UNIVERSITÀ DEGLI STUDI DI CAGLIARI
Facoltà di Scienze Matematiche, Fisiche e Naturali
Scuola di Dottorato in Scienze e Tecnologie Fisiche
XXIV Ciclo

Search for Millisecond Pulsars for the Pulsar Timing Array project

Ph.D. Thesis by
Sabrina Milia

Tutor
Prof. Luciano Burderi

Supervisors
Dr. Andrea Possenti – Dr. Marta Burgay
Dr. Willem van Straten

Anno Accademico 2010-2011

To my family and my friends,
always a light(house) in the darkness

Contents

Introduction

Pulsars are rapidly rotating highly magnetised neutron stars (i.e. ultra dense stars, where about one solar mass is concentrated in a sphere with a radius of ~ 10 km), which irradiate radio beams in a fashion similar to a lighthouse. As a consequence, whenever the beams cut our line of sight we perceive a radio pulses, one (or two) per pulsar rotation, with a frequency up to hundred of times a second.

The pulsar population splits in two main families: the so-called ordinary pulsars (with spin periods between tens of milliseconds and about 8 seconds, and spin period derivatives between 10^{-18} s s^{-1} and 10^{-11} s s^{-1}) and the so-called millisecond pulsars (MSPs, Backer et al. 1982, Nature, 300, 615), having a spin period smaller than the conventional limit of 30 ms and a spin period derivate smaller than about 10^{-18} s s^{-1}. The latter are also dubbed "recycled pulsars" due to their evolutionary mechanism: in fact, they are old and once slowly spinning neutron stars, which have been re-accelerated to short rotational periods via transfer of matter and angular momentum from a stellar companion in a binary system (e.g. Alpar et al, 1982, Nature, 300, 728)

Owing to their compact nature, rapid spin and high inertia, pulsars are in general fairly stable rotators, hence the Times of Arrival (TOAs) of the pulses at a radio telescope can be used as the ticks of a clock. This holds true in particular for the sub-class of the MSPs, whose very rapid rotation and relatively older age (likely implying a fully relaxed neutron star structure), provide better rotational stability than the ordinary pulsars. Indeed, some MSPs rotate so regularly that they can rival the best atomic clocks on Earth over timespan of few months or years (e.g. Lorimer, 2008, Living Review in Relativity, 8).

This feature allows us to use MSPs as tools in a cosmic laboratory. This is performed by exploiting a procedure called *timing* (see e.g. Lyne & Graham-

Smith 2005, Pulsar Astronomy,Cambridge Un Press), which consists in the repeated and regular measurement of the TOAs from a pulsar and then in the search for trends in the series of the TOAs over various timespans, from fraction of seconds to decades. After adopting a timing model, one can measure the differences between the observed TOAs and those predicted by the model and, by minimizing those differences (called *residuals*), one can infer the parameters of the timing model. The values of these parameters represent in turn the physical quantities opening the possibility of using the pulsars for a variety of investigation in fundamental physics and astrophysics (e.g. Lorimer & Kramer 2004, Handbook of Pulsar Astronomy, Cambridge Un Press). They range from very sensitive tests of General Relativity and alternate gravity theories to the discrimination between various proposed Equations of State for the high density matter, from constraining the properties of electromagnetic waves in ultra highly magnetised plasma to the study of the origin and shape of the magnetic field in the Galaxy, from the determination of the distribution of the ionized component in the interstellar medium to the investigation of the latest stages in the stars and binary stars evolution, from the discovery of Earth-mass planets to the measurement of the gravitational potential well of the globular clusters.

As anticipated above, the study of the binary pulsars has already provided the most stringent tests to date of General Relativity in strong gravitational fields and has unambiguously showed the occurrence of the emission of gravitational waves from a binary system comprising two massive bodies in a close orbit (Taylor & Weisberg, 1989, ApJ, 345, 434). In last decades a new exciting perspective has been opened, i.e. to use pulsars also for a *direct* detection of the so far elusive gravitational waves and thereby applying the pulsar timing for cosmological studies (Hellings & Downs 1983, ApJ, 265, L39).

In fact, the gravitational waves (GWs) going across our Galaxy pass over all the Galactic pulsars and the Earth, perturbing the space-time at the pulsar and Earth locations, as well as anywhere along the line-of-sight from the Earth and each of the pulsars. This in turn produces a modulation in the rhythm of the TOAs of the pulses from all the pulsars, with the variation in the TOAs having a strength which is proportional to the amplitude of the GW and a periodicity related to the frequency of the GW. Of course if they are caused by a common physical phenomenon (like a passing-by GW), these variations of the TOAs are

expected to be somehow correlated between the various pulsars, allowing us to disentangle this effect from other effects which could mimic the occurrence of such modulation, like intrinsic irregularities in the rotation of a pulsar, changing interstellar medium along the line of sight, error in the reference clocks used for determining the TOAs and so on.

The consideration of the aforementioned possible sources of additional effects which could mask the signature of a genuine GW shows that a safe direct detection of a GW cannot involve the observation and timing of a single pulsar. Instead, it has been theoretically shown that high precision timing over a 5-10 years data-span of a network of suitable MSPs forming a so-called Pulsar Timing Array (PTA) - in which the pulsars are used as the endpoints of arms of a huge cosmic GW detector - would allow us to overcome the previous problems and open the possibility of a direct detection of GWs (Foster & Backer 1990, ApJ, 361, 300). In particular such apparatus is able to detect GWs in the frequency range between 10^{-9} and 10^{-7} Hz, with the best sensitivity around the nanoHz. Therefore the PTAs nicely complement the other already active or planned GW detectors (e.g. the review Manchester 2011, AIP, 1357, 65).

Given the frequency range of operation, the most favorable source of GWs for a PTA appears to be the cosmological background of GWs produced by the coalescence of supermassive binary black-holes in the early stages of the Universe evolution, at redshift around 1-2 (Sesana et al. 2008, MNRAS, 390, 192). Also a single merging event of a supermassive binary black-hole could be detected, provided it is relatively close and involves very high massive black-holes of about 10^9 solar masses. However, given these constraints, the latter event appear less likely as the source of the first clear detection of GWs by a Pulsar Timing Array. As already mentioned, a problem that arises when considering the detection of the stochastic GW background is to distinguish it from other effects, since the background should yield a TOA perturbation, visible in the timing residuals, similar to that imprinted by the so-called *red timing noise* (e.g. Cordes & Shannon 2011, ArXiv e-prints, arXiv:1106.4047C). This is a kind of noise having power strongly concentrated at lower fluctuation frequencies and that can be yielded by intrinsic irregularities in the pulsar spin and by interstellar scintillation. The challenge is hence to separate the GW background contribution from the red noise one, in order to safely detect the

former. According to Cordes & Shannon 2011, a PTA made up of a sample of 20 super-stable MSPs, each with timing residuals from red noise contribution less than 20 ns over time spans of 5 years and with negligible white noise (i.e. radiometer noise and pulsar jitter), should allow a plausible detection of the GW background. However, if no such super-stable objects exist (so far we know only two MSPs having timing residuals less than 50 ns over 5 years, and only one other less than 100 ns) and hence MSPs have timing noise larger than 20 ns over 5 years, it will be necessary to time many more MSPs, maybe 100 or even more. Therefore it is of *fundamental importance* that more MSPs with timing noise substantially less than 100 ns in a 5-year span are discovered.

From what said, in order to set up a suitable PTA it is necessary on one hand to search for new MSPs having the required clock stability and signal intensity, and on another hand to perform regular high-precision timing observations of the available sample, combining the results from all the pulsars with the use of a solid and well tested software, capable of revealing the genuine GW signal which is searched for.

This work focuses on the first task, in an attempt to enlarge the number of suitable MSPs, in the framework of the *High Time Resolution Universe* (HTRU) survey for pulsars and fast radio transients, that is currently underway at the 64-m Parkes Radio Telescope (NSW, Australia). This experiment (Keith et al. 2010, MNRAS, 409, 619) has been designed in 2007 and started three years ago, with the main scope of largely increasing (possibly doubling) the total number of MSPs known in the Galactic Field (there were only about 40 of them until 2009). The enlarged sample may provide some very good *MSP-clocks* to be added to the still relatively poor list of objects well suited for belonging to a PTA.

In the first chapter of this thesis an overview of the pulsar phenomenon is given, with also a description of the timing technique and its physical applications. The search methods that can be used to analyse the data in order to find isolated and binary pulsars are reported in the second chapter. The third chapter describes part of the work performed by me in the framework of the HTRU survey; in particular the search for MSPs in the HTRU data with a data reduction pipeline sensitive also to highly relativistic systems (i.e. to binary pulsars in close orbits). While performing the aforementioned search,

it emerged the issue of the inspection of the hundreds of thousands of pulsar candidates produced by the adopted pipeline, the vast majority of them being the result of radio interferences. Therefore, a new approach has been explored for making manageable the human intervention in the procedure of selection of the trustable candidates, namely the use of an Artificial Neural Network on the pulsar candidates. The fourth chapter is devoted to report on that. At the end, a brief summary of this thesis work is given, as well as a list of the publications, in preparation and resulting from the HTRU collaborative effort.

Chapter 1

Pulsars: an overview

In this chapter an introductory description of the pulsar phenomenon will be given, together with the explanation of the *timing* procedure that can be employed to study several physical issues, making pulsars a powerful tool for scientific invetigations. One of these applications concerns Pulsar Timing Arrays, where suitable pulsars of the sub-class of millisecond pulsars should allow, owing to their exceptional rotational stability, to detect Gravitational Waves.

This chapter has been largely inspired by the following books: *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005, *Pulsar Astronomy, 3rd ed.* by Lyne and Smith 2005, *An introduction to Radio Astronomy* by Burke and Smith 2002, *An introduction to modern astrophysics* by Carroll and Ostlie 1996 and the useful course on Essential Radio Astronomy by J. J. Condon and S. M. Ransom on the NRAO website (www.cv.nrao.edu/course/astr534/ERA.shtml). So, they are meant to be the bibliographic references for the content of this chapter, whenever an alternate explicit reference is not indicated.

1.1 Pulsar physics

1.1.1 Nature and energetics

Pulsars (pulsa*ting* r*adio sources*) are fascinating celestial objects discovered in 1967 by Jocelyn Bell (Hewish et al. 1968) through their pulsed emission in the radio band of the electromagnetic spectrum. Actually, they have been identified as highly magnetised neutron stars (NSs) which rotate very fast and emit radio beams by a still not completely understood mechanism; so far, about 2000

pulsars have been found, and a number of them have been revealed also at other wavelengths, such as in the optical, X-ray and γ-ray bands.

To understand the origin of a NS, and thus of a pulsar, we must consider the final evolutionary stage of a massive star, specifically of a star with a mass between 8 and 25 solar masses ($M_\odot \sim 2 \times 10^{33}$ g). At the end of its life (i.e. when it has burned all its fuel) such a star undergoes a supernova explosion, during which the majority of its matter is ejected in space and only the central part (the *core*) of the original star is left, in the form of a *compact*[1] object made up of degenerate neutrons, whose pressure avoids a further collapse of the core. Typically, in the end, $\sim 1.2 - 2.0\ M_\odot$ (this is the range of NS masses measured so far) are concentrated in a sphere with a radius of only ~ 10 km, so that a NS is an incredibly dense object: its central density reaches in fact $10^{14} - 10^{15}$ g/cm^3.

Fast spinning and highly magnetic fields heuristically[2] result from conservation of angular momentum and magnetic flux during the supernova explosion (and the related implosion of the core). In particular radio-pulsars have a very small spin period (from milliseconds to seconds) and a very large magnetic field. Young pulsars have magnetic polar field strenghts of about 10^{12} G, value that falls to 10^{10} G for old pulsars and to 10^8 G for the so-called *recycled pulsars*, explained in section 1.1.7; for another type of NSs, namely the *magnetars* (two of them also emitting radio pulses), the magnetic field approaches 10^{15} G, so that they are the most magnetic objects known in the Universe.

The presence of a radio pulsar is revealed by a sequence of pulses observed in the radio band at regular intervals. That happens because the pulsar irradiates radio beam(s) from one (or both) magnetic poles, and the magnetic axis is misaligned with respect to the spin axis. Of course we can observe the emission, in the form of a pulse, only when a beam cuts our line of sight (*lighthouse effect*, see fig.1.1). Since we usually see only one of the two beams, we receive pulses at the same frequency of the pulsar rotation; so the time interval between two pulses is equal to the pulsar spin period, which is very stable since pulsars are

[1] The compactness of an object is defined as the ratio M/R between its mass and radius; a compact object has hence a very large mass concentrated in a small volume.

[2] The details of the process leading to the building of high magnetic fields in a NS are not yet assessed.

extremely massive (and hence with a high inertia) rotators.

However, from the observations it is clear that the pulse period P increases, although very slowly (the value of $\dot{P}$ ranges[3] from something of the order of 10^{-21} s/s for fast spinning millisecond pulsars to 10^{-10} s/s for very slow long-period pulsars and magnetars), i.e. the pulsars *spin down*. The rotational kinetic energy loss is called *spin-down luminosity* and usually represents the total energy output of the pulsar; its value can be calculated by measuring the spin period and its variation with time, according to the formula:

$$\dot{E} = -\frac{dE_{rot}}{dt} = -\frac{d(I\Omega^2/2)}{dt} = 4\pi^2 I \dot{P} P^{-3}, \tag{1.1}$$

where I is the pulsar moment of inertia and $\Omega = 2\pi/P$ is its spin angular frequency. Since $I = kMR^2$ (where k is a constant that depends on the density profile of the object and that, in first approximation, can be set equal to 0.4, i.e. the value for a uniform density sphere), for the typical values of $M = 1.4$ $M_\odot$ and $R = 10$ km the moment of inertia is $I \sim 10^{45}$ g $\cdot$ cm^2 and therefore the spin-down luminosity is:

$$\dot{E} \simeq 3.95 \times 10^{31} \text{erg s}^{-1} \left(\frac{\dot{P}}{10^{-15}}\right) \left(\frac{P}{\text{s}}\right)^{-3}. \tag{1.2}$$

However, from the observations we know that only a very small fraction of this energy is emitted in the radio band between 100 MHz and 100 GHz; most of it is converted into high-energy emission, acceleration of charged particles (originating the so-called pulsar wind), and probably into magnetic dipole radiation (i.e. a monochromatic electromagnetic waves, with frequency equal to the pulsar spin frequency, i.e. of the order of Hz, which is likely quickly absorbed by the interstellar medium).

1.1.2 Magnetic field

As mentioned in the previous section, at first approximation and without accounting for any detailed physical mechanism of emission, the energy budget can be explained considering the fact that the pulsar is a highly magnetised rotating NS having the magnetic axis not aligned with the spin axis, and therefore we can apply to it the *rotating magnetic dipole* model (see fig.1.1).

[3] *www.atnf.csiro.au/people/pulsar/psrcat*; Manchester et al. 2005

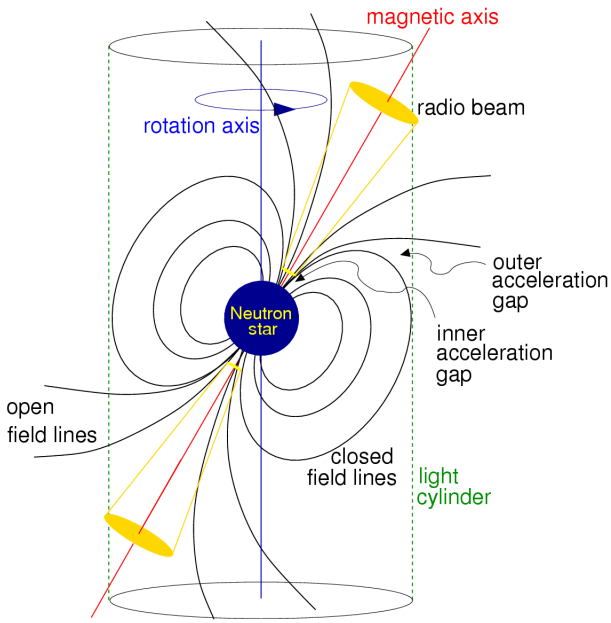

Figure 1.1: Rotating magnetic dipole model of a pulsar. From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

According to such a model, the radiated power is given by the analogue of Larmor's formula from electrodynamics (see Jackson 1962), adapted to a magnetic dipole $\vec{m}$ inclined with respect to the spin axis:

$$\dot{E}_{dip} = \frac{2}{3c^3}|\ddot{\vec{m}}_{\perp}|^2 = \frac{2}{3c^3}|\vec{m}|^2\Omega^4 sin^2\alpha, \tag{1.3}$$

where c is the speed of light, α is the angle between the magnetic dipole and the spin axis, $\vec{m_{\perp}} = \vec{m}\, sin\alpha$ is the component of the magnetic dipole moment perpendicular to the spin axis, and $\vec{m} = \vec{m_0}\, exp(-i\Omega t)$. Since the dipole is inclined it undergoes a variation with time, that causes the emission. Equating the previuos expression to the spin-down luminosity (eq. (1.1)) we obtain the variation of the rotational frequency with time:

$$\dot{\Omega} = -\left(2\frac{|\vec{m}|^2 sin^2\alpha}{3Ic^3}\right)\Omega^3. \tag{1.4}$$

More in general, we can say that Ω evolves according to a power law:

$$\dot{\Omega} = -k\Omega^n, \tag{1.5}$$

where k is a constant and n is called *braking index*; from eq. (1.4) we have that $n = 3$ for a magnetic dipole in the vacuum. Deriving $\dot{\Omega}$ with respect to time we can find an expression for the braking index in terms of Ω or $\nu = \Omega/2\pi$:

$$n = \frac{\Omega\ddot{\Omega}}{\dot{\Omega}^2} = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}. \tag{1.6}$$

Again from classical electrodynamics, we have that $|\vec{m}| \approx B_S R^3$ for a uniformly magnetised sphere with radius R and surface magnetic field B_S; replacing $|\vec{m}|$ obtained from this expression in eq. (1.4), expressing Ω and $\dot{\Omega}$ in terms of P and $\dot{P}$, and rearranging, we can obtain B_S:

$$B_S = \left(\frac{3c^3}{8\pi^2}\frac{I}{R^6 sin^2\alpha}P\dot{P}\right)^{1/2} \simeq 3.2 \times 10^{19}\sqrt{P\dot{P}} \quad \text{G}, \tag{1.7}$$

where the last term has been calculated for the canonical values of $I = 10^{45}$ g $\cdot$ cm^2 and $R = 10$ km, and assuming $\alpha = 90°$.

Thanks to this important relation, it is possible to infer the value of the magnetic field at the surface of the pulsar by the observations of P and $\dot{P}$.

1.1.3 A model for the radio emission mechanism

What said in the previous section can not explain why pulsars emit also in the radio band. Goldreich and Julian 1969 elaborated one of the most popular models, called *polar-cap model*, in an effort to explain that emission on the basis of plasma extracted from the NS surface, for the simplest case in which magnetic and spin axes are aligned. Any attempts to extend the results to the more realistic case of magnetic axis inclined with respect to the spin axis failed (Mestel and Pryce 1992) and nowadays numerical simulations seem to be the best way to solve the problem, but anyway the polar-cap model is still useful to describe some basic concepts related to pulsar observations, like the *magnetosphere*.

According to that model, since the high dipole magnetic field of the NS rotates with the star, changing rapidly at any point in space, the few charged particles inside the NS experience strong Lorentz's forces, and a high induced

electric field is yielded; owing to the fact that the NS electrical conductivity is very high ($\sigma \to \infty$), the particles assume, in the shortest possible time, a configuration to perfectly compensate such forces with electrical forces (and hence balance the induced field with a static field $\vec{E}$ inside the NS):

$$q\vec{E} = -\frac{q}{c}\vec{v} \wedge \vec{B}, \tag{1.8}$$

(in cgs units), from which:

$$\vec{E} + \frac{1}{c}\left(\vec{\Omega} \wedge \vec{r}\right) \wedge \vec{B} = 0, \tag{1.9}$$

where q is the particle charge, $\vec{r} = (r, \theta)$ is the positional vector of the particles in polar coordinates with respect to the centre of the NS, $\vec{v} = \vec{\Omega} \wedge \vec{r}$ is the particle velocity, and $\vec{B}$ is the NS magnetic field. The charge density ρ_i of the particles after the redistribution satisfies Maxwell's equation, so that:

$$\begin{aligned} \rho_i(r,\theta) &= \frac{1}{4\pi}\nabla \cdot \vec{E} = \frac{1}{4\pi}\nabla \cdot \left[-\frac{1}{c}\left(\vec{\Omega} \wedge \vec{r}\right) \wedge \vec{B}\right] \\ &= -\frac{1}{4\pi c}\left[\left(\vec{\Omega} \cdot \vec{B}\right)\nabla \cdot \vec{r} - \left(\vec{r} \cdot \vec{B}\right)\nabla \cdot \vec{\Omega}\right], \end{aligned} \tag{1.10}$$

and, since $\nabla \cdot \vec{r} = (\partial r/\partial r) + (\partial\theta/\partial\theta) = 2$ and $\vec{\Omega}$ is constant in space:

$$\rho_i(r,\theta) = -\frac{1}{2\pi c}\vec{\Omega} \cdot \vec{B}. \tag{1.11}$$

From this equation we can infer two things: first of all, the absolute value of the charge density is higher in the regions where $\vec{B}$ is parallel to the spin axis (poles and equator, in this case); secondly, on the NS surface there is a charge separation: for example, if at the poles $\vec{\Omega} \cdot \vec{B} > 0$, i.e. $\vec{B}$ and $\vec{\Omega}$ point in the same direction, then the charge density is negative in the poles, while it is positive in the equator since $\vec{B}$ has opposite direction; vice versa if at the poles $\vec{\Omega} \cdot \vec{B} < 0$.

If outside the NS there is a vacuum, this surface charge separation induce a strong external electric field (which can be found obtaining the surface electric field from eq. (1.9) and then the surface electric potential, and using the latter as a boundary condition to solve Laplace's equation in the vacuum), which is a quadrupole field. Evaluating its value on the surface of the NS, and considering the component parallel to $\vec{B}$, we find:

$$E_{\parallel} = \left(\frac{\vec{E} \cdot \vec{B}}{B} \right)_{r=R} = -\frac{\Omega R B_S}{c} \cos^3\theta, \tag{1.12}$$

so that, except from the equator where $\theta = 90°$ and $E_{\parallel} = 0$, it always exists a component of $\vec{E}$ parallel to $\vec{B}$, that has a maximum at the poles and decreases towards the equator, where $\vec{E}$ is perpendicular to $\vec{B}$. The value of such a component at the poles is so high that the corresponding electric force ($F = qE_{\parallel}$) acting on the surface charged particles exceeds gravity force by several orders of magnitude so that, provided that the charges are not heavily bound, they can be extracted from the surface. The particles are then channelised along the magnetic field lines and create a plasma which surrounds the NS, forming the so-called *pulsar magnetosphere*. The charge distribution ρ_e of the particles in the magnetosphere has the same expression as the one inside the NS (eq. (1.11)), since the plasma outside the star experiences the same induced electric field as the pulsar interior due to the rotating magnetic field, and hence the charges arrange to balance it.

Since the magnetic field lines rotate with the NS, also the particles in the magnetosphere are forced to co-rotate with it, but that is possible only up to a distance where the tangential velocity v_t along the field lines is equal to c; for longer distances in fact such a velocity should exceed the speed of light in order to maintain the magnetic field and the plasma in co-rotation with the NS. Consequently the field lines, that are closed up to that maximum distance, are forced to open beyond it; we can indicate such a distance as the radius R_{LC} of the so-called *light cylinder* (represented in fig.1.1), a surface that contains all the closed field lines (the last one, for which $v_t = c$, is tangential to it) and outside which the co-rotation is not possible anymore. So we have:

$$\Omega R_{LC} = c \quad \Longrightarrow \quad R_{LC} = \frac{c}{\Omega} = \frac{cP}{2\pi} \approx 5 \times 10^4 \left(\frac{P}{\mathrm{s}} \right) \quad \mathrm{km}. \tag{1.13}$$

The last closed field line define the *polar cap* on the NS surface, that is the region with boundary (R, θ_p), centred on the magnetic pole, that contains only the open field lines; θ_p is hence a critical angle, since for $\theta < \theta_p$ the charges flow along the open field lines giving rise to the radio emission, while for $\theta > \theta_p$ the charges are forced along the closed lines to co-rotate with the NS and will

not leave the magnetosphere ever. From the equation that describe a field line: $\sin^2\theta/r = \text{const}$, we can obtain an estimate of θ_p:

$$\frac{\sin^2\theta_p}{R} = \frac{\sin^2 90^\circ}{R_{LC}}, \tag{1.14}$$

from which, for the canonical value of $R = 10$ km:

$$\sin\theta_p = \sqrt{\frac{R}{R_{LC}}} = \sqrt{\frac{R\Omega}{c}} = \sqrt{\frac{2\pi R}{cP}} \approx 1.4 \times 10^{-2} \left(\frac{P}{\text{s}}\right)^{-1/2}. \tag{1.15}$$

The radius of the polar cap on the NS surface, if it is small, is given by:

$$r_p \simeq R\sin\theta_p = 140 \left(\frac{P}{\text{s}}\right)^{-1/2} \quad \text{m}. \tag{1.16}$$

Electrons in the polar cap are quickly accelerated to relativistic speeds by the induced electric field; as they move along the open magnetic field lines, which are curved, these particles experience a centripetal acceleration and hence emit *curvature radiation* in the form of high-energy (γ-ray band) photons. These energetic photons interact with the strong magnetic field or with lower-energy photons creating electron-positron pairs ($\gamma \rightarrow e^- + e^+$), which in turn are accelerated and emit more high-energy photons, but with lower energy. These photons follow the same destiny as the previous ones, so that a cascade process is yielded and pairs and photons with progressive lower energies are created, until emission in the radio band is produced.

Since the particles are relativistic, the radio emission is extremely collimated in the direction of their motion, i.e. inside the cone defined by the polar cap. The emission is hence strongly anisotropic and can be seen by an observer only if his line of sight crosses the direction of the emission cone.

The NS thus emits two narrow radio beams from its magnetic poles; if, as in a realistic case, the magnetic axis is inclined with respect to the spin axis, the NS acts as a lighthouse and, as mentioned in section 1.1.1, we observe a sequence of pulses corresponding to our line of sight cutting the beam. The NS has therefore become a *pulsar*.

The radio emission of a pulsar is *coherent* (see later in the text). In fact, the brightness temperatures for pulsar radio emission, i.e. the temperature that a black body should have to radiate the observed brightness, are very

high, typically between $10^{25} - 10^{35}$ K; we know that the youngest pulsars have surface temperatures of only $\sim 10^7$ K and central temperatures of only $\sim 10^9$ K (and older pulsars even lower), thus the radio emission can not be thermal. Furthermore, neither other incoherent processes like emission of curvature or synchrotron radiation can directly produce the pulsar radio emission, since the kinetic energy of the emitting particles must always be higher than or equal to the brightness temperature, but in that case the particles would have a very high energy and would emit in X-ray or γ-ray bands, not in the radio band.
The high brightness temperatures can be explained if we take into account *collective* processes of emission, i.e. if we suppose that, for some mechanism still unknown, bunches of N electrons in volumes whose dimensions are smaller than λ (the radiation wavelength) radiate as a unique charge Ne, where e is the charge of each electron. The generated electromagnetic fields are then *coherent* or *in phase*. Since, according to Larmor's formula, the power radiated by a charge q is proportional to q^2, it is clear that in this case the radiated power is N times the power radiated by N electrons that emit incoherently, accounting for the high brightness temperatures.

1.1.4 Emission features

As already said, pulsars are very weak radio sources. Consequently, in general we are not able to observe single pulses, unless the source is uncommonly strong; therefore we need to sum coherently (to *fold*) hundreds or even thousands of subsequent pulses, obtaining the so-called *integrated pulse profile* (see fig.1.2 for some examples) which has a high enough signal-to-noise ratio (S/N), i.e. can be well identified above the background noise. While the shapes of the single pulses can vary considerably from pulse to pulse, the integrated profile has the property of being very stable for a certain pulsar observed at the same radio frequency, i.e. it shows always the same shape. Nevertheless, from pulsar to pulsar such a profile can assume very different shapes, for example it can show only one peak and a gaussian form, or can be double-peaked, or have even more complicated shapes. This can be explained by means of two competing models: the first one is the *nested cone* model (Rankin 1993; Gil et al. 1993), according to which the radio beam has a multiple cone structure nested around a central core component; the second one is the *patchy beam* model (Lyne and

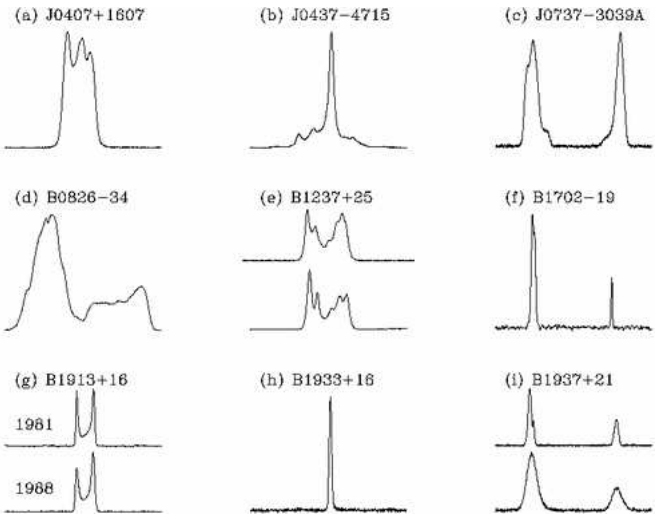

Figure 1.2: Examples of integrated pulse profiles. From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

Manchester 1988), according to which in the interior of the radio beam there are discrete emitting regions whose location are variable in time (likely according to some stochastic process) and whose integrated (in time) contribution produces the shape of the pulse profile. In both cases, the number of components in the pulse profile depends on the portion of the beam that is cut by our line of sight (see fig.1.3).

In some cases, we observe two pulses separated by about 180°, i.e. besides the main pulse also a secondary pulse called *interpulse* is visible. One hypothesis is that the latter is emitted by the opposite magnetic pole with respect to the pole emitting the main pulse, and the geometry of the system allow us to see both the poles (that is, the magnetic axis is perpendicular to the spin axis). Another possible explanation suggests that inter-pulses are the emission from extreme edges of a single wide beam.

As mentioned above, the integrated pulse profile is usually very stable, but it makes an exception for a number of pulsars, for which instead it can vary in time, probably due to a precessional effect that makes the radio beam to change its orientation with respect to our line of sight, or due to a transition between two competing profiles (*mode-changing* phenomenon).

Furthermore, the integrated profile varies with the observing frequency. In many *ordinary* pulsars (see section 1.1.7) the pulse width is larger and the

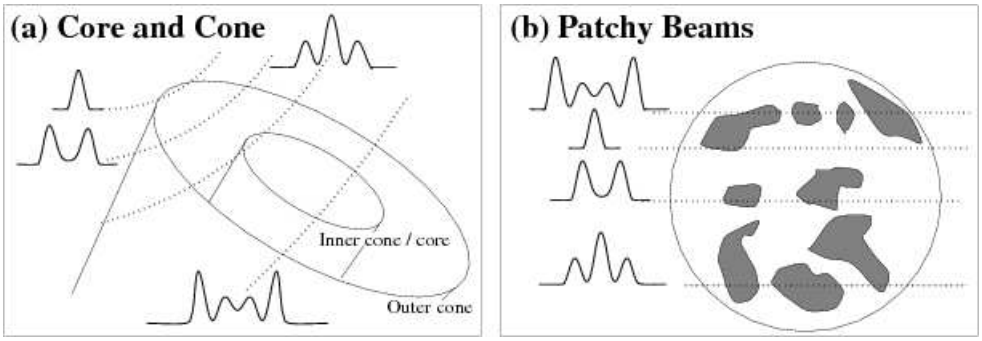

Figure 1.3: Beam models to explain the presence of different components in the pulse profiles: a) *nested cone* model (Rankin 1993; Gil et al. 1993); b) *patchy beam* model (Lyne and Manchester 1988). Figures from *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

profile components are more separated if we observe at low frequencies (these effects are instead very weak for *millisecond* pulsars, section 1.1.7 again). Also, for both the types of pulsars the number of profile components and/or their relative intensity can change with the observing frequency.

The *equivalent width* W_{eq} of the pulse is defined as the ratio between the pulse energy and the peak flux density (i.e. as the width of a top-hat pulse having the same area and peak flux as the true pulse profile), and for integrated profiles is often $\sim 10^\circ$ in *longitude* (or few hundredths in *rotational phase*)[4] ϕ, so that the pulse *duty cycle* $\delta = W_{eq}/P$, where P is the pulse period, is about 3%.

If the S/N is high enough it is possible to observe the individual pulses, that show a rich diversity of behaviour, i.e. a variety of structures and components; the statistical distribution of these components over a range of longitudes, together with the width and the probability distribution of the intensities of the pulses, determine the stability of the integrated profile. Generally, the individual pulses are formed by a number of characteristic components called *sub-pulses*, having a typical width in longitude of 1° to 3°, which are interpreted as radiation emitted by different regions of the polar cap within the distribution of locations covered by the integrated profile. A very common observed phenomenon is the so-called *pulse drifting*, i.e. the progressive change

[4]The longitude runs from 0° to 360° along a pulsar rotation, whereas the phase conventionally goes from 0 to 1.

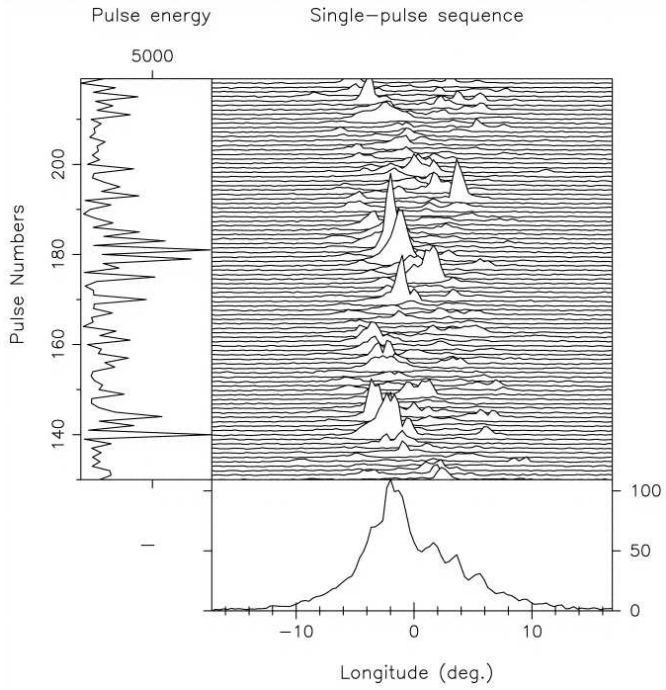

Figure 1.4: Individual pulses from the pulsar B0943+10, with their energy and, in the bottom panel, the integrated profile; it can be noted the drift of the sub-pulses (Deshpande and Rankin 1999).

in longitude of the sub-pulses in successive pulses (see fig.1.4), so that the sub-pulses cross the 'window' of the integrated profile; this is regarded as a lateral movement of an area of excitation across the polar cap, and the track of this movement can be possibly closed giving a pattern of excitation rotating round the polar cap (see fig.1.5).

Another structure of the individual pulses is the *microstructure*, which appears on a much shorter time scale (μs) with respect to the sub-pulses, is broad-band and often quasi-periodic; it appears to be a modulation of sub-pulse radiation rather than a distinct component of radiation, although it takes the form of intense micropulses.

Sometimes the emission from some pulsars switches off (i.e. it decreases to a level well below one per cent of the mean pulsar power) for many pulse periods, a phenomenon called *nulling*. A possible interpretation correlates the nulling with the pulsar age, so that it should happen mainly in old pulsars, which should hence approach the end of their lives and of their radio emission

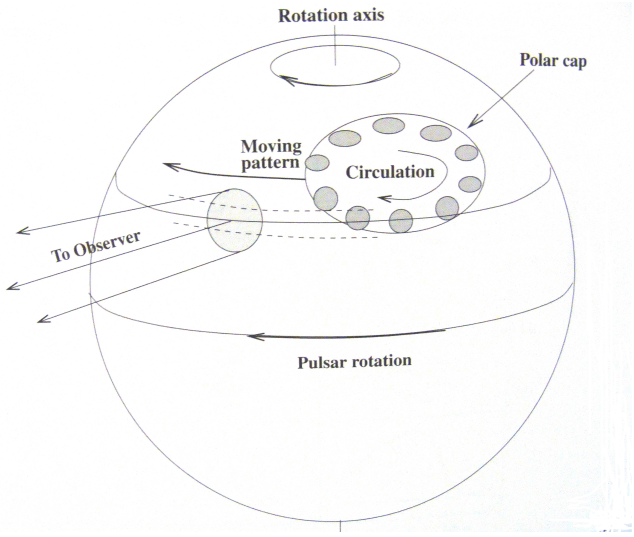

Figure 1.5: Circulation of regions that gives rise to the drifting of the sub-pulses. Figure from *An introduction to Radio Astronomy* by Burke and Smith 2002.

(see section 1.1.7) through progressive larger stages of nulling. Anyway, from observations a clear nulling-age effect is not obvious yet.

In some pulsars, for instance in the Crab pulsar, short bursts of emission can be observed, called *giant pulses*, which have an intensity up to 1000 times that of an individual pulse; in the Crab their lasting is of the order of nanosecond. They could be connected to the high-energy emission of these pulsars.

