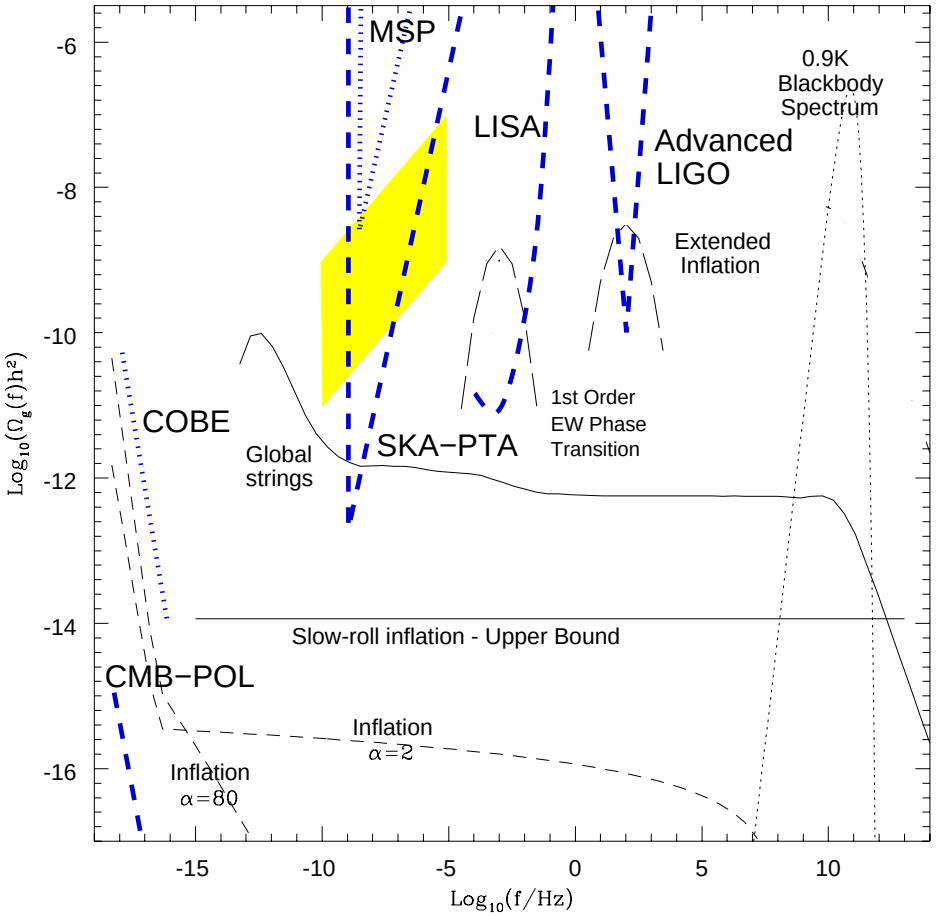

Figure 31: *Summary of the gravitational wave spectrum showing the location in phase space of the pulsar timing array (PTA) and its extension with the Square Kilometre Array (SKA). Figure updated by Michael Kramer [202] from an original design by Richard Battye.*

relativity discussed in Section 4.4, along with further effects, are reviewed elsewhere in this journal by Will [402] and Stairs [347]. Further discussions on the realities of using pulsars as gravity wave detectors can be found in other review articles [326, 17, 148].

5 Summary and Future Prospects

The main aim of this article was to review some of the many astrophysical applications provided by the present sample of binary and millisecond radio pulsars that are relevant to gravitational physics. The topics covered here, along with the bibliography and associated tables of observational parameters, should be useful to those wishing to delve deeper into the vast body of literature that exists.

Through an understanding of the Galactic population of radio pulsars summarised in Section 3 it is possible to predict the detection statistics of terrestrial gravitational wave detectors to nearby rapidly spinning neutron stars (see Section 3.3), as well as coalescing relativistic binaries at cosmic distances (see Section 3.4). Continued improvements in gravitational wave detector sensitivities should result in a number of interesting developments and contributions in this area. These developments and contributions might include the detection of presently known radio pulsars, as well as a population of coalescing binary systems which have not yet been detected as radio pulsars.

