

**Southern Galactic Pulsars with MeerKAT**  
**Surveys, Timing, and Fun with**  
**Old and New Discoveries**

Dissertation  
zur  
Erlangung des Doktorgrades (Dr. rer. nat.)  
der  
Mathematisch-Naturwissenschaftlichen Fakultät  
der  
Rheinischen Friedrich-Wilhelms-Universität Bonn

von  
Miquel Colom i Bernadich  
aus  
Barcelona, Catalunya

Bonn, 27.11.2024

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen  
Friedrich-Wilhelms-Universität Bonn

Gutachter/Betreuer:	Prof. Dr. Michael Kramer
Gutachter:	Prof. Dr. Norbert Langer
Tag der Promotion:	10.09.2025
Erscheinungsjahr:	2025

---

## Funding acknowledgements

---

My PhD research was funded by the Max-Planck-Gesellschaft (MPG), of which the MPIfR is a member. The research uses data from the MeerKAT and Parkes/Murriyang radio telescopes. The MeerKAT telescope is operated by the South African Radio Astronomy Observatory (SARAO), which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. The Parkes Observatory is part of the Australia Telescope National Facility (ATNF), which is funded by the Australian Government for operation as a National Facility managed by the Commonwealth Scientific and Industrial Research Organisation (CSIRO). This research uses data products from the following MeerKAT User-Supplied Equipments: the Filterbanking Beamformer User-Supplied Equipment (FBUSE), the Accelerated Pulsar Search User-Supplied Equipment (APSUSE), and the Pulsar Timing User-Supplied Equipment (PTUSE). FBUSE, APSUSE, and one of the four PTUSE machines are funded and provided by the MPIfR and the MPG. This research also uses data from the MeerKAT S-band receivers. The MeerKAT S-band receivers were built at the MPIfR, and funded by the MPIfR and the MPG.





---

## English summary

---

Where to start explaining this thesis? It is a thesis about astronomy. But instead of using optical telescopes in the mountains, it uses large radio antennas as observatories. I observe pulsars, highly magnetized fast-rotating neutron stars, not by looking at images but by studying the time evolution of their signal. The main topics of my thesis are searching new pulsars with the MeerKAT and Parkes/Murriyang telescopes and measuring their rotation and orbits when found in binary systems. By measuring their orbits we can also measure the masses of the stars, which provide us with very valuable information about the astrophysics of the systems and their evolution. Being very aware that these words may seem incomprehensible, I will go step by step.


### Neutron stars and pulsars

Stars like our Sun are kept stable by the balance between the force of gravity and the radiation generated by nuclear fusion in their cores. When the nuclear fuel runs out after millions or billions of years, gravity is left without a counterweight, and stars collapse on themselves. In stars with more than eight times the mass of the Sun, this collapse is catastrophic: the stellar core contracts uncontrollably and its density grows exponentially. During this process, the free electrons of the plasma are absorbed by the protons in atomic nuclei, forming neutrons. This way, all matter in the stellar core is quickly converted into neutrons. But neutrons, following the Pauli exclusion principle, cannot occupy the same location and energy states as other neutrons, producing a strong repulsion. Suddenly, the collapse is halted, and a shock wave and the neutrino wind emitting the process impacts onto the outer layers of the star, ejecting them in a violent explosion. It's a supernova, and it leaves behind a neutron star, the star's old core, naked and exposed to space.

Neutron stars have between one and three solar masses and a radius no larger than 15 kilometers. This makes them the second most compact objects in the Universe, just behind stellar-mass black holes. And while they are not black holes, their density and gravity are so extreme they would collapse into one if they were only slightly more massive. They are also the most magnetized objects in the Universe, with magnetic fields on their surfaces of intensities between  $10^8$  and  $10^{15}$  Gauss. For reference, the intensity of the Earth's magnetic field is one Gauss, and that of a refrigerator magnet 50 Gauss. Finally, and most relevant to this thesis, they can spin over their own axis in a matter of seconds or milliseconds, a property that turns them into pulsars.

Pulsars are the observational manifestation of the spin of neutron stars. Their magnetic fields, anchored to the surface, rotate with it. From its magnetic poles, two rotating beams of radio waves are emitted, producing pulses that we can detect if they reach Earth. Radio telescopes observe these periodic pulses in the same way that sailors perceive the periodic signal of a lighthouse on the horizon.

## Lighthouse:



## Pulsar:




Figure 1: Diagram comparing lighthouses and pulsars. Although the mechanisms of light emission are very different, for practical purposes, they behave in very similar and completely analogous way.

In fact, the arrival rate of the pulses is even more stable and regular than a lighthouse, because they originate in the rotation of a neutron star in the vacuum of space, without friction or impediments. In this way, a pulsar is also a clock that we use as a reference to measure experiments.

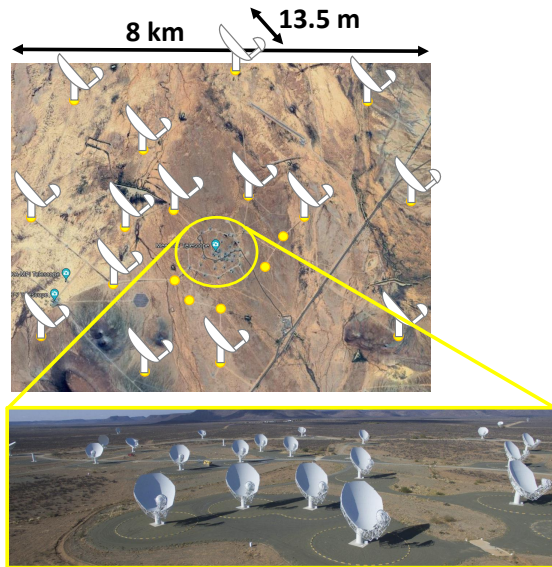
As of today, approximately 4,000 pulsars are known in the Milky Way and its immediate neighbourhood. The population of pulsars is very diverse, but it can be divided into two main groups. The vast majority, 90%, are regular or *canonical* pulsars that were born and spin in isolation. Their magnetic fields are stronger than  $10^{10}$  Gauss and their rotation periods of about 50 milliseconds or greater. The remaining 10% are *recycled* pulsars that have formed in binary systems. These pulsars have or have had a companion star in orbit, from which they have captured matter during their evolution. This *recycling* process accelerates its spin and reduces its magnetic field below  $10^{10}$  Gauss, producing *millisecond pulsars* with rotation periods as short as just a few milliseconds.

## Radio telescopes

I have mentioned that radio pulsars are observed with radio telescopes. A radio telescope is a large radio antenna, similar to parabolic TV antennas, but tens of meters in size. They capture radio waves, light with wavelengths on the order of centimetres, and record the signal as a function of time. When we observe pulsars, their signals are detected as periodic pulses.

The telescopes I have used in this thesis are the 64 antennas of the MeerKAT radio telescope in the Karoo desert, South Africa, and the 64-meter diameter Murriyang parabolic antenna at Parkes Observatory, Australia. Both telescopes are located in the Southern Hemisphere, which is key to their success. Murriyang was built in 1961 and has been one of the best tools for discovering pulsars, with

## MeerKAT (South Africa)



## Parkes/Murriyang (Australia)



Figure 2: Images of the MeerKAT and Parkes telescopes. In the case of MeerKAT, since it is an array of antennas, a map of the distribution of antennas (image taken from Google Maps) and a photograph of the antennas at the center of the array are shown. Image credits: SARAO and David McClenaghan/CSIRO

more than a thousand discoveries in regions of the sky that powerful radio telescopes in the Northern Hemisphere cannot observe. MeerKAT, built in 2017, is the state-of-the-art of radio astronomy in the Southern Hemisphere, in a way the successor to Murriyang, being able to detect signals four times fainter than Murriyang in the same observation time. It is also an interferometer that uses multiple detectors in different positions. Its 64 antennas of 13.5 meters each are distributed in an area of 8 kilometers in size, allowing it to triangulate the position of the signals in the sky with great precision.