The spectra of the pulsar radio emission follow a power law trend for most pulsars:

$$S_{mean}(\nu) \propto \nu^{-\alpha} \quad 0 \leq \alpha \leq 4, \tag{1.17}$$

where ν is the observing frequency, $S_{mean}(\nu)$ is the integrated flux density of the pulse profile averaged over the pulse period, and α is the *spectral index*. The spectra are hence quite steep, especially at higher frequencies. There are nevertheless some deviations from a single power-law behavior, for examples for some pulsars it is necessary a two-component power-law model, or there is a roll over in the spectra at low frequencies or a turn-up at millimeter wavelengths. The mean flux density of the current samples is about 0.8 mJy (1 Jy = 10^{-26} W m^{-2} Hz^{-1}) at a frequency of 1.4 GHz, with a range between 20 μJy and 5 Jy; hence, as already said, pulsars are extremely weak radio sources.

Radio emission from pulsars is strongly polarised; the average degree of linear polarisation is about 20% of the total intensity, while for circular polarisation the average degree is about 10%. The position angle of linear polarisation can vary considerably throughout the pulse or from pulse to pulse, so that even if there are strongly polarised individual pulses (that can reach even 100% of polarisation), the integrated profile can be poorly polarised. Usually the polarisation angle varies with longitude through the pulse, describing a S-shape curve.

1.1.5 The interstellar medium effects

The radio signals that we receive can be very different from those that left the pulsars, since to reach our telescopes these signals have to travel through the interstellar medium (ISM), that is partly made up of ionised particles. The signal propagation is then influenced by the interaction with the ionised component of the ISM, by means of three different effects: *dispersion*, *scattering* and *scintillation*. The dispersion can be explained by propagation through a homogeneous medium, while for scattering and scintillation it is necessary to consider propagation through a more realistic turbulent ISM.

Dispersion

As the pulses travel through the ISM, considered as a homogeneous medium, the time-varying electric field of the radio waves causes the electrons that are encountered along the way to vibrate; this process slows the radio waves below the speed of light in a vacuum, with a greater retardation at lower frequencies. So the group velocity of the wave propagating in the medium is:

$$v_g = c\mu = c\sqrt{1 - \left(\frac{\nu_p}{\nu}\right)^2} \quad < c, \tag{1.18}$$

where μ represents the frequency-dependent index of refraction of the wave in the medium, ν is the wave (observing) frequency and ν_p is the *plasma frequency*, i.e. the frequency below which a signal is totally absorbed by the medium (from last equation, if $\nu < \nu_p$ the wave can not propagate):

$$\nu_p = \sqrt{\frac{e^2 n_e}{\pi m_e}}, \tag{1.19}$$

where e, n_e and m_e are respectively the charge, number density and mass of an electron. For the ISM typically $n_e \sim 0.03$ electrons per cm^3, so the plasma frequency is $\nu_p \simeq 1.58$ kHz. Since pulsar radio signals have frequencies much higher than ν_p (and this also make them not to be absorbed by the ISM), we can approximate eq. (1.18) in this way:

$$\frac{1}{v_g} \simeq \frac{1}{c}\left[1 + \frac{1}{2}\left(\frac{\nu_p}{\nu}\right)^2\right]. \tag{1.20}$$

The signal propagating from the pulsar to the observer is delayed, with respect to a signal of infinite frequency (for which $v_g = c$), by a time:

$$t = \left(\int_0^d \frac{dl}{v_g}\right) - \frac{d}{c}, \tag{1.21}$$

where d is the lenght of the path; replacing eq. (1.20) we obtain:

$$t = \frac{e^2}{2\pi m_e c}\frac{\int_0^d n_e dl}{\nu^2} = D \times \frac{DM}{\nu^2} \simeq 4.15 \times 10^3 \frac{DM}{\nu_{MHz}^2} \quad \text{s}, \tag{1.22}$$

where ν_{MHz} is the frequency in MHz,

$$DM = \int_0^d n_e dl \tag{1.23}$$

is the *dispersion measure*, that is usually expressed in pc cm^{-3} and represents the column density of the ISM free electrons along the observer line of sight to the pulsar, and

$$D = \frac{e^2}{2\pi m_e c} \simeq 4.15 \times 10^3 \,\text{MHz}^2\text{pc}^{-1}\text{cm}^3\text{s} \tag{1.24}$$

is the *dispersion constant*.

Thus a sharp pulse emitted at the pulsar, with all its component frequencies peaking at the same time, is gradually drawn out or *dispersed* as it travels to Earth (see fig.1.6), since the lower frequencies are slower than the higher ones, and the delay depends on the DM; in fact, the delay between two frequencies ν_1 and ν_2, both in MHz, is:

$$\Delta t \simeq 4.15 \times 10^3 \left(\frac{1}{\nu_1^2} - \frac{1}{\nu_2^2}\right) DM \quad \text{s}. \tag{1.25}$$

For a known pulsar, by measuring the pulse arrival time at different frequencies it is then possible to calculate the DM value, and assuming a model

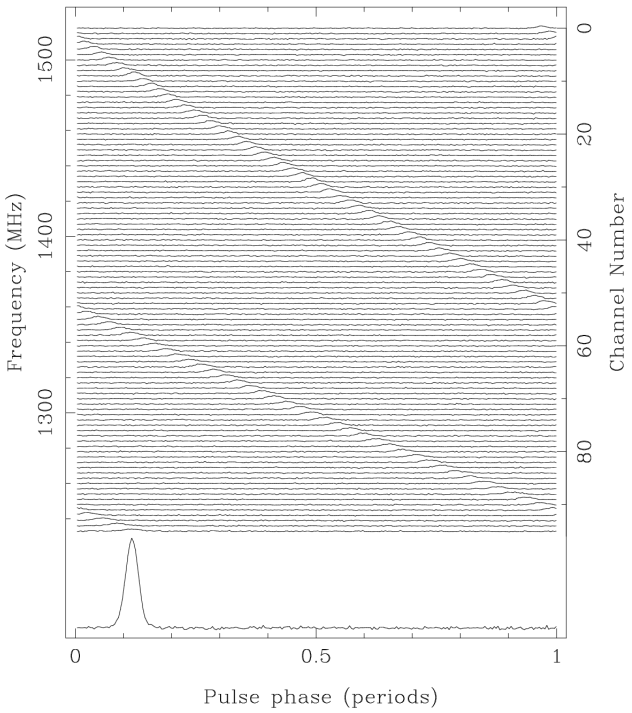

Figure 1.6: Example of the dispersion of a signal due to the ISM; the recovered profile is showed in the bottom. Figure provided by Andrew Lyne for the *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

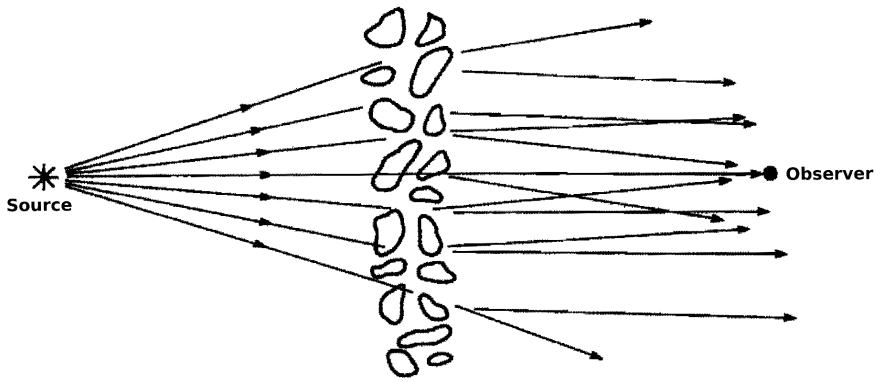

Figure 1.7: The thin screen model. From *Pulsar Astronomy, 3rd ed.* by Lyne and Smith 2005.

for the distribution n_e of ISM free electrons in the Galaxy (for instance, Cordes and Lazio 2002) from eq. (1.23) one can find an estimation of the distance of the pulsar.

Since the pulse profiles of the dispersed signals are totally smeared and can not be discerned above the background noise, we need to *dedisperse* in order to recover the original profile. This technique will be illustrated in chapter 2.

Scattering

As mentioned above, in reality the ISM is not homogeneous, i.e. the electron density shows changes in concentration that distort and scatter the pulse shape. This effect can be described in a simplified way by examining the *thin screen* model (Scheuer 1968), in which the random irregularities of n_e between the source and the observer are effectively concentrated into a thin screen roughly midway along the propagation path. Consequently, the photons emitted by the source that encounter those irregularities are scattered by them and arrive to the observer at different times (see fig.1.7); if Δt is the time delay of a scattered ray with respect to a ray that travelled undeflected, the observed intensity of the source is:

$$I(t) \propto e^{-\Delta t/\tau_s}, \tag{1.26}$$

where τ_s is the *scattering timescale*:

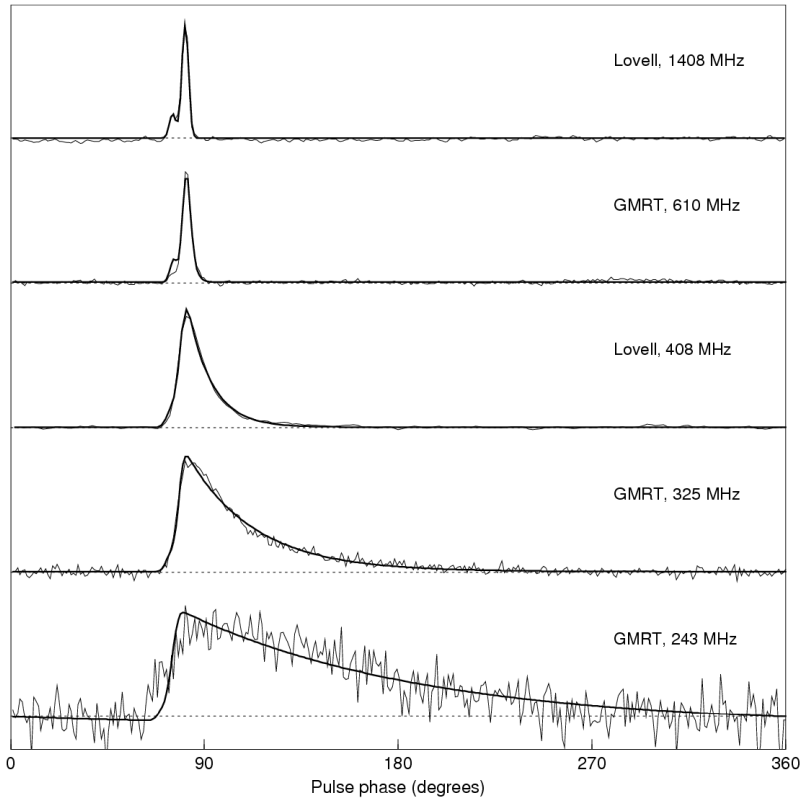

Figure 1.8: Broadening of the pulses due to scattering for the pulsar B1831−03, observed at different frequencies. Figure provided by Oliver Löhmer for the *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

$$\tau_s \propto \frac{d^2}{\nu^4}, \tag{1.27}$$

d is the distance between the pulsar and the observer, and ν is the wave frequency. Hence, eq. (1.26) explains the observed asymmetric time broadening of the pulses (*scattering tails*), which strongly depends on the frequency (see fig.1.8). That causes a stretching of the true pulse shape with a reduction in the S/N, especially for lower frequencies[5].

In order to remove the scattering effects, the only possibility is observing at high frequencies.

[5] Another consequence of scattering is an increasing in the apparent angular size of the source.

Scintillation

The interstellar scintillation is an effect similar to the optical 'twinkling' of stars due to the Earth atmosphere. In fact, it causes short-term variations in the observed intensity of many pulsars, and can be accounted for by using the same thin screen model for a non-homogeneous ISM described in the previous section.

Due to scattering, the deflected rays arrive at the observer randomly delayed with respect to the undeflected rays, hence showing different wave phases ϕ which, for signals received over the time τ_s, cover a range of phases $\delta\phi \sim 2\pi\nu\tau_s$ (ν and τ_s defined as in the previous section). If the phases of the waves (both scattered and unscattered) do not differ by more than about 1 radian, there can occur *interference* of these waves at the observation point, more in general a patchy interference pattern with maxima and minima of intensity is created at the *plane* of the observer. Due to relative motions between the pulsar, the ISM and the observer this plane 'moves', so that in the observation point we experience a change in the intensity of the signals, that can therefore enhance or reduce.

Since the phase depends on frequency, the limitation in the phase difference to have interference implies a limitation in the range of frequencies of the interfering waves, i.e. only the waves with frequencies inside the so-called *scintillation bandwidth* $\Delta\nu$ will contribute:

$$2\pi\Delta\nu\tau_s \sim 1 \quad \Longrightarrow \quad \Delta\nu \propto \frac{1}{\tau_s} \propto \nu^4. \tag{1.28}$$

Therefore we see an intensity changing both in time and in frequency, and the scintillation bandwidth is larger if we observe at high frequencies.

Scintillation can be very useful, for example in the searching for new pulsars, since sometimes we can luckily observe an otherwise undetectable (because too faint) pulsar whose intensity is enhanced by scintillation, which makes the S/N rise and hence allows us to discover the pulsar.

1.1.6 Age

We can calculate an estimate of the pulsar age from the relation of the pulsar slowdown in the rotating magnetic dipole model, eq. (1.5): $\dot{\Omega} = d\Omega/dt = -k\Omega^n$,

by integrating it:

$$\int_0^t dt = -\frac{1}{k}\int_{\Omega_i}^{\Omega} \Omega^{-n} d\Omega,$$

where Ω_i is the initial (i.e. at birth) angular velocity of the pulsar, and Ω is its angular velocity at the present time t; solving the integral we find:

$$t = -\frac{\Omega}{(n-1)\,\dot{\Omega}}\left[1 - \frac{\Omega^{n-1}}{\Omega_i^{n-1}}\right]. \tag{1.29}$$

Assuming that $n \neq 1$ and that the angular velocity at birth was much higher than the present value ($\Omega_i \gg \Omega$, according to the core-collapse theories of massive stars that suggest pulsars are born with very small spin periods, of the order of ms), we can approximate in this way:

$$\tau = -\frac{1}{n-1}\,\frac{\Omega}{\dot{\Omega}} = \frac{1}{n-1}\,\frac{P}{\dot{P}}, \tag{1.30}$$

where P is the present value of the pulsar spin period. The time τ is called *characteristic age*, and for purely magnetic dipole braking ($n = 3$) its expression is:

$$\tau = \frac{P}{2\dot{P}}. \tag{1.31}$$

Hence, from a measurement of its P and $\dot{P}$ we can infer the characteristic age of a pulsar.

Nevertheless, this estimate should be considered with some care. In fact, comparing the value of the characteristic age with the value of the age obtained in other ways (such as for example from the association of a pulsar with a supernova of known age, or from the proper motion[6] of the pulsar, see section 1.2.1), the correspondence is good for some pulsars, like the Crab, but is very bad for others. This can be due to the assumptions we made of a very high initial angular velocity and that $n = 3$ is the opportune value of the breaking index[7].

[6] The proper motion of a celestial object, usually expressed in arcseconds or milli-arcseconds per year, is a slow, angular change in the object equatorial coordinates due to its intrinsic transverse velocity along the celestial sphere.

[7] In fact, recent estimates suggest a wide range of initial spin periods, up to 140 ms; and in some pulsars there are significant deviations from the value of 3 for the braking index, which may indicate that part of the torque on the pulsar that causes its spin-down is due to outflow of particles.

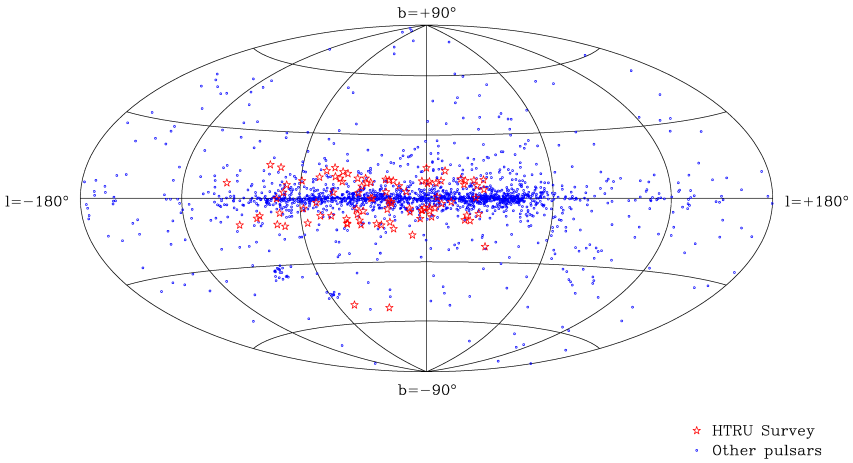

Figure 1.9: Distribution of pulsars on a Hammer-Aitoff projection in Galactic coordinates. The red stars are those discovered so far by the High Time Resolution Universe (HTRU) survey.

1.1.7 Population

As already mentioned, about 2000 pulsars are known so far. Their distribution (fig.1.9) shows that most of them are concentrated on our Galaxy plane, that is in agreement with the hypothesis that pulsars are the final stage of massive (O and B) stars, which lie on the Galactic disk too.

While the massive stars show a radial distribution around the Galactic centre, the known pulsars, from their projection onto the Galactic plane (fig.1.10), seem to be clustered around the Sun, but this is just an observational bias, a selection effect due to the fact that pulsars are weak sources and can not be detected if they are too far away from us; moreover, also the propagation effects due to the ISM, described in section 1.1.5, that distort the pulses, contribute to the selection effect. Therefore, it is also expected that in the Galaxy there are many more active pulsars with respect to the known ones, maybe about 10^5 in total (Lorimer et al. 2006).

If we consider the height above and below the Galactic plane, z, the distribution $N(|z|)$ of the known pulsars is approximately exponential:

$$N(|z|) = N_0 e^{-|z|/h}, \tag{1.32}$$

where $h \sim 300 - 350$ pc (e.g. Mdzinarishvili and Melikidze 2004) is the scale

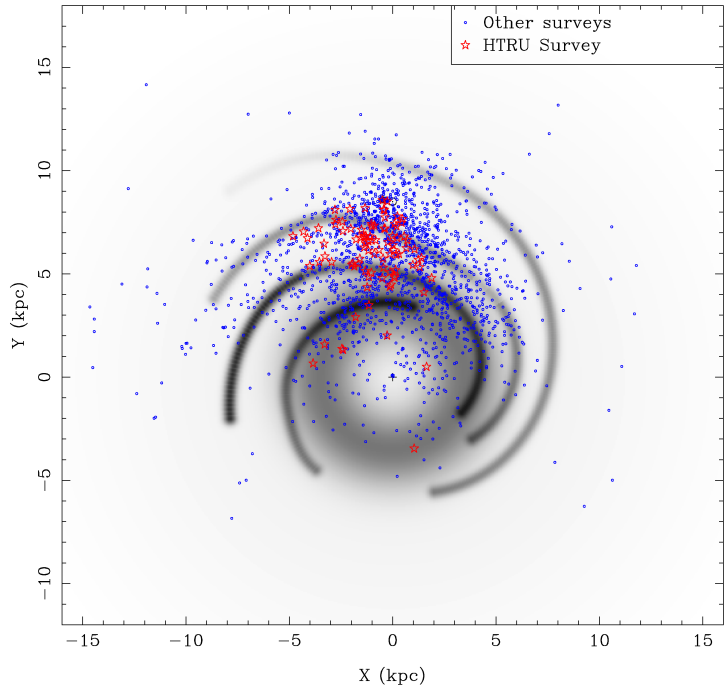

Figure 1.10: The observed pulsar distribution (circles and stars) and model electron density distribution (grey scale) projected onto the Galactic plane. In these coordinates, the Sun is at (0.0, 8.5) kpc and the Galactic centre is at the origin. The red stars are the pulsars discovered so far by the High Time Resolution Universe (HTRU) survey. The electron density distribution is the NE2001 model by Cordes and Lazio 2002. Darker areas correspond to regions of enhanced electron density.

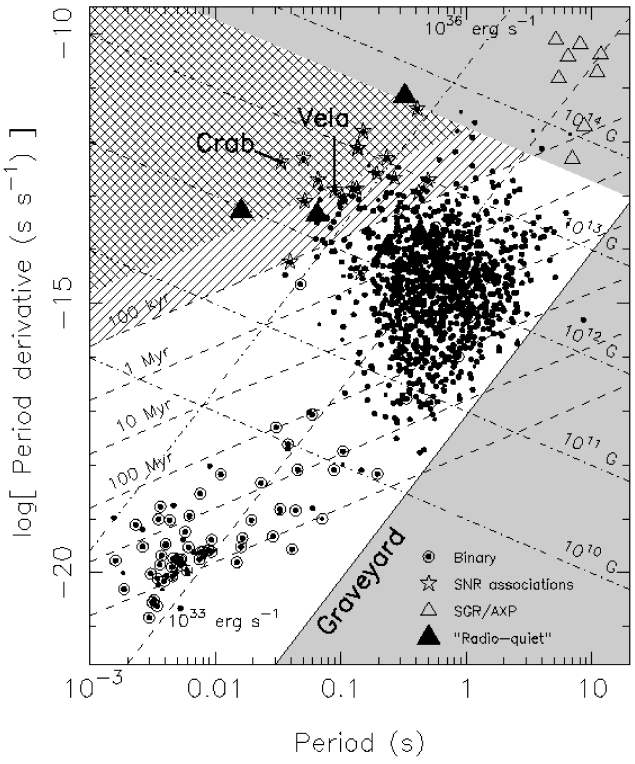

Figure 1.11: The P-$\dot{P}$ diagram, which represents the evolution of the spin during the pulsar life. From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

height corresponding to a decrease of $1/e$ in the number of pulsars with respect to the value N_0 on the plane. This value of h is much higher than the value for the O-B stars ($h \sim 80$ pc) and for the supernova remnants ($h \leq 100$ pc), but that can be explained by the fact that if pulsars, as we believe, were born in supernova explosions (so very close to the plane), during these they probably received a 'kick' of several hundred km s^{-1} (maybe due to small asymmetries in the explosions) that made them start to go away from the plane. This is consistent with the high velocity of pulsars, that can reach and even exceed[8] 1000 km s^{-1}, calculated by measuring their proper motions (see section 1.2.1).

Pulsar evolution and recycling model

The P-$\dot{P}$ diagram in fig.1.11 helps describing the evolution of the pulsar spin with time. It is evident that two quite different populations of pulsars exist:

[8]The millisecond pulsars, described in section 1.1.7, are an exception since they have on average lower velocities, about 100 km s^{-1}.

the *ordinary* pulsars, with periods ranging from some hundredths of a second to few seconds and with $\dot{P} \sim 10^{-17} - 10^{-11}$ s s^{-1}, and the *millisecond* pulsars (MSPs), with spin periods of the order of millisecond and $\dot{P} \sim 10^{-18} - 10^{-21}$ s s^{-1}.

Since for the magnetic dipole model the characteristic age $\tau \propto P/\dot{P}$ (section 1.1.6) and the magnetic field $B \propto \sqrt{P\dot{P}}$ (section 1.1.2), we obtain that the ordinary pulsars are younger and have very high magnetic field strenghts (typically $\tau \sim 10^7$ yr and $B \sim 10^{12}$ G), while millisecond pulsars are older and have lower magnetic field (typically $\tau \sim 10^9$ yr and $B \sim 10^8$ G). Those facts - as well as the much more frequent occurrence of binary systems in the population of millisecond pulsars with respect to that of the ordinary pulsars - suggest a possible evolutionary scenario for the two families of radio pulsars.

According to this scenario, a pulsar is born with a very small spin period, of the order of tens or hundreds ms, and a high magnetic field, heuristically as a result of the conservation of angular momentum and magnetic flux during the collapse of the core of the massive star. This implies that $\dot{P}$ is very high for young pulsars, so that they rapidly spin down and quickly ($10^5 - 10^6$ yr) reach the ordinary pulsar region in the P-$\dot{P}$ diagram. At some point the pulsars cross a region called *death valley* entering in the *graveyard*, i.e. the region in the diagram where the mechanism of radio emission 'switches off'.

This is the destiny of an *isolated* pulsar. Nevertheless, from the observations we know that a pulsar (ordinary or millisecond) can be part of a *binary system*, where the companion can be either a main sequence star, a white dwarf or another neutron star[9]. According to the so-called *recycling model* (Alpar et al. 1982), if an ordinary pulsar is in a binary its final destiny can be different from the fate of an isolated pulsar because of accretion processes from the companion, that can result in the production of a millisecond pulsar.

To explain what happens, we need to describe the evolution of a binary system since its formation (see fig.1.12). Initially the system is made up of two main sequence stars. The primary more massive star evolves first; if it satisfies the opportune conditions, at the end of its life it undergoes a supernova explosion and becomes a neutron star, with fast spin period and high magnetic

[9]The companion could be even a black hole, due to processes like, for example, the capture of an isolated pulsar by the black hole gravitational field. Anyway the pulsar-black hole system is still a 'Holy Grail' of astronomy since none of these systems has been found yet.

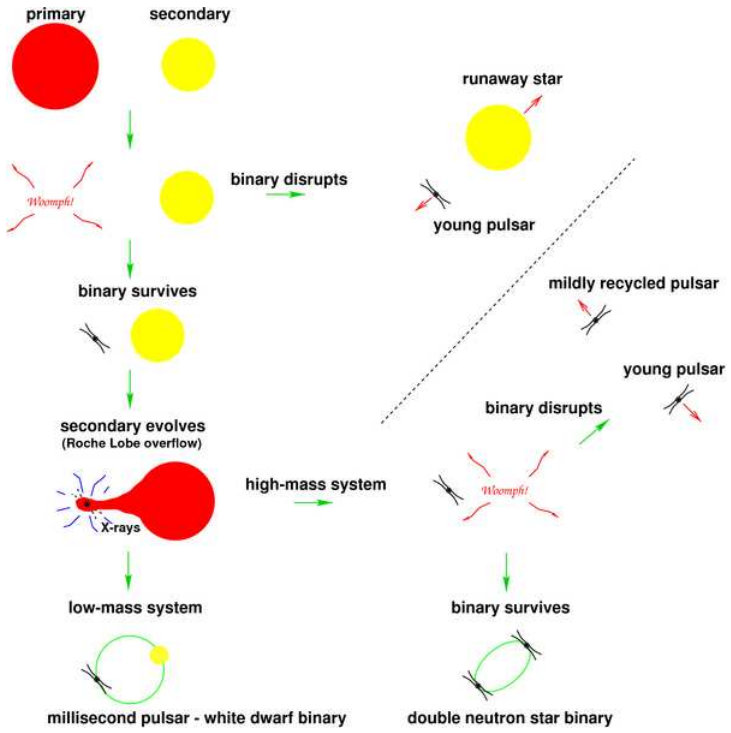

Figure 1.12: Possible evolutionary scenarii for a binary system. From Lorimer 2001.

field, which starts to emit as a radio pulsar and slows down, quickly approaching the ordinary pulsar region in the P-$\dot{P}$ diagram.

In a few cases the binary system can survive the supernova explosion, so that the pulsar and its companion remain bound (*ordinary pulsar-main sequence star system*). At some point during its evolution the companion reaches in turn the stage of red giant, filling or almost filling its Roche Lobe; at that point an accretion process can start, in which matter and angular momentum from the companion are transferred to the NS either by Roche Lobe Overflow, via an accretion disk, or via a plasma wind. At this stage the NS is visible as X-ray source, while its spin period decreases (due to the accreted angular momentum) so that the pulsar is 'spun up', and its magnetic field decays dramatically (maybe due to the accretion itself, see for example Jahan Miri and Bhattacharya 1994). It can be shown that higher the accretion rate $\dot{M}$ is, lower the achievable spin period is; since there is an upper limit to the accretion rate (the *Eddington limit*, $\dot{M}_{Edd} = 1.5 \cdot 10^{-8} R_6 \quad M_\odot \ yr^{-1}$, where $R_6 = R_{NS}/(10^6 cm)$ and R_{NS} is the NS radius), there is consequently a lower limit to the spin period that can

be achieved, given by $P_{min} \propto B^{6/7}\dot{M}^{-3/7}$.
If the companion mass (which the duration of the mass transfer phase depends on) is suitable, the pulsar can cross the death line again, moving toward the bottom-left part of the P-$\dot{P}$ diagram, and at the end of the accretion it turns on again as a radio pulsar. For this reason these pulsars are called *recycled pulsars*; from what said, they have very fast spin period and low magnetic field.

The evolution of the companion depends on its mass. If it is sufficiently massive, after the red giant phase (and hence the accretion phase, during which the system is called *High Mass X-ray Binary*, HMXB) the companion will undergo a supernova explosion in turn; if the binary survives again, at the end we will have a system made up of a young neutron star and a recycled pulsar (*double neutron star binary*)[10], where the latter has a spin period of $P \geq$ 20 ms and the orbit of the system is quite eccentric ($0.1 \leq e \leq 0.9$, where e is the orbit eccentricity). If the companion is not massive enough, it will not undergo a supernova explosion, and the accretion phase (during which the system is called *Low Mass X-ray Binary*, LMXB) will last much longer (up to 10^8 yr); consequently, the pulsar spin period can reach much lower values, of the order of few milliseconds, i.e. the LMXBs are considered to be the progenitors of the millisecond pulsars. The final system will be therefore typically made up of a millisecond pulsar and a white dwarf, with almost circular orbits ($10^{-5} \leq e \leq 10^{-1}$) (*millisecond pulsar-white dwarf binary*).

Nevertheless, from the observations we know that, although almost 80 per cent of the millisecond pulsars are in binary systems, some of them are isolated. The reason is still unknown, a possibility could be the ablation of the companion by the pulsar, but the time scales for ablation seem to be too long. The recent discovery of a system made up of a MSP and an ultra-low mass Carbon white dwarf (Bailes et al. 2011), where the latter has a planet mass and can be the core of a white dwarf that narrowly avoided complete destruction, could give some insights.

[10]The only one double-pulsar system known so far, J0737-3039 (Burgay et al. 2003, Lyne et al. 2004), is made up of a 22.7 ms pulsar and a 2.8 s long-period pulsar, in excellent agreement with this scenario.

1.2 Pulsar timing

When a new pulsar is discovered, it is submitted to a procedure called *timing* in order to obtain the values of its spin and astrometric parameters, and to get also other information about the binary motion (if the pulsar is part of a binary system) and the propagation of the pulses throughout the ISM.

Timing a pulsar consists in measuring and 'phase-connecting' (see later in the text) the Times Of Arrival (TOAs) of the pulsar radio pulses to the telescope. A TOA is defined as the time of arrival of a *fiducial point* on the integrated profile of an observation (remembering that this profile is very stable in shape). This TOA is obtained by a cross-correlation between the integrated profile typically obtained folding few hundreds or thousands pulses and a 'template' with high S/N (called *standard profile*, and obtained summing in phase many earlier observations at the particular observing frequency). The fraction of period at which we have the best cross-correlation between the two profiles is added to the starting or to the middle time of the observation, giving the time of arrival of the fiducial point, i.e. the TOA.

After measuring, we compare these TOAs with the ones predicted by a given pulsar model, through a multi-parametric fit where we minimize the differences between the measured and the predicted TOAs (called *timing residuals*); in this way we obtain the best fit values for the searched parameters. On a operational ground, if we consider a model in which the pulsar has spin frequency ν in a reference frame comoving with the pulsar (described in the next section), and expand the pulsar spin evolution in a Taylor series around a reference epoch t_0, we obtain at a time t in this frame :

$$\nu(t) = \nu_0 + \dot{\nu}_0(t - t_0) + \frac{1}{2}\ddot{\nu}_0(t - t_0)^2 + ..., \tag{1.33}$$

where $\nu_0, \dot{\nu}_0, ...$ are the values calculated in t_0.

Whence, remembering that the rotational phase $\phi(t)$ results from the integral of the spin frequency in time

$$\phi(t) = \phi_0 + \nu_0(t - t_0) + \frac{1}{2}\dot{\nu}_0(t - t_0)^2 + \frac{1}{6}\ddot{\nu}_0(t - t_0)^3 + ..., \tag{1.34}$$

where ϕ_0 is a reference phase calculated in t_0 (and usually assumed to be 0).

If we know the evolution of $\nu(t)$ accurately enough, and suppose that both t_0 and t correspond to the arrival of a pulse, our expansion for $\phi(t)$ should give an integer number n at t, i.e. we would be able to predict the phase of each pulse that will arrive after t_0 (*phase connection*). In other words we can unambiguously account for every single rotation of the pulsar even over long periods (years or decades) of time.

In order to obtain the best values for the rotational parameters in the simple model above, one typically starts with a first guessed set of parameters and then iteratively proceeds by minimizing the expression:

$$\chi^2 = \sum_i \left(\frac{\phi(t_i) - n_i}{\sigma_i} \right)^2 , \tag{1.35}$$

where n_i is the nearest integer to the phases $\phi(t_i)$ associated to the observed TOAs t_i in the assumed model for the given parameters, and σ_i is the TOA uncertainty.

Once the fit is good over a long enough dataspan (see section 1.2.1), one is confident to have obtained significant values for the model parameters.

The procedure above would hold for TOAs observed in a system comoving with the pulsar. However, that is never the case and the real observational situation is that we are observing from a rotating Earth which orbits the Sun. Moreover the space between us and the source is not empty and the source itself is moving with respect to us. Therefore, one needs to correct the observed TOAs before applying to them the simple model above.

The corrections to be applied to the so-called *topocentric* TOAs (i.e. the observed TOAs) will be described in the next two sections, taking into account separately isolated and binary pulsars.

1.2.1 Isolated pulsars

As mentioned in the previous section, if we want to perform the timing of an isolated pulsar we must make some corrections to the measured TOAs.

First of all, we must consider that our reference frame is not inertial, so that the pulses arriving at our telescopes are biased by the motion of the Earth both around its axis and around the Sun. Not correcting for these motions, we would see a sinusoidal variation of the observed arrival times during one year

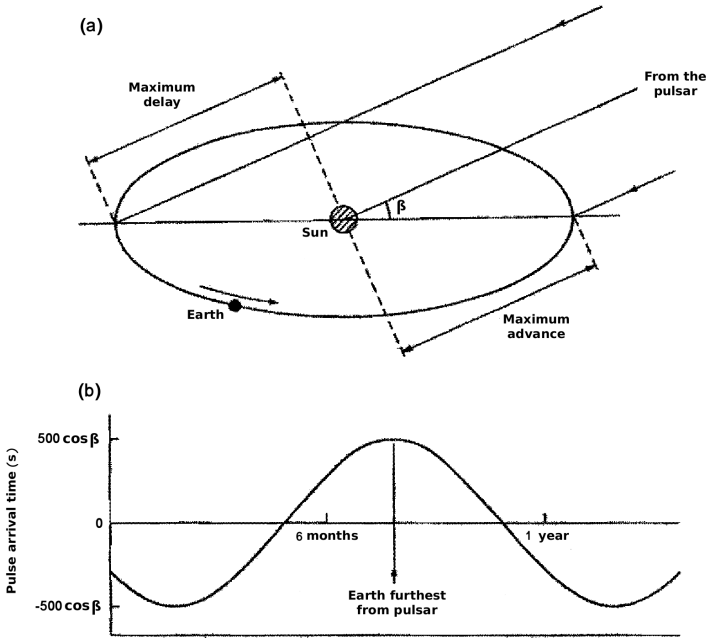

Figure 1.13: The Earth orbital motion causes a sinusoidal variation of the pulse TOAs in one year. From *Pulsar Astronomy, 3rd ed.* by Lyne and Smith 2005.

(see fig.1.13).

To remove this effect, we need to transform our *topocentric* TOAs to *barycentric* TOAs, i.e. calculate their values with respect to the barycenter of the Solar System (SSB), which is to a very good approximation an inertial reference frame.

Nevertheless, other effects besides the Earth motion must be taken into account in making this transformation, so that the expression for a barycentric TOA is:

$$t_{SSB} = t_{topo} + t_{clock} - \frac{\Delta D}{\nu^2} + \Delta_R + \Delta_S + \Delta_E, \tag{1.36}$$

where t_{topo} is the topocentric TOA, t_{clock} represents clock correction that accounts for differences between the observatory clocks and terrestrial time standards, $\Delta D/\nu^2$ accounts for the dispersion delay caused by the ISM (section 1.1.5: $\Delta D = D \times DM$, where D is the dispersion constant, eq. (1.24), and DM the dispersion measure, eq. (1.23));

- Δ_R is the *Römer delay*,

$$\Delta_R = -\frac{\vec{r}\cdot\vec{s}}{c} + \frac{(\vec{r}\cdot\vec{s})^2 - |\vec{r}|^2}{2cd}, \tag{1.37}$$

where $\vec{r}$ is the vector connecting the SSB to the Earth and $\vec{s}$ is a unit vector pointing from the SSB to the pulsar, positioned at a distance d.
The first term in this equation is the light travel time between the Earth and the SSB, that undergoes the annual sinusoidal variation showed in image (b) in fig.1.13) due to the Earth motion around the Sun; its maximum value is about $500\cos\beta$ s, where β is the ecliptic latitude. It's worth remarking that the occurrence of this annual modulation opens the possibility of deriving very accurate astrometric positioning of a pulsar in celestial coordinates.
The second term in the equation is the so-called *timing parallax*, another annual effect that causes a variation in the pulse arrival time having an amplitude of $l^2\cos\beta/2cd$, where l is the distance between the Earth and the Sun. This parallax is a measurement of the curvature of the wavefronts in different positions of the Earth's orbit around the Sun. Unfortunately this effect is very small and is measurable only for a few nearby pulsars, providing, in this case, a very good trigonometric measurement of the pulsar distance.