The phenomenal timing stability of radio pulsars leads naturally to a large number of applications, including their use as laboratories for relativistic gravity (see Section 4.5), constraining the equation of state of superdense matter (see Section 4.6) and as natural detectors of gravitational radiation (see Section 4.7). Long-term timing experiments of the present sample of millisecond and binary pulsars currently underway appear to have tremendous potential in these areas and perhaps detect the gravitational wave background (if it exists) within the next decade.

These applications will benefit greatly from the continuing discovery of new systems by the present generation of radio pulsar searches which continue to probe new areas of parameter space. Based on the results presented in Section 3.3, it is clear that we are aware of only a few per cent of the total active pulsar population in our Galaxy. In all likelihood, then, we have not seen all of the pulsar zoo. More sensitive surveys using a multibeam system on the Arecibo telescope [277] are now being carried out. Future surveys with the Square Kilometre Array [158] will ultimately provide a far more complete census of the Galactic pulsar population. Three possible "holy grails" of pulsar astronomy which could soon be found are:

An X-ray/radio millisecond pulsar While the link between millisecond pulsars and X-ray binaries appears to be compelling, despite intensive searches [50], no radio pulsations have been detected in these binaries. This could be a result of free-free absorption of any radio waves by the thick accretion disk, or perhaps quenching the accelerating potential in the neutron star magnetosphere by infalling matter. The recent discovery of radio pulsations from two magnetars [65, 64] demonstrates that radio emission can be found in seemingly unlikely sites. A millisecond pulsar in transition between an X-ray and radio phase would be another such remarkable discovery.

A pulsar–black hole binary system: Following the wide variety of science possible from the new double pulsar J0737−3039 (see Sections 2.6, 3.4.1 and 4.4), a radio pulsar with a black-hole companion would without doubt be a fantastic gravitational laboratory. Excellent places to look for such systems are globular clusters and the Galactic centre [300].

A sub-millisecond pulsar: The most rapidly spinning radio pulsar currently known is the 1.39-ms PSR J1748−2446ad in Terzan 5 [145]. Do neutron stars with kHz spin rates exist? Searches now have sensitivity to such objects [49] and a discovery of even one would constrain

the equation of state of matter at high densities. As mentioned in Section 2.2, the R-mode instability may prevent neutron stars from reaching such high spin rates [41, 69]. Of course, this does not stop observers from looking for such objects! The current generation of surveys summarized in Section 2.11 continues to be more and more sensitive to this putative source population.

The discovery of the eccentric millisecond binary pulsar J1903+0327 [71] (see Sections 2.9 and 4.6) was completely unprecedented within the framework of our theories for binary pulsar formation. It therefore serves as a stark reminder that we should "expect the unexpected"! Even though the study of radio pulsars is now 40 years old, these experiences tell us that many exciting discoveries lie ahead of us.

Pulsar astronomy remains an extremely active area of modern astrophysics and the next decade will undoubtedly continue to produce new results from currently known objects as well as new surprises. Although it is often stated that we are bound by available computational resources, In my opinion, pulsar research is currently limited by a shortage of researchers, and not necessarily computational resources. Keen students more than ever are needed to help shape the future of this exciting and continually evolving field.

6 Acknowledgements

Many thanks to Chunglee Kim and Evan Kean who read and provided very useful comments on this revised review, and to Scott Ransom and Matthew Bailes for providing unpublished data given in Table 4. I am indebted to a number of colleagues who, as noted in many of the figure captions, kindly gave permission to use their figures in this article. Frequent use was made of NASA's magnificent Astrophysics Data System [139], the arXiv e-Print service [90] and of course the *Google* during the literature searches. My research is partially supported by West Virginia EPSCoR through a Research Challenge Grant. Finally, I would like to thank the *Living Reviews* editors for their support and patience during the completion of this long-overdue update.