## Measuring pulsars

The spin of millisecond pulsars is even more stable than that of canonical pulsars, and so they are the best clocks available for our experiments. In binary systems, their orbital motions are measurable thanks to the Doppler effect, which induces apparent changes in the periodicity of their pulses. Just as the siren of an ambulance sounds higher-pitched when running towards us and lower-pitched when moving away, the rate of rotation of a pulsar in a binary system is observed to be higher when the pulsar is hurtling towards us and lower when it goes away. This can be measured very precisely with the method of *pulsar timing*. With this method, we build an orbital and spin model of the pulsar that predicts the expected times of arrival of each pulse at the telescope. The expected arrival times are then compared to the measured ones, and corrections are made to the orbital and spin model until the predictions match the measurements.

Timing enables many experiments. Its precision is so high that, in the most compact binary systems, we cannot describe the pulsar's orbit with simple Newtonian physics but must incorporate elements of




Figure 3: Doppler effect illustrated for both an ambulance moving in a straight line and a pulsar moving in an orbit. Both the speed of sound and the speed of light are constant, but the movement of the emitter contracts or spaces out the waves or pulses, changing the perception of the listener/observer.

Albert Einstein's general relativity. These elements depend on the orbital parameters and the masses of the stars in the binary system. Therefore, measuring them allows us to obtain an estimate of the masses. In this way, we can measure the masses of neutron stars and the white dwarfs they orbit with, which allows us to study the distribution of stellar masses or check if binary evolution theories are correct. For example, there has been speculation on whether millisecond pulsars are more massive than canonical pulsars because of the mass accreted during recycling. Measuring pulsars masses is therefore essential to verify our understanding of this process. There are also theoretical predictions that relate the orbital parameters and the pulsar spin to the mass of the companion stars. These theories are as diverse as the categories of binary systems with pulsars in existence. Therefore, it is important to discover and measure many of them for a complete understanding. Finally, precise measurements also result in experiments of fundamental physics. The very precisely measured masses of neutron stars can be used to test models of condensed matter and test whether general relativity is correct in extreme gravity situations. For example, the stellar masses should be consistent with the orbital decay due to the emission of gravitational waves as predicted by general relativity, a phenomenon that is also measurable by the timing of pulsars.

## The projects in this thesis

This thesis has six chapters. Chapters 1 and 2 present a scientific and methodological introduction to the topics covered so far in greater detail, while Chapters 3, 4, and 5 describe an array of experiments performed with pulsars. I detail each of these experiments below.

## Pulsar searching with MeerKAT

To aid in discovering new pulsars with the great sensitivity of MeerKAT, I joined the MPIfR-MeerKAT Galactic Plane Survey (MMGPS), a collaborative project that had been in the planning for a decade before the construction of MeerKAT. The goal of this survey is to observe the southern Galactic plane and discover pulsars too faint for Murriyang. For this, Fourier transforms are performed on time series recorded by the observations, and the characteristic periodic signals of pulsars are searched in that domain. With observations starting in 2020, we have discovered 82 new pulsars so far, 17 of which are recycled ones in binary systems. I was involved in the inspection of candidates, the localisation of discoveries in the sky using MeerKAT's multiple antennas, and the timing several of these discoveries.

I have localised and timed up to six pulsars in binary systems discovered in this survey. Three of them are millisecond pulsars in circular binary systems: **PSR J1306–6043**, **PSR J1108–6329**, and **PSR J1543–5439**, with spin periods of 5.7, 4.3, and 4.3 milliseconds, respectively. Their companion stars are light white dwarfs, with masses between 0.2 and 0.4 solar masses. The precision of the timing has not been sufficient to measure the stellar masses precisely, but future measurements could help test theories of stellar evolution that relate the orbital period to the mass of the companion white dwarf.

I have also analyzed more massive systems. **PSR J1708–4843** and **PSR J1015–5359**, with spin periods of 17 and 21 milliseconds, are partially recycled pulsars in circular orbits with companions of at least 0.5 and 1.2 solar masses, respectively. The companion of **PSR J1015–5359** is the most massive we have discovered in the MMGPS, and it may belong to a class of very massive pulsar-white dwarf binary systems of which only four were known so far, with only two with mass measurements via pulsar timing. Furthermore, with an orbital period of eight hours, it would be the most compact known binary system of this type. This fact amplifies the relativistic effects of its orbit and it could result in experiments on extreme gravity in the future.

## Timing a double neutron star system

Another big project of this thesis has been dedicated to **PSR J1208–5936**, the sixth MMGPS discovery I worked with. Instead of a white dwarf companion, this 28-millisecond recycled pulsar has a neutron star companion, in an orbit just over 12 hours long and a large orbital eccentricity. The eccentricity is indicative of the occurrence of a second supernova after the recycling process, which resulted in a kick to the second-born neutron star that broke the circularity of the orbit. Only about twenty of these systems were known in the Milky Way before its discovery, so it was a very significant contribution.

The timing of **PSR J1208–5936** with MeerKAT has yielded many interesting measurements and science. Thanks to its eccentricity, we have measured the relativistic effects very quickly. The mass of the system adding both stars is  $2.582 \pm 0.006$  solar masses, and the individual masses are  $1.3 \pm 0.3$  solar masses each. Its orbital eccentricity is higher than that of other similar systems with similar orbital periods, and this could be correlated with a massive companion, but the precision of individual mass measurements is not enough to verify this hypothesis. Even so, the accuracy is enough to predict a merging time of 7.2 billion years due the emission of gravitational waves, producing one of the neutron star merger events observed by the gravitational-wave detector LIGO. With this discovery, we have updated the rate of neutron star mergers in the Milky Way to a value of  $25^{+19}_{-9}$  mergers per million years, consistent with past estimates but slightly more reduced. The reduction shows that radio surveys are scanning the sky than we are discovering merging double neutron star systems.

## Timing of a massive system

Besides participating in the MMGPS, I have also done an independent timing experiment on the 34-millisecond recycled pulsar **PSR J1227–6208**. With an orbital period of one week and a very massive binary companion of more than 1.2 solar masses, it is one of the few four exceptionally massive systems known before the discovery of **PSR J1015–5359**. This pulsar was discovered in 2011 with Murriyang, but the telescope’s sensitivity was not sufficient to measure the relativistic effects of its orbit. Therefore, when MeerKAT was inaugurated, it was observed again for more precise timing and mass measurements.

Timing was complicated by interferences from the interstellar medium between the pulsar and Earth, which introduce random variations in the times of arrival of the pulses at the telescope, a phenomenon called *timing noise*. We modelled the characteristic frequencies of these variations with Bayesian methods and created different noise models to obtain appropriate uncertainties in the mass measurements. Combining observations from the Murriyang and MeerKAT telescopes from 2012 to 2023, we have measured the mass of the companion white dwarf to be between 1.2 and 1.5 solar masses, confirming its extremely massive nature. This white dwarf could be very close to the Chandrasekhar limit of 1.38 solar masses, the theoretical maximum mass above which white dwarfs collapse into a neutron stars. This timing measurement is only the third one in systems of this type, and it is therefore a very significant result.

We also found that the mass of the pulsar is probably very high, between 1.2 and 1.7 solar masses. However, the spin period of 34 milliseconds is long compared to other recycled pulsars with spin periods below 10 milliseconds, which have thus acquired much more mass than their companions. Therefore, this tells us that massive neutron stars can exist without deep recycling, a result consistent with the measurement of high masses in other pulsars that are also only partially recycled.