- Δ_S is the *Shapiro delay* (Shapiro 1964),

$$\Delta_S = -\frac{2GM_\odot}{c^3}\ln(1+\cos\theta), \tag{1.38}$$

where G is Newton's gravitational constant, $M_\odot$ is the mass of the Sun, and θ is the angle subtended from the pulsar to the Earth and the Sun at a certain epoch.
This is a relativistic (General Relativity) effect due to the fact that the pulses are delayed by travelling through the gravitational field of the solar system, i.e. due to the space-time curvature created by the presence of masses (essentially the Sun mass) in it.

- Δ_E is the *Einstein delay*, whose derivative is given by:

$$\frac{d\Delta_E}{dt} = \sum_i \frac{GM_i}{c^2 r_i} + \frac{v_\oplus^2}{2c^2}, \tag{1.39}$$

where M_i are the masses of all bodies in the solar system except for the Earth, r_i is the distance of the Earth from the mass M_i, and $v_\oplus$ is the Earth velocity with respect to the SSB.
This is a relativistic (Special Relativity) effect due to a combination of time dilation due to the Earth motion and gravitational redshift caused by the other bodies in the solar system.

By applying the formula 1.36, we obtain the barycentric TOAs.
If the pulsar is moving with respect to the SSB, we must consider also other effects due to the transverse component v_t of the pulsar velocity, related to its proper motion μ, according to:

$$v_t = 4.74 \frac{\mu}{\text{mas yr}^{-1}} \frac{d}{\text{kpc}} \quad \text{km s}^{-1}, \tag{1.40}$$

where $\mu = \sqrt{\mu_\alpha^2 + \mu_\delta^2}$, $\mu_\alpha = \dot{\alpha}\cos\delta$, $\mu_\delta = \dot{\delta}$, where μ_α and μ_δ are the proper motions in right ascension, α, and declination, δ, respectively.

Although it does not directly enter the timing formula (i.e. the following corrections are usually applied after having obtained a good set of timing parameters) it is important to remember another effect of the transverse motion of a pulsar which impacts on the determination of its intrinsic parameters, i.e. the so-called *Shklovskii effect* (Shklovskii 1970) or 'secular acceleration'. It affects all the time derivatives of a given quantity, e.g. the derivative of the spin period, $\dot{P}$. Since this contribution is inversely proportional to d it is usually very small and can be neglected, but it must be accounted for the case of recycled pulsars (having small values of $\dot{P}$) at close distances.

Once we have corrected our measured TOAs, keeping into account all the mentioned effects, we can fit the TOAs to our model (eq. (1.35)) using the already mentioned least-square fit procedure.
If the fit is good, the residuals are randomly distributed around zero, i.e. there are no unmodeled effects (case (a) in fig.1.14). On the contrary, if trends are presents in the residuals they are an indication for a wrong determination of some parameter of the model. For example, if $\dot{P}$ is underestimated, the residuals show a parabolic trend (case (b) in the figure), while if there is an error in the

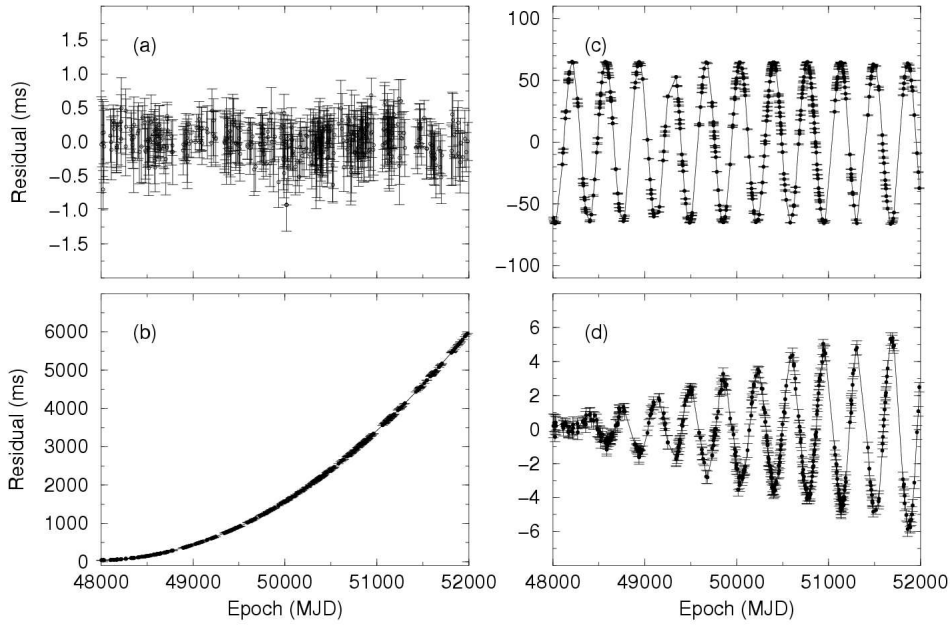

Figure 1.14: Examples of timing residuals. See the explanation of the different cases in the text. From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

pulsar position (i.e. in the $\vec{s}$ vector), the residuals are sinusoidal with a period of 1 year (case (c)). In case (d) the proper motion has been neglected, and the residuals show an annual sinusoidal trend with an increasing amplitude.

1.2.2 Binary pulsars

If the pulsar is part of a binary system, an observer also see the occurrence of a periodic modulation in the TOAs as the pulsar moves along the orbit, since they anticipate when the pulsar is approaching us and they delay when it is going away from us (orbital Doppler effect).
Therefore we need to add new terms in the expression (eq. (1.36)) in order to take into account the motion of the pulsar around the centre of mass of the binary system:

$$\begin{aligned} t_{SSB} &= t_{topo} + t_{clock} - \frac{\Delta D}{\nu^2} + \Delta_R + \Delta_S + \Delta_E \\ &+ \Delta_{RB} + \Delta_{SB} + \Delta_{EB} + \Delta_{AB}, \end{aligned} \tag{1.41}$$

where Δ_{RB} is a further Römer delay caused by the orbital motion of the pulsar, Δ_{SB} and Δ_{EB} are Shapiro and Einstein delays due to the companion gravitational field, and finally Δ_{AB} is a delay caused by the aberration of the electromagnetic waves due to the orbital motion.

If the binary is *non-relativistic* the pulsar motion follows Kepler's laws, so that we can describe it by means of seven independent *Keplerian parameters*: the orbital period, P_b; the semi-major axis a of the pulsar orbit around the common centre of mass; the orbital inclination i, i.e. the angle between the orbital plane and the plane of the sky; the orbit eccentricity, e; the periastron longitude, ω; the epoch of periastron passage, T_0; the position angle of the ascending node, Ω_{asc}.
Five parameters are usually obtained by fitting the TOAs (P_b, $x = a\ \sin i$, that is the projection of the semi-major orbital axis, e, w, T_0), whereas a sixth one, Ω_{asc}, is measurable only in particular cases.
From these parameters it is possible to derive the *mass function*:

$$f(M_p,\ M_c) = \frac{(M_c \sin i)^3}{(M_p + M_c)^2} = \frac{4\pi^2\,(a \sin i)^3}{G P_b^2}, \tag{1.42}$$

where M_p is the pulsar mass and M_c the companion mass. Owing to this relation we can put some costraints on the mass of the companion. In fact, assuming a typical value of $M_p = 1.4\ \mathrm{M}_\odot$ and an edge on orbit ($i = 90°$), from this equation we obtain the minimum possible value for M_c.

If the pulsar is part of a *relativistic binary*, i.e. it is in a close orbit with another massive compact object, like a white dwarf, another neutron star or even a black hole, the orbital motion can not be described simply by Kepler's laws, and other parameters called *post-Keplerian parameters* need to be added in our fit, in order to consider the relativistic effects due to the strong gravitational fields and the high orbital velocities. The most commonly measured parameters are the relativistic periastron advance $\dot{\omega}$, the parameter γ which takes into account time dilation and gravitational redshift, the orbital decay $\dot{P}_b$ which measures the orbital period decreasing rate due to the emission of gravitational waves, and the range r and shape s of the Shapiro delay caused by the space-time deformations around the companion star.

It is worth concluding this introduction to the timing of pulsar, noticing

that unfortunately not all the pulsars are equally good for applying to them the procedure above. In fact, in few cases the timing residuals show the signatures of un-modeled variations of the pulsar rotation with time, called *timing noise*. The exact origin of that is still debated. It could be due, for example, to changes in the interior of the NS related to the slow down and concerning for instance the occurrence of an independent motion of the crust and of the fluid interior. Alternatively it has been recently proposed that it can be due to random changes in the magnetic field or in the structure of the magnetosphere. As a matter of fact, timing noise is more prominent in the younger pulsars, while millisecond pulsars are, on the contrary, not heavily affected by it and hence are intrinsically better clocks (see next section).

1.2.3 A Pulsar Timing Array

Owing to the fact that the recycled pulsars (i) are rapidly spinning; (ii) are old neutron stars (where the internal structure of the star is usually not affected by the glitches shown in young neutron stars) and (iii) are not heavily affected by timing noise (section 1.2), they can be *exceptional clocks*, in some cases rivaling our best atomic clocks over time spans of months and years: in fact, currently the spin period of MSPs is known with a very high precision (the period of J1909-3744 is the one that so far has been measured with the best precision[11]: 0.0029471080681076401 $\pm$ 0.0000000000000000009 s, i.e. with 19 digits after the decimal point, Verbiest et al. 2009) and their rotational stability, measured by the σ_z parameter described in Matsakis et al. 1997, can be as good as 10^{-15} on a 10-year time span (Verbiest et al. 2009, Verbiest et al. 2008). Precise timing of these objects is then a powerful tool to investigate several celestial phenomena.

One of the possible physical uses of the recycled pulsars is in the framework of the so-called *Pulsar Timing Array* (PTA), for the detection of gravitational waves (GWs), whose existence has been verified so far only in an indirect way.

In fact, when a GW passes over our Galaxy pulsars and the Earth, it perturbs the space-time and hence modulates the signals from the pulsars with the variation in the TOAs of a strength which is proportional to the amplitude of the GW and a periodicity related to its frequency.

[11] *www.atnf.csiro.au/people/pulsar/psrcat*; Manchester et al. 2005

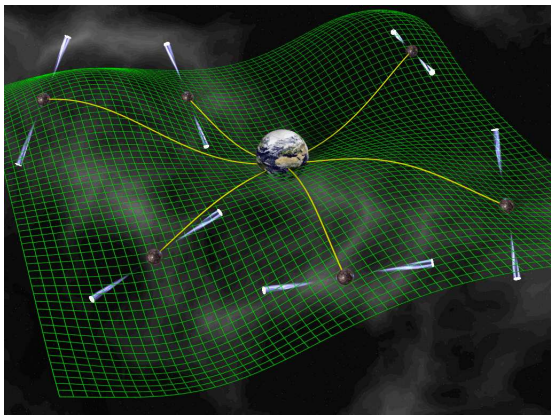

Figure 1.15: A Pulsar Timing Array to detect low-frequency GWs. Credit: NANOGrav.

The effect of course affects to TOAs of any pulsar, but it is expected to be tiny (even in the best cases) and therefore if one observes only one pulsar it turns out to be almost impossible to isolate this effect from many other possible causes for a periodic variation in the TOAs: for example low-level residual timing noise, fluctuations in the density of the interstellar medium along the line of sight, errors in the reference clocks at the observatory, errors in the planetary ephemeris which are necessary for performing the conversion from the topocentric to the barycentric TOAs, etc.

Most of these issues can be addressed and solved if one uses not a single but a 'network' of many suitable recycled pulsars placed in different directions of the sky (that is, a PTA). In fact, it has been theoretically shown that by means of simultaneous high precision timing of many pulsars belonging to a PTA one could really obtain a direct detection of GWs (see fig.1.15). In particular, given the typical interval between timing observations of pulsars (from months to decades, $\sim 10^7$ s to 10^9 s) a PTA is sensitive to GWs of frequency smaller than 10^{-7} Hz and higher than 10^{-9} Hz, with in practice a better sensitivity for the latter. This range of nanoHz frequencies is not covered by any other detector and thus PTAs nicely complements the other (mostly interferometric) detectors (either ground-based or in space) which have been (or will be) built in the aim of directly detecting GWs.

In particular the range of the nHz corresponds to the frequency range of the

predicted stochastic cosmological GW background resulting from the merging of supermassive BHs in the early phases of the Universe. So, detection and measurement of this background would result a breakthrough step not only for gravitational physics, but also for cosmological studies. The possibility of detecting gravitational signals from single monochromatic sources (like merging events in massive black-hole binaries, being relatively close to us and involving very high massive BHs of about 10^{-9} solar masses) standing above the root-mean-square (rms) value of the background has been recently investigated by Sesana and Vecchio 2010, but such signals have a much lower probability than the background of being detected by a PTA. However a problem arises concerning a clear detection of the GW background. In fact this has a dimensionless strain amplitude spectrum $h_c(f) = Af^{-2/3}$ (e.g. Cordes and Shannon 2011), where $h_c(f)$ is the characteristic amplitude of the GW signal as a function of the frequency f and we can use a fiducial value $A = 10^{-15}yr^{-2/3}$; the spectrum of timing residuals $\propto f^{-13/3}$ and the rms residual scales as $\sigma_{gw}(T) \propto T^{5/3}$, where T is the time span of the observations. Unfortunately red timing noise (see also the previous section), which is a low-frequency non-Gaussian noise present in the timing residuals (common in ordinary pulsars but also expected in MSP at low levels; see e.g. Kaspi et al. 1994 for the first detection of such a kind of noise in a MSP, B1937+21; Verbiest et al. 2008 for red timing noise in the closest and brightest MSP known, J0437−4715) due to irregularities in the pulsar rotation or to interstellar scintillation, yields a rms residual $\sigma_r(T) \propto T^{2\pm0.2}$ corresponding to a spectrum of timing residuals $\propto f^{-5\pm0.4}$. Hence the red noise spectrum is similar to the GW background spectrum, and the challenge will be to separate the red noise contribution from the GW background one, in order to safely detect the latter. According to Cordes and Shannon 2011, a minimum of 20 super-stable MSPs, each with rms timing residuals from red noise contribution less than 20 ns over time spans of 5 years and with negligible white noise (i.e. radiometer noise and pulsar jitter), should allow a plausible detection of the GW background. However, if such super-stable MSPs will not be found (currently we only know two MSPs, J1713+0747 and J1909−3744, having rms timing residuals less than 50 ns over a 5-year time span, Demorest et al. 2009, and one other less than 100 ns, J0437−4715, Manchester 2010) a much larger set of MSPs will be necessary,

maybe 100 or even more, having timing noise substantially less than 100 ns in a 5-year span.

Osłowski et al. 2011 have recently discussed the limits in the achievable precision of pulsar timing. They describe the intrinsic variability of the pulsar radio emission, which causes stochastic fluctuations in the TOAs, as *stochastic wideband impulse modulated self-noise* (SWIMS), a pulsar-intrinsic noise (also called in other ways, like jitter noise, intermittent emission or simply self-noise) likely due to the stochastic subpulse structure observed in single pulses, which yields pulse profile variations. They conclude that this self-noise may be a limiting factor for timing precision of every MSP currently observed by PTAs, even when larger and more sensitive antennae will be built. For example, they found that for MSP J0437−4715 a timing precision better than 30 - 40 ns in one hour at an observing wavelenght of 20 cm is highly unlikely, owing to the intrinsic variability of the pulsar signal and hence to SWIM. Also in this case it would be necessary to significantly increase the number of MSPs in the PTA.

From what said so far, it is clear that the discovery of more MSPs having the mentioned required features is of *fundamental importance* for the detection of the cosmological background of GWs.

As it can be inferred from what said above, the observational problem is to set up a suitable PTA, and in order to do that two things are necessary:

1. to find a large sample of recycled pulsars randomly distributed in space and having the required clock stability and signal intensity;

2. to perform regular high precision (100 ns) timing of the available sample.

In order to implement the second point, 3 large collaborations have been set up: the PPTA (Parkes Pulsar Timing Array), the EPTA (European Pulsar Timing Array) and NANOGrav, which are already exploiting the largest radiotelescopes in the world for observing a set of suitable targets.

The first point highlights the importance of the *pulsar surveys* (see section 1.3) since, as already said, we know only a handful of recycled pulsars having the suitable features to be part of a PTA and need to largely increase their number. This Thesis work focused on this task.

1.3 Why pulsar surveys?

A *survey* is the observation of a portion of the sky to blindly search for new objects, in this case pulsars.

The importance of finding new pulsars has been highlighted in the previous section for the case of the PTAs and the related scientific aim of detecting gravitational waves and performing investigations in cosmology.

However, recycled pulsars, owing to their often exceptional period stability and hence to the possibility of high precision timing (see the previous section), have already shown to be unique physical tools for a wide range of studies in fundamental physics and astrophysics, some of which are described in the following.

- Testing General Relativity: as said in section 1.2.2, in a relativistic binary there are significant deviations from Kepler's laws of motion; furthermore, the pulse path from the pulsar to Earth is modified due to the distortion of the space-time around the pulsar for the presence of the other compact object. Hence, by comparing the values of the system parameters obtained from the measured TOAs with the predictions of GR we can perform precision tests of this theory. The only known double-pulsar, J0737-3039 (Burgay et al. 2003, Lyne et al. 2004), provides so far the most precise test of GR in the strong field regime (see fig.1.16); discovering additional of these systems may allow us to carry out even more stringent tests of relativistic gravity.

- Studying super-dense matter: the behavior of super-dense matter, whose density can not be reached in terrestrial laboratories but only in the interior of a NS, can be studied through the measurement of the structural and/or rotational parameters of a NS. In fact, for a given mass, radius, internal density and structure of a NS depend on the Equations of State (EoS) for the nuclear matter at ultra high density. So far, this EoS is still unknown and about 30 models have been proposed for that to date, grouped in two main categories of EoS, the *soft* EoS and the *stiff* EoS. The *soft* EoS predicts a high compressibility of the super-dense matter, so that the NS would have a small radius, a high central density and a thin

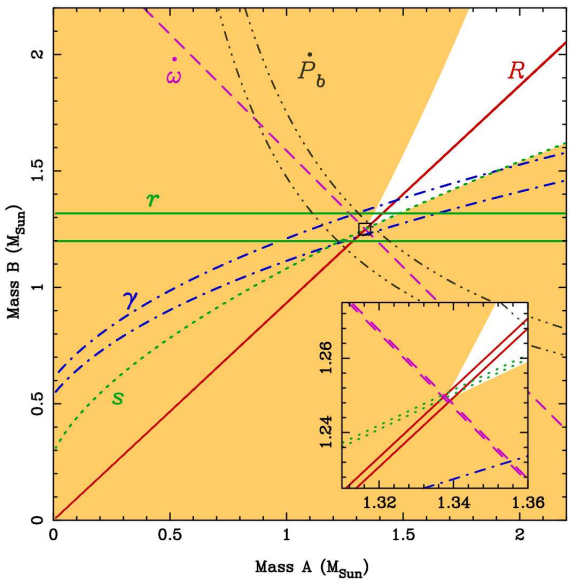

Figure 1.16: Kramer et al. 2006 measured all five post-Keplerian parameters and also the mass ratio R of the two pulsars, finding that GR is correct at the 0.05 per cent level.

crust; the *stiff* EoS instead predicts (for the same NS mass of the case of the soft EoS) a lower compressibility, and hence the NS would have a larger radius, a lower central density and a thicker crust. Moreover the stiff EoS predicts a larger maximum mass for a NS than the soft EoS. So, on one hand, from the measurement of the mass of the NS (as it is possible in the relativistic pulsar binaries) one can test (and possibly discriminated) between the EoSs. On another hand, since the minimum spin period reachable by a rotating NS is[12] $P_{min} \propto \sqrt{R^3/M}$, where R and M are respectively radius and mass of the NS, the two kinds of EoS predict different values of P_{min}. The discovery of a pulsar with $P \leq 1$ ms (*sub-millisecond pulsar*) would give credit to the *soft* EoSs, leading to the rejection of most of the stiff EoSs.

- Studying the plasma physics under extreme conditions: the plasma in the pulsar magnetosphere experiences super-strong magnetic fields; these

[12] P_{min} has been obtained equating the centrifugal force with the gravitational force at the equator at the breakup limit. This dependence of P_{min} from R and M is roughly reproduced also when accounting for the relativistic corrections to the neutron star structure.

extreme physical conditions can be probed if the pulsar is part of a binary system, as it can be done for example for the double-pulsar system J0737-3039, where as the pulsars move along their orbits the line of sight from pulsar A to us passes through the magnetosphere of pulsar B: any change in the transmission properties of the signals can give information about the plasma density and the structure of the magnetic field of pulsar B.

- Finding extra-solar planets: if a body having a planet mass is orbiting around a pulsar, we can observe variations in the TOAs that lead to reveal the presence of the planet and its mass. This pulsar-planet systems are not expected to be common, due to the violent conditions of NS formation, however the first planetary system outside the Solar System (and the only Earth-mass and Moon mass planet known until very recently) was discover to orbit a pulsar: in fact the 6.2 ms pulsar B1257+12 is orbited by three planets (one planet has a Lunar mass, while the other two are Earth-mass planets). Another system is a pulsar orbited at the same time by a planet and by a white dwarf: it is the 11 ms pulsar B1620-26 in the globular cluster M4 (the planet mass is in the range 1-2 Jupiter masses). Finally there is the recently found 'diamond planet' in orbit around the 5.7 ms pulsar J1719-1438 (the planet has a mass near that of Jupiter, is made up of Carbon and is incredibly dense, see Bailes et al. 2011; see section 1.1.7 for details), which have been discovered in the survey for pulsar which is in the focus of the work of this Thesis.

Besides the ones mentioned above, many other astrophysical fundamental issues can be addressed by means of the study of the young and ordinary pulsars, such as investigating the magnetic field of the Galaxy and the interstellar medium.

Chapter 2

Search methods

This chapter describes the different search methods that can be used to process data in order to find isolated and binary pulsars; in the former case, the so-called *standard search* is employed, while in the latter case, especially for close (short orbital period) binaries, several kinds of *acceleration search* can be used.

For more details about the content of this chapter see the *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

2.1 The sensitivity issue in the search for pulsars

In general, pulsars are very weak radio sources and the mean flux density decreases at higher frequencies according to a power law (eq. (1.17)); therefore, to detect them it is very important the *sensitivity* of the radiotelescope that we use. The sensitivity S_{min} is the minimum flux that the antenna is able to detect at a particular observing frequency ν and corresponding to a particular threshold $(\mathrm{S/N})_{min}$, i.e. the minimum S/N that we choose to consider (for example often the value of 8 is chosen, so that it is less probable that what we think is a signal is instead only a noise fluctuation). Hence a signal detected with that $(\mathrm{S/N})_{min}$ has a flux density S_{min} , given by:

$$S_{min} = \epsilon \, \frac{(S/N)_{min} \, T_{sys}}{G\sqrt{n_p \, \Delta t \, \Delta\nu_{MHz}}} \sqrt{\frac{W_e}{P - W_e}} \quad \mathrm{mJy}, \tag{2.1}$$

where $\epsilon \geq 1$ is a factor that accounts for the loss in sensitivity due to the digitalisation and the subsequent transmission of the signal, G is the telescope gain (in K/Jy), n_p is the number of observed polarisations of the signal ($n_p = 1$ for single-polarisation observations, $n_p = 2$ if two orthogonal polarisations are

summed), Δt is the integration time in seconds, $\Delta\nu_{MHz}$ represents the observing bandwidth in MHz, P is the pulsar period. The *system temperature* T_{sys} (in K) is given by:

$$T_{sys} = T_{rec} + T_{spill} + T_{atm} + T_{sky}, \tag{2.2}$$

where T_{rec} is the noise temperature of the receiver, T_{spill} is the contribution of the 'spillover noise' due to the ground, T_{atm} is caused by the Earth atmosphere, and T_{sky} represents the noise temperature of the sky background (3 K from the cosmic microwave radiation, that is everywhere, and a strong contribution from electrons in the Galactic plane emitting synchrotron radiation, that makes the T_{sky} to be a function of sky position and observing frequency, $\propto \nu^{-2.7}$).
The *effective width* W_e of the pulse in eq. (2.1), which we want to be as small as possible, is given by:

$$W_e = \sqrt{W_i^2 + (\beta\delta t)^2 + \delta t_{DM}^2 + \delta t_{scatt}^2}, \tag{2.3}$$

where W_i is the intrinsic width of the pulse, $\beta \sim 2$ is an instrument factor, δt is the *sampling time* (i.e. the rate at which we record the signals arriving during the integration time), δt_{scatt} is the broadening of the pulse due to scattering ($\delta t_{scatt} \propto \nu^{-4}$), and δt_{DM} is the pulse broadening caused by dispersion, due to the fact that the radio pulsar signal is broad-band; if in our telescope we have a receiver with bandwidth $\Delta\nu_{MHz}$, and ν_{MHz} is the central observing frequency, deriving eq. (1.22) we can obtain the rate at which the pulse traverse the radio spectrum:

$$\dot{\nu} = \frac{d\nu}{dt} = \frac{\nu_{MHz}^3}{8.3 \times 10^3\, DM} \quad \text{MHz s}^{-1}, \tag{2.4}$$

and hence the broadening of the pulse traversing the receiver bandwidth:

$$\Delta t_{DM} = 8.3 \times 10^3 \frac{\Delta\nu_{MHz}}{\nu_{MHz}^3} DM \quad \text{s}. \tag{2.5}$$

We can however minimise the dispersion effects by splitting the total bandwidth into several channels having a frequency width $\delta\nu_{MHz}$, so that in each channel the broadening will be smaller:

$$\delta t_{DM} = 8.3 \times 10^3 \frac{\delta\nu_{MHz}}{\nu_{MHz}^3} DM \quad \text{s}. \tag{2.6}$$

From what said so far, it is clear that for our aim we need to reach a value of S_{min} as low as possible. Observing at high frequencies can help since T_{sky}, δt_{scatt} and δt_{DM} decrease with increasing frequency, but as already mentioned also the pulsar flux density decreases; hence a compromise is necessary between these two things.

To reduce S_{min} we can also play with the other parameters in the previuos equations (if technically possible), for example increasing the integration time Δt, reducing the sampling time δt, using a larger observing bandwidth $\Delta\nu_{MHz}$ and/or dividing the latter into a high number of channels with width $\delta\nu_{MHz}$; nevertheless, we must keep in mind that in all these cases the quantity of collected data can become huge, and therefore we need to have the resources to store it up and the computing power to analyse such a mass of data.

To decide the strategy of a survey the first thing to do is defining the goals to be achieved, i.e. for example what kind of pulsars (section 1.1.7) we want to find.

If our target are the ordinary pulsars we must observe mainly on the Galactic plane, since they are younger and did not go far away from it yet; in this direction, however, the ISM density, and hence the DM, is high. In general that is not a problem for ordinary pulsars, because the broadening of the pulse induced by the ISM is normally small with respect to P ($P_{mean} \sim 0.5$ s) also at small radio frequencies. Nevertheless, if we want to observe beyond the Galactic Centre, or find also the youngest pulsars, that have periods of only few tens of milliseconds, it is opportune to observe at high frequencies ($1 \div 2$ GHz)[1], so that the ISM effects, especially scattering, for which there are no technical expedients as for the dispersion, are lower.

If instead we want to find millisecond pulsars, since they are old and hence already moved we must search above or below the Galactic plane, up to a height $z \sim 500$ pc. Here the ISM density is lower so that its effects should be less important, but since MSPs have short periods even a small pulse broadening caused by ISM can drastically reduce the signal strenght. For this reason it is opportune to observe also MSPs at high frequencies.

[1]Pulsars can be observed at wavelenghts from about 15 m to 3 cm, i.e. at frequencies from 20 MHz to 10 GHz.

2.2 Search techniques

Once the signals have been acquired by the radiotelescope, digitalised and recorded on magnetic tapes, they must be analysed in order to single out possible periodic signals of a pulsar, immersed in the background noise.

Data analysis is made up of several stages, and different techniques need to be used depending on what we are searching for (e.g. isolated pulsars, binary pulsars, single pulses, etc.). The initial two stages are common to all the techniques, since they 'clean' the data from spurious signals and eliminate the dispersion suffered by the signals when they cross the ISM.

In fact, the first step of processing is identifying Radio Frequency Interference (RFI), i.e. man-made terrestrial radio signals emitted for example by radars, satellites, etc. (see also section 4.1 for the RFI problem), which can cover the weak signals from pulsars and/or even have a periodic pulsar-like behaviour, that can deceive us inducing to consider them as pulsar candidates. Nevertheless, a first 'cleaning' can be done owing to the fact that most sources of RFI, being produced on Earth, are undispersed and hence detectable in the so-called *zero-DM* data set (i.e. the data set summed in frequency without de-dispersion). To do that we can use two main approaches, one in the frequency domain and one in the time domain; for example, as described in Keith et al. 2010, in the *HTRU survey* pipeline (see chapter 3) we remove both 'bad' spectral channels and 'bad' time samples.

'Bad' spectral channels are the channels affected by periodic RFI (usually appearing many times in different sky positions), that can be identified doing an individual Fourier transform of the zero-DM data set for each channel, so that the power spectra (see section 2.2.1) are obtained; if we find in them any frequencies corresponding to a spectral S/N of the detected signals higher than 15, the corresponding frequency channels are rejected[2].

'Bad' time samples are the ones, in our series of recorded signal samples of an observation (*time series*), affected by sporadic *bursts* of RFI (i.e. non-periodic), that can be identified by comparing them with the expected mean and standard deviation of the single zero-DM time series obtained by summing the data in

[2]It could be possible also to create a spectral 'mask', i.e. a list of identified RFI frequencies, for RFI occurring frequently.

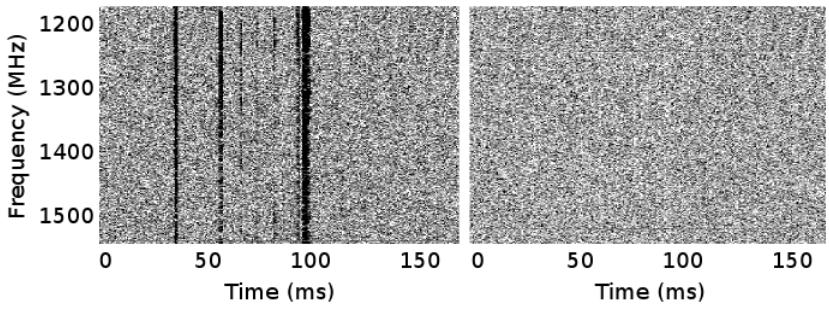

Figure 2.1: RFI excision: replacement of time samples affected by bursts of RFI (vertical lines in the left panel) by random noise (right panel). From Keith et al. 2010.

frequency. If a time sample differs from the mean value by more than five standard deviations (5σ), it is removed from the time series and replaced with random noise, as in fig.2.1.

Unfortunately, this cleaning is not able to remove *all* RFI signals, and most of the candidates indicated by the Fourier domain search in the de-dispersed time series (section 2.2.1) are RFI.

As already mentioned, the second step of data analysis is the so-called *de-dispersion*: this is a technique to eliminate the broadening of a pulse due to dispersion, i.e. to the different delays experienced by the several frequency components of the signal as the latter crosses the ISM (section 1.1.5). There are different algorithms to remove dispersion, the simplest one consists in the following procedure: after dividing the total observing bandwidth into several channels with width $\delta\nu_{MHz}$, as said in section 2.1, we consider the time series as it appears in each channel, assign a value to the DM and calculate the delay of the signal in every channel with respect to a reference frequency, according to eq. (1.25); then we apply the calculated delays to the time series in all the channels (see fig.2.2) and sum all in frequency, obtaining a single de-dispersed time series.

Since we do not know the origin of a signal and hence the right DM value, we have to perform many trials producing many de-dispersed time series, each corresponding to a different DM value. We must be careful with the choice of the trial DM values though, since if the interval between two values is too

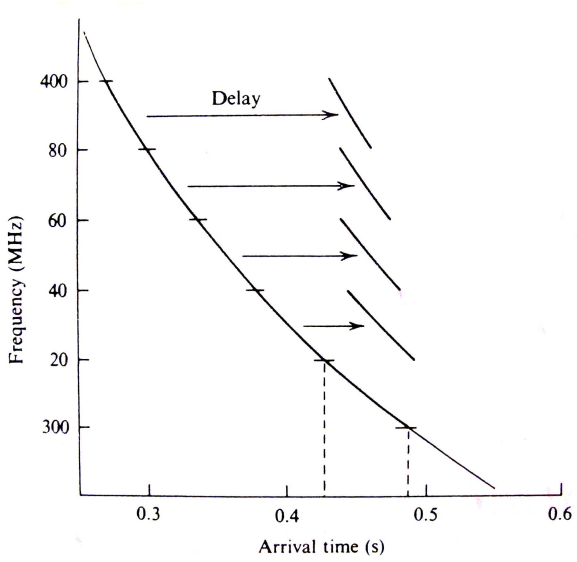

Figure 2.2: De-dispersion criterion (see the text). From *Pulsar Astronomy, 3rd ed.* by Lyne and Smith 2005.

large the true DM value of an unknown pulsar could lie within the interval, with a loss of sensitivity (see fig.2.3); on the other hand, if the interval is too small we waste computing time to produce de-dispersed time series that for neighbouring DM values will be almost identical. The minimum value of DM for which doing de-dispersion makes sense is the one corresponding to a delay, and hence to a broadening, of the signal over the whole bandwidth[3] (eq. (2.5)) equal to the sampling time (since it is the smallest delay that we are able to measure):

$$\Delta t_{DM} = 8.3 \times 10^3 \, \frac{\Delta\nu_{MHz}}{\nu^3_{MHz}} \, DM_{min} = \delta t, \tag{2.7}$$

from which:

$$DM_{min} = 1.2 \times 10^{-4} \, \frac{\nu^3_{MHz}}{\Delta\nu_{MHz}} \, \delta t. \tag{2.8}$$

The value of DM_{min} is also the step δDM that we use to increment the DM values from DM_{min} to a certain value DM_{diag} (see below), producing at every step an increment of δt in the delay over the whole bandwidth, Δt_{DM}, since again it makes no sense using a DM step corresponding to an increment in the

[3]It is therefore the delay between the time of arrival of the signal highest frequency in the highest channel and the one of the signal lowest frequency in the lowest channel, that is the slowest component.

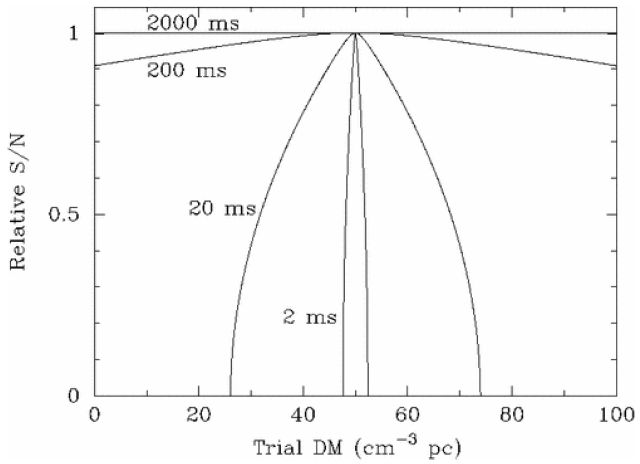

Figure 2.3: S/N versus trial DM relative to the S/N value for a true DM of 50 pc cm^{-3}, spanning an 8 MHz band centred at 430 MHz, plotted for a variety of pulse periods. The right choice of the DM step is critical for the detection of pulsars with periods below a few hundred ms, since they suffer a strong decreasing of the S/N for wrong DM values, while pulsars with periods above that value do not experience such an issue. From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

delay smaller than the sampling time (and hence not measurable). Therefore the second trial DM value is $DM_{min} + \delta DM = 2DM_{min}$, and corresponds to a delay over the whole bandwidth $\Delta t_{DM} = 2\delta t$; and so on. We maintain this value of δDM up to the DM value corresponding to a delay of the signal equal to the sampling time in *each* channel (eq. (2.6)) with respect to the previous channel:

$$\delta t_{DM} = 8.3 \times 10^3 \, \frac{\delta \nu_{MHz}}{\nu^3_{MHz}} \, DM_{diag} = \delta t, \tag{2.9}$$

from which:

$$DM_{diag} = 1.2 \times 10^{-4} \, \frac{\nu^3_{MHz}}{\delta \nu_{MHz}} \, \delta t. \tag{2.10}$$

Hence in this case the delay over the whole bandwidth will be $\Delta t_{DM} = n \, \delta t$, where n is the number of channels; this DM value is called *diagonal DM*. For values larger than DM_{diag} the delay, and hence the broadening, in each channel becomes larger than δt; if it is smaller than $2\delta t$, and we vary the DM again with a step $\delta DM = DM_{min}$, we produce again an increment in the delay over the whole bandwidth equal to δt, but this time it is not good, since in this way we de-disperse pulses broader than δt with a time resolution smaller than the pulse width itself. Therefore in this case we choose a step δDM such that the

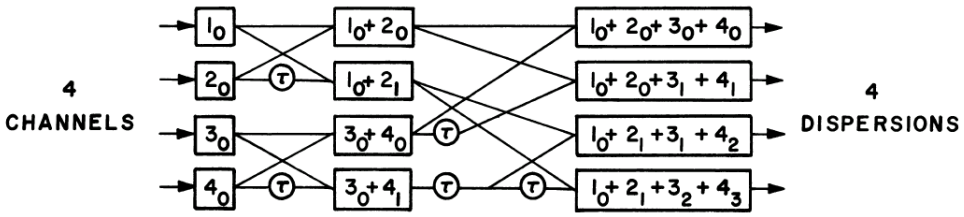

Figure 2.4: Scheme of the tree algorithm for de-dispersion, in the four-channel case. Rectangles represent summations; the numbers within them represent the channel numbers in order of increasing frequency, while the subscripts on the channel numbers indicate the accumulated delay of a particular datum, in units of the sampling time τ. Circles indicate the introduction of one unit of delay. From Taylor 1974.

increment in Δt_{DM} is equal to $2\delta t$. We do the same if the delay in each channel lies between $2\delta t$ and $3\delta t$, choosing for the same reason a DM step such that the increment in Δt_{DM} is $3\delta t$; and so on.

The maximum value of DM to reach must be determined depending on the periodicity that we expect to find. For example, for millisecond pulsars we de-disperse up to a DM value corresponding to a delay in each channel equal to half the average period of a MSP (i.e. $\delta t_{DM} \sim 2$ ms), so that the signal can be still identified.

A problem with the method described above is that it requires n^2 operations, so that it is computationally expensive. An alternative method is the so-called *tree-algorithm* (Taylor 1974), that requires instead only $n\log_2 n$ operations. It is built as a 'tree', where each 'branch' starts from a different couple of channels (so that the total number of channels must be an integral power of 2); the procedure consists in introducing progressive delays into the paths of the higher frequency channels, and then combining the delayed signals to produce different de-dispersed output signals (i.e. different de-dispersed time series), whose number is equal to the number of frequency channels (see fig.2.4, where the algorithm provides dispersion filters 'tuned' to four different dispersions).

In this method we assume that the delay across the whole bandwidth varies as a linear function of frequency: $\Delta t_{DM} \propto \nu$, while the exact form of the

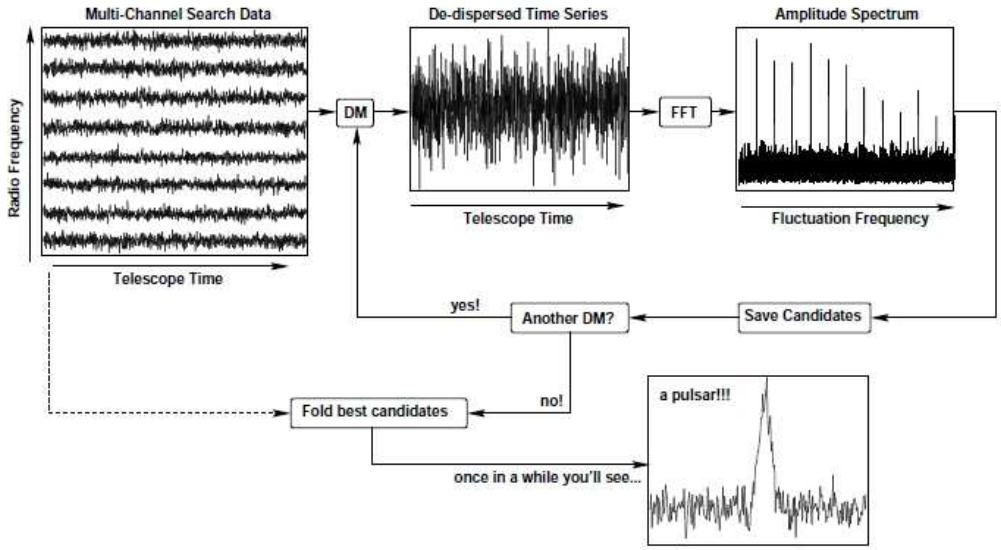

Figure 2.5: Scheme of the main steps of the standard search. From Lorimer 2001.

delay is quadratic (eq. (1.25): $\Delta t_{DM} \propto \nu^{-2}$); this assumption is good if the employed total bandwidths are narrow, but if they are large, as in many recent surveys, the approximation is not good. In that case a linearisation approach can be used, for example padding the data with empty channels so that the delay approximates a straight line; however, the cost of linearising the data may exceed the gain from using this algorithm.

In the following sections we describe some of the different techniques used for the remaining steps of data analysis. Search methods for single-pulses and for long-period ($P_{spin} > 2$ s) pulsars will not be described since they are not subjects of this thesis.

2.2.1 Standard search

The most common and general procedure to analyse the data searching for pulsars with unkown period and DM, especially if isolated, is the *standard search*, whose main steps are illustrated in fig.2.5.

The first step after de-dispersion is searching the data for the presence of periodic signals, i.e. for trains of regularly spaced pulses that may be hidden within the background noise; this can be done moving from the time domain to the frequency domain, since in this way the periodicity of the signals is brought

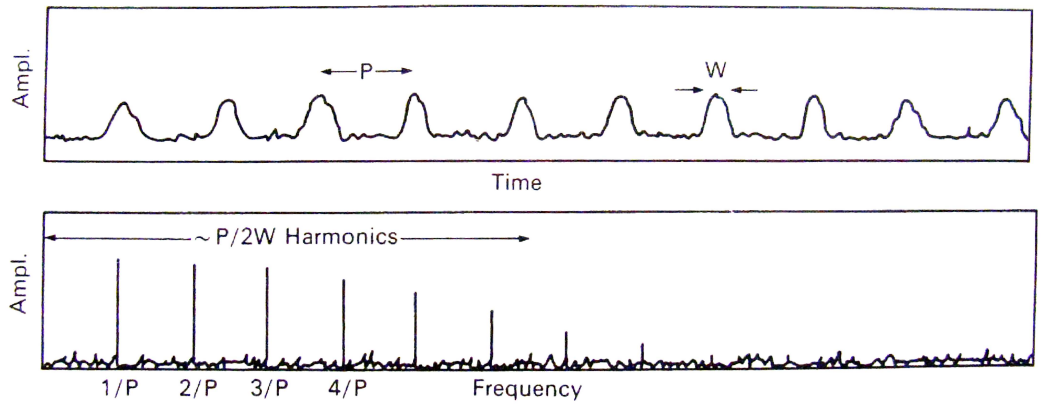

Figure 2.6: A pulse train with period P and pulse width W is Fourier transformed, giving a spectrum where its periodicity is brought out. From *Pulsar Astronomy, 3rd ed.* by Lyne and Smith 2005.

out. In fact, if we have a pulse train with period P and pulse width W, like in fig.2.6, and we Fourier transform it, we produce in this way a spectrum like the one in the figure, that shows how much power is associated to each frequency corresponding to such pulse train. In fact, assuming that the background noise is purely Gaussian, after the Fourier transformation the noise is still Gaussian, while the pulse train, being periodic (but not purely sinusoidal) and hence expressible as the sum of sinusoidal components, produces a series of frequencies, where $\nu = 1/P$ is the fundamental, and 2ν, 3ν, and so on, are the harmonics, corresponding to the various components. The power of the pulse train is then distributed among the fundamental and the harmonics, whose number is calculated to be proportional to P/W. Therefore, pulsars with a large duty cycle ($\delta = W/P$) like the MSPs, hence having quasi-sinusoidal pulses so that also the train is quasi-sinusoidal, have a small number of harmonics and the power is mainly concentrated in the fundamental feature; pulsars with a small duty cycle like the ordinary pulsars instead, having narrow pulse widths so that the pulses and hence the train are not quasi-sinusoidal, have a large number of harmonics with amplitude comparable to the fundamental.

In order to recover the power distributed among the harmonics, especially in the last case, we use a technique called 'incoherent harmonic summing' illustrated in fig.2.7, where the original spectrum is stretched by a factor of 2, and the resulting spectrum is added to the original one. In this way the

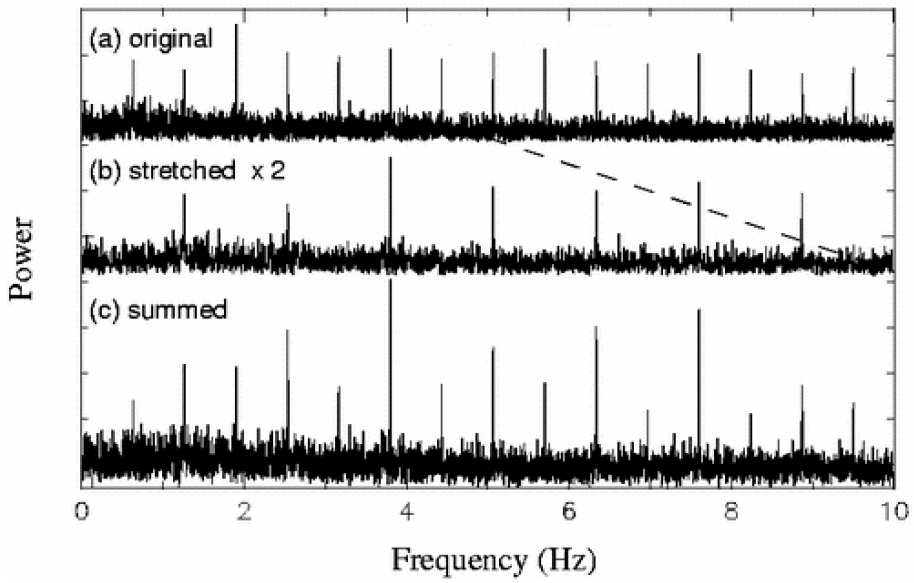

Figure 2.7: The harmonic summing technique. Figure provided by Dipankar Bhattacharya for the *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

second harmonic is summed to the fundamental, so that we recover its power and the S/N increases, also because the noise increases by a factor of $\sqrt{2}$ while the amplitudes of the fundamental and the harmonic add directly. The process can be repeated by stretching the obtained summed spectrum by a factor of 2 again, so that the fourth harmonic can be added to the sum of the fundamental and the second harmonic; and so on. Taking care to add in odd-numbered harmonics, we can obtain the sum of the first two, four, eight, sixteen, etc. harmonics, producing a large increment in the S/N of narrow width pulses.

In reality, however, what we have is a de-dispersed time series made up of N data points sampled at a time δt, rather than a continuous data stream; hence to move from the time domain to the frequency domain we compute the *Discrete Fourier Transform* (DFT) of the time series. The k^{th} Fourier component yielded by the DFT has a frequency $\nu_k = k/(N\delta t) = k/T_{obs}$, where T_{obs} is the lenght of the observation and $1 < k < N/2$. These ν_k represent therefore the frequencies that we are able to reveal for a time series of that lenght T_{obs} and sampled at that time δt: the lowest frequency is $1/T_{obs}$, the highest one is $1/(2\delta t)$ and the interval between each frequency and the next one is $1/T_{obs}$. We can then define a *Fourier bin* as the frequency interval having ν_k

as central frequency and lenght $1/T_{obs}$, i.e. as our *spectral resolution*. Therefore the fundamental and the harmonic frequencies of our pulse train hidden in the time series, once the latter is Fourier transformed, that are spaced by $1/P$, will lie in that frequency range, i.e. in some of the bins, and each of them may coincide with the central frequency of a bin or not. In the latter case we will have a not precise knowledge of it.

A problem with using a DFT algorithm is that it requires N^2 operations; given that, for modern applications, time series can have even more than 2^{25} samples, this can become very time consuming. Hence what we use is the *Fast Fourier Transform* algorithm (FFT), that is computationally convenient since it requires only $N\log_2 N$ operations.

The next step of processing is the so-called *candidate sorting*. The result of data analysis up to this stage is a usually long list of candidate periods (found by the FFT) from each DM value used to de-disperse, with their spectral S/N ratios for all harmonic summings. A pulsar may appear many times in this list at different S/N values, but the maximum S/N should correspond to the DM value that is closest to the true DM. Hence in the sorting procedure we group the candidates with the same period into a single candidate and combine candidates that are harmonically related; then we select the candidates with the highest spectral S/N ratios. In this way we obtain a manageable number of pulsar candidates.

At this stage, we de-disperse and fold the raw data (i.e. the original dispersed time series) at the candidate DM and period respectively, for each selected candidate. The *folding* technique is used to increase the S/N so that the signal may hopefully be identifiable above the noise in the time domain. It consists in dividing the de-dispersed time series of lenght T_{obs} into a number N_{subint} of identical time intervals called *subintegrations* (usually 32 or 64); each subintegration is then divided into intervals of lenght equal to the candidate period P, and each interval is in turn divided into n_{bins} equally spaced elements (called *time bins*, for example we can choose $n_{bins} = 128$) across the candidate period. Hence every time bin corresponds to a particular rotational phase of the candidate, and summing the signals in corresponding bins (e.g. the k^{th} bin) over all the intervals within a subintegration we are doing a coeherent (i.e. in phase) sum, obtaining an increased S/N of the integrated pulse profile (since

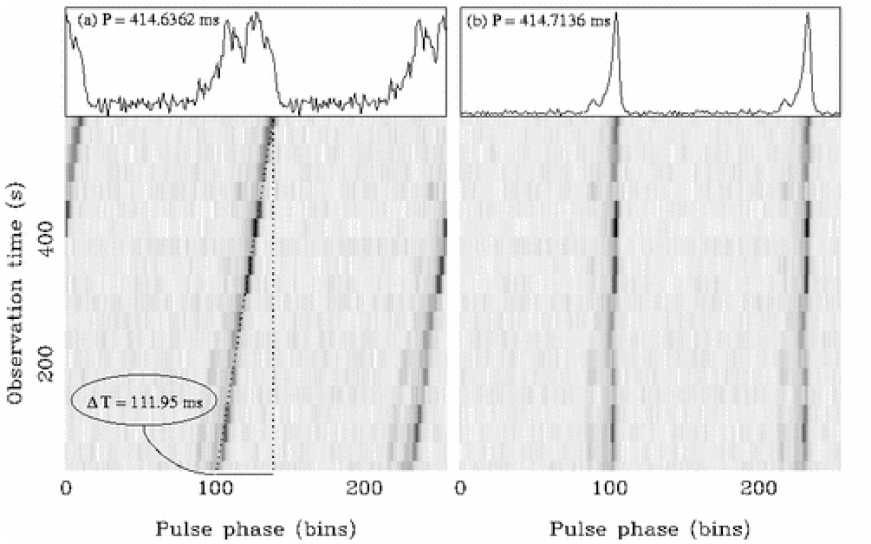

Figure 2.8: A time series for a pulsar folded at a slightly incorrect period (a) and at the correct period (b). From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

the signal increases with the number of summed elements, i.e. of intervals, while the noise increases only with the square root of it).

Nevertheless, this technique requires a precise identification of the period of the signal. If the candidate period does not match it exactly (i.e. if the signal frequency does not coincide with the central frequency of the concerned Fourier bin, which is the candidate frequency), we fold the time series with an imprecise period and the difference is evident among the subintegrations, since the phase of the folded signal drifts slightly from subintegration to subintegration, and therefore the integrated pulse will appear significantly broadened and with reduced S/N with respect to the correct period case. Hence, making a plot of the folded signal in all the subintegrations versus the rotational phase, if the period is correct we see a vertical line across the subintegrations, while if it is imprecise we see a diagonal line (see fig.2.8). In the latter case then, in order to find the correct value of the period, we use a procedure called *pdm optimisation*, which consists in performing period trials, using different values around the candidate period (i.e. adding trial delays to the period) to sum all the subintegrations, until the highest folded S/N is found. This refined value is called *best* or *optimised* period. By analogy, if we de-disperse the data into a

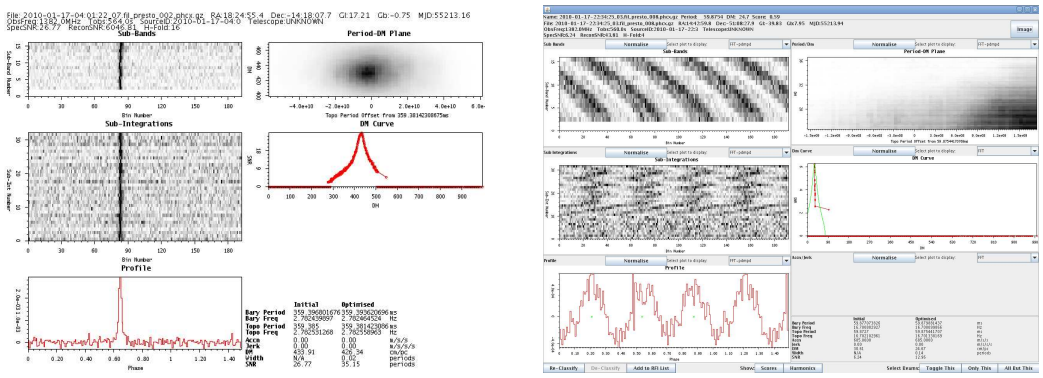

Figure 2.9: Examples of JREAPER diagnostic plots. Left: plot of the pulsar J1824-1423; right: plot of a candidate clearly identifiable as RFI. Clockwise from the bottom left, the sub-plots show the integrated pulse profile, the S/N vs pulse phase for 32 subintegrations, the S/N vs pulse phase in 16 sub-bands, the dependence of S/N on DM and on offset from the nominal period, and the spectral S/N vs trial DM (S/N - DM $curve$), where the red curve is the observed one, while the green curve is the theoretical prediction for that candidate.

number $N_{subbands}$ of identical *sub-bands* (into which we choose to divide the total bandwidth; for example 16) and fold at the best period obtained above, if the pulse profile is delayed across the de-dispersed sub-bands we can improve our estimate of the DM performing DM trials around the candidate value, until again the highest folded S/N is found. This is part of the pdm optimisation procedure too, and the refined DM value is called *best* or *optimised* DM.

The last step of data analysis is the visual inspection of the candidates. In fact the best P and DM, together with the related parameters, are then used to make the final diagnostic plots for the candidates, like the ones in fig.2.9, obtained using the graphical tool JREAPER (Keith et al. 2009). We then inspect the plots of the candidates having an optimised S/N, i.e. corresponding to both best P and best DM, higher than a particular value, usually 8 or 9 (since below that value it is more probable that the candidate is not a pulsar), and we classify them as possible new pulsars (that need to be confirmed by re-observation) or non-pulsars (noise, RFI)[4]. The classification is quite subjective, anyway several features typical of a pulsar can help us to make a decision

[4]The candidate may also be an already known pulsar, but usually candidates are checked by a software program during the processing to identify possible known pulsars, to avoid wasting time during the visual inspection; however, sometimes not all the known pulsars present among the candidates are identified, hence it is always opportune to double-check at least the most probable candidates with the pulsar catalogue.

about a candidate: a high optimised S/N, together with a well defined peak in the spectral S/N vs trial DM plot (*S/N - DM curve*), a defined spot in the gray-scale showing the dependence of the folded S/N on the DM and on the offset from the *nominal* (i.e. returned from the Fourier domain search) period, signal continuity across the sub-bands (visible in the sub-bands vs phase plot, where the signal should be linear), signal continuity in time (visible in the subintegrations vs phase plot, where the signal should be linear, or parabolic in case of pulsars in binary systems, see section 2.2.2), and a well peaked pulse profile (fig.2.9).

In general, what said is useful to make a distinction between pulsars and non-pulsars; sometimes, however, a kind of RFI may show these same features, deceiving and inducing us to re-observe it too.

2.2.2 Acceleration search: pulsars in binary systems

If a pulsar is part of a binary system, we observe an apparent change in the frequency of the pulse trains from it as the pulsar moves along its orbit: the time of arrivals of the pulses appear closer each other (*blue-shifted*) if the pulsar is moving toward us, while they are more separated (*red-shifted*) if it is moving away from us (*orbital Doppler effect*). Hence the pulsar spin frequency (and period) seems to vary during the observation time, and this effect is heavier for close (short orbital period) binaries; that makes the Fourier analysis of the standard search be less sensitive to binary pulsars, especially if in close binary systems and observed for a long time. In fact, due to this effect the signal power in the Fourier spectrum is *smeared* over a large frequency range (i.e. many Fourier bins), with a great decrease of the spectral S/N. This could make a candidate not to be selected during the sorting stage; or, if it is selected, it could be rejected anyway after folding, since the apparent change in its period causes a related change in its rotational phase across the folded subintegrations, visible as a parabolic trend of the signal in the subintegrations vs phase plot, producing a broadened integrated profile with an even further reduced S/N (see case (a) in fig.2.10).

However, several techniques can be used to partially, or sometimes fully, recover the loss of sensitivity due to binary motion; they all are grouped under the name of *acceleration search*. In the following a description of some of them

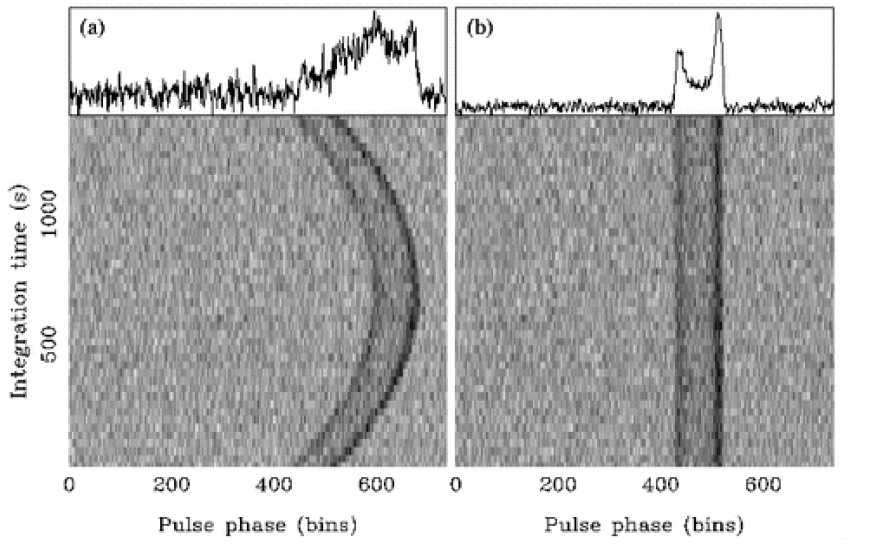

Figure 2.10: Example of folding results for the time series of a binary pulsar, (a) folded without any corrections for the orbital motion, (b) folded assuming an opportune value for its orbital acceleration (note that the integrated pulse profiles are normalised at the peak hence are not to scale). From *Handbook of Pulsar Astronomy* by Lorimer and Kramer 2005.

is given.

Time domain resampling

This technique works in the time domain and consists in resampling the de-dispersed time series in order to transform it to the rest frame of an inertial observer with respect to the pulsar, i.e. removing the pulsar orbital motion. This is done by 'stretching' or 'compressing' the data using different trial values for the orbital acceleration, as described below; the new samples in the time series are created by linear interpolation over the original time series, taking care that the total number of samples in the two series is the same (see for example Camilo et al. 2000).

If τ is a time interval in the pulsar frame corresponding to a time interval t in the non-inertial observer frame, these intervals are related by the formula of the Doppler effect:

$$\tau(t) = \tau_0 \left(1 + \frac{v_l(t)}{c}\right), \tag{2.11}$$

where $v_l(t)$ is the observed radial velocity of the pulsar along the line of sight, and terms in (v_l/c) of order higher than the first have been neglected since $v_l << c$. The choice of the constant τ_0 is explained below.

If we know the orbital parameters of the binary systems we can calculate $v_l(t)$ from Kepler's laws, but if we are doing a blind search, like in this case, we need to search over all the orbital parameter space, that is very computationally expensive. Hence, an economical solution is assuming a constant orbital acceleration a_l during the observation, so that $v_l(t) = a_l t$. In this way, what we do is resampling our time series using the time intervals calculated by eq. (2.11) for a range of trial values of the acceleration, producing a new time series for each of those values.

The value of τ_0 is chosen so that τ is equal to the original sampling interval, δt, at the midpoint of the observation, because this guarantees that the number of samples in the corrected time series is the same as in the original one; so we have:

$$\tau(t = T_{obs}/2) = \tau_0 \left(1 + a_l \frac{T_{obs}/2}{c}\right) = \delta t \quad \Longrightarrow \quad \tau_0 = \frac{\delta t}{1 + a_l T_{obs}/2c}, \tag{2.12}$$

where T_{obs} is the observation time.

The acceleration step between two trial values must be chosen with care, not to over-sample (increasing uselessly the processing time) nor under-sample (losing sensitivity) the acceleration space. In order to conveniently choose this step, first of all we calculate the number of Fourier bins an accelerated signal will occupy if no correction for the orbital motion is applied, i.e. how many bins are *drifted* by its frequency over the course of an observation; applying the formula for the Doppler effect again, we have:

$$\nu(t) = \nu_0 \left(1 - \frac{v_l(t)}{c}\right) = \nu_0 \left(1 - \frac{a_l t}{c}\right), \tag{2.13}$$

where $\nu(t)$ is the observed pulse frequency and ν_0 is the true spin frequency of the pulsar in its rest frame. Therefore the frequency drift will be:

$$\mid \dot{\nu} \mid = a_l \frac{\nu_0}{c}. \tag{2.14}$$

Then the number z of Fourier bins drifted during the observation time T_{obs}, remembering that the width of a bin is $\Delta\nu = 1/T_{obs}$, will be:

$$z = \frac{\mid \dot{\nu} \mid T_{obs}}{\Delta\nu} = \frac{\nu_0}{c} a_l T_{obs}^2. \tag{2.15}$$

If we define Δa_l as the amount by which an assumed trial acceleration is different from the true value, the corresponding number of bins drifted by the signal (i.e. the residual drift after our correction of the time series) is:

$$\Delta z = \frac{\nu_0}{c} \Delta a_l T_{obs}^2, \tag{2.16}$$

from which:

$$\Delta a_l = \frac{cP_0}{T_{obs}^2} \Delta z, \tag{2.17}$$

where $P_0 = 1/\nu_0$ is the true spin period of the pulsar. Since we want $\Delta z < 1$ so that the residual drift lies within our spectral resolution, we obtain:

$$\Delta a_l < \frac{cP_0}{T_{obs}^2}. \tag{2.18}$$

Hence if we choose Δa_l as our acceleration step, this choice guarantees that any pulsar period or harmonic with a period greater than P_0 will not drift by more than one Fourier bin during the observation time, after our correction of the time series, even if the value of acceleration used to correct it does not coincide exactly with the true value (but lies within a lenght Δa_l from it). For example, Camilo et al. 2000 chose $\Delta a_l = 0.3$ m s^{-2} to correct a 17.5 min time series from the globular cluster 47 Tucanae, and this choice guarantees a correction within one Fourier bin for any pulsars with $P > 2$ ms.

Once the de-dispersed time series has been corrected, producing a number of new time series equal to the number of trial acceleration values, we perform on each of these new series the FFT analysis and all the subsequent steps described for the standard search (section 2.2.1); for a pulsar corrected with an opportune acceleration value, we obtain the results showed in fig.2.10 case (b), where the effects of the orbital motion have been fully removed.

This resampling method has however some problems. First of all, since we made the assumption of a constant acceleration during the observation, this is valid only if we observe the pulsar for a time corresponding to a small part of its orbit, i.e. for $T_{obs} << P_{orb}$, where P_{orb} is the pulsar orbital period.

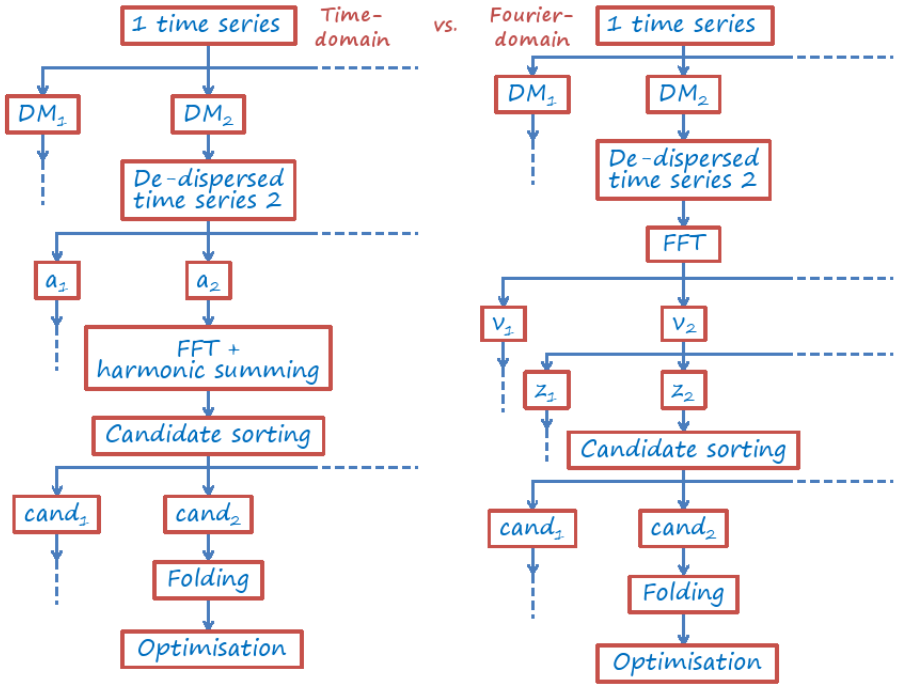

Figure 2.11: Comparison between time-domain and Fourier-domain acceleration searches.

Another issue is that for long time series ($N > 2^{23}$) the computational time becomes very long. In fact, each of the n_{DM} de-dispersed time series yielded for all DM trials is searched for n_a acceleration trial values, producing in total $n_{DM} \times n_a$ time series; each of them will be then Fourier analysed, i.e. a huge number of FFTs must be performed, that for long time series is very computationally expensive.

Correlation technique

The *correlation technique* developed by Ransom et al. 2002 is a Fourier-domain version of the time-domain acceleration search to detect a binary pulsar if the orbital period is much longer than the observation time; with respect to the time-domain technique, this method has great computational advantages when the time series is long. In fact, as said in the previous section, time-domain techniques require a full-lenght FFT for each trial acceleration, while the correlation technique requires only a single full-lenght FFT for each de-dispersed time series, i.e. for each trial DM (see fig.2.11).

The method that this technique uses to remove the orbital effects consists in applying a filter to the FFT of the time series to 'sweep up' the power of the signal, spread over a number of Fourier bins, into a single bin. In fact, if we consider the sinusoidal components of a not sinusoidal signal (like the one from a pulsar, in its rest frame), the Fourier transform of each of them is a sinc function; but due to the orbital motion of the pulsar, the Fourier transform that we obtain is a convolution of the sinc function with a *finite impulse response (FIR) filter*, i.e. a filter that spreads the sinc response (hence the power) over a certain finite number (say m) of bins of the frequency range analysed. Hence, applying an opportune *inverse* FIR filter to the Fourier transform in the appropriate bins we can recover the sinc function (i.e. removing the orbital motion). Therefore we need to do only one FFT of the whole data set, and then work with 'local' Fourier amplitudes, i.e. the ones in the m bins around the true frequency of the considered signal component.

For a signal with a constant frequency derivative, if F_{k-r_0} is the Fourier response of a signal component biased by the orbital motion (where $\mid k - r_0 \mid$ is the frequency offset of the k^{th} bin with respect to a reference bin r_0), response which goes to zero as $\mid k - r_0 \mid$ approaches some number of bins $m/2$, the corrected Fourier response is:

$$F_{r_0} = \sum_{k=r_0-m/2}^{r_0+m/2} F_k F^*_{r_0-k} \,, \tag{2.19}$$

where we have applied the inverse filter $F^*_{r_0-k}$ (the complex conjugate of F_{k-r_0}, 'frequency-reversed') to the Fourier response in every bin k within the m involved bins around r_0.

The form of $F^*_{r_0-k}$ is a function of both r_0 and the drifted bins $\dot{r}$, to be precise $\dot{r}$ is the number of closest bins to r_0 where most of the signal power has been spread (relatively evenly); in fact, in reality an additional small amount of signal has 'leaked' into bins further away, but however eq. (2.19) takes into account this fact, and the correct Fourier amplitude will be well-approximated if m is chosen such that $\dot{r} < m \leq 2\dot{r}$.

For unknown pulsar signals in our time series, since we do not know their frequencies nor the orbital motions we have to calculate the form of the filter (and the sum in eq. (2.19)) searching the data both in a range of r and in a

range of $\dot{r}$, i.e. covering portions of the $\nu - \dot{\nu}$ plane. From what said, the search in $\dot{r}$ is equivalent to the time-domain acceleration search but is computationally much cheaper.

Using this technique, Ransom et al. 2001 discovered the 1.7 h binary MSP J1807-2459 (with a spin period of 3.06 ms) in the globular cluster NGC 6544. The search was conducted choosing the value of 100 as the maximum number of Fourier bins that could be drifted during the observation (i.e. the maximum trial value); from eq. (2.15), this choice of $z_{max} = 100$ (i.e. +100 and -100 around the frequency) corresponds to an acceleration of 500 m s^{-2} for a signal with $\nu_0 = 10$ Hz or 5 m s^{-2} for a signal with $\nu_0 = 1000$ Hz, T_{obs} being 28.9 mins. Since most known binary pulsars show maximal accelerations of only a few m s^{-2}, in this way the observers were sensitive to all but the most exotic binaries or pulsars with spin periods ($P_0 = 1/\nu_0$) much less than ~ 2 ms.

Stack search

Another acceleration search that works in the frequency domain is the *stack search*, which is even faster than the correlation technique but, differently from it and from the time-domain resampling, does not perform a *coherent* search in acceleration space, i.e. the information on the phase of the signal is lost, implying a lower sensitivity.

This method consists in splitting a de-dispersed time series into a number n_s of identical segments, each of which is Fourier-transformed separately, yielding n_s spectra. In this way, the width of a Fourier bin in each segment is larger than for the entire observation, and furthermore during each segment the drift in frequency of an accelerated signal is smaller; as a consequence, the number of bins drifted over all the observation is reduced by n_s^2.

Then, to compensate for the drift, linear or parabolic frequency shifts (*slides*) are applied to each spectrum so that the peaks of an accelerated signal in each spectrum are alligned; then the spectra are added together (*stacked*). The resulting stacked spectrum hence does not show any orbital motion effects, and the S/N of the signal in it is considerable higher with respect to the standard search, in spite of the lost of the phase information.

Phase-modulation search

The *phase-modulation* technique (Ransom et al. 2003) is another computationally cheap but incoherent search for binary pulsars, for the case of orbital periods $P_{orb} \leq \frac{2}{3}T_{obs}$, where the approximation of a constant acceleration during the observation time is not applicable anymore.

Fortunately, if we have an observation covering several orbits of a binary system, the Fourier power spectrum of the time series shows a characteristic shape imprinted by the *modulation*, i.e. the periodic change, in the phase of the signal due to the orbital motion. In fact, differences in light travel time across the projected orbit advance or delay the pulse phase in a periodic fashion.

Ransom et al. 2003 demonstrated that the imprint can be described by a set of sidebands around the spin frequency of the pulsar, and the sidebands are regularly spaced by a number of Fourier bins $r_{orb} = T_{obs}/P_{orb}$. In fact, the sidebands have periodicities related to P_{orb}, therefore we can obtain the orbital period, possibly together with its harmonics, by doing an FFT of the region of the sidebands (like when we obtain the spin periodicities by doing the FFT of the time series). In this way it is possible to obtain the features of the binary system.

This technique is being used in a number of searches of globular clusters.

Chapter 3

The acceleration search in the HTRU survey

This thesis work focuses mainly on the search for recycled pulsars (most of which are in binary systems) hopefully suitable to be part of a Pulsar Timing Array, in the context of the *High Time Resolution Universe Pulsar Survey* (HTRU); for this purpose the data have been analysed by using an acceleration search technique, exploiting the processing power of the supercomputers at the Cagliari Astronomical Observatory (Sardinia, Italy) and, in the last six months of this PhD, at the Swinburne University of Technology (Melbourne, Australia).

In this chapter I will give a description of the survey and of such a work, which yielded the discovery of a 2.7 ms recycled pulsar, that so far seems very promising for the PTA project.

3.1 The HTRU survey

This section describes the outlines of the HTRU survey of the southern sky (a twin survey is currently underway at the Effelsberg Radiotelescope, in Germany, for the northern sky); for a more detailed explanation see Keith et al. 2010.

The HTRU survey is the result of an international effort involving several pulsar groups in Australia (Swinburne University of Technology in Melbourne and CSIRO Astronomy and Space Science in Sydney) and Europe (Max Planck Institut für Radioastronomie in Bonn, Jodrell Bank Centre for Astrophysics in Manchester and Cagliari Astronomical Observatory in Sardinia). It consists in a uniform digital all-sky survey of the Galaxy for pulsars and radio transients that is being carried out at the Parkes 64-m Radiotelescope (NSW, Australia)

since November 2008. Owing to the survey parameters and strategy (see below), we expect that, once it is concluded, about a Petabyte of digital data will be available.

The problem of storage has been solved by recording the data on 800GB LTO tapes, more than 400 used to date (November 2011), while supercomputers in Melbourne, Manchester and Cagliari are currently being employed to process the data.

The HTRU survey aims at achieving two main goals: the first one is the study of the Galactic population of pulsars and in particular MSPs, the second one is the search for transient radio signals on timescales down to tens of microseconds. For this purpose, the strategy of the survey consists in dividing the sky into three parts (see fig.3.1). The first part is the *low-latitude* survey, which covers the sky region included between $-80° < l < 30°$ and $|b| < 3.5°$ on the Galactic plane, with pointings of 4300 s that allow to observe more in-depth through the disk and reach more distant systems, with respect to the previuos surveys; in this region the presence of faint ordinary pulsars, binary pulsars and pulsar-BH systems is expected. The second portion is the *mid-latitude* survey, which covers the region $-120° < l < 30°$ and $|b| < 15°$ with pointings of lenght 540 s, where we look for bright MSPs that could be in binary systems with a white dwarf companion and that could be suitable for a Pulsar Timing Array. Finally, the third part is the *high-latitude* survey, which consists of 270 s pointings covering the sky area with $\delta < +10°$ (excluding the area overlapping with the mid-latitude one) and that will allow us to have an overview of the transient events over the entire southern sky.