---

# Contents

---

<b>Funding acknowledgements</b>	<b>iii</b>
<b>English summary</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.2 Pulsars: What are they? . . . . .	3
1.2.1 Formation I: stars and white dwarfs . . . . .	4
1.2.2 Formation II: stars and supernovae . . . . .	5
1.2.3 From neutron stars to pulsars I: basic neutrons star properties . . . . .	7
1.2.4 From neutron stars to pulsars II: the magnetic dipole model . . . . .	9
1.2.5 Radio emission I: properties of radio emission . . . . .	10
1.2.6 Radio emission II: interstellar medium effects . . . . .	13
1.2.7 The pulsar population I: properties from spin parameters . . . . .	18
1.2.8 The pulsar population II: a visit to the pulsar zoo . . . . .	19
1.3 Binary pulsars: Why are they so interesting? . . . . .	22
1.3.1 Binary parameters: Keplerian orbits and the mass function . . . . .	23
1.3.2 Binary populations: a visit to the binary zoo . . . . .	25
1.3.3 Binary evolution I: mass transfer and pulsar recycling . . . . .	25
1.3.4 Binary evolution II: the making of the binary zoo . . . . .	29
1.4 Beyond just pulsars: what do we get from binary systems? . . . . .	33
1.4.1 The neutron star mass distribution . . . . .	34
1.4.2 The physics of ultra-dense matter . . . . .	35
1.4.3 Supernova physics and kicks . . . . .	35
1.4.4 The white dwarf mass distribution . . . . .	36
1.4.5 Pulsars and gravitational-wave astronomy . . . . .	37
<b>2 Methods: radio data and pulsars</b>	<b>41</b>
2.1 Before we start . . . . .	41
2.2 Radio astronomy, radio telescopes, and data acquisition . . . . .	41
2.2.1 From radio light to radio data: how do we record radio waves? . . . . .	42
2.2.2 Radioastronomy 101: what does a radio telescope see? . . . . .	43
2.2.3 The telescopes of this thesis: Parkes/Murriyang and MeerKAT . . . . .	46
2.3 Pulsar searching . . . . .	47
2.3.1 Filterbanks: your frequency- and time-resolved friends. . . . .	48

2.3.2	De-dispersion: make many time series . . . . .	48
2.3.3	Single pulse searches: bright pulsars want to be found . . . . .	50
2.3.4	Periodicity searches: time domain, or Fourier domain? . . . . .	50
2.3.5	Acceleration searches: chase the binary pulsars! . . . . .	54
2.3.6	Identification and inspection of candidates . . . . .	56
2.3.7	Radio-frequency interference: the nightmare of radio astronomy . . . . .	58
2.4	Pulsar timing . . . . .	60
2.4.1	Pulsar archives: the backbone of pulsar timing . . . . .	61
2.4.2	Orbital solutions: the low-precision cousins of timing solution . . . . .	61
2.4.3	The principles of pulsar timing I: measuring ToAs . . . . .	66
2.4.4	The principles of pulsar timing II: the timing formula . . . . .	69
2.4.5	Advanced pulsar timing I: the binary timing formula . . . . .	74
2.4.6	Advanced pulsar timing II: biting timing solutions with <i>dracula</i> . . . . .	78
2.4.7	Mass measurements I: physics beyond Isaac Newton . . . . .	81
2.4.8	Mass measurements II: getting proper mass constraints . . . . .	86
<b>3</b>	<b>New Galactic discoveries with MeerKAT: searching, finding, and solving.</b>	<b>89</b>
3.1	Before we start . . . . .	89
3.2	The MPIfR-MeerKAT Galactic Plane survey . . . . .	90
3.2.1	Science objectives: on the hunt for faint binary pulsars . . . . .	90
3.2.2	The MMGPS pipeline: from observations to discoveries . . . . .	91
3.2.3	My contribution to RFI excision during search . . . . .	95
3.2.4	Managing discoveries: confirmation, localisation, and follow-up . . . . .	100
3.3	MMGPS discoveries: an isolated and binary pulsar jamboree . . . . .	105
3.3.1	Details about isolated MMGPS-L discoveries . . . . .	107
3.3.2	The MMGPS-L binary jamboree . . . . .	111
3.3.3	PSR J1306–6043: the tutorial pulsar . . . . .	114
3.3.4	PSR J1208–5936: the pulsar that built CandyCracker and <i>dracula2</i> . . . . .	115
3.3.5	PSR J1708–4843: the vanishing pulsar . . . . .	118
3.3.6	PSR J1108–5936 and PSR J1543–5439: soft-looking but hard-to-bite pulsars . . . . .	122
3.3.7	PSR J1015–5358: the evil twin of PSR J1708–4843 . . . . .	124
3.4	SNR G340.6+0.3 and the MeerKAT S-band receivers: the search for the missing pulsar . . . . .	127
<b>4</b>	<b>The eccentric DNS PSR J1208–5935 and a neutron star merger rate update</b>	<b>131</b>
4.1	Context and contributions . . . . .	131
4.2	Summary of the publication . . . . .	132
4.2.1	Orbital campaign . . . . .	132
4.2.2	Timing and mass measurements . . . . .	133
4.2.3	Science from PSR J1208–5936 . . . . .	133
4.2.4	Updating the neutron star merger rate . . . . .	135
4.2.5	Final remarks and future prospects . . . . .	139
<b>5</b>	<b>PSR J1227-6208: the most massive known PRP–ONeMg WD system?</b>	<b>141</b>
5.1	Context and contributions . . . . .	141



5.2	Summary of the publication	142
5.2.1	Data used in the experiment	142
5.2.2	Profile measurements	143
5.2.3	Timing models and timing noises	143
5.2.4	PK parameters and mass constraints	146
5.2.5	Orbital geometry constraints	146
5.2.6	Astrophysical implications	147
5.2.7	Future prospects	147
<b>6</b>	<b>Conclusions and future prospects</b>	<b>149</b>
6.1	Conclusions	149
6.2	Future prospects	150
6.2.1	For the MMGPS surveys	150
6.2.2	For PSR J1208–5936	151
6.2.3	For PSR J1227–6208	152
6.2.4	For pulsars in the era of MeerKAT+ and the SKA	153
	<b>Epilogue</b>	<b>155</b>
<b>A</b>	<b>The many angles of pulsar astronomy</b>	<b>157</b>
A.1	Sky coordinate systems	157
A.2	Pulsar binary angles	159
A.2.1	Orbital phases: Running orbital angles	159
A.2.2	Orbital geometry: static orbital angles	161
<b>B</b>	<b>Paper on PSR J1208–5936</b>	<b>163</b>
<b>C</b>	<b>Paper on PSR J1227–6208</b>	<b>183</b>
<b>D</b>	<b>Music to spin to</b>	<b>205</b>
D.1	Music to spin to, the list	205
	<b>Bibliography</b>	<b>211</b>
	<b>List of Figures</b>	<b>229</b>
	<b>List of Tables</b>	<b>231</b>



---

# Introduction

---

## 1.1 Motivation

In 1982, Donald Backer and his collaborators announced an unusual discovery. They reported the Arecibo observations of **PSR B1937+21**, a radio pulsar with a period of 1.588 ms (Backer et al., 1982). This implied that the underlying neutron star, the collapsed core of a former massive star that underwent supernova, rotates on its axis once every 1.588 ms. For comparison, this is 842 times faster than the first pulsar discovery **PSR B1919+21** (Hewish et al., 1968), and 21 times quicker than the fastest pulsar known at the time, **PSR B0531+21** at the core of the **Crab Nebula supernova remnant** (the Crab pulsar, Comella et al., 1969). The discovery raised many questions in the heads of pulsar astronomers. How could a neutron star achieve such an incredible rotation rate? This pulsar is rotating so fast that matter on its equator moves at more than 10% of the speed of light, putting it on the edge of tearing itself apart with its own centrifugal force. Was it perhaps an extremely young specimen? But this explanation was not satisfactory because, unlike the Crab pulsar, also the youngest pulsar known at the time, it was not surrounded by an energised supernova remnant nebula (Radhakrishnan and Srinivasan, 1982).

Instead, the explanation turned out exactly opposite. It was not a young pulsar, but an ancient one. The key to this puzzle was a previous Nobel Prize-worthy Arecibo discovery. A bit earlier, in 1975, Russell Hulse and Joseph Taylor discovered **PSR B1913+16**, a 59.03-ms pulsar in an eccentric compact orbit with an unseen companion star (Hulse and Taylor, 1975). This discovery, nicknamed the *Hulse–Taylor pulsar*, was a whole beast in itself. It was special from the start because of the compact nature of the unseen companion, which could only be a neutron star or a black hole. The study of this system led to the confirmation of the existence of gravitational waves predicted by Albert Einstein in 1916 ago via the observation of a shrinking of the orbit (Taylor et al., 1979; Taylor and Weisberg, 1982), leading to a Nobel Prize in 1993, and convincing governments around the world to fund gravitational-wave detectors (such as *LIGO*) whose science fruits we enjoy today.