Observations are being done by using the Parkes 20cm Multibeam receiver (Staveley-Smith et al. 1996) and the new digital filterbank Berkeley-Parkes-Swinburne Recorder (BPSR), which allows a very high frequency and time resolution and therefore a great sensitivity towards short-term signals. The observing bandwidth is 400 MHz, centered on 1382 MHz and subdivided into 1024 spectral channels of width 390 kHz (in order to minimise the pulse broadening due to dispersion in the ISM, see sections 1.1.5 and 2.1); unfortunately, the high-frequency part of this band is badly affected by RFI from a geostationary satellite, hence the effective observing bandwidth is reduced to 341 MHz, centered on 1352 MHz, and the effective number of frequency channels

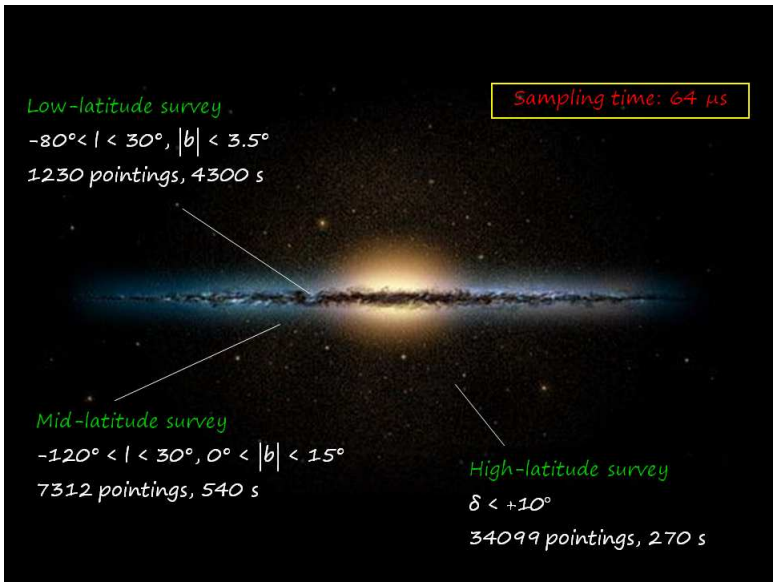

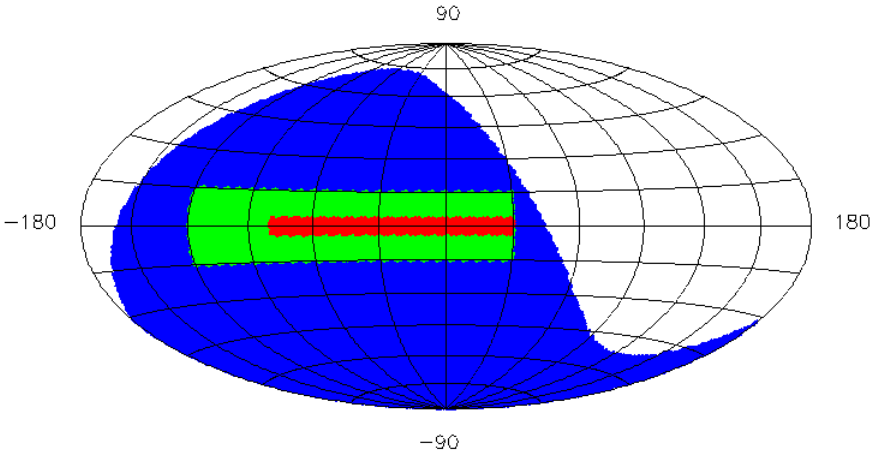

Figure 3.1: The HTRU survey. Top: the survey strategy (the picture of the Galaxy is a computer artwork by CHRIS BUTLER/SCIENCE PHOTO LIBRARY). Bottom: map of the sky areas covered by the three parts of the survey, on a Hammer-Aitoff projection in Galactic coordinates; the red region is the low-latitude survey area, the green one is the mid-latitude survey area, and the blue zone is the high-latitude survey area.

Beam	Central	Inner Ring	Outer Ring
Telescope gain (K/Jy)	0.735	0.690	0.581
Half-power beam width (arcmin)	14.0	14.1	14.5
Beam ellipticity	0.0	0.03	0.06
Coma lobe (db)	none	-17	-14

Table 3.1: Multibeam receiver feed parameters for the central feed, the inner and the outer rings of feeds. From Manchester et al. 2001.

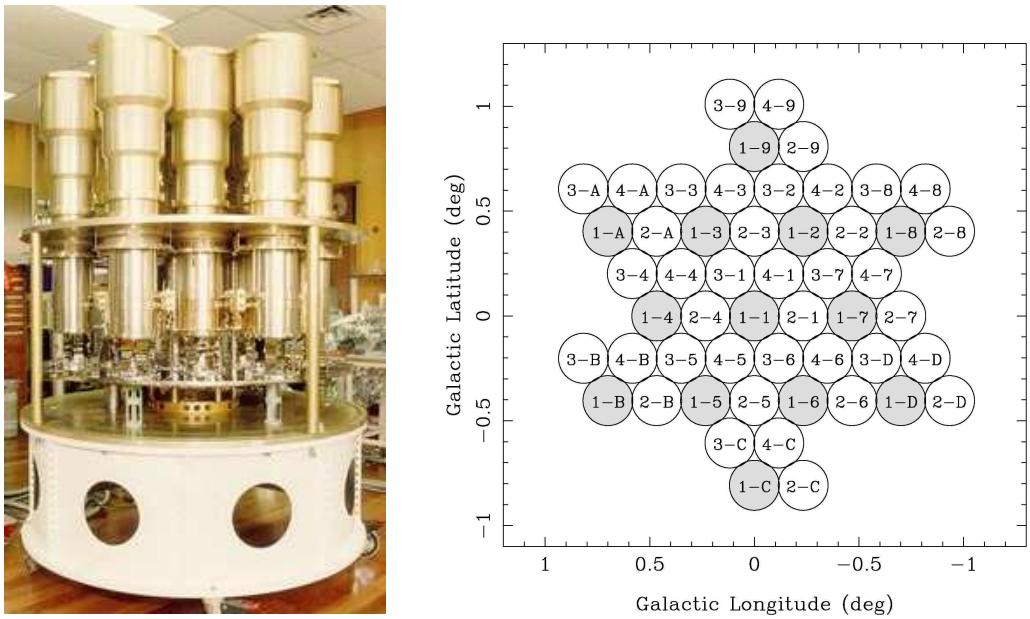

Figure 3.2: Left: the Parkes Multibeam receiver. Right: beam pattern of the Multibeam receiver for a cluster of four pointings, with the feed oriented at a Galactic position angle of 30°; the first number in each beam is the pointing number, while the second one is the beam number. From Manchester et al. 2001.

is 874. The data are 2-bit sampled every 64 μs.

The Multibeam receiver is mounted in the prime focus of the telescope and is made up of 13 feeds, one is in the central position and the others form two hexagonal rings around it; the 13 corresponding beams on the sky are spaced by approximately two beamwidhts, and have a FWHM (full width at half maximum) of 14 arcminutes. The feed parameters are listed in table 3.1. This beam configuration allows a survey to be completed faster, using interleaved pointings to cover a given region; for example, in fig.3.2 it is shown how a cluster of only four pointings covers a region about 1.5° across, for the Multibeam receiver oriented at a Galactic position angle of 30°.

Survey	Mean (K)	Max (K)	Min (K)
Low	7.6	36.0	1.6
Mid	2.5	9.1	0.6
High	1.0	2.4	0.6

Table 3.2: The mean, maximum and minimum values of the noise temperature of the sky in the regions covered by the three parts of the HTRU survey. From Keith et al. 2010.

The minimum detectable flux density (*sensitivity*) for the centre of the central beam of a given survey observation, and for a pulsar of period P, can be computed by using eq. (2.1):

$$S_{min} = \frac{\epsilon\,(S/N)_{min}\,T_{sys}}{G\sqrt{n_p\,\Delta t\,\Delta\nu_{MHz}}}\sqrt{\frac{W_e}{P-W_e}} \quad \text{mJy},$$

where in our case $\epsilon \sim 1$, $(\mathrm{S/N})_{min} = 8$, $G = 0.735$, $n_p = 2$, Δt is equal to 4300, 540 or 270 depending on the considered part of the survey (low, mid or high latitude respectively), $\Delta\nu_{MHz} = 341$, and T_{sys} is measured to be 23 K + T_{sky} at 1400 MHz, where T_{sky} depends strongly on the observed area of the sky (see table 3.2).

The first term on the right-hand side of the equation, calculated by using the mean values for T_{sky}, is equal to 0.20, 0.47 and 0.61 mJy for the low, mid or high latitude surveys respectively; the second term depends on the pulsar period and DM: in fig.3.3, it is shown the value of S_{min} as a function of P and DM for the mid-latitude survey.
However, we must consider that the sky coverage is not uniform, since the outer beams of the receiver are less sensitive than the central beam, and furthermore in each beam the sensitivity reduces towards the edges; hence the mean sensitivity over the mid-latitude survey is calculated to be ~ 0.25 mJy.

By a comparison with previous surveys, as for instance the Parkes Multibeam Pulsar Survey (PMPS), it resulted that the mid-latitude HTRU survey is significantly more sensitive to short period pulsars at large DMs, as shown in Keith et al. 2010, owing to a combination of several factors, such as the narrower frequency channels.

The observations of the low-latitude and high-latitude surveys are still in progress, while the mid-latitude survey has been completed, and the processing

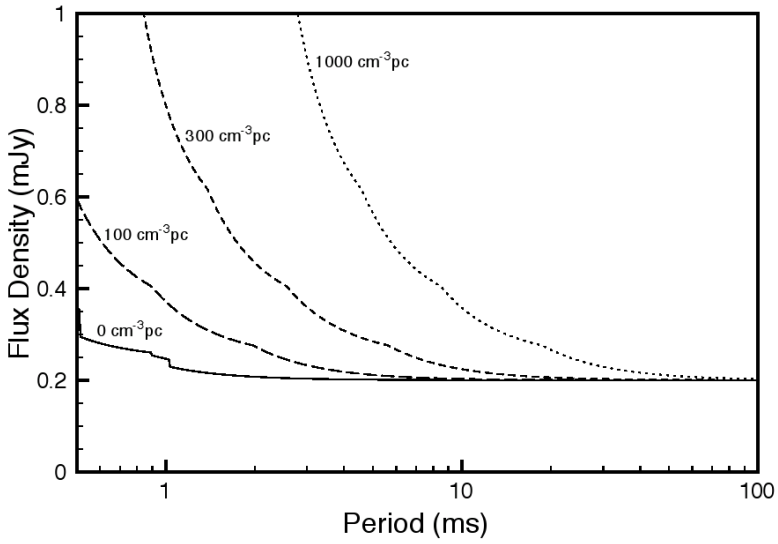

Figure 3.3: The mean sensitivity for the mid-latitude survey as a function of pulsar period and DM, assuming a pulse duty cycle (W/P) of 5%. It can be noted that this survey is sensitive to MSPs out to DMs of a few hundred pc cm^{-3}. From Keith et al. 2010.

of the related data with the HTRU standard pipeline, that analyses the data by using the standard search method (sections 2.2 and 2.2.1), has also been just completed.

The processing of the mid-latitude data is being carried out also with the single-pulse search pipeline and with the acceleration pipeline (i.e. exploiting an acceleration search technique, see section 2.2.2); the latter was the main task of this thesis work and will be described in the next section.

The effectiveness of the HTRU survey has been already proved by the numerous discoveries made so far by analysing the mid-latitude part; those already published include the first magnetar discovered by its radio emission (Levin et al. 2010), 13 transient objects (Burke-Spolaor et al. 2011) and 34 pulsars, 12 out of which are MSPs and number among them a 'black widow', a good timer already selected to be part of the Parkes Pulsar Timing Array (Hobbs et al. 2009) and a MSP with a planet-companion. The discovered MSPs are reported in Keith et al. 2010, Bates et al. 2011, Keith et al. 2011 and Bailes et al. 2011.

3.2 The search for recycled pulsars

Recycled pulsars, especially if in relativistic binary systems (section 1.2.2), are powerful tools to investigate several physical and astrophysical issues, as explained in section 1.3; as already mentioned, the main aim of this thesis work was to find recycled pulsars hopefully suitable for a Pulsar Timing Array, i.e. a network of this kind of pulsars that would allow to detect Gravitational Waves in the nanoHz frequency range.

Since we expect to find this type of pulsars mostly at intermediate latitudes of our Galaxy, the present work focused on analysing the data of the mid-latitude survey.

In order to choose the opportune analysis method for reaching our goal, some steps had been necessary, as described in the following sections.

3.2.1 Standard search

The first step was to implement the HTRU pipeline (called HITRUN) for the standard search on the CYBERSAR supercomputer at the Cagliari Astronomical Observatory, in order to test that the steps common to all the search methods (RFI removal, de-dispersion, FFT) worked properly. CYBERSAR (Porceddu and D'Amico 2009) is a 75-node cluster, shared by different research groups in Cagliari; we had at our disposal 30 dual-core and dual-CPU nodes, so 4 CPUs per node, each node having 2 scratch disks with 4 GB RAM, and 250 GB of disk storage.

The first stage of our pipeline is RFI removal. First of all, we always remove from the data a list of channels that we know to be affected by RFI. Then, as described in section 2.2, we remove both 'bad' spectral channels and 'bad' time samples in the zero-DM (pre-dedispersion) data set, i.e. channels where we found periodicities (by an individual FFT of the data in each channel) with spectral S/N > 15, and time samples containing bursts of RFI (in the single zero-DM time series obtained summing the data in frequency) which are then replaced by random noise (fig.2.1). In the latter case our threshold for removal is 5σ. In order to remove longer duration RFI bursts, we repeat the process adding together 2, 4, 8, 16 and 32 consecutive time samples before creating the single time series.

The second stage of the pipeline is dedispersion (section 2.2); the chosen algorithm is a ν^{-2} dedispersion algorithm, since in our case the 'tree algorithm' would require a too high cost for linearising the data with respect to the gain in efficiency. Our DM trial values range from 0 to 1000 pc cm^{-3} with a total of 1196 dedispersion steps; in order to perform the process in a shorter time, our algorithm is a multi-threaded code, i.e. it carries out 4 dedispersion threads simultaneously, exploiting the feature of 4 CPUs per node in our cluster.

The next steps of our pipeline are those described in section 2.2.1. All the obtained dedispersed time series are searched for the presence of periodic signals by doing a FFT of each of them. From the resulting power spectra we remove the so-called *red noise* (i.e. a usually significant low-frequency noise due to fluctuations in the receiver and/or data acquisition systems) at frequencies below $\sim$ 10 Hz by subtracting a running mean and dividing by a running variance, that also yields the normalisation of the spectra. Then any peaks with a S/N $>$ 6 in the spectra are recorded for later sorting; to recover the power in the harmonics of narrow pulses, we also perform the harmonic summing of 2, 4, 8 and 16 harmonics. The SEEK[1] software is used for both the searching and the harmonic summing processes.

The candidate sorting is then performed by a software that is part of the PULSAR HUNTER[2] package, which groups detections with the same frequency into a single candidate and combines harmonically related candidates, yielding a manageable number of pulsar candidates promoted to the next analysis steps.

We then dedisperse and fold the raw data at the candidate DM and period for all the sorted candidates, by using the DSPSR[3] software; for the folding, we chose to use 32 subintegrations and 128 time bins.

The PSRCHIVE package (Hotan et al. 2004) is used at this point to remove the bad channels and then scrunch in frequency obtaining 16 subbands, and then to perform the optimisation procedure for each candidate, searching for the best values of both DM and period. These values are then used to produce diagnostic plots for each candidate.

The entire processing of the data of one beam, for an observation of lenght 540 s (mid-latitude survey), took about 3.5 hours on a dual processor dual core

[1] http://sigproc.sourceforge.net/seek/
[2] http://pulsarhunter.sourceforge.net/
[3] http://dspsr.sourceforge.net/

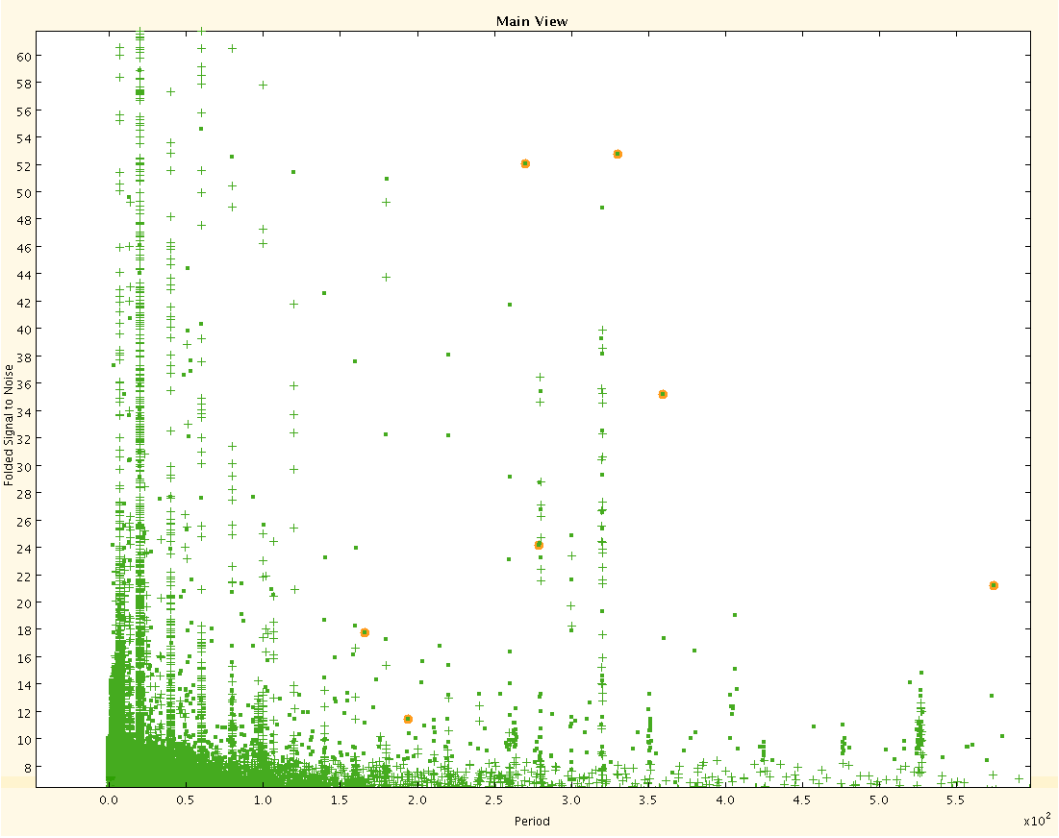

Figure 3.4: Example of the JREAPER main screen, where the candidates are plotted as points on the S/N vs period plane. Selecting a point, the candidate diagnostic plots are shown. The main screen is zoomed to show a period range up to 6×10^2 ms. The orange points are known pulsars. When a candidate has been visualised, the corresponding green cross turns into a green square.

server of the CYBERSAR cluster.

In order to select pulsar candidates to be re-observed (in fact, to confirm that they actually are pulsars we need to re-detect them), we visually inspect the candidates by using the JREAPER graphical tool (Keith et al. 2009), which is part of the PULSAR HUNTER package. In fig.3.4 it is shown its main screen, where the user can choose to display the candidates, as points on a plane, in different ways, for example with their S/N as a function of period, as in the figure, or the DM, and so on. This is useful because we can avoid to look at the diagnostic plots of many candidates just owing to their positions on the display: for instance, often RFI signals are disposed on a vertical column in the S/N vs period plane, or they have a very small DM value since they have a terrestrial

origin; hence we can ignore those points, saving time. Selecting instead a point on the plane, JREAPER will show the corresponding diagnostic plots, as shown in fig.2.9; the user can then make a decision about it and classify the candidate as a possible pulsar to be re-observed or not.

After implementing the pipeline on our cluster, I used it to analyse several survey tapes (each 800 GB LTO tape contains about 7 hours of survey for all the 13 beams of the MB receiver), in order to test if the codes and the cluster were working properly (and this also yielded some candidates for re-observation).

In order to check the performance of the standard search on accelerated MSPs, I then tested it on the millisecond pulsar (A) of the double pulsar system J0737−3039, that is the most exotic object (higly relativistic) known so far and that is made up of a 22.7 ms recycled pulsar and a 2.8 s ordinary pulsar in a close binary, with an orbital period of only 2.4 hr.

The result of the test was that the pulsar was not among the final candidates selected by the pipeline; therefore, looking for it in the list of the candidates yielded by the FFT step the pulsar was found with a spectral S/N lower than 6, that is the reason why it had not been selected in the sorting step. Hence by folding the time series by hand using the pulsar period resulting from the FFT, we obtained the output shown in fig.3.5, which is produced by *pdmp*, the optimisation software of the PSRCHIVE package. It can be noted that, since no corrections for the orbital motion have been performed, the signal trend in the subintegrations vs phase plot is parabolic, due to the orbital Doppler effect, as explained in section 2.2.2; as a consequence, the integrated profile is broadened and the S/N is highly reduced: in this case, in fact, the S/N of pulsar A is only ~ 7 (as we will see in section 3.2.3, removing the orbital motion the S/N will rise to ~ 27). Hence, even if the pulsar had been selected at the sorting step, it would be discarded at this stage, since in our analysis of the survey we inspect by eye only the candidates having S/N from 8 (or 9) up.

From what said, this test showed that the standard search does not allow us to find this kind of pulsar.

3.2.2 Simulations

In order to test more in depth the success of the standard search in finding out recycled pulsars in close binaries, in the mid-latitude HTRU survey, and hence

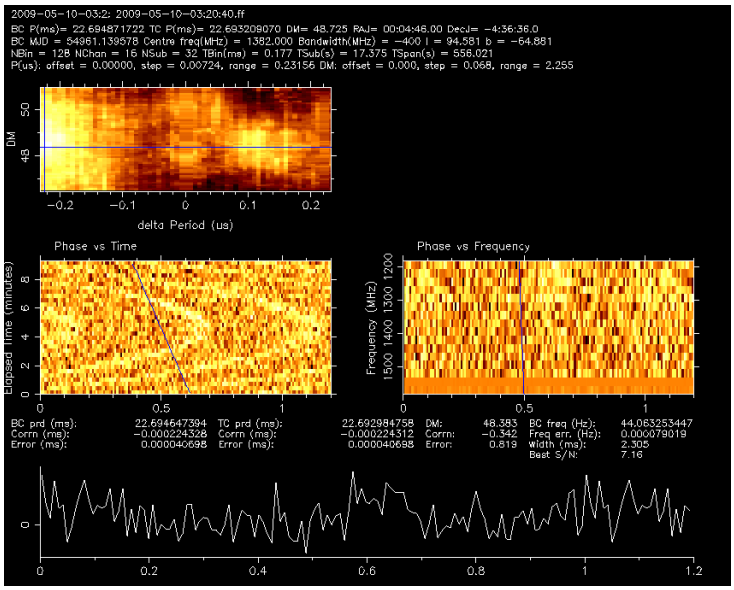

Figure 3.5: Output plot produced by *pdmp* of pulsar A of the double pulsar system J0737−3039, analysed without correcting for the orbital motion. From the top: the dependence of S/N on DM and on offset from the nominal period, the S/N vs pulse phase for 32 subintegrations (left), the S/N vs pulse phase in 16 sub-bands (right) and the integrated pulse profile (bottom).

to understand which processing method was the best to find different classes of them, I made some simulations, imagining to observe in our survey three different pulsar populations:

1. known binary pulsars;
2. the same sample as in point 1. but with each pulsar at a double and a triple distance from us;
3. fake samples of interesting close binaries with orbital period shorter than 2 hours and different kinds of companion (white dwarf, neutron stars, black holes) for the pulsar.

Essentially I simulated the selection effects of our survey, in order to see how many of the pulsars in the different populations would be detected with the standard search.

The results of the simulations therefore represent our possibilities to find recycled pulsars in close binaries without any corrections for the orbital motion.

Simulation 1: known binary pulsars

When I made this first simulation 56 Galactic MSPs with a measured flux density were known[4], 40 out of them were in binaries. One of the latter was pulsar A of the double pulsar system J0737−3039, tested before.

The first part of the simulation consisted in calculating the minimum flux that we can detect in the mid-latitude survey (T_{obs} = 540 s) for each MSP of the sample, at our observing frequency (1400 MHz). For this pourpose I used the sensitivity formula (eq. (2.1)), entering the values of the survey parameters as described in section 3.1 (except for the T_{sky}, calculated by the code TT408, which is based on the paper by Haslam et al. 1982, at 408 MHz and then scaled at 1400 MHz according to $\nu^{-2.7}$), and those of the pulsar parameters. For the pulse effective width, I used the corresponding values on the catalogue, except for those pulsars whose width value was absent; in the latter case, I used the eq. (2.3):

$$W_e = \sqrt{W_i^2 + (\beta\delta t)^2 + \delta t_{DM}^2 + \delta t_{scatt}^2},$$

giving a value equal to 15% P for the pulse intrinsic width W_i (where P was the period of the pulsar), $\beta = 2$, the sampling time $\delta t = 64\ \mu s$, δt_{DM} given by eq. (2.6) with $\delta\nu_{MHz} = 0.39$ (the width of our channels), and δt_{scatt} calculated by the code NE2001 (Cordes and Lazio 2002) at 1 GHz and then scaled at 1400 MHz according to $\nu^{-4.4}$.

For the pulsars of the sample that were in binary systems, I also had to consider the broadening ΔW of the pulse width due to the Doppler effect. I assumed that the orbital acceleration of the pulsar along the line of sight was constant during the observing time (since the latter was always much shorter than the orbital period), and used a code that from the orbital parameters of the pulsar, through the Kepler's laws, is able to calculate the minimum and maximum values assumed by the acceleration during the orbit, and also the corresponding values for the pulse broadening ΔW_{min} and ΔW_{max} (i.e. the 'best' and the 'worst' case respectively), according to the lenght of the observation (the longer the observation the larger the broadening, for a certain value of acceleration).

[4] *www.atnf.csiro.au/people/pulsar/psrcat*

Hence, for this kind of pulsars, instead of using W_e in the sensitivity equation I used W_{TOT}:

$$W_{TOT} \sim \sqrt{W_e^2 + \Delta W^2}, \tag{3.1}$$

calculated for both ΔW_{min} and ΔW_{max}.

The last step was comparing the calculated minimum detectable flux to the true pulsar flux taken from the catalogue (and scaled with $\nu^{-1.7}$ if it was the flux at 400 MHz instead of that at 1400 MHz) for each MSP: if the pulsar flux was higher than the minimum flux, this meant that the pulsar was detectable in our mid-latitude survey without any corrections for the orbital motion. The comparison was made for both the best and the worst smearing case.

One of the results was that $\sim 20\%$ of the known sample was undetectable just for flux reasons, since the true pulsar flux was very low; but the important result of this simulation for our aim is that, considering the pulse smearing for the binary pulsars, the double pulsar was undetectable in about 3/4 of its orbit, because the smearing was greater than 50% for a 9 min observation.

Other 4 binary pulsars could have been undetectable due to the pulse broadening at least in a fraction of the orbit, but they were bright enough to be visible anyway despite the smearing.

Simulation 2: double and triple distance

Since, owing to the features of our survey, we are able to detect MSPs with higher DMs and hence at larger distances, I did the same test as in the first simulation but doubling and tripling the distance d of each pulsar of the known sample.

First of all I changed the flux (that is $\propto 1/d^2$), but keeping DM and smearing corresponding to the real distance, in order to test just the effects of the variation in distance.

One of the results was that $\sim 50\%$ of the known sample became undetectable at double distance and $\sim 75\%$ at triple distance; 2 of the 4 binary pulsars detectable at the real distance despite the Doppler smearing disappeared at double distance and minimum Doppler smearing, while the other 2 at triple distance and only at maximum Doppler smearing.

Again, however, the important result for us is that the double pulsar was undetectable in more than 3/4 of its orbit at double distance, while became undetectable in *any* fraction of the orbit at triple distance.

Using DM and smearing corresponding to the new distances, the results did not change; hence in this case the decisive factor is the decrease in flux due to the increased distance.

Simulation 3: fake samples

In the third simulation I did the same analysis again, but on three different samples of fake pulsars, which I created through the *psrpop*[5] code and then put in a close binary system with a white dwarf (WD), a neutron star (NS) or a black hole (BH) companion (taking into account the pulsar evolution described in section 1.1.7). Hence this simulation was divided into three parts, as described below.

a) MSP + WD case

In this part I created a sample of 50,000 MSPs with the following distributions (from Cordes and Chernoff 1997 and Lorimer et al. 2006):

- luminosity: power law $\mathrm{dlog}N/\mathrm{dlog}L = -2$, and with cutoffs $L_{min} = 1.1$ and $L_{max} = 1600$ mJy kpc^2;
- spin period: power law $\mathrm{d}N/\mathrm{d}P \propto P^{-2}$, with a lower cutoff at 1 ms and an upper one at 30 ms;
- radial distribution: Gaussian with $\sigma = 6.5$ kpc;
- height distribution: exponential with z-scale height = 0.5 kpc.

First of all I checked if these pulsars would have been detectable in the mid-latitude survey if they were isolated, repeating the same test on the flux described in the first simulation and checking if the pulsar position fell in the sky area covered by the survey. The result was that $\sim$1% out of the 50,000 pulsars would have been detectable.

[5] http://psrpop.sourceforge.net

Then I placed the detectable ones in close ($P_{orb} \leq 2$ hr) binary systems with a WD companion, using a flat distribution for the orbital period ranging from 0.5 hr to 2 hr with a step of 15 mins; for the mass of the companion, I also used a flat distribution ranging from 0.01 to 1.2 $M_\odot$ but considering three intervals:

- $0.01 - 0.05$ $M_\odot$ (step: 0.01 $M_\odot$): low mass WDs;
- $0.1 - 0.3$ $M_\odot$ (step: 0.05 $M_\odot$): typical He - WDs;
- $0.8 - 1.2$ $M_\odot$ (step: 0.1 $M_\odot$): typical CO - WDs.

At this point I checked the detectability of the pulsars in these binary systems using the usual method, considering the maximum orbital smearing in the formula for the minimum flux (hence my estimates are lower limits to the visibility of the binaries with the standard analysis).

The result was that only $\sim$ 10% of the pulsars detectable if isolated remained detectable without any corrections for the orbital motion when placed in this kind of system. About 5% of these detectable binary pulsars had a He - WD as a companion, none had a CO - WD and all the others had low mass companions.

b) PSR + NS case

In the second part of the simulation I created a sample of 50,000 recycled pulsars with the same distributions as the previous case but the spin period, for which I used a Gaussian with a mean of 85 ms and $\sigma = 65$ ms to simulate a population of mildly recycled MSPs, and the height distribution, for which the z-scale height was 350 pc.

In this case $\sim$ 2.4% of 50,000 would have been detectable by the mid-latitude survey.

Then I placed the detectable ones in a close binary with a NS (creating a double NS system), hence with a companion mass of $M_c = 1.4$ $M_\odot$, using the same orbital period distribution as in the MSP + WD case and considering again the maximum orbital smearing.

In this way I obtained that $\sim$ 17% of the previously detectable pulsars remained detectable by the standard pipeline when put in a double NS system.

	Case (a) MSP + WD	Case (b) PSR + NS	Case (c) PSR + BH
Companion mass	$0.01 - 1.2$ $M_\odot$	1.4 $M_\odot$	$5 - 10$ $M_\odot$
Detectable pulsars*	10%	17%	60%

*(% with respect to the detectable isolated pulsars of the corresponding case)

Table 3.3: Percentages of pulsars still detectable by the standard search in the mid-latitude survey, with respect to those detectable if isolated in the corresponding case, when placed in three different kinds of binary system with $P_{orb} \leq 2$ hr and considering the maximum orbital smearing.

c) PSR + BH case

In the last part of the simulation the fake sample was made up of 50,000 ordinary pulsars, with distributions:

- luminosity: power law with slope $= -0.6$, $L_{min} = 0.1$ and $L_{max} = 1000$ mJy kpc^2;
- spin period: Gaussian with a mean of 550 ms and $\sigma = 450$ ms;
- radial distribution: Gaussian with $\sigma = 6.5$ kpc;
- height distribution: exponential with z-scale height = 350 pc.

This time $\sim 3.5\%$ of 50,000 would have been detectable by the mid-latitude survey if isolated.

Placing the detectable ones in a close binary with a BH, using a flat distribution for the BH mass from 5 to 10 $M_\odot$ (with a step of 0.25 $M_\odot$), the same orbital period distribution as in the previous cases and the maximum orbital smearing, the result was that $\sim 60\%$ of the pulsars detectable if isolated were still detectable by our standard pipeline if put in this type of binaries.

The results of the three parts of this simulation are summarised in table 3.3 and show that, without removing the orbital motion effects, a significant fraction of the potentially detectable pulsars in orbit shorter than 2 hours would be missed.

Global results

The global result of all the simulations is that the standard search is able to detect in the mid-latitude survey only a small fraction of pulsars in close binaries, due to the effects of the orbital motion.

As a consequence, in our case the data analysis method suitable to find this type of pulsars is the acceleration search, which removes those effects.

3.2.3 Acceleration search

Once the opportune analysis method had been established, I worked on implementing the pipeline for the acceleration search on the CYBERSAR supercomputer; I did the same work also on the Green Machine supercomputing facility at Swinburne University, in Melbourne, which consists of 145 Dell Power Edge 1950 server-class nodes, each with dual quad-core Clovertown processors (2.33 GHz), 16 GB RAM and 1 TB of disk storage.

The difference in the acceleration pipeline with respect to the standard one is that, after doing the FFT of the single de-dispersed time series for a certain value of the trial DM, we use the code PRESTO[6] to perform the acceleration search in the Fourier domain with harmonic summing (up to 16 harmonics, in our case), yielding the candidates for the sorting step.

To be precise, we used the routine *accelsearch* in PRESTO, which performs the correlation technique described in section 2.2.2 and hence searches the data covering portions of the $\nu - \dot{\nu}$ plane; this technique is suitable for the mid-latitude observations since even in case of quite close binaries the observing time would be much shorter than the orbital period.

In our pipeline we chose to give as the maximum number of bins drifted by the frequency during the observing time, due to the orbital motion, $z_{max} = 100$ (i.e. +100 and −100 around the frequency), while for the search in frequency $\nu_{min} = 1$ Hz and $\nu_{max} = 10000$ Hz (both for the *highest* harmonic). The acceleration step is calculated through eq. (2.18).

The other steps of the pipeline are the same as the standard pipeline, from the candidate sorting to the visual inspection (section 3.2.2).

After implementing the pipeline, I did several tests with different known

[6] http://www.cv.nrao.edu/~sransom/presto/

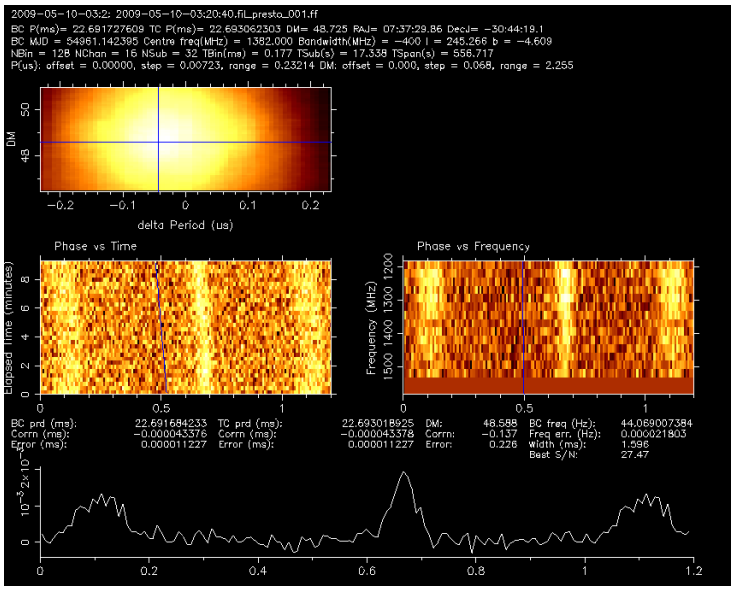

Figure 3.6: Output plot produced by *pdmp* of pulsar A of the double pulsar system J0737−3039, analysed with the acceleration search, which removed the orbital motion correcting for an acceleration value of 222.5 m s^{-2}. From the top: the dependence of S/N on DM and on offset from the nominal period, the S/N vs pulse phase for 32 subintegrations (left), the S/N vs pulse phase in 16 sub-bands (right) and the integrated pulse profile (bottom).

binary pulsars; for example, testing the double pulsar I obtained that the pulsar was present among the final candidates selected by the pipeline, and the orbital motion was correctly removed by the code (for an acceleration value of 222.5 m s^{-2}) as shown in fig.3.6, where the signal trend in the subintegrations vs phase plot is linear and the integrated profile is good, with a S/N of $\sim$ 27 (to be compared with the result without orbital motion corrections, section 3.2.2 and fig.3.5, where the S/N was only $\sim$ 7). Hence, this pulsar would not be missed by our pipeline anymore.

Once the code had been tested, the processing of the mid-latitude data with the acceleration search could start. To analyse one beam, on one cluster node, the pipeline takes a long time, from $\sim$ 20 to $\sim$ 24 hours, depending on the amount of RFI present in the data; in case of large quantities of RFI, the number of candidates yielded per beam can be very high ($>$ 600).

Due to these reasons, despite the use of supercomputers the percentage of the survey processed so far both in Cagliari and in Melbourne is $\sim$ 3% of the

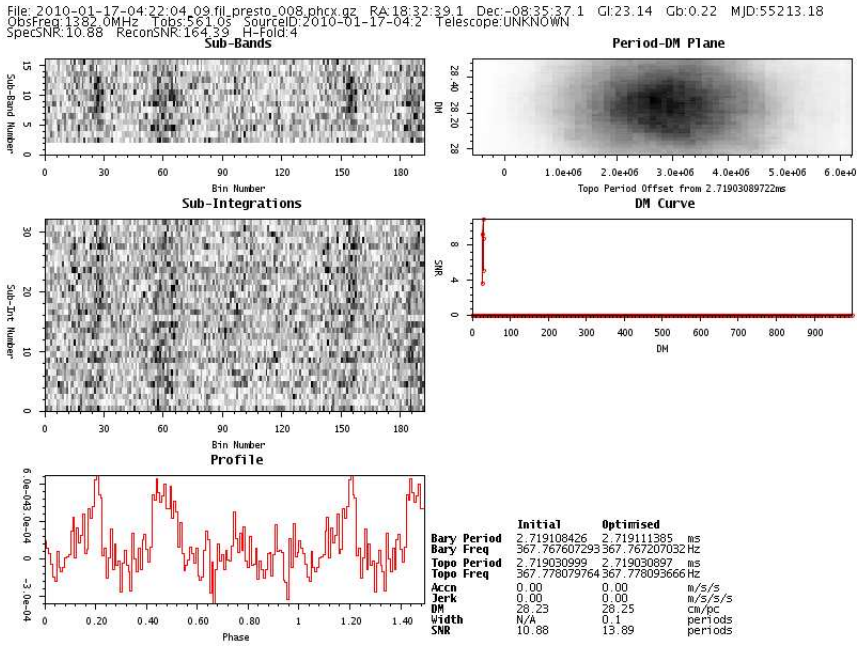

Figure 3.7: Discovery plot of the 2.7 ms pulsar J1832−0835.

total (which, for the mid-latitude survey, is 95,056 beams, i.e. 7312 pointings), and about one third of the produced candidates has been inspected by eye.

The main result of this work is that one of the candidates that I proposed to date for re-observation has been confirmed to be a 2.7 ms pulsar with some interesting features, as described in the next section, that could make it suitable for a Pulsar Timing Array. Ironically, despite it has been identified through the acceleration search, this pulsar is isolated (the only isolated MSP found to date in the HTRU survey!).

3.2.4 PSR J1832−0835

In fig.3.7 it is shown the discovery plot of the pulsar J1832−0835, found with the acceleration search and having a spin period of only 2.7 ms, no orbital acceleration and an interesting integrated profile made up of three peaks.

Since in the context of timing (section 1.2 and subsections) usually MSPs are very stable (except for those very few examples with timing noise), and furthermore this pulsar shows a sharp profile with several components that makes more efficient the cross-correlation from which we obtain the TOAs of the pulses, J1832−0835 could turn out to be a good timer for the PTA project.

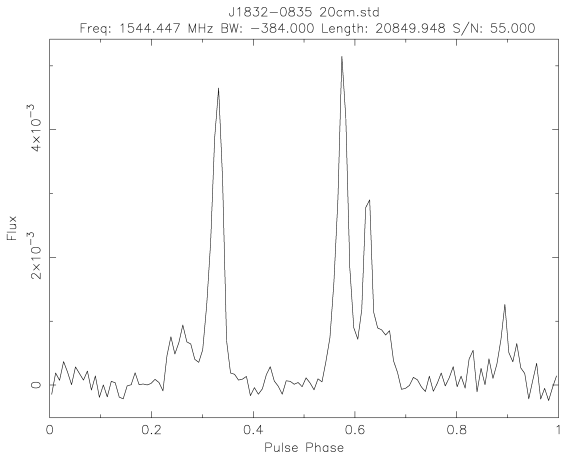

Figure 3.8: Standard profile used for the timing of the 2.7 ms pulsar J1832−0835.

Moreover, since this pulsar typically has a S/N of about 10-15 for an observing time of 15-20 mins (the variation of the S/N with the observations is due to the fact that the pulsar scintillates slightly, which confirms that the DM value is quite low), it is bright enough to allow to obtain quite small TOAs uncertainties with a reasonable integration time; hence, although it is slightly fainter than the average of the pulsars of the Parkes PTA, owing to these features it would be suitable for being part of a timing array. Due to its position in the sky, it could be a candidate to be observed with the Sardinia Radio Telescope (SRT), which is part of the European PTA.

The timing of this pulsar started at the Parkes Radio Telescope and is being performed now at Jodrell Bank, by using the 76-m Lovell Telescope. In fig.3.8 the standard profile used for the timing is shown. To date, we have only 7 months of timing observations, that are not enough to be sure about its actual intrinsic stability.

However, so far that seems to be quite good: the timing solution in fig.3.9 has a root mean square of the post-fit residuals between measured and model arrival times of ~ 1.8 μs over 7 months of observation (the TOA uncertainties σ_i in eq. (1.35) are ~ 2.5 μs on average over an observation 15-20 min long), that corresponds to a rotational stability[7] of 9.4 $\times 10^{-14}$.

[7]This value of the stability has been obtained as the ratio between the RMS and the time span, and is a figure of merit, that is a good estimate since the time span is short. With a longer time span we will be able to calculate it more precisely by using the σ_z parameter from Matsakis et al. 1997.

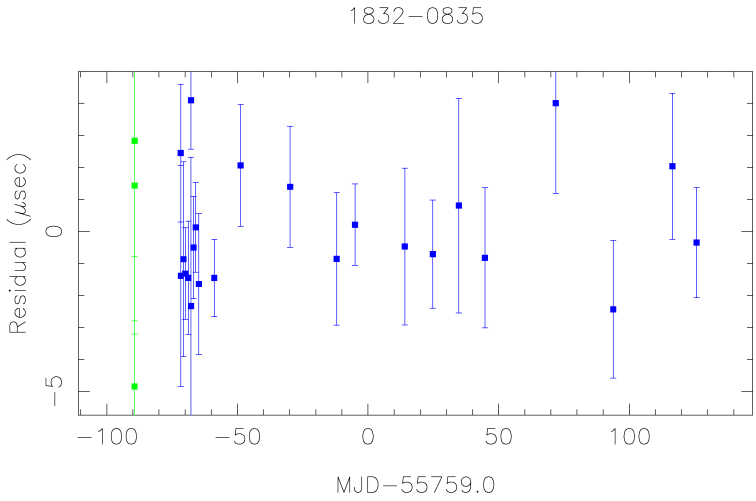

Figure 3.9: Present timing residuals of the 2.7 ms pulsar J1832−0835 for a data span of 7 months. The green points correspond to timing observations carried out by using the Parkes Radio Telescope, while the blue points by using the Lovell Radio Telescope at Jodrell Bank.

The question now is if it will maintain such a good stability over longer data span (of the order of years).

In table 3.4 the values of the pulsar parameters obtained through the timing are summarised; since the data span is shorter than one year, the errors (at 2σ) in the determination of right ascension, declination and frequency derivative, which co-vary, are underestimated.

This pulsar will be published soon in a paper of the HTRU collaboration currently in preparation at the Cagliari Astronomical Observatory, together with 4 binary MSPs: J1431−5736, J1546−4552, J1825−0322 and J2236−5526, as soon as we have one year of timing data for all the 5 pulsars.

The timing of these binary pulsars is mainly being performed at the Parkes Radio Telescope, except for J1825−0322, for which it is mainly being performed at Jodrell Bank. In fig.3.10 they are shown the standard profiles used for the timing of the 4 binary pulsars. Their timing solutions are illustrated in figures 3.11 (J1431−5736), 3.12 (J1546−4552), 3.13 (J1825−0322), and 3.14 (J2236−5526), while in table 3.5 a summary of the values of the parameters obtained through the timing (with errors at 2σ), for all the four pulsars, is reported. For the pulsar for which we do not have approximately one year of timing data, no value is reported for the first period derivative because it is

Parameter	Value
Right ascension (J2000)	$18^{\rm h}32^{\rm m}27^{\rm s}.5949(8)$
Declination (J2000)	-08°36′54″.99(5)
Spin frequency (ν)	367.767115561(1) Hz
Frequency derivative ($\dot{\nu}$)	-8(4) $\times 10^{-16}$ Hz s^{-1}
Epoch	MJD 55759
Dispersion measure (DM)	28.185(1) pc cm^{-3}
Number of TOAs	27
RMS of fit	1.751 μs
TOA range	MJD 55670-55885

Table 3.4: Positional and rotational parameter values of J1832−0835 obtained for a data span of 7 months. The two-sigma errors on the last digit of each parameter are reported in parentheses. The last three rows report the number of TOAs, the root mean square of the timing residuals and the time span covered by the timing data used to obtain the reported ephemerides.

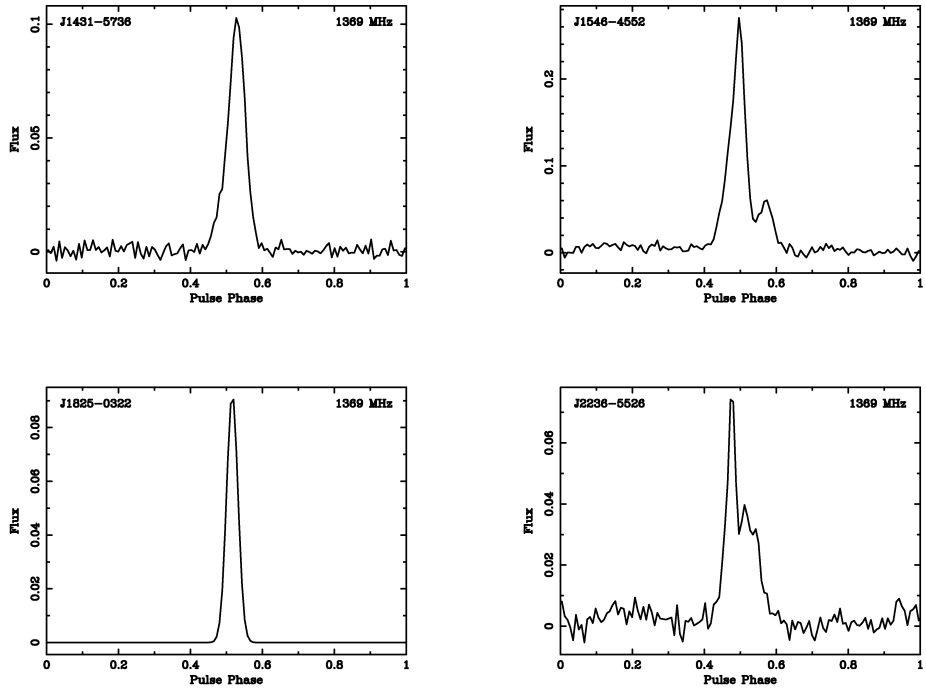

Figure 3.10: Standard profiles used for the timing of the 4 binary MSPs J1431−5736, J1546−4552, J1825−0322 and J2236−5526.

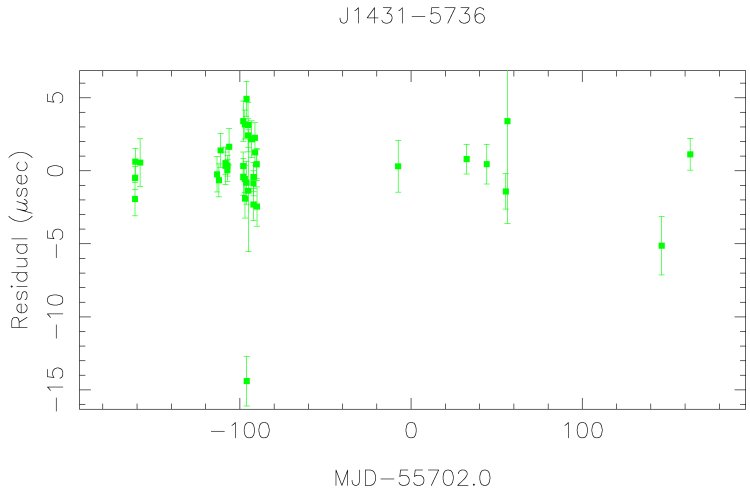

Figure 3.11: Present timing residuals of the 4.1 ms pulsar J1431−5736 for a data span of 325 days.

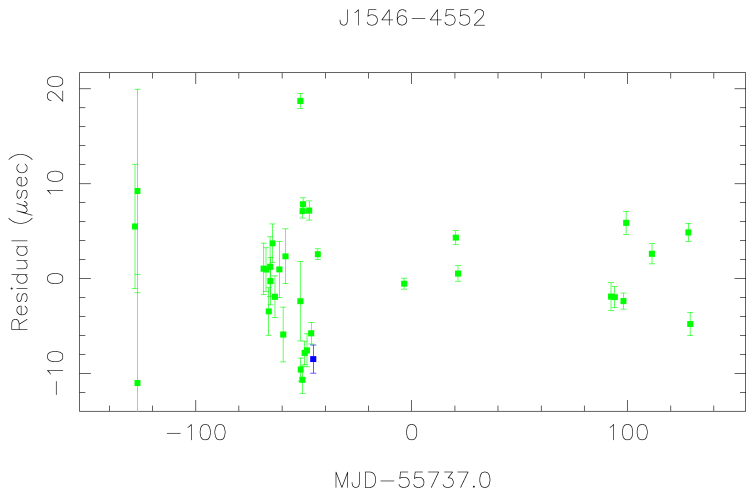

Figure 3.12: Present timing residuals of the 3.6 ms pulsar J1546−4552 for a data span of 257 days. The green points correspond to timing observations carried out by using the Parkes Radio Telescope, while the blue point by using the Lovell Radio Telescope at Jodrell Bank.

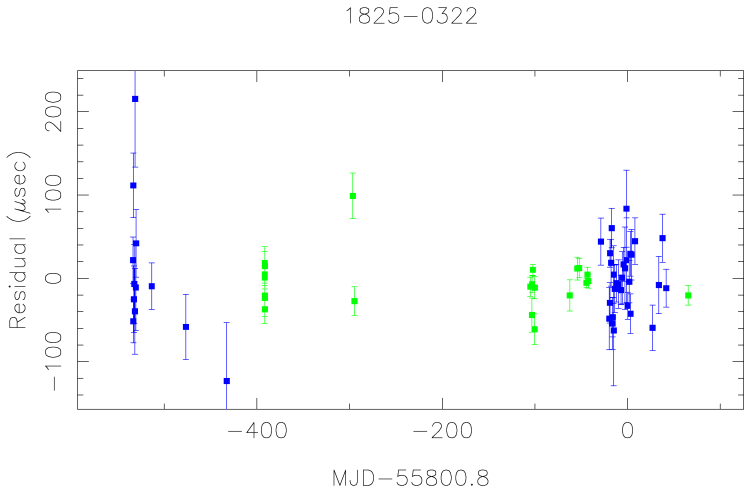

Figure 3.13: Present timing residuals of the 4.6 ms pulsar J1825−0322 for a data span of 599 days. The green points correspond to timing observations carried out by using the Parkes Radio Telescope, while the blue points by using the Lovell Radio Telescope at Jodrell Bank.

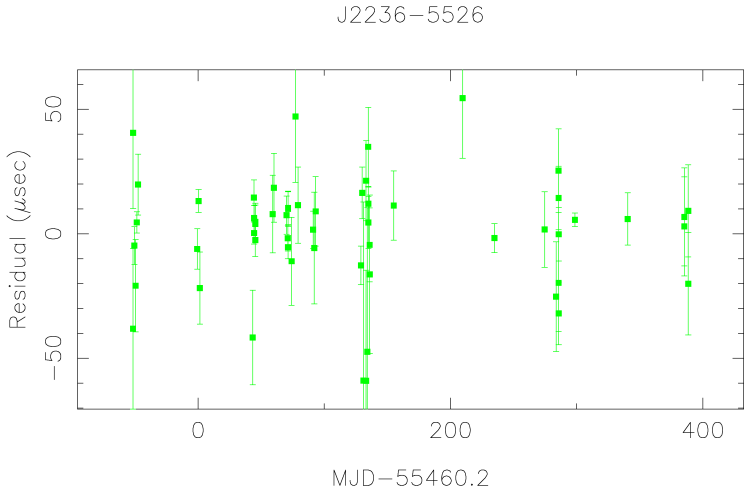

Figure 3.14: Present timing residuals of the 6.9 ms pulsar J2236−5526 for a data span of 458 days.

PSRJ	J1431-5736	J1546-4552	J1825-0322	J2236-5526
RAJ	14:31:03.4959(4)	15:45:55.948(3)	18:25:55.9523(5)	22:36:51.8515(6)
DECJ	-57:40:11.657(9)	-45:50:37.52(7)	-03:19:57.55(3)	-55:27:48.837(7)
P0 (ms)	4.1105439566762(17)	3.57528861788(6)	4.553527919749(13)	6.90754939267(2)
P1	5.7(11)E-21	–	7.0(7)E-21	9.5(15)E-21
EPOCH (MJD)	55702.00000	55737.00000	55800.82910	55460.19044
DM (pc/cm^3)	131.39(3)	68.390(8)	119.47(4)	20.09(4)
PB (days)	2.726855837(18)	6.20306488(13)	52.6305024(15)	12.6891870(3)
A1 (lt-s)	2.2698868(13)	3.846905(5)	18.266400(12)	8.775870(8)
T0 (MJD)	55461.0699555(10)	55611.38(16)	55805.06(5)	55472.56(7)
ECC	–	1.1(2)E-05	0.000194(1)	4.91(18)E-05
OM (deg)	–	220(9)	93.2(3)	351(2)
SPAN (days)	325.128	257.172	599.195	457.666
RMS (μs)	2.134	6.295	28.389	9.753

Table 3.5: Positional, rotational and orbital parameters for the four binary millisecond pulsars soon to be published together with the isolated MSP J1832-0835 discovered in this work. The two-sigma errors on the last digit(s) of each parameter are reported in parentheses. For the pulsar for which we do not have approximately one year of timing data, no value is reported for the first period derivative because it is either unconstrained and/or covaries with the positional parameters. The last two rows report the time span covered by the timing data used to obtain the reported ephemerides and the root mean square of the timing residuals.

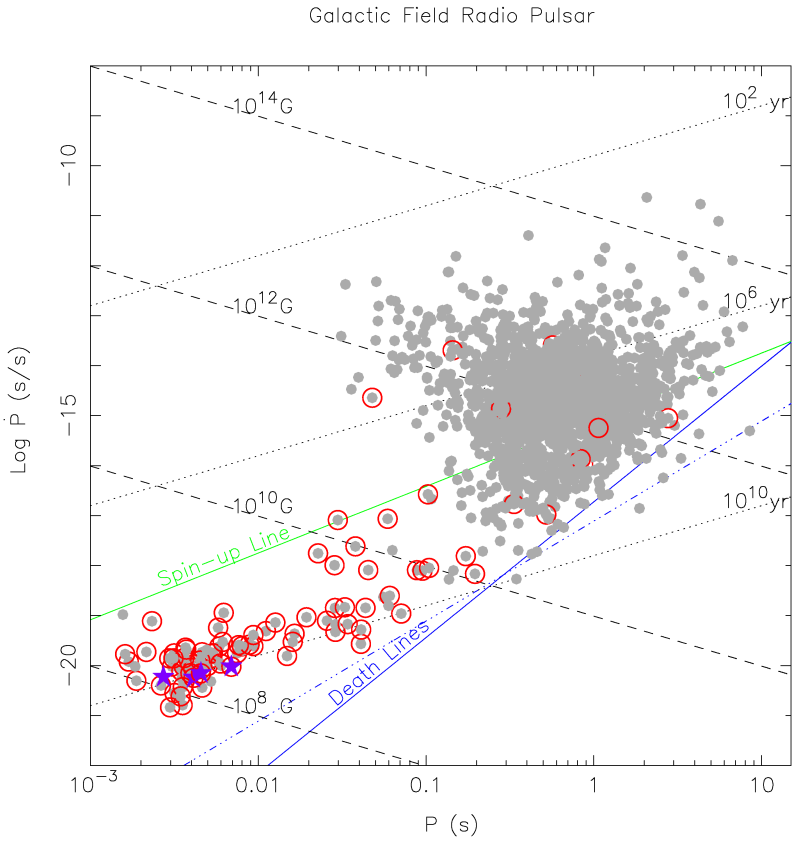

Figure 3.15: $P - \dot{P}$ diagram showing four of the five MSPs described in the text, represented by the stars. The grey points are the Galactic radio-pulsars in the pulsar catalogue, while the circles represent the binary pulsars. The only star not surrounded by a circle, since isolated, is J1832−0835, for which the determination of the period derivative is still uncertain due to the fact that the data span is much less than one year.

either unconstrained and/or covaries with the positional parameters.

The rotational stabilities of the 4 binary MSPs are respectively 7.6 $\times 10^{-14}$ (J1431−5736), 2.8 $\times 10^{-13}$ (J1546−4552), 5.5 $\times 10^{-13}$ (J1825−0322) and 2.5 $\times 10^{-13}$ (J2236−5526). Hence J1431−5736 seems to be slightly better than J1832−0835, whose stability is 9.4 $\times 10^{-14}$, but this is due to the fact that we have more data for J1431−5736 than for J1832−0835 and therefore we could obtain a timing solution for the former by using only the most recent data, and of the highest quality.

Fig.3.15 shows the $P - \dot{P}$ diagram with four of the five mentioned MSPs (those for which the period derivative has been estimated).

3.2.5 Issues

As already mentioned, the acceleration pipeline presents some issues.
In fact, it requires a very long computing time to analyse one beam, about 1 day, that is 6-7 times the time required by the standard search; this is essentially due to the long time necessary to PRESTO to perform the acceleration search. Considering, as already said, that the total number of beams in the mid-latitude survey is 95,056, it is clear why only a small percentage of the survey has been processed so far.

Furthermore, the code yields a huge number of candidates for visual inspection, usually of between 300 and 600 per beam but even more depending on the amount of RFI in the data. For a large-scale survey as HTRU, this means a number of candidates that is not human manageable, even inspecting only the candidates with $S/N > 9$.

The computing time problem is going to be solved owing to the implementation of the version for graphic processing units (GPUs) of the acceleration pipeline, that is now underway at the Swinburne University and that will be many times faster than the version for CPUs.

For the candidate problem, I worked on an automation approach by the use of an Artificial Neural Network to distinguish between pulsars and non-pulsars among the candidates; this work and its results will be described in chapter 4.

Chapter 4

A Neural Net approach to the candidate problem

As mentioned in the previuos chapter, one of the big issues I dealt with during the analysis of the HTRU survey data has been the huge number of candidates to be visually inspected yielded by the pipeline. This was due to both the mass of data and the strong presence of Radio Frequency Interference (RFI) in it, which produced lots of false pulsar candidates (most of the total candidates). Such an issue is common to all the large-scale recent surveys and will be critical also for the future (for example for surveys carried out with the Square Kilometer Array); therefore a solution is urgently necessary.

In this chapter I will describe a new approach to the problem consisting in the use of Artificial Neural Nets (ANNs), in an attempt to make a machine able to discern between pulsars and non-pulsars among the candidates. The innovations with respect to previous similar attempts are in our goal and our method. In fact, in addition to the potential for the discovery of new pulsars, we aimed to produce a large and novel set of classified candidates that can be used to better train future machine learning algorithms; to reach this aim, we used a method consisting of (i) training and then applying to all the HTRU survey candidates a 'pseudo-committee' of ten ANNs, rather than a single neural net, to explore the consequences of random number generation, and (ii) exploring the use of different ratios of pulsars and non-pulsars in the training data set of the nets, and using this ratio to reduce the rate of RFI, among the candidates, uncorrectly classified as pulsars by the ANNs (although this also led to an increase in the rate of pulsars wrongly classified as non-pulsars, however in

an acceptable percentage). Owing to this method, the number of candidates selected by the nets, and hence to visually inspect, reduced to something that is human manageable. This set could hopefully include new pulsars to be discovered; furthermore, once the inspection is completed, the set will be used to produce the mentioned new training set by which improving, in different possible ways, the training of the nets.

This work has been carried out at the Swinburne University of Technology (Melbourne, Australia) under the supervision of Dr. Willem van Straten, and is going to be reported in a paper of the HTRU collaboration (Milia, van Straten et al., in preparation).

4.1 The RFI problem

As described in section 3.1, the HTRU survey, once it is concluded, is expected to have yielded about a Petabyte of digital data. Such a mass of data brings several practical issues with it. Besides the problems of storage and data processing, respectively solved by using several hundred 800GB LTO tapes and 3 supercomputers, a further big issue is the huge amount of pulsar candidates resulting from our processing pipeline.

In fact, another key matter that must be taken into account is the increasing problem of Radio Frequency Interference (RFI). Due to the extraordinary development of modern technologies and the consequent increase, both in number and power, of satellites, radars, telecommunication systems and electronic devices which emit signals in the radio frequency band, radioastronomers have to face the constant presence of a large amount of RFI in their data. This can both cover the weak signals coming from pulsars and behave like pulsars, i.e. showing periodicity that is typical of the latter, thereby incrementing the number of candidates that must be evaluated in the search for real pulsars. Although the first step of the HTRU processing pipeline is RFI removal, which removes both ‘bad’ spectral channels and ‘bad’ time series (see the first paper of the HTRU series, Keith et al. 2010, and section 2.2), the presence of RFI among the pulsar candidates is nevertheless overwhelming.

The combination of the RFI problem with the amount of data collected so far yielded more than 19 million candidates from the beginning of the HTRU

survey, only for the mid-latitude part analysed with the standard pipeline. Since all these candidates need to be visually inspected to decide if they are:

1. possible new pulsars (to be re-observed to confirm the identification);
2. already known pulsars;
3. non-pulsars (noise, RFI),

it is evident that such a task is becoming impossible, at least in a 'reasonable' time scale. In fact, making a decision about a candidate can require from a fraction of a second to several minutes, depending on how clearly it belongs to one of the three mentioned categories; this means that the time necessary to look at all the candidates from the HTRU survey can range from something of the order of a year to an entire human lifetime or even more. As a consequence, a better approach is urgently required.

4.2 Previous automation attempts

An interesting possibility to face the candidate problem is to exploit machine resources, making the selection process as automatic as possible. This means that the machine must be able to recognise a possible pulsar, distinguishing it from noise and RFI signals.

As a first step towards this purpose Keith et al. 2009, using the candidates resulting from the Parkes Multi-beam Pulsar Survey reprocessing made in 2002 (Faulkner et al. 2004), implemented in the graphical tool JREAPER a ranking method based on assigning to each candidate a numerical score between 0 and 1 (the latter indicating a pulsar-like candidate), according to the results of several tests performed on the candidate features (like spin period stability and variation of the signal-to-noise ratio, S/N, with trial dispersion measure, DM). This approach yielded very good results, with the discovery of 28 previously unknown pulsars. One problem with the approach outlined in Keith et al. 2009 is that it is able to identify only typical pulsars, due to the particular settings of the scoring functions (for example, the spin period stability required to score high is not good for binaries, that can face a big variation of their spin period with time as they go along their orbit).

Reprocessing the same survey, Eatough et al. 2010 (see also Eatough 2009; Eatough et al. in prep.) used a development of the scoring method consisting in a *machine learning* approach (for a good review about the current state of data mining and machine learning in astronomy, see Ball and Brunner 2010), more specifically in the employment of an *Artificial Neural Network* (ANN), i.e. a software that tries to reproduce the decision-making faculty of the neural net of the human brain. The method consists in training an ANN to distinguish between pulsars and non-pulsars according to some criteria (that is the values of suitable parameters, called 'input scores'), using a training dataset made up of known pulsars and certain non-pulsars. The trained net is then launched on the pulsar candidates, among which it makes a distinction indicating, through the so-called 'output scores', which ones are more probable to be real pulsars (and that will be then double-checked by eye) and which ones RFI or noise. From the results obtained by Eatough et al. 2010, this approach seems to be quite encouraging, both for the discovery of a further new pulsar and since, testing the trained ANN on a random sample of about 2.5 million candidates resulting from the survey processing, the net selected as possible pulsars only 13000 out of them, among which 92 per cent of the known pulsars present in the testing sample was found (with an average of one pulsar every 30 candidates selected by the net). The fact that some pulsars have not been recovered can be due, according to the authors, to different reasons, such as a too small number of pulsars of a certain type in the training set, abnormal plots (from which the candidate input scores are taken) generated by the used search software for a kind of pulsars, and finally the so-called 'over-training' of the net. The latter means that the ANN could have been excessively trained with RFI and/or noise signals, present in the training set in much greater quantities than pulsars; as a consequence, the net classification ability might be biased towards the class of non-pulsars, i.e. the ANN might have considered the missing pulsars as RFI or noise, since it is more trained to recognise the latter.

Despite all these problems, in part related to the fact that they were the first attempts to apply an ANN to the pulsar candidate selection, the machine learning approach looks very promising. For this reason, this method has been already used for a couple of years for the HTRU survey, submitting to an ANN the portion of the survey data processed by the HYDRA supercomputer in

Manchester (Bates 2010). In this case, the net has been trained to discern on the base of the values of 22 parameters, conveniently chosen so that they well represent the differences between pulsars and non-pulsars, for example in pulse profile, DM curve, best candidate parameters (period, S/N, and DM), and several other features (see section 4.5.1 and subsections for a more detailed description). The so trained ANN rejects as non-pulsars 99.7 per cent of the candidates, making the quantity of data to be inspected by eye manageable. Owing to the use of this net, several new pulsars, among which some MSPs, have been discovered. It has also been confirmed a strong dependence of the new and the recovered pulsars on the kind of known pulsars present in the training set. From this it can be inferred the need to include in the latter as much variety of pulsars as possible, although for some kind of pulsars, like the MSPs, the known available sample is quite small; and attention should be paid to avoid the over-training of the net in favour of a particular sub-class.

4.3 Outline of our approach

Taking cues from the aforementioned works, we started to set up a neural net to be used for the candidate selection; but, unlike the previous cases, we chose a different approach.

First of all, besides the discovery of new pulsars, our goal was to produce a data set to use in the future in order to improve the training of the net. For this reason, rather than making only a distinction between pulsars and non-pulsars, we decided to subdivide the 'positive' output of the ANN (i.e. the candidates indicated as pulsars), when inspecting it by eye, in three categories: (i) pulsars, (ii) non-pulsars and (iii) 'confusion', where the latter are the candidates which it is difficult to make a decision about, and that therefore might equally be pulsars or non-pulsars. These categories can then be used in different ways (still to be tested) to create an improved training set. If we choose to keep trying to recognise pulsars, i.e. we still want to train the net to distinguish between pulsars and non-pulsars, we can use the obtained pulsar and non-pulsar categories for the training set, without including the confusing sources, in order to avoid errors in the training. If we decide to change our approach instead, training the net to recognise RFI, i.e. to distinguish between 'RFI'

and 'non-RFI', it might prove useful to include both pulsars and confusion in the non-RFI category for the training set, hence giving the neural net more examples of 'what might not necessarily be RFI'; with this method, both the number of pulsars wrongly identified as RFI and the number of RFI wrongly classified as pulsars by the net should reduce (see section 4.5.4).

Another new feature of our approach, with respect to previous techniques, consists in the method we used to achieve our aim, that was made up of two parts:

1. Rather then train a single neural net, we trained and then applied to all the HTRU survey candidates a 'pseudo-committee' (see section 4.5.3) of ten ANNs, in order to overcome the intrinsic randomness of neural nets (only the candidates indicated as pulsars by all the 10 ANNs are currently being inspected by eye);

2. We explored the use of different ratios of pulsars and non-pulsars in the training set, and used this ratio to fine tune the percentage of non-pulsars misclassified as pulsars by the ANNs, reducing it to a quite small value (although this also produced an increment in the rate of pulsars wrongly classified as non-pulsars, however in an acceptable percentage).

Moreover, the number of parameters used by the nets has been increased to 27 (the same 22 used in Manchester and 5 new, see section 4.5.1 and subsections).

This method led to a huge reduction of the number of candidates to visually inspect among all the HTRU survey candidates. The inspection is still in progress, with both the aims of finding new pulsars and classifying the candidates in the pulsar, non-pulsar and confusion categories; once a good training set has been obtained, the intention is to make it public, so that even the people who do not have at their disposal a quantity of data sufficient to train an ANN can be enabled to do it.

One of the problems we dealt with was the presence of some outliers among the scores of many parameters of the training set (for both pulsars and non-pulsars); outliers are values that lie well outside the range in which the values of the scores lie for all the other elements of that class, and they may be correct or might also correspond to bad data and/or as-yet-undetected software bugs. Their presence aroused the suspicion that the entries with such outliers might

be eschewing the training process, resulting in inaccurate neural nets. Section 4.5.1 describes the tests performed to investigate this issue.

Before describing this work more in detail, a brief introduction to the ANN concepts will be given; to study ANNs in depth see for example Bishop 1995.

4.4 Artificial Neural Networks

As already said, ANNs are softwares that try to mimic the brain and its decision-making faculty; although they are not 'intelligent' in the true sense of the word, they are able to recognise patterns and making simple rules for complex problems, owing to their learning capability when submitted to a training. It is especially important the fact that they can generalise from a set of training data, being able to make predictions also on elements outside the original set.

4.4.1 History of ANNs

The study of the human brain dates back thousands of years but only recently, owing to the modern day electronics, man has begun to try and emulate the human brain and its thinking processes.

The preliminary theoretical base for contemporary neural networks was indipendently proposed by Alexander Bain (Bain 1873) and William James (James 1890); in their work, both toughts and body activity resulted from interactions among neurons within the brain.

The modern era of neural network research started in 1943 with the work done by Warren McCulloch, a neurophysiologist, and Walter Pitts, a logician (McCullock and Pitts 1943). They developed models of neural networks based on several assumptions about how neurons worked; their networks were based on simple neurons which were considered to be binary devices with fixed thresholds. The results of their model were simple logic functions such as "a or b" and "a and b". They also designed and built a primitive artificial neural network using simple electric circuits.

The next major development in neural network technology arrived in 1949 with a book, "The Organization of Behavior" (Hebb 1949) written by the psychologist Donald Hebb, which supported and reinforced McCulloch-Pitts's theory about neurons and how they work, and described how neural pathways

are strengthened each time they were used (this principle is the basis of the training of an ANN).

During the 1950s both corporations and Academia worked on developing NNs, and the outcome was research stimulation in Artificial Intelligence (AI) and Neural Networks. In 1954 Farley and Clark first used computational machines, then called calculators, to simulate a network based on the Hebbian learning hypothesis at MIT (Farley and Clark 1954); in 1956 other neural network computational machines were created (Rochester et al. 1956).

Two years later Frank Rosenblatt, a neuro-biologist at Cornell University, created the Perceptron (Rosenblatt 1958), which was the first 'practical' artificial neural network, consisting in a two-layer learning computer network using simple addition and subtraction. In that way this system could learn to connect or associate a given input to a random output unit.

Between 1959 and 1960, Bernard Wildrow and Marcian Hoff of Stanford University, in the USA, developed the ADALINE (ADAptive LINear Elements) and MADELINE (Multiple ADAptive LINear Elements) models (Widrow and Hoff 1960). These were the first neural networks that could be applied to real problems: they are still used as a filter to remove echoes from telephone lines. They consisted in an analogue electronic device made from simple components; the method used for learning was different to that of the Perceptron, and employed the Least-Mean-Squares (LMS) learning rule.

In 1969 Marvin Minsky and Seymour Papert proved that both Perceptron and ADALINE/MADELINE models had limited capabilities (Minsky and Papert 1969). The significant result of their work was to eliminate funding for research with neural network simulations, and the conclusions supported the disenhantment of reserchers in the field. Furthermore, the capabilities of artificial neural networks had been completely blown out of proportion by writers and producers of books and movies, so that people believed that such neural networks could do anything, resulting in disappointment when people realised that this was not so. These factors contributed to large-scale critique of AI and neural networks.

However, although public interest and available funding were minimal, several researchers continued working to develop neuromorphically based computational methods for problems such as pattern recognition.

In 1974 Paul Werbos developed and used the back-propagation learning method (Werbos 1974), however several years passed before this approach was popularised. In essence, the back-propagation network is a Perceptron with multiple layers, a different threshold function in the artificial neuron, and a more robust and capable learning rule.

One year later Fukushima developed a step wise trained multilayered neural network for interpretation of handwritten characters, called the Cognitron (Fukushima 1975). This kind of networks can propagate information in one direction only, or they can bounce back and forth until self-activation at a node occurs and the network settles on a final state.

The ability for bi-directional flow of inputs between neurons/nodes was produced with the Hopfield's network in 1982, which consisted of only one layer whose neurons were fully connected with each other. Hopfield presented a paper (Hopfield 1982) to the scientific community in which he stated that the approach to AI should not be to purely imitate the human brain but instead to use its concepts to build machines that could solve dynamic problems. He showed what such networks were capable of and how they would work, convincing scientists and researchers at the National Academy of Sciences of the USA to renew interest into the research of AI and neural networks.

At about the same time at a conference in Japan about neural networks, Japan announced that they had again begun exploring the possibilities of neural networks. The United States feared that they would be left behind in terms of research and technology and almost immediately began funding for AI and neural network projects.