Putting aside my enthusiasm for the Hulse–Taylor binary, these two pulsars have a characteristic missing in other “normal” pulsars: a highly stable rotational period. Pulsars are not eternal. Instead, their spin periods present subtle but often erratic evolution, with the overall tendency of slowing down over time as they lose energy until their emission turns off millions of years after their birth. However, **PSR B1937+21** and the Hulse–Taylor pulsar had a much more stable spin, which slowed down orders of more slowly than those of other pulsars by orders of magnitude. The difference? Binary evolution:

it was quickly understood that these two pulsars shared a common past involving the death of their orbital companions (Radhakrishnan and Srinivasan, 1982). This kind of special pulsars are now called *recycled* pulsars. The word "recycled" makes sense, as these pulsars received a second life by feeding upon the gas of their ageing companions, bringing them to significantly faster and more stable spin periods (Bhattacharya and van den Heuvel, 1991). In the case Hulse–Taylor pulsar, like in many others, the stripped companion remains as an unseen neutron star, while **PSR B1937+21** drove its unfortunate companion into full evaporation.

As exciting as this story sounds already, this is only the beginning. As already hinted for the Hulse–Taylor binary, recycled pulsars in binary systems are very promising for testing fundamental physics. Owing to their highly stable rotations, astronomers can use the highly precise technique of *pulsar timing* to model their rotation, turning them into laboratories for gravity and binary evolution. While we can not set up experiments with mountain-crushing gravity or with matter with the densities equivalent to those of atomic nuclei without obliterating planet Earth, the Universe has been kind enough to serve us with hundreds of very stable naturally-made clocks in binaries in the radio sky. A few notable examples are the Hulse–Taylor binary, of which I have already enthused enough, the pulsar **PSR J0337+1715** in a triple star system, with which we have tested the strong-equivalence principle of Albert Einstein’s general relativity (Ransom et al., 2014; Voisin et al., 2020), the double pulsar system **PSR J0737–3039A/B**, with which we have made the most precise neutron star mass measurements (with  $90,000\sigma$  significance!), the most stringent tests of gravity up to first post-Newtonian order, and probed the propagation of light in the highest space-time curvature (Lyne et al., 2004; Kramer et al., 2021a), the pulsar–white dwarf binary **PSR J1141–6545**<sup>1</sup>, with which the dragging of space-time due to the rotation of the white dwarf has been measured (Kaspi et al., 2000; Antoniadis et al., 2011; Venkatraman Krishnan et al., 2020), and even just the most and least massive pulsar masses measured in **PSR J0740+6620** and **PSR J0453+1559** (Fonseca et al., 2021; Martinez et al., 2015).

So, with all of this in mind, what are the contributions of the researcher writing this thesis? Well, I have had the opportunity to participate in and lead some projects using MeerKAT pulsar data. MeerKAT, located at the Karoo Desert in South Africa and administered by the South African Radio Astronomy Observatory (**SARAO**), is a Square-Kilometer Array (SKA) precursor and the most sensitive radio interferometer and pulsar facility in the Southern Hemisphere (Jonas and MeerKAT Team, 2016; Bailes et al., 2020). For comparison, it has three times the collecting area of the former title holder, the Parkes/Murriyang telescope from the Australia Telescope National Facility (**ATNF**), which I have also had the fortune to use in my research.

Therefore, MeerKAT is a great opportunity for untapping the unexplored depths of the radio sky in the Southern Galactic plane, a region with great potential for new pulsar binary discoveries due to the high stellar density. For reference, Parkes/Murriyang already discovered more than 1,000 pulsars in this sky region across a variety of surveys: 742 in the Parkes Multibeam Pulsar Survey (Manchester et al., 2001; Morris et al., 2002; Kramer et al., 2003; Hobbs et al., 2004; Faulkner et al., 2004; Lorimer et al., 2007), with 24 further discoveries from re-processing of data from said survey (Knispel et al., 2013), 69 from the Swinburne Intermediate-Latitude Pulsar Survey (Edwards et al., 2001), 18 from the Parkes High Latitude Pulsar Survey (Burgay et al., 2006), 26 from another high latitude survey (Jacoby et al., 2007; Jacoby et al., 2009), and 289 from the High Time Resolution Universe pulsar survey and its multiple re-processing with ever-improving search techniques (Keith

---

<sup>1</sup> J1141–6545 is a bit of an exception: this pulsar is not recycled. Nonetheless, I will show later that this has a very good explanation.

et al., 2010; Bates et al., 2011; Burke-Spolaor et al., 2011; Keith et al., 2012; Bates et al., 2012; Burgay et al., 2013; Ng et al., 2014; Ng et al., 2015; Bates et al., 2015; Cameron et al., 2018; Morello et al., 2019; Burgay et al., 2019; Cameron et al., 2020; Sengar et al., 2022; Wongphechauxsorn et al., 2024). That is not to be joked about: at some point more than half of all known pulsars had been discovered with Parkes/Murriyang, with the other half having been discovered by multiple Northern Hemisphere observatories. The large number of discoveries allowed us to understand the pulsar population to a precision much higher before (Lorimer et al., 2007; Lorimer et al., 2015; Levin et al., 2013; Burgay et al., 2013). Now, it is MeerKAT's turn to explore exploring the depths of the Southern Galactic plane, from the same observing latitude, but with the much improved sensitivity.

MeerKAT is also a great tool for improving timing measurements of already-known pulsars of which observations with less sensitive telescopes struggle to place meaningful constraints. Therefore, in this thesis, I show how MeerKAT has enabled new and exciting pulsar science with the discovery of more than 80 new pulsars in the southern Galactic plane, 17 of which in binaries, during the MPIfR-MeerKAT Galactic Plane Survey (MMGPS). To this end, I describe the techniques used to solve and extract scientific information from the discoveries, and how they fit with the current paradigm of binary pulsar science, with a particular focus on **PSR J1208–5936**, a new eccentric double neutron star system which is merging due to the emission of gravitational waves. I also describe how I used MeerKAT data to significantly improve the precision of mass measurements of the massive pulsar–white dwarf binary **PSR J1227–6028**, leading to a very interesting discussion on how this system fits with the scarce population of pulsars with white dwarf binaries close to the Chandrasekhar limit.

## 1.2 Pulsars: What are they?

If you just read the introduction and are new to pulsars, you must be thinking: “What is a pulsar? How do they form? Can I eat one?” No, you can not eat pulsars. This is because they are fast-rotating, highly-magnetized neutron stars, the collapsed cores of former massive stars that underwent a violent supernova event, for the most part<sup>2</sup>. Containing  $\sim 1.4$  solar masses within a radius of  $\sim 15$  km, neutron stars are the most extreme confirmed objects in the universe besides black holes. Right after the discovery of the neutron particle in 1932 by James Chadwick (Chadwick, 1932), neutron stars were first hypothesised in 1934 by Walter Baade and Fritz Zwicky when they realised that the gravitational collapse of an entire star was needed to explain the large energies observed in cosmic rays reaching Earth. This would occur in a supernova event, and the remaining object would be a star made entirely of neutrons, a *neutron star* (Baade and Zwicky, 1934).

Observational confirmation of this stellar remnant was a rather lucky event: Jocelyn Bell discovered the first pulsar in 1967 while looking for the scintillation of quasars with the Mullard Radio Astronomy Observatory<sup>3</sup>. Pulsar **PSR B1919+21** (back then labelled **CP 1919**) was observed as a very bright periodically pulsating radio source, which turned on every 1.337 s with extreme regularity<sup>4</sup> (Hewish et

<sup>2</sup> Less than a handful of pulsars are actually highly-magnetized, rotating white dwarfs (e.g. Pelisoli et al., 2023). But they exist only in very specific conditions, have very different spin properties and emission mechanisms, and are a topic for another thesis altogether.

<sup>3</sup> Despite this, it was her supervisor and British astronomer Anthony Hewish who received the Nobel Prize for the discovery in 1974. However, Jocelyn Bell is nowadays recognized as the original discoverer.

<sup>4</sup> The regularity of the source was so surprising that the discoverers thought of a possible alien origin, calling it “Little

al., 1968). Not much later, in 1969, the discovery of the 33.392-ms Crab pulsar inside of its eponymous supernova remnant nebula (Comella et al., 1969) settled the idea that such regular periodicity could only originate in the rotation of neutron star born in supernovae, the only astrophysical objects capable of reaching such spin periods without breaking themselves apart due to centrifugal forces.

During this introduction, I have already used concepts such as “white dwarfs”, “neutron stars”, “recycling” or even *pulsars* without going into detail about their meaning. In the following sections, I will show how all of this makes any sense at all, starting with how to produce neutron stars from regular dying stars. For this, we will begin with a brief overview of the life cycle of regular stars and, especially, how they end their lives.

### 1.2.1 Formation I: stars and white dwarfs

Before delving into the formation of neutron stars and pulsars, we must understand the origin of their cousin stellar remnants: *white dwarfs*. This exploration is a necessary, as many aspects of binary stellar evolution in later sections of this chapter will depend on a good understanding of stellar evolution and the formation of white dwarfs. On top of this, pulsars are essentially white dwarfs brought to the extreme.

Stellar evolution is the result of a titanic balance between two opposing forces: their own gravity pushing for collapse against the outward pressure arising from their interiors (e.g. Carroll and Ostlie, 1996). In the Sun and all the stars that we see in the night sky, the outward pressure originates in the release of energy via nuclear fusion, producing starlight as a side-product. As shown in Fig. 1.1, stars spend the majority of their life span in the *main sequence*, fusing hydrogen into helium in their cores (hydrogen core-core burning). As hydrogen exhausts and helium accumulates, nuclear fusion moves to the other shells of the core (hydrogen-shell burning) and the other layers of the star expand in size by a couple orders of magnitude, turning it into a *giant star*. If stars are massive enough, they can move into stages of helium burning into carbon and oxygen, then into carbon burning into neon and magnesium, and so forth with increasingly heavier elements.

However, the star can also be not massive enough. For stars that started less than 8 *solar masses* in the main sequence ( $M_{\text{MS}} < 8 M_{\odot}$ <sup>5</sup>), when the core exhausts its nuclear fuel but is unable to move into the next stages of nuclear fusion, the other layers continue to expand asymptotically until they become gravitational unbound and are expelled to form *planetary nebulae*, leaving the exposed the inert core behind. This is now a *white dwarf*. With the size of planet Earth and with masses of  $M_{\text{WD}} < 1.38 M_{\odot}$ , white dwarfs are already amongst the densest of objects in the Universe. The only force keeping them from collapsing under their own extreme gravity is the electron degeneracy pressure arising from the Pauli exclusion principle, and while they are born very hot and bright ( $\sim 100,000$  K at the surface), they will slowly radiate away all of their excess heat and cool down over billions of years.

As shown in Fig. 1.1, stars in different mass ranges produce white dwarfs of different compositions. The lightest ones are *helium white dwarfs* (He WD,  $M_{\text{WD}} < 0.33 M_{\odot}$ ). However, since light stars burn their fuel slowly, their formation takes longer than the current age of the Universe (13.8 Gyr, Planck Collaboration, 2020), so no isolated He WD has formed yet as of today. Key emphasis on “isolated”, as we actually see plenty of He WD in orbit with recycled pulsars! But let us not get ahead of ourselves. On the other hand, *carbon-oxygen white dwarfs* (CO WD,  $0.33 < M_{\text{WD}} \lesssim 1.1 M_{\odot}$ ) are

---

Green Men”.

<sup>5</sup>  $M_{\odot}$ : 1 solar mass, equal to 333,060 Earth masses or  $1.988 \times 10^{30}$  kg. Stars are pretty massive.

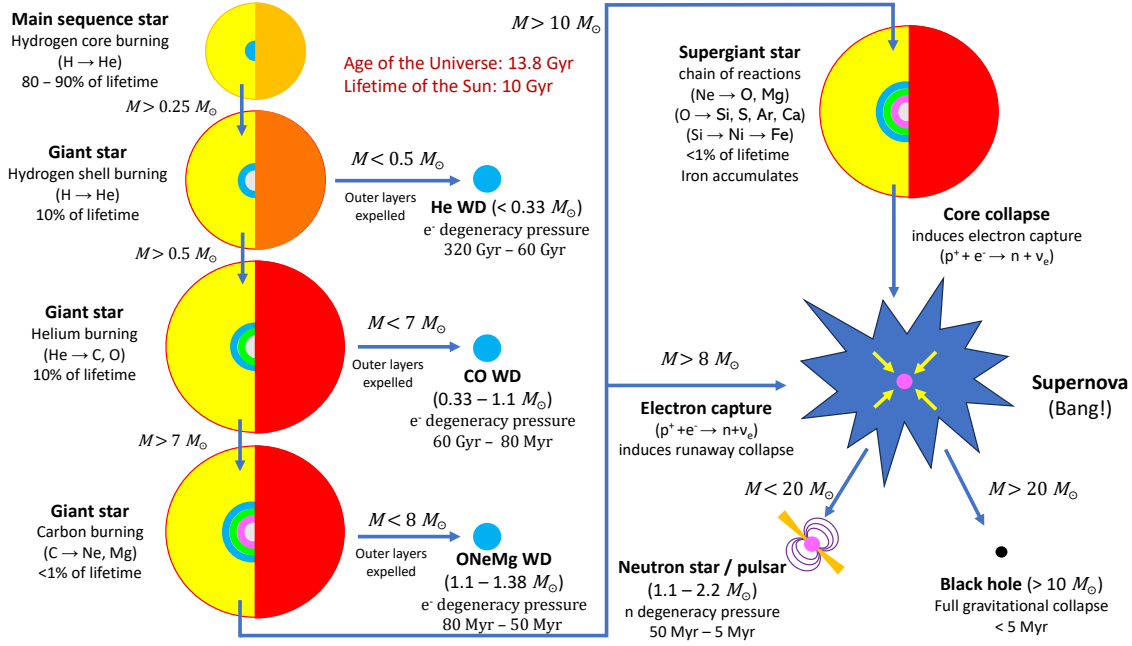


Figure 1.1: Diagram depicting the evolution of isolated stars from the main sequence to the formation of compact stellar remnants, with branching paths according to stellar masses. The typical mass ranges of the final stellar remnants and the typical formation times (lifetime of the progenitor star) are indicated. The approximate fraction of the lifespan spent by stars on each stage is also indicated. The helium/carbon core-burning and shell-burning stages are depicted jointly for simplicity, but one should mind that they happen one after the other, just like the hydrogen core-burning and shell-burning stages. Also, the drawings are not at scale, as giant stars are  $\sim 100$ – $1000$  larger than main-sequence stars.

abundant in the Milky Way, with Sirius B, the companion of Sirius A (the brightest star in the sky), being the closest example (Bond et al., 2017; Brosch, 2008). Finally, *Oxygen-neon-magnesium white dwarfs* (ONeMg WD,  $1.1 \lesssim M_{WD} < 1.38 M_{\odot}$ ) would be the most extreme and rare types of white dwarf (Truran and Livio, 1986), with some candidates being discovered only very recently (e.g. Kepler et al., 2016).

### 1.2.2 Formation II: stars and supernovae

In theory, non-rotating and maximally-rotating<sup>6</sup> white dwarfs do not exist above the  $1.38 M_{\odot}$  (e.g. Nomoto et al., 1984; Takahashi et al., 2013; Chanlaridis et al., 2022) and  $1.48 M_{\odot}$  (e.g. Yoon and Langer, 2004), respectively, the point where gravity overcomes the electron pressure and resumes the gravitational collapse. This is the *Chandrasekhar limit*, owing the name to Subrahmanyan Chandrasekhar, who made the prediction based on theoretical calculations of white dwarf equations of state (Chandrasekhar, 1931). So what, happens when a star's core is more massive than the said limit? The fireworks begin now.

As shown in Fig. 1.1, stars with originally  $M_{MS} > 8 M_{\odot}$  continue to fuse heavier elements until the electron degeneracy pressure can no longer keep the core from collapsing into itself. The result is a

<sup>6</sup> Centrifugal forces from rotation can also prevent the gravitational collapse of a star, but up to a maximum mass



sudden runaway gravitational collapse at the end of their life. Courtesy of the inverse  $\beta$ -decay reaction, electrons ( $e^-$ ) are absorbed by protons ( $p^+$ ), producing neutrons ( $n$ ) and releasing electron neutrinos ( $\nu_e$ ). In less than a day, the core becomes entirely composed of neutrons and reaches a radius of only  $\sim 15$  km, at which point the neutron degeneracy pressure suddenly halts the collapse. The release of gravitational energy, the violent rebound pressure shocks and the out-flowing wind of neutrinos blow the outer layers away (Burrows and Vartanyan, 2021) in an explosion that outshines the whole galaxy, a *supernova*.