A general re-emergence of interest in the neural network field was due to several factors, such as books and conferences, that provided a forum for people in diverse fields with specialised technical languages, the news media picked up on the increased activity and tutorials, that helped disseminate the technology, academic programs and courses, that were introduced at most major Universities (in US and Europe). 1986 saw the first annual Neural Networks for Computing conference that drew more than 1800 delegates.

The rediscovery of Werbos's backpropagation algorithm by Rumelhart, Hinton and Williams in 1986 was another important factor behind the repopularisation of neural networks (Rumelhart et al. 1986); they succeeded

in making the method widely known, so that the back-propagation learning emerged as the most popular learning set for the training of multilayer Perceptrons.

Today significant progress has been made in the field of neural networks, enough to attract a great deal of attention and fund further research. The tasks to which artificial neural networks are currently applied fall within the following broad categories:

- function approximation, or regression analysis, including time series prediction and modeling;
- classification, including pattern and sequence recognition, novelty detection and sequential decision making;
- data processing, including filtering, clustering, blind signal separation and compression.

Application areas of ANNs include system identification and control (vehicle control, process control), game-playing and decision making (backgammon, chess, etc.), pattern recognition (radar systems, face identification, object recognition, etc.), sequence recognition (gesture, speech, handwritten text recognition), medical diagnosis, financial applications, data mining (or knowledge discovery in databases, 'KDD'), visualisation and e-mail spam filtering.

The challenge today lies in finding ways to electronically implement the principles of neural network technology. Electronics companies are working on three types of neuro-chips, namely digital, analog and optical. With this prospect, the future of neural network technology looks very promising.

4.4.2 The human brain Neural Network

ANNs are modelled on the human brain structure. The brain is a part of the central nervous system and is made up of around 10^{11} neurons higly interconnected to form a complex network (see fig.4.1).

In a simplified explanation, each neuron consists of a *nucleus*, that is its central part, *dendrites*, an *axon* and *synaptic connections* (see fig.4.2). Dendrites are the sensing part of the neuron; through them it receives electric

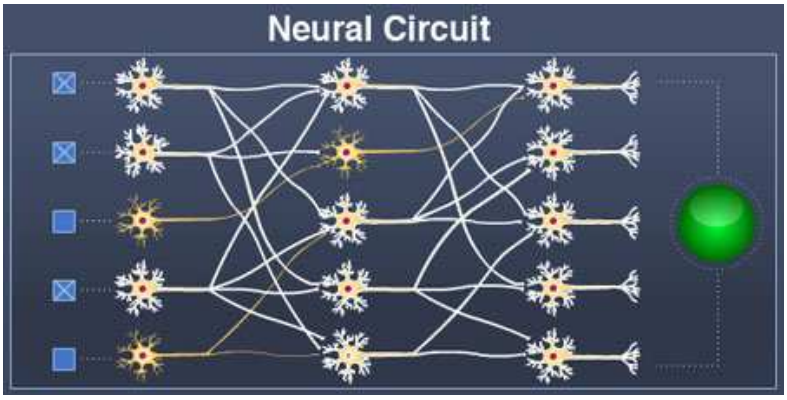

Figure 4.1: Scheme of a 'Neural Circuit', whose lighting up depends on the number of neurons that are passing on the signal. In the first layer on the left, the activated neurons are indicated by a cross. Figure taken by M. Arauz from an interesting and funny interactive experiment with a virtual neuron realized by Children' s Hospital Boston (http://www.childrenshospital.org/research/_neuron/index.html, on November 2011).

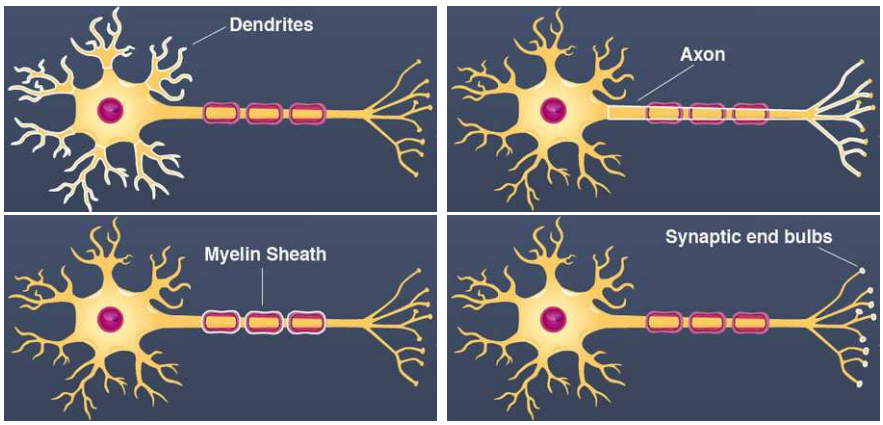

Figure 4.2: Scheme of a neuron. Dendrites are the sensing part, which receives the signal from other neurons; if the neuron activates, the axon carries the signal out to any other neurons connected to the neuron through the synaptic connections. The Myelin Sheath protects the axon, and also increases the speed and strength of the signal. Figures taken by M. Arauz from an interesting and funny interactive experiment with a virtual neuron realized by Children' s Hospital Boston (http://www.childrenshospital.org/research/_neuron/index.html, on November 2011).

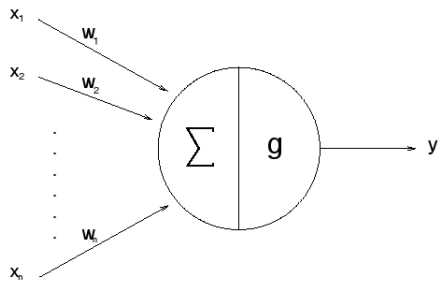

Figure 4.3: Scheme of an artificial neuron. From Nissen 2003.

pulses from other neurons, and if these pulses are enough (i.e. they exceed a certain threshold) the neuron activates and fires a pulse (called 'action potential') through its axon. The pulse will reach any other neurons that the neuron is connected to, through the synaptic connections. In this way the information (signal) can propagate across the Neural Net; in fig.4.1 it is shown a 'Neural Circuit', that only lights up if enough of the individual neurons are passing on the signal.

Both the synaptic connections and the threshold to activate the neuron change during its life; this allows the Neural Net to learn.

4.4.3 The artificial neuron

An artificial neuron is a mathematical function which tries to reproduce the behaviour of biological neurons. It takes the form:

$$y(\mathbf{x}) = g\left(\sum_{i=1}^{n} w_i x_i\right), \tag{4.1}$$

where $\mathbf{x} = (x_1 \dots x_n)$ is the *input vector* (or *pattern*), representing the inputs ('signals') received through the n 'dendrites', $y(\mathbf{x})$ is the output 'axon' and $(w_1 \dots w_n)$ are weights determining how much the inputs should be weighted (in a real neuron, the weights are given by the number of pulses a neuron sends out, the strenght of the pulses and how closely connected the neurons are). In fact, the weighted sum is then passed through a (usually non-linear, see below) function g known as an *activation function*, which on the basis of the sum decides if there should be any output from the neuron and, in that case, how powerful it should be (see fig.4.3).

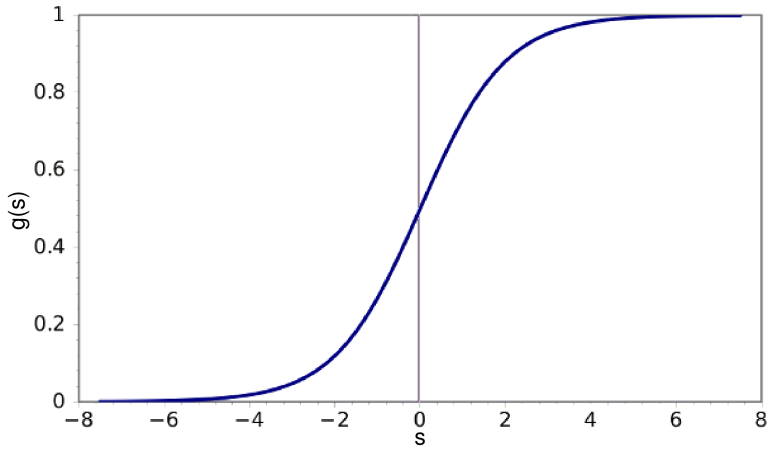

Figure 4.4: Graph of a logistic sigmoid function.

The easiest form of the activation function is a simple threshold function returning 1 or 0 (depending on whether the neuron is firing or not respectively). However, in the case of artificial neurons it is useful to have a smooth function, which returns an output between 0 and 1 (or -1 and 1, depending on the function used), so that also outputs not exactly equal to 1 can be considered as positive (we can for example fix an output threshold of 0.5 and say that all the outputs higher than that are positive). Hence, one of the most commonly used activation function takes the form of a 'logistic sigmoid function', having a characteristic S-shape (see fig.4.4):

$$g(s) = \frac{1}{1 + \exp(-s)} \ , \tag{4.2}$$

where $s = \sum_{i=1}^{n} w_i x_i$ is the weighted sum in eq. (4.1). This function assumes real values, is differentiable (that will be necessary for the 'back-propagation', see section 4.4.5) and in general is non-linear, even if for small s it shows a linear behaviuor.

The inputs and the weights are not restricted and can in principle be between $-\infty$ and $+\infty$, but they are often small values centered around zero.

4.4.4 Creation of an ANN

Artificial neurons can be combined to create an ANN; the most common kind of net is the *fully connected multilayer feedforward ANN*, where the neurons are

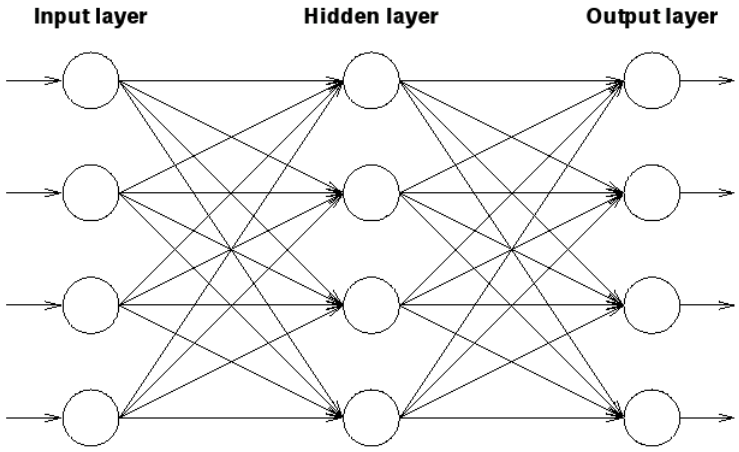

Figure 4.5: Scheme of a fully connected multilayer feedforward ANN. The circles represent the neurons, arranged in this case in three layers with the same number of neurons. The arrows between the layers indicate that the connections only go forward from one layer to the next. From Nissen 2003.

ordered in layers and each neuron of a layer is connected with all the neurons of the previuos and of the next layers (see fig.4.5). However, the connections only go forward from one layer to the next, i.e. the information only propagates in one direction, without going back to the previous layers (there are no loops).

The first layer is the *input layer*, whose number of neurons is equal to the number n of components of the input vectors (see previuos section) that we want to give to the ANN; the last one is the *output layer*, which contains a number of neurons equal to the number of desired outputs (for example, if we just want a 'yes' or a 'no' answer about our input candidate, the number of output neurons will be 2: one of them will 'fire' only if the answer is yes, the other only if the answer is no). Between the input and the output layers we can have a number of *hidden layers*, made up of any number of neurons, and the connections (and weights) to and from these layers determine how well the ANN performs.

That is how the ANN works: as said above, each neuron in the input layer represents one component of the input vector $\mathbf{x} = (x_1 \dots x_n)$, and sends its information x_i to all the neurons in the hidden layer (assuming for instance that we have only one of this kind of layers); then each neuron y_j in the hidden layer takes the form:

$$y_j = g(s_j^y) = \frac{1}{1 + \exp(-s_j^y)} \ , \tag{4.3}$$

where s_j^y is the weighted sum calculated by that neuron in that layer, for its interaction with all the neurons of the input layer:

$$s_j^y = \left(\sum_{i=1}^{n} w_{ij} x_i \right), \tag{4.4}$$

where w_{ij} is the weight associated between the i^{th} neuron in the input layer and the j^{th} neuron in the hidden layer. The index $j = 1, 2, 3, \ldots, m$, where m is the number of neurons in the hidden layer.

Hence also the hidden layer can be represented by a vector $\mathbf{y}$, with components $(y_1 \ldots y_m)$. The activation function of each neuron in this layer will decide, on the basis of the sum, about the neuron output; then all these outputs are sent to the neurons of the next layer, that in this case is the last one. Here again the weighted sum $s_k^z = \sum_{j=1}^{m} w_{jk} y_j$ and the function $z_k = g(s_k^z)$ are calculated for each neuron ($k = 1, 2, 3, \ldots, l$, where l is the number of neurons in the output layer); w_{jk} is the weight associated between the j^{th} neuron in the hidden layer and the k^{th} neuron in the output layer. Again, also this layer can be represented by a vector $\mathbf{z}$, with components $(z_1 \ldots z_l)$. Since it is the last layer, the outputs of these neurons will give the final output of the ANN.

4.4.5 Training an ANN

It is clear from what said in the previous section that the weigths are the crucial part of an ANN, since they determine, through the activation function, the output of each neuron; hence it is very important that they have the opportune values to make our ANN work well.

To find these opportune values, the ANN is submitted to a process called *training* or *learning*; in a *supervised* learning, the ANN is trained to return a specific output when given a specific input, i.e. its internal weights, that initially assume random values, are slowly adjusted (through a continuous training on a set of input and output data) in order to produce the desired output. The hope is that, once the ANN is trained, when the net is shown a new set of input variables it will give a correct output.

A common algorithm to train ANNs is the so-called 'backpropagation' algorithm, that works in reverse from the output to the input layer in this way: after propagating an input through the net, the error between the obtained output and the desired output is calculated, and the error is propagated back through the net while the weights are adjusted in order to make the error smaller. This is done for all the input vectors, and then repeated for a number of iterations (also called *epochs*) until a certain stop criterion is reached.

It can be shown that the weights must be adjusted on the basis of some δ values, that for the k^{th} neuron in the output layer are calculated according to this equation:

$$\delta_k = g'(s_k^z)\ (z_k - d_k)\,, \tag{4.5}$$

where g' is the derived activation function and d_k is the desired output of neuron k, while for the j^{th} neuron in the hidden layer:

$$\delta_j = \eta\, g'(s_j^y) \sum_k \delta_k w_{jk}, \tag{4.6}$$

where η is the learning rate parameter, which determines how much the weight should be adjusted, and the sum runs over all neurons k to which neuron j sends connections (since in our case we have only one hidden layer, and also the net is fully connected, the neurons k are all the ones in the output layer).

Hence we calculate the δ_j values for preceding layers from the δ_k values of successive layers. Then we can calculate the Δw values that the weights should be adjusted by:

$$\Delta w_{jk} = \delta_j z_k\,; \tag{4.7}$$

so the weight w_{jk} is adjusted in this way: $w_{jk}(t+1) = w_{jk}(t) + \Delta w_{jk}(t+1)$, where t indicates one training iteration, and $(t+1)$ is the next one. The backpropagation algorithm then moves on to the next input vector of the training set, repeating the same process of adjusting the weights.

The δ values are related to the derivatives, with respect to each of the weights in the ANN, of the so-called *error functions* E^p, where p labels the input vectors; these functions can be defined for example as the standard sum-of-squares functions:

$$E^p = \frac{1}{2} \sum_{k=1}^{l} (z_k - d_k)^2 , \tag{4.8}$$

or as the mean square error functions:

$$E^p = \frac{1}{l} \sum_{k=1}^{l} (z_k - d_k)^2 , \tag{4.9}$$

the latter meaning that each E^p is the average of all the square errors in each of the training pairs (input-output).
A total error function can be then defined, for example as the sum of the error functions of all the inputs in the training set:

$$E = \sum_{p} E^p ; \tag{4.10}$$

the stop criterion to the training process is typically determined by measuring the value of E during the training: when it reaches a certain limit the training is stopped. More advanced stopping criteria involves both training and testing data (where the latter is another set of data, different from the training set, used to test if the output given by the ANN is good); for example, if we do the testing during the training we should stop when the error of the testing data is not improving anymore, i.e. it has reached a minimum, in order to avoid the *overtraining* of the net, that occurs when too much training is applied to a set of data and the error in the testing set begins to increase again. In fact in that case the ANN looses generalisation and becomes able to recognise only input vectors that match precisely those in the training set.

Another stop criterion than the value of the combined mean square error could be that each of the training pairs should have a mean square error lower than a given value.

4.5 Our techniques

4.5.1 Scores

The first thing to do to use an ANN is to generate the components of the input vectors, usually called *scores*, i.e. numbers that describe the characteristics of the input.

In our case the input is a pulsar candidate (or, for the training and testing sets, known pulsars and certain non-pulsars, i.e. RFI or noise), and we can create its scores on the basis of the features shown in its diagnostic plot. In the HTRU pipeline, the values of the basic candidate parameters used to make the plot like period, S/N, DM, etc., are stored in a text file in the Pulsar Hunter[1] Candidate XML (PHCX) format, where XML is the well-defined eXtensible Markup Language[2] format, that can be read both by humans and by softwares. In this work some Phyton[3] scripts have been used to take the information about the features from the phcx file of the candidate and generate the relative scores.

In this case we chose to create 27 scores, adding 5 to the 22 scores used by the ANN in Manchester, that should hopefully allow the ANN to well distinguish the features of a pulsar from those of a non-pulsar and hence to make a correct decision about a candidate.

In the following I will give a brief description of some of the used scores, whose summary is reported in table 4.1 (see also Bates 2010 for the 22 scores used in Manchester).

Profile fitting scores

Since RFI profiles are often described by a sine or a sine squared function, the candidate profile is fitted with both these functions, and the χ^2 values of the two fits are used as scores. Then the profile is fitted with Gaussian functions with one or two peaks, that are instead good descriptions of 'typical' pulsar profiles, and we take as scores the value of several parameters from the two fits, such as the FWHM (full width at half maximum) and the χ^2 values. Furthermore, a Gaussian fit on the derivative of the profile is performed, and the difference between the expectation value of this fit and the expectation value of the Gaussian fit on the profile is taken as a score.

DM curve fitting scores

From eq. (2.1) we know that the S/N of a pulse follows this relation:

[1] http://pulsarhunter.sourceforge.net/
[2] http://www.w3.org/XML
[3] http://www.python.org/

Table 4.1: Summary of the scores for the ANN used in this work.

#	Score description
Sinusoid Fitting	
1	χ^2 value from fitting the pulse profile with a sin curve
2	χ^2 value from fitting the pulse profile with a $\sin^2$ curve
Gaussian Fitting	
3	Difference between expectation value of pulse profile and fitted Gaussian
4	Max. value of pulse profile / Max. value of Gaussian
5	FHWM of Gaussian fit
6	χ^2 value from Gaussian fit
7	Average FHWM from fitting two Gaussians to pulse profile
8	χ^2 value from fitting two Gaussians to pulse profile
9	Fit Gaussian to dy/dx of profile. Difference between expectation values
Candidate Parameters	
10	Best period (ms)
11	Best S/N value
12	Best DM value (DM_{best})
13	Pulse width
DM Curve Fitting	
14	S/N / $\sqrt{(P - W_e)/W_e}$
15	After fitting the DM curve, calculate $(\mathrm{S/N})_{fit}$ / $\sqrt{(P - W_e)/W_e}$
16	mod($DM_{fit} - DM_{best}$)
17	χ^2 value from DM curve fit
Sub-band scores	
18	RMS of peak positions in all sub-bands
19	Average correlation coefficient for each pair of sub-bands
20	Sum of correlation coefficients
Pulse Profile Tests	
21	Number of peaks in the pulse profile
22	Integrated area under the pulse profile
Subintegration scores	
23	RMS of peak positions in all subintegrations
24	Average correlation coefficient for each pair of subintegrations
25	Sum of correlation coefficients
26	Mean of correlation coefficients
27	Entropy

$$S/N \propto \sqrt{\frac{P - W_e}{W_e}}, \tag{4.11}$$

where P is the pulse period and W_e is the effective pulse width, given by eq. (2.3). Since W_e is a function of DM, we can evaluate the behaviour of the S/N for our trial DM values (see fig.2.3), i.e. predict the shape of the *DM curve*. As a score to represent this theoretical curve we use the proportionality constant between the two terms in eq. 4.11, calculated for the peak of the curve, i.e. for the value of DM at which the pulse would be correctly de-dispersed. Then we take the actual values of the spectral S/N at each trial DM from the phcx file of the candidate, and fit these data with the relation above; for a pulsar, the fit will be good. The χ^2 value of this fit is hence another score, together with the value of the proportionality constant after the fit, $(\mathrm{S/N})_{fit}$ / $\sqrt{(P - W_e)/W_e}$. The last DM fitting score is the deviation of the fitted DM from the best DM as recorded in the phcx data.

Sub-band and subintegration scores

Since a pulsar is visible right across all the observing bandwidth, not just within a few of the sub-bands (section 2.2.1), and the shape of the pulse is consistent over the entire bandwidth, we can test our candidate running a window of a width equal to the pulse width (taken from the phcx file) along the phase axis of each sub-band, and integrate within the window. The center of the window is then identified with the maximum value for each sub-band, so that we can calculate how close the maxima are to each other. The returned value, that will be our first sub-band score, is the standard deviation of the positions of the maxima, that hence should be as low as possible to have a good candidate. Another score is the correlation coefficient of the amplitudes across the whole pulse between sub-band pairs, averaged across all the pairs; since the pulse from a pulsar is strongly correlated across the whole spectrum, in that case the correlation coefficient is high (i.e. closed to 1). We find the value of one more score measuring the correlation of the whole profile to each sub-band and adding all correlation coefficients together (this score is quite similar to the previuos one).

To test the persistence of the signal over time, in order to be able to identify

and reject short bursts of RFI, we introduced four new scores related to how the signal changes over time. Three of them are the equivalent of the three sub-band scores, but using subintegrations (section 2.2.1) rather than sub-bands; the fourth score is a mean of the correlation coefficients obtained from the correlation of the whole profile to each subintegration.

Entropy

We introduced a last new score on the basis of the concept of entropy, that is a measure of the degree of order of a system: more the latter is ordered, lower its entropy is. If we have a periodic signal and Fourier transform it, its degree of order will depend on the quantity of its power being in its harmonics. For example, a sinusoidal RFI signal will have a very low entropy, because all of its power will be in one harmonic. A pulsar should have power in multiple harmonics and therefore have a higher value of entropy. White noise (like the Gaussian noise) has its power distributed uniformly over the entire frequency range, i.e. it has equal power in all harmonics and hence maximum entropy.

The entropy score is then calculated by doing an FFT of the candidate pulse profile, computing the power of each resulting element and finally calculating the sum:

$$-\sum(k_i \log(k_i)) \,, \tag{4.12}$$

where k_i is the fractional power in each harmonic.

Outlier and dimensionality issues

Training and testing the ANN, as will be described in the following sections, we dealt with the presence of some outliers among the values of the scores of the training set; as already said, the outliers are values that lie well outside the range in which the other values of the scores lie for the same class of inputs (e.g. in our case pulsars or non-pulsars). There was a suspicion that the inputs with such outliers might be eschewing the training process, possibly making the ANN inaccurate. Hence some tests were performed using some 'filters', i.e. it was decided to discard any entry for which one of the scores was more than four sigma outside the mean for that score. Nevertheless, this resulted in no better accuracy than when using all entries.

Then this test was repeated with some variations, like using median rather than mean; calculating separate mean (or median) for each class of inputs and filtering accordingly; running the test using only the best performing pair or triplet of score fields; training with the filtered dataset and testing with the unfiltered one. None of these attempts resulted in measurable improvements in the ANN performance.

Another issue that can arise when dealing with a generic ANN is related to dimensionality, i.e. to the number of score fields. In fact, even if in general a large number of input scores is desirable since it can better represent the input data, and furthermore the ANN is able to make more complicated decisions, in many cases it may worsen the ANN performance. This is a problem called the *curse of dimensionality* (see for example Bishop 1995): each of the training entry corresponds to a point in the input space, whose dimension is equal to the number of score fields (i.e. the scores are the coordinates of the point in this space); if we imagine to divide the input space in a large number of cells, each of the training entries correspond to a point in one of the cells, and carries an associated value of the output variable y. If more than one training point fall in the same cell, the average value of y for all these points is returned. Therefore, if we have a new point, we can determine a corresponding value for y by finding which cell the point falls in (this is the generalisation ability of the ANN); nevertheless, if we increase the dimensionality of the input space, the number of training points necessary to have the complete mapping of the space (i.e. to have at least one training point in each cell) grows *exponentially* with the dimensionality. Since in practice the quantity of training data is limited, increasing dimensionality leads to the point where this data is very sparse, with many low density or even empty regions. This makes it difficult to the ANN to generalise from the training data and produce useful new results. However, the feed-forward ANN is able to exploit correlations in the data and reduce the dimensionality of the problem.

The dimensionality can also be reduced by eliminating scores that are irrelevant or redundant, i.e. that are correlated with other scores. For this reason different kind of elimination tests have been performed on our ANN, for instance doing training/testing rounds using all available input fields but one, with a different field left off on each round, in order to see how the performance

of the ANN varied and understand the significance of the field left off. In another test, input fields were eliminated one by one but not added back. Finally, another round of tests was executed performing training/testing rounds with all possible combinations of two and three input fields.

In all these cases, no significant improvements in the ANN performance were observed. This confirms that we had sufficiently many training data points to fully sample the dimensionality of the problem space.

4.5.2 The (Fast)ANN used in this work

The kind of ANN used in this work is the *Fast Artificial Neural Network* (FANN)[4], a multilayer ANN library in C language created and actively maintained by the Computer Scientist Steffen Nissen, also author of the report *Implementation of a Fast Artificial Neural Network Library (FANN)* (2003). Moreover, a reference manual is available online[5].

We chose to use a FANN because it presents numerous valuable features, for example the fact that it is an open source, easy to use and fast (up to 150 times faster execution than other libraries), that it can be trained with different kind of backpropagation algorithms, and that we can choose among several different activation functions; but many other features could be enumerated.

Many tests were performed to find the combination of parameters giving the best performance of the FANN in our case; the result was a FANN with an architecture 27:18:2 (i.e. 3 layers, with the input layer made up of 27 neurons, the only one hidden layer with 18 neurons and the output layer with 2 neurons), a learning algorithm called iRPROP (improved Resilient backPROPagation, Ingel and Hüsken 2000), that is an improved version of the standard RPROP (Riedmiller and Braun 1993), and a fast symmetric sigmoid-like activation function (Elliott 1993), that is a higher-speed approximation of the sigmoid activation function and takes the form:

$$y = \frac{x \cdot s}{1 + \mid x \cdot s \mid}, \tag{4.13}$$

where x is the input to the activation function, y is the output and s is a steepness parameter (in our case fixed to 0.1); the span of this function is

[4] http://leenissen.dk/fann/wp/
[5] http://leenissen.dk/fann/html/files/fann-h.html

$-1 < y < 1$. The stop criterion used during training was the mean square error value (eq. (4.9)). A future work will be to implement in the FANN a stop criterion based on testing during training.

4.5.3 Application of the FANN to the candidate problem

Our first step to apply the described FANN to the candidate problem was creating a training and a testing set for the net, both made up of known pulsars and non-pulsars. Hence we took all the known pulsars re-detected by our pipeline at the Swinburne University since the beginning of the HTRU survey, whose number was 644, and a group of 28444 candidates previously classified by eye, obtained by discarding the known pulsars and the most probable pulsars from the candidates already visually inspected; then we calculated the scores described in section 4.5.1 and subsections.

Then we created the training set so that it contained 50% pulsars and 50% non-pulsars, using half of the known pulsar sample (322) and an equal number of non-pulsars. The remaining half and the remaining non-pulsars were then used to create the testing set, so that all the elements contained in the latter were different from the elements in the training set.

The next step was training the FANN by using the described training set, where the desired output for each input vector was a vector $\mathbf{d} = (1, -1)$ for pulsars and a vector $\mathbf{d} = (-1, 1)$ for non-pulsars (because of the chosen activation function, see previous section). Then we performed the testing by using the described testing set; the result was quite good, since the FANN correctly identified as pulsars 97.52% of the known pulsars present in the testing set, and as non-pulsars 98.45% of the non-pulsars in the same set.

In order to reduce the effects of the randomness implicit in the training process, and also any potential biases caused by the selection of the training and testing sets, a *cross-validation* system was developed that randomises the datasets and performs repeated tests while averaging the results. In fact, using the full available sample of pulsars and non-pulsars, it creates temporary training sets containing again equal numbers of randomly selected pulsars and non-pulsars (again half of the known pulsars and an equal number of non-pulsars) and testing sets containing all the remaining elements; those temporary sets are then used to do ten rounds of training/testing, the results are logged,

and the whole process (including the selection of the elements for the two sets) is repeated nine more times. The result was still quite good, with an average of 97.98% correctly identified pulsars and 97.78% correctly identified non-pulsars.

However, for a large-scale survey like HTRU which yields million candidates most of which are non-pulsars, even 2% non-pulsars wrongly identified as pulsars means a very high number of candidates to visually inspect (since all the candidates indicated as pulsars by the net are then inspected by eye). Hence we would like to increment the percentage of non-pulsars correctly identified by the FANN, but maintaining at the same time a high percentage of correctly identified pulsars.

Two useful concepts in the classification of objects are the *completeness* and the *efficiency* (see for example Ball and Brunner 2010), defined in terms of true and false positives (TP and FP) and true and false negatives (TN and FN). In our case:

- TP are the pulsars correctly classified as pulsars by the net;
- FP are the non-pulsars wrongly classified as pulsars;
- TN are the non-pulsars correctly classified as non-pulsars;
- FN are the pulsars wrongly classified as non-pulsars.

The completeness is defined as:

$$completeness = \frac{TP}{TP + FN}, \tag{4.14}$$

i.e. in our case is the fraction of correctly identified pulsars with respect to the total of pulsars (correctly or wrongly classified); the efficiency is defined as:

$$efficiency = \frac{TP}{TP + FP}, \tag{4.15}$$

i.e. in our case is the fraction of correctly identified pulsars with respect to the total of objects (correctly or wrongly) classified as pulsars.

What we want is that both these quantities be high, since a high completeness means that many pulsars have been correctly classified as pulsars and only a few have been wrongly identified as non-pulsars, while a high efficiency means that most of the objects that have been classified as pulsars

Percentage and number of pulsars in the training set	Percentage and number of non-pulsars in the training set	Total number of objects in the training set	Number of non-pulsars in the testing set (number of pulsars = 322)	Total number of objects in the testing set	Testing results for pulsars, averaged on 10 trials and 10 rounds per each trial	Testing results for non-pulsars, averaged on 10 trials and 10 rounds per each trial
3% = 322	97% = 10412	10734	18032	18354	92.60%	99.71%
5% = 322	95% = 6118	6440	22326	22648	92.74%	99.68%
10% = 322	90% = 2899	3221	25545	25867	95.01%	99.48%
20% = 322	80% = 1289	1611	27155	27477	96.51%	99.11%
30% = 322	70% = 752	1074	27692	28014	97.21%	98.79%
40% = 322	60% = 483	805	27961	28283	97.13%	98.43%
50% = 322	50% = 322	644	28122	28444	97.72%	97.77%
60% = 322	40% = 215	537	28229	28551	98.36%	97.17%
70% = 322	30% = 138	460	28306	28628	98.94%	95.43%
80% = 322	20% = 81	403	28363	28685	98.84%	93.07%
90% = 322	10% = 36	358	28408	28730	98.81%	91.05%

Figure 4.6: Cross-validation results for nets trained with different percentages of pulsars and non-pulsars in the training set.

are actually pulsars and only a few non-pulsars have been wrongly identified as pulsars.

Nevertheless, there is generally a tradeoff involved since it is not possible to increment both, but a choice needs to be made depending on the application.

In our case, we decided to train several nets with different percentages of pulsars and non-pulsars in the training set (i.e. not just 50% - 50%), to investigate the change in the net performance and hence in completeness and efficiency, in order to choose the best case for our purpose; the cross-validation results are reported in fig.4.6.

The number of pulsars in the training set was maintained equal to 322, but the percentage changed according to the change in the number of non-pulsars, that were added or removed; all the remaining elements of the full sample were used in the testing set.

As shown in the table, the average percentages of correctly identified pulsars and non-pulsars (last two columns), with respect to the total of pulsars and to the total of non-pulsars present in the testing set respectively, are similar for the 50% - 50% case, as we found in the previuos test, but change in an opposite way when the percentages of pulsars and non-pulsars in the training set are modified

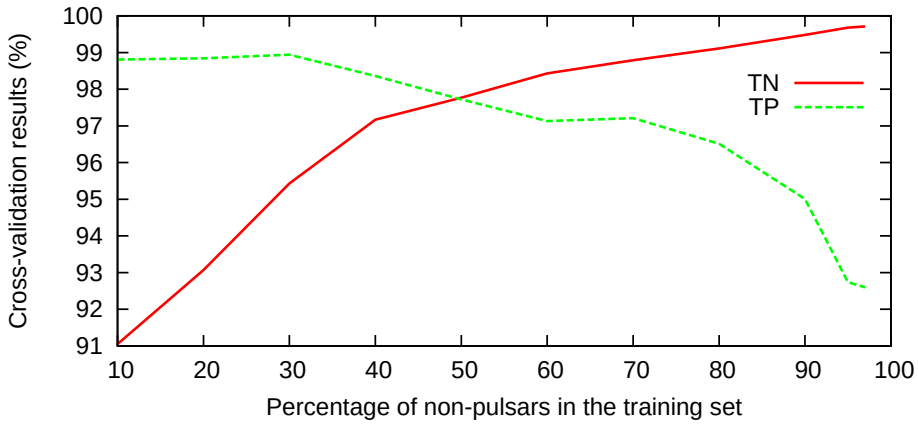

Figure 4.7: Cross-validation results as functions of the percentage of non-pulsars in the training set, for pulsars (TP, green curve) and non-pulsars (TN, red curve).

(hence confirming the need of a tradeoff between completeness and efficiency). In fig.4.7 these cross-validation results as functions of the percentage of non-pulsars in the training set are shown, both for pulsars and non-pulsars, while in fig.4.8 completeness and efficiency for each row of the table in fig.4.6 are plotted, also in this case as functions of the percentage of non-pulsars in the training set. It is evident from the last plot that incrementing such a percentage we have significant gains in efficiency with a small penalty in completeness.

At this point we made a particular choice. Since, as already said, our intention was to cut the number of non-pulsars wrongly classified as pulsars, in order to have a human manageable number of candidates to visually inspect, we decided to favour the correct classification of non-pulsars rather than that of pulsars, choosing to apply to our survey candidates the net trained with 3% pulsars and 97% non-pulsars, for which the completeness was 0.926 and the efficiency ~ 0.851. In fact, such a net was able to correctly identify on average 99.71% non-pulsars and 92.60% pulsars; therefore our FP percentage was only 0.29%, that for million non-pulsars is just a few thousands. It is true that in this way we have a percentage of FN of 7.4%, i.e. we miss such a percentage of pulsars since they are wrongly classified as non-pulsars, but considering the small number of pulsars present among the candidates, this means that we are missing only a few pulsars: we think that this is quite acceptable, with respect to the huge number of non-pulsars eliminated.

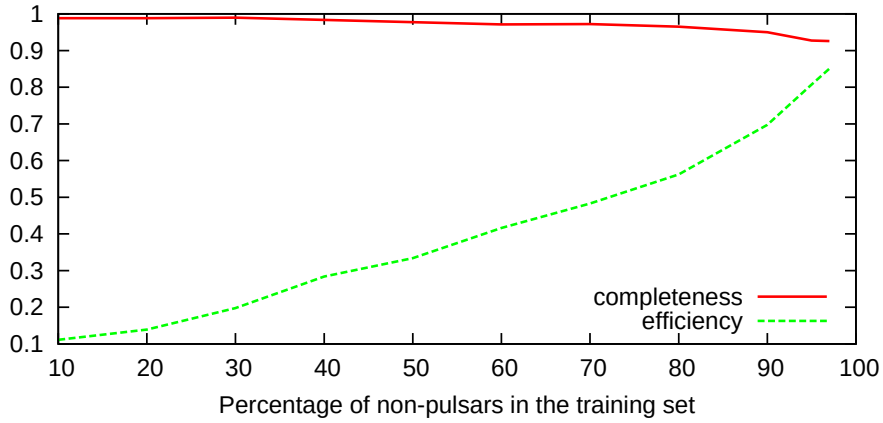

Figure 4.8: Completeness (red curve) and efficiency (green curve) plotted for each row of the table in fig.4.6, as functions of the percentage of non-pulsars in the training set.

Since neural nets are trained using random numbers and there may not be a unique minimum in the error function, we wanted to know the consequences of random training; hence we decided to train, with the mentioned percentage of pulsars and non-pulsars, a 'pseudo-committee'[6] of 10 nets, and in the end we exploited it to reduce the number of FP for the visual inspection, choosing to only look at the candidates classified as pulsars by *all* the nets (hence that received 10 'yes' votes). The training sets were made up of elements taken randomly from the full available sample, hence they should be quite different for each net.

The testing results for the 10 nets were of between 89.13% and 93.48% for the correct classification of pulsars (with respectively 99.76% and 99.78% for the correct identification of non-pulsars), while they were of between 99.68% and 99.88% for the correct classification of non-pulsars (with respectively 90.99% and 90.06% for the correct identification of pulsars). Hence, in the worst case, the percentage of FP was 0.32%.