In stars with  $8 < M_{\text{MS}} < 10 M_{\odot}$ , the process of electron capture begins in the ONeMg core, leading to an *electron-capture supernova* (Miyaji et al., 1980; Nomoto, 1984; Nomoto, 1987; Takahashi et al., 2013) which leaves behind the exposed core of neutrons, a *neutron star*. While this type of supernova has remained hypothetical for a long time, the recent discovery of an electron-capture supernova candidate (Hiramatsu et al., 2021) has given a lot of strength to this formation channel. Even earlier, it has long been speculated that the famous Crab nebula supernova remnant and its pulsar are the result of an electron-capture supernova due to its chemical composition (e.g. Nomoto, 1982; Smith, 2013; Yang and Chevalier, 2015).

On the other hand, stars with  $M_{\text{MS}} > 10 M_{\odot}$  will continue to fuse higher elements until they produce iron, the element with the largest nuclear binding energy, and therefore the most stable element in the Universe. With iron being unable to release further energy via nuclear fusion, an *iron core-collapse supernova* occurs (Carroll and Ostlie, 1996). If the mass of the core is less than  $\sim 2 M_{\odot}$ , the neutron degeneracy pressure will halt the collapse and it will produce a neutron star. Otherwise, if the core is more massive (only in stars with originally with  $M_{\text{MS}} > 20 M_{\odot}$ ), not even neutrons can save the day and the remnant becomes a black hole instead.

Supernova events are not only the start of a pulsar's life, but they can also have a significant evolution on their influence. Young pulsars are often observed with sky positions and proper motions that diverge from that of their supernova remnant, and some pulsars such as the 98.81-ms **PSR J1747–2958** have even been *kicked out of their supernova remnant* (Camilo et al., 2002; Gaensler et al., 2004). This is because supernovae are not spherically symmetric, and a combination of the turbulent asymmetric collapse and the post-supernova fallback lead to birth velocity kick that can be even higher than  $1,000 \text{ km s}^{-1}$  (e.g. Janka et al., 2022). This extremely relevant for pulsars forming in binary systems, as the supernova kick may lead to the disruption of the systems, or to the introduction of a significant orbital eccentricity (e.g. Tauris et al., 2017). In fact, from binary systems, there is tentative evidence that more massive stars are born with larger supernovae kicks, resulting perhaps from more massive progenitors (Andrews and Mandel, 2019). However, for isolated pulsars (90% of the population), this relationship can not be corroborated, as their masses are unknown and the connection to their supernova remnant is lost to time.

An interesting side-effect of supernovae kicks is that neutron stars travel far from their birthplace, making them a run-away population. The Galactic height<sup>7</sup> ( $z$ ) distribution of a given stellar population can be described with a negative exponential

$$\rho(z) \propto \exp\left(-\frac{z}{z_{\text{scale}}}\right) \quad (1.1)$$

defined by with a scale factor  $z_{\text{scale}}$ . For observed isolated radio pulsars,  $z_{\text{scale}}$  is approximately 350 pc,

---

<sup>7</sup> Galactic height referees to "height" above or below the base Galactic plane. For instance, the Sun is at  $z = 5 - 30$  pc above the Galactic plane (e.g. Humphreys and Larsen, 1995; Siebert, 2019). A parsec (pc) is equal to 3.26 light years.



while for their progenitors (main sequence stars with masses between 8 and 20  $M_{\odot}$ ) it takes values between 50 and 100 pc (Mdžinarishvili and Melikidze, 2004), meaning that pulsars tend to escape the Galactic plane due to the supernova kicks.

### 1.2.3 From neutron stars to pulsars I: basic neutrons star properties

We have seen that neutron stars are the end-product of the evolution of stars with masses between 8 and 20  $M_{\odot}$ . However, what does a neutron star actually look like, and what are the physical properties that allow them to become periodic radio pulsars? In this section, we revise some basic neutron star properties and lay the groundwork for understanding pulsars.

Neutron stars are extremely compact and dense objects, overcome in these properties only by stellar-mass black holes. The density of degenerate neutron matter is similar to that of atomic nuclei: from  $10^{17}$  to  $10^{18}$   $\text{kg m}^{-3}$  or, equivalently, 8 billion humans per teaspoon. **More than 50 neutron stars masses** have been measured via pulsar timing in binary systems, with values ranging from 1.1 to 2.2  $M_{\odot}$  (Özel and Freire, 2016). Modelling of the X-ray emission of some neutron stars has constrained their radii to between 11 and 14 km (e.g. Riley et al., 2019; Riley et al., 2021). For comparison, the Schwarzschild radius

$$r_s = \frac{2GM}{c^2}, \quad (1.2)$$

which defines the event horizon of a black hole, is just over 4 km for a  $M = 1.4 M_{\odot}$  object, meaning that neutron stars are barely holding themselves from collapsing into black holes, with a surface gravity that would crush anything that exists on Earth, including mountains. These extreme conditions result in unique forms of matter that can only exist in ultra-dense environments, such as *nuclear pasta* or the hypothetical quark-gluon plasma (Ascenzi et al., 2024).

The fastest pulsar is **PSR J1748–2446ad**, with a spin period of 1.396 ms (Hessels et al., 2006), while the slowest pulsar is **PSR J0901–4046** with 75.886 s (Caleb et al., 2022). While there is no theoretical limit to how slow can a neutron star rotate, the shortest spin period before it breaks apart due to centrifugal forces is expected to be  $\sim 0.5$  ms. However, their birth periods are  $\lesssim 100$  ms (e.g. Janka et al., 2022), while typical rotation periods of their progenitor stars are a few days. Also, neutron stars are also the most magnetized objects in the universe. While the strength of their magnetic fields can not be directly measured, the estimated surface strengths range between  $10^8$  and  $10^{15}$  G. In other words, the strongest neutron star magnetic fields could rip the iron out of your blood cells.

How can neutron stars achieve such fast rotation periods from such slow-spinning progenitors? Where do they super-strong magnetic fields come from? Let us begin with the spin. There are many proposed processes that explain the spin-up of the collapsing core, but they all agree on one common dominant factor: the conservation of the star's angular momentum,

$$L \propto M \times R^2 \times \Omega_s, \quad (1.3)$$

during the collapse, where  $M$  is the stellar mass,  $R$  is the stellar radius, and  $\Omega_s = 2\pi/P_s$  the angular velocity, where  $P_s$  is the spin period. As is shown in Fig. 1.2, assuming a constant  $M$  and  $L$  during the collapse of a star, the reduced  $R$  implies a larger  $\Omega_s$ . If we take the solar values as a proxy ( $R \approx 700,000$  km,  $P_s \approx 30$  days) and a typical neutron star radius of  $R \approx 15$  km, this results in a post-collapse spin period of  $P_s \approx 1.2$  ms. That is very close to the  $P_s = 1.588$  ms spin period of **PSR B1937+21**, albeit only by chance, because as discussed at the start of this chapter it achieved




Figure 1.2: Diagram depicting how neutron stars get their highly magnetised and rapidly spinning properties from the gravitational collapse of a progenitor star. While the angular momentum (red vector) and magnetic flux (purple lines) are preserved, a more compact mass distribution and a smaller surface area result in a faster spin and a stronger magnetic field on the surface.

such fast spin much later than its birth. Of course, this is a ballpark and unrealistic estimate, but it is nonetheless an enlightening one. In reality, a star's density distribution changes through its history, the majority of the mass is lost before and during the supernova event, and many kinds of interactions can occur between outer layers and the collapsing core, pre- and post-supernova (Janka et al., 2022, and references therein). Thus, most neutron stars are born with  $P_s = 10\text{--}100$  ms.

Similar to their spin periods, the super-strong magnetic fields of neutron stars can be heuristically explained via the inheritance of the magnetic field of their progenitor. In this case, we look at the preservation of magnetic flux through their surfaces,

$$F_B \propto B_s \times A, \quad (1.4)$$

during the collapse, where  $B_s$  is the magnetic strength at the star's surface, and  $A \propto R^2$  the surface area of the star. As is shown in Fig. 1.2, the reduced post-collapse surface area results in a stronger magnetic field in the surface. If, like before, we take the solar values as a proxy, its current surface magnetic field strength of  $B \approx 1$  G becomes  $B \approx 2 \times 10^9$  G after the collapse, which is consistent with the weakest estimated magnetic fields in pulsars. In addition, some authors have proposed the existence of some amplification processes during the collapse, leading to even higher magnetic fields (e.g. Spruit, 2009, and references therein).

### 1.2.4 From neutron stars to pulsars II: the magnetic dipole model

For all of you who are thinking, “Neutron stars are cool, but when will we talk about actual pulsars?”, the answer is *now*. Now that we understand the physical conditions of neutron stars, I will *try* to answer the question of “How does pulsar emission work?”. Short answer: nobody knows the details. Long answer: many useful models have been built to explain the characteristics of observed pulsar emission with various degrees of success (for reviews, see Harding, 2018; Philippov and Kramer, 2022). Here I will present a very simple one: the *magnetic dipole model*, based on the rapidly rotating and highly magnetised nature of neutron stars.

The model, depicted in Fig. 1.3, assumes a dipole magnetic field centred on the neutron star. The magnetic field attempts to co-rotate rigidly with the star, but there is a distance

$$d_c = c \times \frac{P_s}{2\pi} \quad (1.5)$$

from the rotation axis where it would move at the speed of light  $c$  to follow. This distance defines the *light cylinder*. For a  $P_s = 1$  s neutron star, it lies at  $d_c = 0.16$  *light seconds*<sup>8</sup> (ls)  $\approx 3000$  radii, and for a  $P_s = 1$  ms neutron star lies at  $d_c \approx 3$  radii. Outside of this cylinder, the magnetic field no longer co-rotates with the neutron star, and thus the field lines detach from its surface. This creates two regions within the magnetosphere of the neutron star: one where the magnetic field lines are closed and another where they are open.

According to classical electrodynamics, a rotating magnetic field induces a Lorentz force on charged particles. In a conducting medium such as the magnetosphere of a neutron star, charged particles will distribute themselves until they produce an electric field that counteracts this Lorentz force, reaching a state of equilibrium (Goldreich and Julian, 1969). However, since a pulsar’s magnetosphere has open magnetic field lines, this state of equilibrium can not be reached. Instead, charged particles (electrons or protons, depending on the pole) are constantly accelerated along the open magnetic field lines, from the base of the open field lines on the magnetic poles of the star, the *polar acceleration gap*, towards the border between the open and closed regions at the light cylinder, the *outer acceleration gap* (Sturrock, 1971).

The beam emission is produced in the polar acceleration gap, where the charged particles accelerated along the magnetic field lines produce very-high-energy  $\gamma$ -ray curvature emission<sup>9</sup>. In turn, these  $\gamma$ -rays split into electron-positron pairs, which are further accelerated into high energies, leading to the production of the next generation of  $\gamma$  radiation and particle pairs. The sequential production of pairs increases the plasma density by several orders of magnitude, until at a certain distance from the polar gaps curvature radiation is produced in the low-energy radio frequency, leading to the emission of highly collimated radio beams from each magnetic pole (Komesaroff, 1970; Sturrock, 1971). If the solid angle covered by a beam includes the Earth, we detect it as a radio pulse that repeats periodically with each rotation of the neutron star. Thus, pulsars are essentially radio-emitting Galactic lighthouses<sup>10</sup> arising from the highly magnetized and fast-rotating nature of neutron stars.

<sup>8</sup> Light-second: the distance that light overtakes in one second. For reference, the Earth-Moon distance is 1.3 ls. As you will see later on, this is a very useful unit for binary pulsars.

<sup>9</sup> Curvature emission: light emitted by a charged particle accelerated along curved magnetic field lines.

<sup>10</sup> Sailors use lighthouses to identify their position relating to shorelines at night. Likewise, NASA once thought pulsars could aid an advanced alien civilisation in locating the position of the Solar System, so they drew a pulsar map on the cover of each copy of *The Golden Record* carried by each Voyager interstellar probe.




Figure 1.3: Diagram of the magnetic dipole model and its observational manifestation as a radio pulsar. The inclination angle of the magnetic dipole with respect to the rotation axis and the size of the light cylinder defines the regions with open and closed magnetic field lines and the acceleration gaps. Charged particles accelerated along the gaps produce radio emission. From afar, the emitted radio beams are seen as pulsations when they sweep the observer.

### 1.2.5 Radio emission I: properties of radio emission

The pulsar dipole model predicts simple and symmetric pulse shapes. However, reality is often much more complicated. Pulsar radio emission is detected across multiple radio frequencies with a diverse array of morphology and characteristics. In this section, I summarize a few characteristics of pulsar emission to lay the groundwork for what kind of signals pulsar astronomers look for.

#### Pulse profiles

Some pulsars like the first discovery **PSR B1919+21** are bright enough that individual pulses are detected<sup>11</sup>, but these are only exceptions. Most pulsars have flux densities below the mJy<sup>12</sup>, meaning that only the addition of several hundreds or thousands of pulses shows a signal above the noise baseline. This integration of pulses produces a *pulse profile*, an average of many individual pulses.

Pulse profiles can take a multitude of shapes. While many pulsars present simple pulse profiles with a narrow component, some look skewed, many have more than one feature, and others look like the outline of **mount Montserrat in Catalonia**. To illustrate this point, two different pulse profiles of

<sup>11</sup> Fun fact: the cover of Joy Division's music album *Unknown Pleasures* depicts individual radio pulses from **PSR B1919+21** as seen by the Arecibo telescope.

<sup>12</sup> Jy: unit of energy flux per unit time, per unit area, per unit frequency, equivalent to  $10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$ . Radio astronomers are used to looking at radio data within an particular bandwidth of frequency, hence the need to express flux over frequency units.

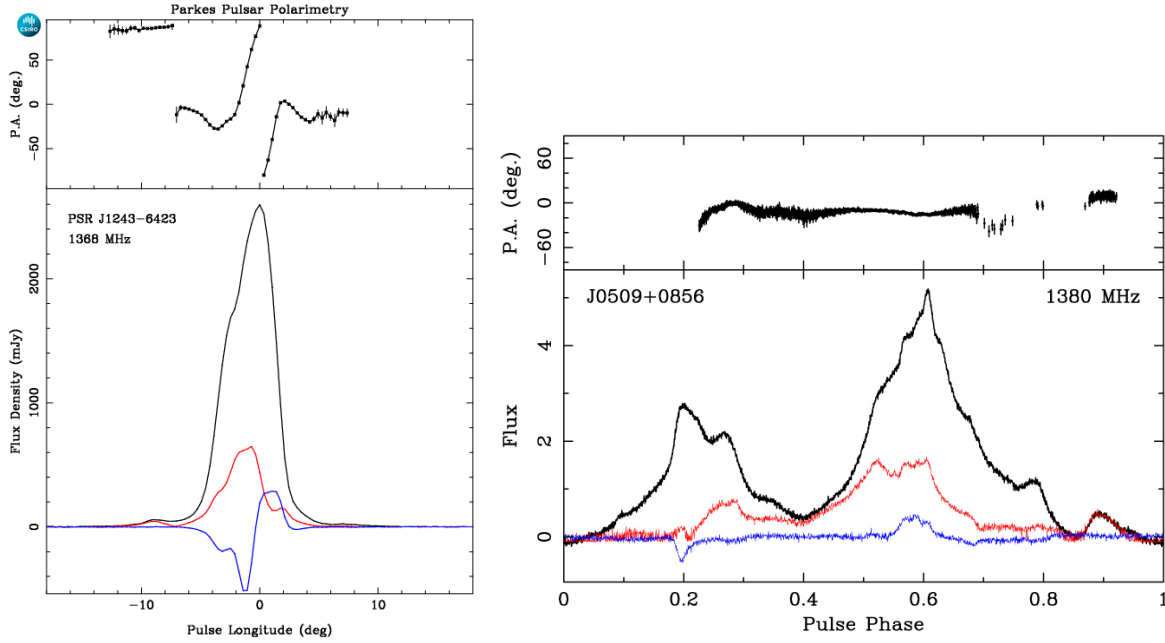


Figure 1.4: Pulse profiles (integrated radio flux against pulsar rotational phase) in total intensity (black lines), linear polarisation (red lines) and circular polarisation (blue lines). The top plots show the position angle of linear polarisation. **Left: PSR B1240–64** ( $P_s = 0.388$  s) as seen by the Parkes telescope (Johnston and Kerr, 2018). The x-axis goes from  $-18$  to  $+18$  degrees out of 360, so it is a narrow profile. **Right: PSR J0509+0856** ( $P_s = 4.056$  ms) as seen by the Arecibo telescope (Martinez et al., 2019).

vastly different complexity are shown in Fig. 1.4. The profile of **PSR B1240–64** ( $P_s = 0.388$  ms) is narrow and relatively simple, with a single main pulse and some minor side-features. On the other hand **PSR J0509+0856** ( $P_s = 4.056$  ms) emits a very complex pulse across almost all of its rotational phase. Indeed, fast-spinning pulsars tend to have wider and more complex pulse profiles (Karastergiou et al., 2024, e.g.), but in general all types of pulsars are susceptible to different degrees of complexity (Kramer et al., 1998).