In addition, we must consider the fact that this time we do not want to average the results of the ten nets, but we want to consider only the candidates

[6] We call our nets a 'pseudo-committee' since they do not constitute exactly a committee, because in the latter the nets should have different numbers of hidden units, or the same architecture but trained to different local minima of the error function, or with different kind of net models; moreover the output of a committee is taken as the *average* of the outputs of the nets in the committee (Bishop 1995).

that will receive 10 yes votes; this means that, among them, the percentage of correctly identified pulsars will be equal to or lower than the lowest testing result among the nets (i.e., even in the case that all the nets correctly identify as pulsars exactly the same candidates, the percentage of these candidates with respect to all the pulsars present among the candidates would be equal to the lowest testing result, 89.13%, because the net yielding this result is not able to correctly identify a higher number of pulsars; but since in general the nets will not classify as pulsars exactly the same candidates, the percentage of correctly identified pulsars among the 10-yes-vote candidates will be even lower than 89.13%, however not too much). Therefore, we should wonder if this percentage is still acceptable: it implies that something more than 10.87% pulsars would be missed. Anyway, as already said, considering the small number of pulsars present among the candidates we still do not miss many pulsars. However, this percentage could be improved by considering also for example the candidates that received *only* 9 yes votes (that is a sub-class of those which received *at least* 9 votes, i.e. 9 or 10, since the 10-yes-vote candidates had obviously received also 9 votes), or even the ones that received *only* 8 yes votes (sub-class of the 8-yes-vote class, for the same reason); but, as we will show soon, in this way the number of FP (and hence of candidates to visually inspect) grows considerably. Anyway this could be something to try in the future.

In order to test our method, we decided to launch the ten trained and tested nets on all the sample of pulsars (644) used for the training/testing process, and then on all the sample of non-pulsars (28444) used for the same purpose. The number of elements that received *at least* a number of yes votes respectively from 0 to 10 are shown in table 4.2 for both the cases. The percentage of pulsars that received 10 yes votes was 86.6%, while the same percentage for non-pulsars was 0%; this can be explained by considering the fact that the number of non-pulsars in the training set of each net was overwhelming (97% of the total of elements) with respect to the number of pulsars, and that each net had more than one third of the 28444 non-pulsars in its training set.

At this point we took all the candidates yielded by the survey pipeline at the Swinburne University since the beginning of the survey, whose impressive number was 19,084,702 (on May 2011), and calculated the scores for all of them. Given the huge number of candidates, this step was going to be very time

Yes votes	Pulsars	Non-pulsars
10	558	0
≥ 9	579	0
≥ 8	589	2
≥ 7	602	4
≥ 6	612	13
≥ 5	622	19
≥ 4	630	36
≥ 3	635	59
≥ 2	639	92
≥ 1	642	193
≥ 0	644	28444

Table 4.2: Number of elements that received *at least* a number of yes votes respectively from 0 to 10 after launching the ten trained and tested nets firstly on all the sample of pulsars (644) and then on all the sample of non-pulsars (28444) used for the training/testing process.

consuming; hence we ran multiple processes in parallel using the Green Machine supercomputing facility at Swinburne University (see section 3.2.3). In fact we divided the candidates into blocks each made up of $\sim$ 35,000 elements and then ran the jobs (i.e. the calculation of the scores for each block of candidates) on different nodes of the supercomputer; then the ten nets were launched on each block. In this way all the scoring/classifying process took only one day and a half.

In table 4.3 it is reported the number of candidates that received *at least* a number of yes votes respectively from 0 to 10, and the percentage with respect to the total of candidates; we can see that $\sim$ 97% of the candidates received 0 yes votes, while 0.21% received 10 yes votes, i.e. $\sim$ 41,000 candidates.

Hence, in this way we managed to reduce the number of candidates to visually inspect from more than 19 millions to only $\sim$ 41,000, i.e. something that is human manageable. From table 4.3 we can notice that if we chose to consider for example also the candidates that received *only* 9 and *only* 8 votes (20,095 and 21,215 respectively), in order to increase the number of correctly identified pulsars, the number of candidates to look at would double.

Yes votes	Candidates	Percentages
10	40864	0.21%
≥ 9	60959	0.32%
≥ 8	82174	0.43%
≥ 7	105113	0.55%
≥ 6	131022	0.69%
≥ 5	162651	0.85%
≥ 4	203728	1.07%
≥ 3	261623	1.37%
≥ 2	352612	1.85%
≥ 1	569911	2.99%
≥ 0	19084702	100%

Table 4.3: Number of candidates that received *at least* a number of yes votes respectively from 0 to 10, and the percentage with respect to the total of candidates, after launching the ten trained nets on all the candidates yielded by the survey pipeline (19,084,702).

4.5.4 Work in progress and future

The work of visually inspecting the 10-yes-vote candidates is in progress. We made an advanced script in C++ to find known pulsars and their harmonics (from 1/20 to 20 times the fundamental frequency) possibly present among these candidates, double-checking the latter with the catalogue of known pulsars. The script found 2,299 known pulsars (since many pulsars were observed several times) and 1,181 harmonics (this is a good result, since it means that the nets are able to recognise harmonics, although no harmonics were in the training sets). In this way the number of candidates to be inspected further reduced to 37,384.

The hope obviously is that among them there be new pulsars to discover; moreover, it could be possible also a re-evaluation of a part of previously wrongly classified data.

Anyway, since most of the 37,000 remaining candidates are surely non-pulsars, this could mean that many pulsar-like RFI deceived the nets. Hence, as already explained in section 4.3, our main aim now is to create a better training set in order to improve the training of the nets.

For this reason, inspecting the candidates, we are dividing them into three categories, pulsars, non-pulsars and 'confusion', where the latter are the

candidates which it is difficult to make a decision about, and that therefore might equally be pulsars or non-pulsars. These categories can be used in different ways (still to be tested) to create the improved training set.

If our approach still consists in making the nets recognise pulsars, i.e. distinguishing between pulsars and non-pulsars, the confusion category should not be included in the new training set, in order to avoid errors in the training. The new training set then should be made up of the newly classified pulsars added to the pulsars of the old training set, and the newly classified non-pulsars (the non-pulsars of the old training set should not be included, since they were obtained by discarding only the most probable pulsars, hence they could include some confusion as well, undermining the work put into separating the classes).

A different approach can consist in training the nets to recognise RFI, i.e. to distinguish between RFI and non-RFI, which may provide better results. In this case, it might be useful that both pulsars and confusion be included in the non-RFI category (together with the pulsars of the old training set), hence feeding the ANNs with more examples of 'what might not necessarily be RFI'. The RFI category should be made up only of the newly classified non-pulsars as in the previous case, for the same reason. Since in this case we have a greater number of classified data points ($\sim$ 40,000) and a better balance between RFI and non-RFI candidates (owing to the inclusion of the confusion category in the non-RFI class) with respect to the old training/testing set (in which there was instead a great asymmetry in the number of pulsars and non-pulsars), the hope is that we can decrease both FN and FP rates, where in this case FN are misclassified RFI and FP are misclassified pulsars, obtaining better trained ANNs.

After obtaining a good training set, we would like to make it public, so that it can be used for instance also by people not involved in large-scale surveys and that hence do not have the possibility to make their own training set. Moreover, we have also the intention to make our scripts an open source.

For the future, besides what said above, we are thinking about other ways to improve the performance of the nets. For example, since almost all the known pulsars found among the 10-yes-vote candidates had S/N $>$ 10, it would be necessary to train the nets with more pulsars having S/N $<$ 10, in order to make the nets able to correctly classify them as well (their percentage was only

5% of the pulsar sample that we used for training/testing). Other possibilities could be exploring other machine learning algorithms (see for example the ones described by Ball and Brunner 2010) and making new kind of scores, especially to better identify RFI; for instance, a new score may be the correlation between the profile of RFI and the folded profile of the candidate: if the correlation is high (~ 1), the candidate is rejected. For this purpose it would be valuable if a catalogue of RFI could be created.

Credits

This work exploited the scripts to calculate the scores made by Christian Nietner for the 22 parameter case at Manchester, and modified by Wilson Afonso for the 27 parameter case at Melbourne. The latter also performed the tests on outliers and dimensionality and on the combination of parameters to have the best performance of a FANN, and made the scripts to train and test it.

Publications

List of the HTRU collaboration papers currently in preparation which are direct results of the work reported in this thesis:

M. Burgay et al. *The High Time Resolution Universe Pulsar Survey - VI. Discovery and timing of further five millisecond pulsars* (provisional). MNRAS in prep.

S. Milia, W. van Straten et al. *The High Time Resolution Universe Pulsar Survey - VII. Use of several Neural Nets to select candidates* (provisional). MNRAS in prep.

Refereed papers produced so far by the HTRU collaboration:

M. Bailes, S. D. Bates, V. Bhalerao, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. DAmico, S. Johnston, M. J. Keith, M. Kramer, S. R. Kulkarni, L. Levin, A. G. Lyne, S. Milia, A. Possenti, L. Spitler, B. Stappers, and W. van Straten. *Transformation of a Star into a Planet in a Millisecond Pulsar Binary.* Science, 333:1717-, September 2011. doi: 10.1126/science.1208890.

S. D. Bates, M. Bailes, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. DAmico, A. Jameson, S. Johnston, M. J. Keith, M. Kramer, L. Levin, A. Lyne, S. Milia,

A. Possenti, B. Stappers, and W. van Straten. *The High Time Resolution Universe Pulsar Survey - II. Discovery of ve millisecond pulsars.* MNRAS, 416:2455-2464, October 2011. doi: 10.1111/j.1365- 2966.2011.18416.x.

S. Burke-Spolaor, M. Bailes, S. Johnston, S. D. Bates, N. D. R. Bhat, M. Burgay, N. DAmico, A. Jameson, M. J. Keith, M. Kramer, L. Levin, S. Milia, A. Possenti, B. Stappers, and W. van Straten. *The High Time Resolution Universe Pulsar Survey - III. Single-pulse searches and preliminary analysis.* MNRAS, 416:2465-2476, October 2011. doi: 10.1111/j.1365-2966.2011.18521.x.

M. J. Keith, A. Jameson, W. van Straten, M. Bailes, S. Johnston, M. Kramer, A. Possenti, S. D. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. DAmico, L. Levin, P. L. McMahon, S. Milia, and B. W. Stappers. *The High Time Resolution Universe Pulsar Survey - I. System conguration and initial discoveries.* MNRAS, 409:619-627, December 2010. doi: 10.1111/j.1365-2966.2010.17325.x.

M. J. Keith, S. Johnston, M. Bailes, S. D. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. DAmico, A. Jameson, M. Kramer, L. Levin, S. Milia, A. Possenti, B. W. Stappers, W. van Straten, and D. Parent. *The High Time Resolution Universe Pulsar Survey - IV. Discovery and polarimetry of millisecond pulsars.* MNRAS, page 1783, November 2011. doi: 10.1111/j.1365-2966.2011.19842.x.

L. Levin, M. Bailes, S. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. DAmico, S. Johnston, M. Keith, M. Kramer, S. Milia, A. Possenti, N. Rea, B. Stappers, and W. van Straten. *A Radio-loud Magnetar in X- ray Quiescence.* ApJ, 721:L33-L37, September 2010. doi: 10.1088/2041- 8205/721/1/L33.

Summary

The work described in this Thesis has been developed in the context of the *High Time Resolution Universe Survey* (HTRU) for pulsars and radio transients, which is in progress at the 64m Radio Telescope of Parkes, in Australia. It is an all-sky survey at 20 cm wavelength, the first part of which (dubbed *med-latitude survey*) looked at a portion of the Southern sky within ± 15 degrees from the Galactic Plane.

The main aim of the survey is to significantly increase the number of known millisecond pulsars in the Galactic Field, hopefully doubling the number of catalogued objects of this category. On a statistical ground, one expect that a fraction of the new pulsars will turn out to have the requested features - rapid spinning, high radio flux, sharp pulse profile, good long term rotational stability - for becoming suitable members of one of the Pulsar Timing Arrays, whose main scope is the direct detection of gravitational waves released by supermassive binary black-holes. Moreover, among the new discoveries, one expect that few objects will turn out to raise interesting scientific cases *per se*.

In fact, the first processing of the data of the *med-latitude survey* has already led to the discovery of ~ 20 millisecond pulsars (plus about 70 ordinary pulsars). Among them few very interesting objects (i.e. the binary millisecond pulsar J1719$-$1438 having a companion with a Jovian mass) and one (the pulsar J1017$-$7156) which is a very good timer and already entered in the list of the targets of the Parkes Pulsar Timing Array (PPTA).

However, since most of the millisecond pulsars are enclosed in binary systems, their periodic radio signal suffers the effects of the orbital motion, and can be missed by the standard codes which are used for searching for isolated radio pulsars and which have been applied for performing the first aforementioned processing of the data of the HTRU survey.

In view of that, within the framework of the HTRU experiment, the part of my job which is reported in this Thesis developed along the following items.

FIRSTLY it consisted in investigating which kind of pulsar binary systems (and quantifying with which probability) may have been missed by the first processing of the data.

From this study, it resulted that $\sim 90\%$ of the binary systems including a millisecond pulsar and a white dwarf in orbit shorter than ~ 2 hours may have not being discovered by a standard pulsar search pipeline. Similar probabilities ($\sim$ 80% of missed objects) applies to the case of mildly recycled pulsars in double neutron star systems. Also long period ordinary pulsars could have escaped discovery (with a probability of $\sim 40\%$) by using the standard codes, if they are orbiting around a stellar mass black-hole companion in a shorter than 2 hours orbit.

SECONDLY my work was devoted to set up and test a binary pulsar search pipeline, suitable for discovering - during a reprocessing of the data of the med-latitude survey - those millisecond pulsars in tight binaries which may have escaped discovery in the first data search.

The chosen pipeline involves using a public available suite of codes (PRESTO) for performing a search for accelerated pulsars (i.e. pulsar in a binary) in the frequency domain, namely the so-called correlation search. The pipeline was tested with few known binary pulsars, and in particular on the famous Double Pulsar PSR J0737−3039, which represents a stereotype for a relativistic binary in a very close orbit. It was checked that this pulsar would have been missed by the standard pipeline, whereas it displayed nicely in the new *accelerated pipeline.*

THIRDLY my study focused on implementing the aforementioned pipeline on two large clusters of CPUs (one located at the INAF-Osservatorio di Cagliari and the other at the Swinburne University of Technology in Melbourne, Australia) and carrying on the initial part of the data reprocessing and inspection of the candidates.

Given the very heavy computational effort required by the new pipeline, up

to now only about 3% of the data have been processed using that, allowing me to discover a new 2.7 ms millisecond pulsar, PSR J1832−0835, which indeed turned out to be solitary (of course the code is sensitive both to binary and isolated pulsars). Follow-up observations of this pulsar provided the positional and rotational parameters of the source, and revealed that it may also be a new good source for the use in the Pulsar Timing Arrays. However this provisional classification will have to be confirmed after few additional months of observations. This new millisecond pulsar, as well as 4 other millisecond pulsars discovered and followed-up in the HTRU survey, will be presented in a paper in preparation.

Furthermore, the data reduction with the new pipeline is keeping on and will be significantly boosted by the availability in few months of new larger clusters of computers at the collaboration sites, as well as the adoption of a new version of the searching codes, suitable for being used on machine based on Graphical Processor Units (GPUs).

However, besides the long timescale necessary for running the new pipeline, a problem emerged in the course of the visual inspection of the pulsar candidates of the med-survey data (both from the standard pipeline and the new pipeline), i.e. their huge number, mainly due to the overwhelming presence of Radio Frequency Interference (RFI), which makes a human selection very long and tedious, if not almost unfeasible in the case of the pipeline for accelerated pulsar candidates. As a consequence of this consideration, my work

FOURTHLY aimed to test the use of an Artificial Neural Network (ANN) in an attempt to solve the aforementioned problem, i.e. trying to make a machine able to distinguish between pulsars and non-pulsars among the candidates.

The possible discovery of new pulsars was not the only goal of such a work: in fact, the aim was also to create a large and novel training set by which improving the training of future machine learning algorithms. In order to reach all these goals, the following method made up of two parts was employed: (i) the investigation of the consequences of random number generation, for which reason a 'pseudo-committee' of ten ANNs, rather than a single neural net, was trained and then applied to all the HTRU survey candidates; (ii) the exploration of the use of different ratios of pulsars and non-pulsars in the training set of

the nets, exploiting this ratio to fine tune the rate of false positives (i.e. the non-pulsars wrongly classified as pulsars by the nets) among the candidates. However, it should be noted that the reduction of such a rate was obtained at the expense increased false negatives (i.e. the pulsars misclassified as non-pulsars), although in an acceptable percentage.

Owing to this method, the number of candidates for visual inspection (among all those yielded by the standard pipeline from the processing of almost all the med-latitude survey data), reduced from more than 19 millions to $\sim$ 41,000, i.e. a human manageable number. These 41,000 candidates are currently still being inspected by eye, with the hope of new pulsar discoveries.

Moreover, the candidates under inspection are being divided into three categories, pulsars, non-pulsars and 'confusion' (i.e. those it is difficult to make a decision about); once the inspection is completed, these categories can be used in different ways to produce an improved training set for the nets. For example, the confusion category could be excluded from it , in order to avoid error in the training. Another possibility, which could maybe produce better results, is to switch to a different approach, choosing to train the nets to recognise RFI instead of pulsars, and hence creating the new training set divided into RFI and non-RFI: the RFI class made up of the non-pulsar category, the non-RFI class made up of the pulsar and confusion categories. This approach could lead to reduce both false positive and false negative rates, yielding better trained ANNs.

Besides the publication (which is in progress) of the results of this investigation, the plan is to make the obtained training set public, so that it can be available to all the pulsar community.

Bibliography

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

M. Bailes, S. D. Bates, V. Bhalerao, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. D'Amico, S. Johnston, M. J. Keith, M. Kramer, S. R. Kulkarni, L. Levin, A. G. Lyne, S. Milia, A. Possenti, L. Spitler, B. Stappers, and W. van Straten. Transformation of a Star into a Planet in a Millisecond Pulsar Binary. *Science*, 333:1717–, September 2011. doi: 10.1126/science.1208890.

A. Bain. *Mind and Body: The Theories of Their Relation*. D. Appleton and Company, New York, 1873.

N. M. Ball and R. J. Brunner. Data Mining and Machine Learning in Astronomy. *International Journal of Modern Physics D*, 19:1049–1106, 2010. doi: 10.1142/S0218271810017160.

S. D. Bates. *Surveys of the galactic plane for pulsars*. PhD thesis, University of Manchester, 2010.

S. D. Bates, M. Bailes, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. D'Amico, A. Jameson, S. Johnston, M. J. Keith, M. Kramer, L. Levin, A. Lyne, S. Milia, A. Possenti, B. Stappers, and W. van Straten. The High Time Resolution Universe Pulsar Survey - II. Discovery of five millisecond pulsars. *MNRAS*, 416:2455–2464, October 2011. doi: 10.1111/j.1365-2966.2011.18416.x.

C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford, 1995.

M. Burgay, N. D'Amico, A. Possenti, R. N. Manchester, A. G. Lyne, B. C. Joshi, M. A. McLaughlin, M. Kramer, J. M. Sarkissian, F. Camilo, V. Kalogera, C. Kim, and D. R. Lorimer. An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system. *Nature*, 426:531–533, December 2003.

B. F. Burke and F. G. Smith. *An introduction to Radio Astronomy*. Cambridge University Press, Cambridge, 2002.

S. Burke-Spolaor, M. Bailes, S. Johnston, S. D. Bates, N. D. R. Bhat, M. Burgay, N. D'Amico, A. Jameson, M. J. Keith, M. Kramer, L. Levin, S. Milia, A. Possenti, B. Stappers, and W. van Straten. The High Time Resolution Universe Pulsar Survey - III. Single-pulse searches and preliminary analysis. *MNRAS*, 416:2465–2476, October 2011. doi: 10.1111/j.1365-2966.2011.18521.x.

F. Camilo, D. R. Lorimer, P. Freire, A. G. Lyne, and R. N. Manchester. Observations of 20 millisecond pulsars in 47 Tucanae at 20 cm. *ApJ*, 535: 975–990, 2000.

B. W. Carroll and D. A. Ostlie. *An introduction to modern astrophysics*. Reading, Mass. : Addison-Wesley Pub., c1996., 1996.

J. M. Cordes and D. F. Chernoff. Neutron Star Population Dynamics. I. Millisecond Pulsars. *ApJ*, 482:971–992, 1997.

J. M. Cordes and T. J. W. Lazio. NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations. *ArXiv Astrophysics e-prints: astro-ph/0207156*, July 2002.

J. M. Cordes and R. M. Shannon. Minimum Requirements for Detecting a Stochastic Gravitational Wave Background Using Pulsars. *ArXiv e-prints*, June 2011.

P. Demorest, J. Lazio, A. Lommen, A. Archibald, Z. Arzoumanian, D. Backer, J. Cordes, P. Demorest, and et al. Gravitational Wave Astronomy Using Pulsars: Massive Black Hole Mergers & the Early Universe. In *AGB Stars and Related Phenomena Astro2010: The Astronomy and*

Astrophysics Decadal Survey, volume 2010 of *Astronomy*, page 64, 2009. arXiv:0902.2968v1.

A. A. Deshpande and J. M. Rankin. Pulsar magnetospheric emission mapping: Images and implications of polar CAP weather. *ApJ*, 524:1008–1013, 1999.

R. P. Eatough. *A search for relativistic binary pulsars in the galactic plane*. PhD thesis, University of Manchester, 2009.

R. P. Eatough, N. Molkenthin, M. Kramer, A. Noutsos, M. J. Keith, B. W. Stappers, and A. G. Lyne. Selection of radio pulsar candidates using artificial neural networks. *MNRAS*, 407:2443–2450, October 2010. doi: 10.1111/j.1365-2966.2010.17082.x.

D. L. Elliott. A better activation function for artificial neural networks. ISR Technical Report TR 93-8, Institute for Systems Research, University of Maryland, 1993.

B. Farley and W.A. Clark. Simulation of Self-Organizing Systems by Digital Computer. *IRE Transactions on Information Theory*, 4:76–84, 1954. doi: 10.1109/TIT.1954.1057468.

A. J. Faulkner, I. H. Stairs, M. Kramer, A. G. Lyne, G. Hobbs, A. Possenti, D. R. Lorimer, R. N. Manchester, M. A. McLaughlin, N. D'Amico, F. Camilo, and M. Burgay. The Parkes Multibeam Pulsar Survey - V. Finding binary and millisecond pulsars. *MNRAS*, 355:147–158, November 2004.

K. Fukushima. Cognitron: A self-organizing multilayered neural network. *Biological Cybernetics*, 20:121–136, 1975. doi: 10.1007/BF00342633.

J. A. Gil, A. Jessner, and M. Kramer. Are there really planets around PSR B1257+12. *A&A*, 271, 1993.

P. Goldreich and W. H. Julian. Pulsar electrodynamics. *ApJ*, 157:869–880, 1969.

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

D. Hebb. *The Organization of Behavior*. Wiley, New York, 1949.

A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a rapidly pulsating radio source. *Nature*, 217:709–713, 1968.

G. B. Hobbs, M. Bailes, N. D. R. Bhat, S. Burke-Spolaor, D. J. Champion, W. Coles, A. Hotan, F. Jenet, and et al. Gravitational-Wave Detection Using Pulsars: Status of the Parkes Pulsar Timing Array Project. *PASA*, 26:103–109, June 2009. doi: 10.1071/AS08023.

J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA, Biophysics*, April 1982.

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

C. Ingel and M. Hüsken. Improving the Rprop learning algorithm. In H. Bothe and R. Rojas, editors, *Proceedings of the Second International ICSC Symposium on Neural Computation (NC 2000)*, pages 115–121, Berlin, 2000. ICSC Academic Press.

J. D. Jackson. *Classical Electrodynamics*. Wiley, 1962.

M. Jahan Miri and D. Bhattacharya. Magnetic evolution of neutron stars in wide low-mass binary systems. *MNRAS*, 269:455–461, 1994.

W. James. *The Principles of Psychology*. H. Holt and Company, New York, 1890.

V. M. Kaspi, J. H. Taylor, and M. Ryba. High-precision timing of millisecond pulsars. III. Long-term monitoring of PSRs B1855+09 and B1937+21. *ApJ*, 428:713–728, 1994.

M. J. Keith, R. P. Eatough, A. G. Lyne, M. Kramer, A. Possenti, F. Camilo, and R. N. Manchester. Discovery of 28 pulsars using new techniques for sorting pulsar candidates. *MNRAS*, 395:837–846, May 2009. doi: 10.1111/j.1365-2966.2009.14543.x.

M. J. Keith, A. Jameson, W. van Straten, M. Bailes, S. Johnston, M. Kramer, A. Possenti, S. D. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. D'Amico, L. Levin, P. L. McMahon, S. Milia, and B. W. Stappers. The High Time Resolution Universe Pulsar Survey - I. System configuration and initial discoveries. *MNRAS*, 409:619–627, December 2010. doi: 10.1111/j.1365-2966.2010.17325.x.

M. J. Keith, S. Johnston, M. Bailes, S. D. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. D'Amico, A. Jameson, M. Kramer, L. Levin, S. Milia, A. Possenti, B. W. Stappers, W. van Straten, and D. Parent. The High Time Resolution Universe Pulsar Survey - IV. Discovery and polarimetry of millisecond pulsars. *MNRAS*, page 1783, November 2011. doi: 10.1111/j.1365-2966.2011.19842.x.

M. Kramer, I. H. Stairs, R. N. Manchester, M. A. McLaughlin, A. G. Lyne, R. D. Ferdman, M. Burgay, D. R. Lorimer, A. Possenti, N. D'Amico, J. M. Sarkissian, G. B. Hobbs, J. E. Reynolds, P. C. C. Freire, and F. Camilo. Tests of General Relativity from Timing the Double Pulsar. *Science*, 314: 97–102, October 2006. doi: 10.1126/science.1132305.

L. Levin, M. Bailes, S. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, N. D'Amico, S. Johnston, M. Keith, M. Kramer, S. Milia, A. Possenti, N. Rea, B. Stappers, and W. van Straten. A Radio-loud Magnetar in X-ray Quiescence. *ApJ*, 721:L33–L37, September 2010. doi: 10.1088/2041-8205/721/1/L33.

D. R. Lorimer. Binary and Millisecond Pulsars at the New Millenium. *Living Reviews in Relativity*, 2001. http://www.livingreviews.org/Articles/Volume4/2001-5lorimer.

D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Cambridge University Press, 2005.

D. R. Lorimer, A. J. Faulkner, A. G. Lyne, R. N. Manchester, M. Kramer, M. A. McLaughlin, G. Hobbs, A. Possenti, I. H. Stairs, F. Camilo, M. Burgay, N. D'Amico, A. Corongiu, and F. Crawford. The Parkes

Multibeam Pulsar Survey - VI. Discovery and timing of 142 pulsars and a Galactic population analysis. *MNRAS*, 372:777–800, October 2006.

A. G. Lyne and R. N. Manchester. The shape of pulsar radio beams. *MNRAS*, 234:477–508, 1988.

A. G. Lyne and F. G. Smith. *Pulsar Astronomy, 3rd ed.* Cambridge University Press, Cambridge, 2005.

A. G. Lyne, M. Burgay, M. Kramer, A. Possenti, R. N. Manchester, F. Camilo, M. A. McLaughlin, D. R. Lorimer, N. D'Amico, B. C. Joshi, J. Reynolds, and P. C. C. Freire. A double-pulsar system: A rare laboratory for relativistic gravity and plasma physics. *Science*, 303:1153–1157, 2004.

R. N. Manchester. The Parkes Pulsar Timing Array Project. *Highlights of Astronomy*, 15:233–233, November 2010. doi: 10.1017/S1743921310008987.

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

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

D. N. Matsakis, J. H. Taylor, and T. M. Eubanks. A statistic for describing pulsar and clock stabilities. *A&A*, 326:924–928, 1997.

W. McCullock and W. Pitts. A Logical Calculus of Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943. doi: 10.1007/BF02478259.

T. G. Mdzinarishvili and G. I. Melikidze. On the z-distribution of pulsars. *A&A*, 425:1009–1012, October 2004. doi: 10.1051/0004-6361:20034410.

L. Mestel and M. H. L. Pryce. On the construction of a simple model pulsar magnetosphere. *MNRAS*, 254:355–360, January 1992.

M. S. Minsky and A. Papert. *An Introduction to Computational Geometry*. MIT Press, 1969.

S. Nissen. Implementation of a fast artificial neural network library (fann). Report, Department of Computer Science, University of Copenhagen (DIKU), 2003.

S. Osłowski, W. van Straten, G. B. Hobbs, M. Bailes, and P. Demorest. High signal-to-noise ratio observations and the ultimate limits of precision pulsar timing. *MNRAS*, 418:1258–1271, December 2011. doi: 10.1111/j.1365-2966.2011.19578.x.

I. Porceddu and N. D'Amico. The Cybersar project . INAF activities. *Memorie della Societa Astronomica Italiana Supplementi*, 13:128, 2009.

J. M. Rankin. Toward an empirical theory of pulsar emission. VI. The geometry of the conal emission region. *ApJ*, 405:285–297, 1993.

S. M. Ransom, L. J. Greenhill, J. R. Herrnstein, R. N. Manchester, F. Camilo, S. S. Eikenberry, and A. G. Lyne. A binary millisecond pulsar in globular cluster NGC 6544. *ApJ*, 546:L25–L28, January 2001.

S. M. Ransom, S. S. Eikenberry, and J. Middleditch. Fourier techniques for very long astrophysical time-series analysis. *AJ*, 124:1788–1809, 2002.

S. M. Ransom, J. M. Cordes, and S. S. Eikenberry. A New Search Technique for Short Orbital Period Binary Pulsars. *ApJ*, 589:911–920, 2003.

M. Riedmiller and H. Braun. A direct adaptive method for faster backpropagation learning: the RPROP algorithm. In Ruspini E. H., editor, *Proceedings of the IEEE International Conference on Neural Networks*, pages 586–591, San Francisco, USA, 1993. IEEE Press.

N. Rochester, J.H. Holland, L.H. Habit, and W.L. Duda. Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Transactions on Information Theory*, 2:80–93, 1956. doi: 10.1109/TIT.1956.1056810.

F. Rosenblatt. The Perceptron: A Probalistic Model For Information Storage And Organization In The Brain. *Psychological Review*, 65:386–408, 1958. doi: 10.1037/h0042519.

D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, and PDP Research Group, editors, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1 of *Foundations*, Cambridge, MA, 1986. MIT Press.

P. A. G. Scheuer. Amplitude variations of pulsed radio sources. *Nature*, 218: 920–922, 1968.

A. Sesana and A. Vecchio. Gravitational waves and pulsar timing: stochastic background, individual sources and parameter estimation. *Classical and Quantum Gravity*, 27(8):084016, April 2010. doi: 10.1088/0264-9381/27/8/084016.

I. I. Shapiro. Fourth test of general relativity. *Phys. Rev. Lett.*, 13:789, 1964.

I. S. Shklovskii. Possible causes of the secular increase in pulsar periods. *Sov. Astron.*, 13:562–565, 1970.

L. Staveley-Smith, W. E. Wilson, T. S. Bird, M. J. Disney, R. D. Ekers, K. C. Freeman, R. F. Haynes, M. W. Sinclair, R. A. Vaile, R. L. Webster, and A. E. Wright. The Parkes 21 cm multibeam receiver. *PASA*, 13:243–248, 1996.

J. H. Taylor. A sensitive method for detecting dispersed radio emission. *A&AS*, 15:367–369, 1974.

J. P. W. Verbiest, M. Bailes, W. van Straten, G. B. Hobbs, R. T. Edwards, R. N. Manchester, N. D. R. Bhat, J. M. Sarkissian, B. A. Jacoby, and S. R. Kulkarni. Precision timing of PSR J0437-4715: an accurate pulsar distance, a high pulsar mass and a limit on the variation of Newton's gravitational constant. *ApJ*, 679:675–680, 2008. doi: 10.1086/529576.

J. P. W. Verbiest, M. Bailes, W. A. Coles, G. B. Hobbs, W. van Straten, D. J. Champion, F. A. Jenet, R. N. Manchester, N. D. R. Bhat, J. M.

Sarkissian, D. Yardley, S. Burke-Spolaor, A. W. Hotan, and X. P. You. Timing stability of millisecond pulsars and prospects for gravitational-wave detection. *MNRAS*, 400:951–968, December 2009. doi: 10.1111/j.1365-2966.2009.15508.x.

P.J. Werbos. *Beyond regression: new tools for prediction and analysis in the behavioral sciences*. PhD thesis, Harvard University, 1974.

B. Widrow and M. E. Jr. Hoff. Adaptive switching circuits. In *1960 IRE WESCON Convention Record*, volume 4, pages 96–104, New York, 1960. IRE.

Acknowledgements

This work made use of results produced by the Cybersar Project managed by the Consorzio CyberSAR, a project co-funded by the Italian Ministry of University and Research (MIUR) within the Programma Operativo Nazionale 2000-2006 'Ricerca Scientifica, Sviluppo Tecnologico, Alta Formazione' per le Regioni Italiane dell' Obiettivo 1 (Campania, Calabria, Puglia, Basilicata, Sicilia, Sardegna) Asse II, Misura II.2 'Società dell'Informazione', Azione a 'Sistemi di calcolo e simulazione ad alte prestazioni'. More information is available at http://www.cybersar.com.

Observations have been carried out at the Parkes Observatory, that is part of the Australia Telescope which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO.

I acknowledge INAF, the Cagliari Observatory and the University of Cagliari for financial support. I also thank the Cagliari Observatory for the hospitality, and in particular Dr. Andrea Possenti and Prof. Nichi D'Amico.

I would like to acknowledge the Australian Government for allowing my stay in Melbourne in the framework of the Endeavour Award program: I am truly grateful for the opportunity this program has provided.

I am really grateful also to all the people who supported me in the last three years.

First of all, I wish to thank my supervisors Dr. Andrea Possenti and Dr. Marta Burgay at the Cagliari Astronomical Observatory, for making me fully live the experience of a PhD: these three years have been very intense, sometimes difficult but also exciting and rich in events, adventures and encounters with interesting people. I learnt so many things, deepening my scientific knowledge and not only. Thank you for making all of this come true, and for your constant encouragement which has been of vital importance to make me achieve my PhD

goals. Many thanks to Andrea also for kindly guiding me in my work choices.

I really want to thank Dr. Willem van Straten of the Swinburne University of Technology, who supervised me while I was in Australia, for all the things he taught me and for his great availability, kindness and patience, and sense of humor as well (I apologise with him for not accomplishing the point number 6 of our list, 'find BH-PSR system'. I am sorry!).

Thanks to my tutor, Prof. Luciano Burderi, for his support at the University of Cagliari and for solving lots of bureaucratic issues.

A huge 'thank you!' to all the members of the HTRU collaboration: it has been an honour and a pleasure to meet and work with so nice people! Thanks for warmly welcoming me at the beginning of my PhD and for your kindness during all these years: I am not able to explain how much this meant for me. In particular, I thank Prof. Matthew Bailes also for allowing me to live my Australian experience at Swinburne: it has been very important for my life, and I will not forget it. Thank you so much. I also thank Dr. Mike Keith, who kindly gave the cue for my work with the Neural Nets.

I would like to acknowledge Prof. Monica Colpi for her kindness and enthusiasm towards my work, and for valuable comments and suggestions.

I am very grateful to Dr. Ralph Eatough for useful discussions about my Neural Net work and for his great kindness.

Many thanks to all the people at the Cagliari Observatory and University, for their friendship and valuable help in many occasions. In particular, I thank my special friends Noemi and Valentina, always by my side in both happy and sad situations, and whose inestimable support helped me to reach the end of my PhD (thank you, girls! You are the best friends one could wish for); thanks to Alessandro for taming my computer lots of times and for teaching me many tricks; to Paolo for being a black belt in finding the funny side of every situation, managing to ease my spirit in many occasions; thanks to Federica, Matteo, Antonella and Silvia for taking care of me when I needed, to Antonella and Riccardo for their help with the cluster and to Antonio for being my Jedi Master of Linux and introducing me to the secrets of the penguin's Force.

Special thanks go to all the Swinburne gang: it was so nice and funny to meet you, guys! Thanks for making me feel a part of your group, and for helping me every time that I had a problem. In particular, I thank my dear

friend Lina for taking care of me on countless occasions, helping me so much both on a working and personal level: you are an angel. Many thanks also to my friend Paul, for his sensitivity and for making me laugh lots of times, at least as soon as I started to understand his English (since he refuses to speak with the Oxford accent). I also thank Lina and Graeme for always helping me to catch the right tram, so that I did not get lost in Melbourne too many times.

I really want to thank my family and all my friends for tolerating to see me next to never the last three years without disowning and/or disinheriting me: your patience and constant support have been of fundamental importance, and contributed to yield this thesis work. Thank you so much.

Many thanks also to the gorgeus staff at the Parkes Observatory, for their warm heart and for taking care of and delightfully spoiling me every time that I went there: I will always remember all of you, with affection and gratitude (I will also remember all the bugs and spiders that were there in summer, but them *without* either affection or gratitude). I wish also to thank all the people I observed with, for interesting chats, card games and falls off bike.

Finally, I thank my friend Marco for granting the outlandish requests of a tired brain and taking the wonderful picture at the beginning of this thesis, in a cold winter night in company with mouflons and wild boars: *Deu ti du paghidi*[7].

In general, I want to acknowledge all the people who helped me to reach this important goal, letting me to add another piece to the puzzle of my life. Thanks all from the heart.

[7]Typical expression of gratitude in Sardinian language, meaning ‘May God repay you for that’.