The process of shaping pulse morphology is not fully understood. Some authors have proposed the existence of nested conical emission regions to explain the multi-component nature of some peaks (Rankin, 1993). However, the lack of symmetry in some pulsars has led some to think that only random patches of the magnetic poles (*beamlets*) are active in some of the pulsars with complex beam morphology (Lyne and Manchester, 1988). The study of pulse morphology is often limited by one-dimensional sampling of the observer’s line-of-sight (LOS) tracing over the pulsar’s pole as it rotates. Nonetheless, some exceptional pulsars such as the Hulse–Taylor one and **PSR J1906+0746** precess<sup>13</sup> due to binary orbital motion, allowing us to do sample two-dimensional patches of their pulses across the years, revealing the complex geography of their poles (Kramer, 1998; Weisberg and Taylor, 2002; Desvignes et al., 2019).

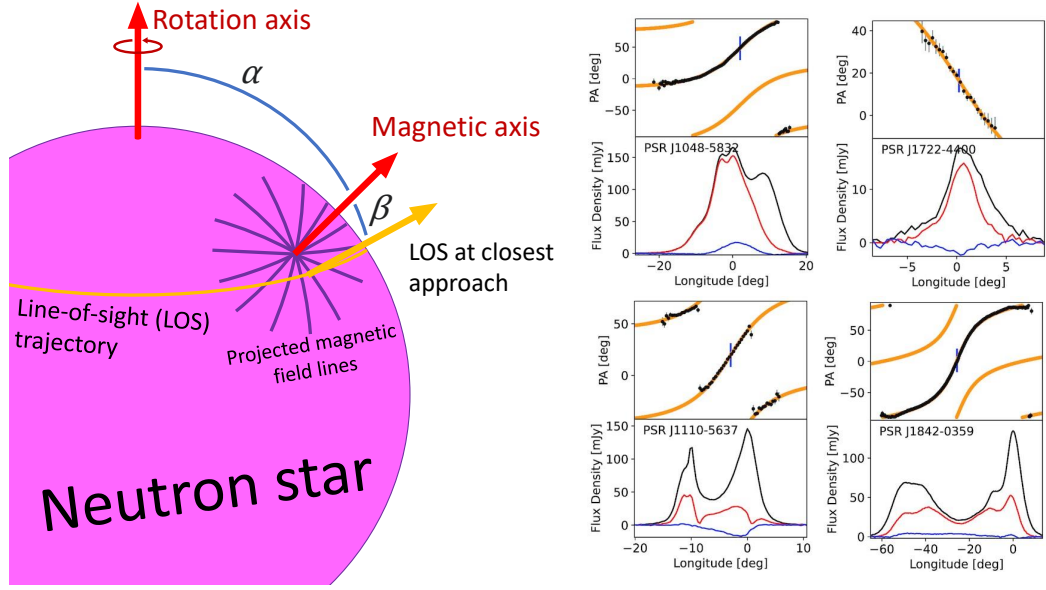


Figure 1.5: **Left:** diagram of the rotating vector model. The line of sight of the observer (yellow arrow) traces a 1-dimensional section of the surface and the beam as the neutron star rotates (yellow trajectory). The direction of the magnetic field lines projected on the surface (purple lines) defines the polarisation angle of light. **Right:** successful fits of the rotation vector model to polarised pulse profiles, from MeerKAT data (Johnston et al., 2023). In the bottom plots, total pulse intensity (black), linear polarisation (red) and circular polarisation (blue) are depicted as a function of the pulse rotational phase (longitude in degrees), while in the upper plots, the measured (black dots) and modelled PA (yellow lines) are shown.

## Polarisation

Pulsar radio light is often highly polarised, as depicted in Fig. 1.4. One of the big successes of the rotating dipole model is that the characteristics of linearly polarised light are consistent with it. For linearly polarised light, the direction of the oscillating electric field of the radio wave is described by the *polarisation angle* (PA), which in turn is determined by the projected direction of magnetic lines along which charged particles are accelerated. The derived *rotating vector model* (RVM) predicts that the PA is a function of the rotational phase of the pulsar  $\phi$ , the angle between the rotation and magnetic axes  $\alpha$  and the angle of closest LOS approach to the magnetic pole or *impact parameter*  $\beta$  (Radhakrishnan and Cooke, 1969, Fig. 1.5):

$$\tan (PA - PA_0) = \frac{\sin \alpha \sin (\phi - \phi_0)}{\sin (\alpha + \beta) \cos \alpha - \cos (\alpha + \beta) \sin \alpha \cos (\phi - \phi_0)}, \quad (1.6)$$

where  $PA_0$  and  $\phi_0$  stand for origin constant. This model predicts a *polarisation swing*, where  $PA - PA_0$  has opposite signs before and after the point of closest approach to the magnetic pole, a phenomenon observed in the majority of pulsars with polarised emission. This is used to determine beam geometry with a high degree of success, as depicted in Fig. 1.5 for four different pulsar from the MeerKAT Thousand Pulsar Array (e.g. Johnston et al., 2023).

<sup>13</sup> Precession: rotation of the rotation axis around a second axis.

On the other hand, the circular polarisation or radio emission is "an embarrassment" (Melrose, 2003)<sup>14</sup>. Despite the fact that pulsars are basically the only circularly polarised sources in the radio sky and that they are identifiable in the imaging domain because of this very reason (e.g. Wang et al., 2022b), circular polarisation is playing hard to get for emission models in general. However, the magnitude degree of circular polarisation is often lower than that of linear polarisation, and modelling efforts attempting to account for it persist to this day (e.g. Gangadhara, 2010; Oswald et al., 2023).

### Frequency evolution

Pulse profiles can change considerably in shape across the radio band, a phenomenon called *frequency evolution*. It is generally observed that pulse widths increase with decreasing frequency, a phenomenon interpreted as higher-frequency emission being produced closer to the acceleration gap (Cordes, 1978). It is also often the case that different pulse components in complex pulse profiles change their relative separation (in rotational phase) and intensities.

Another relevant observational phenomenon is that pulsar brightness increases with decreasing frequency  $f$  quite steeply. This can usually be described with a power law

$$F(f) = F_{\text{ref}} \left( \frac{f}{\text{GHz}} \right)^{\xi}, \quad (1.7)$$

where  $F$  stands for the mean flux density (averaged across the entire rotational phase),  $F_{\text{ref}}$  is the reference flux density at  $f = 1$  GHz and  $\xi$  is the *spectral index*, with an average value of  $\xi \approx -1.6$  with large spread across individual pulsars (Bates et al., 2013; Jankowski et al., 2018). Thus, pulsars are more efficient at generating radio-emission at low frequencies, and naively one would think it is preferable to observe pulsars at low radio frequencies. However, in the following section, I show how this is very often not the case.

#### 1.2.6 Radio emission II: interstellar medium effects

Besides the intrinsic emission properties of pulsars discussed above, there are also extrinsic properties arising from the interaction of the radio waves with the *interstellar medium* (ISM) along the LOS between the source and the observer. Therefore, these effects depend on the observer, meaning that an alien civilisation observing the same pulsar from another corner of the Galaxy will describe completely different ISM-derived pulse properties. Fig. 1.6 depicts the most significant ones are: *dispersion*, *scintillation* and *scattering*. These phenomena are highly correlated and often occur simultaneously, with each of them presenting its own set of observational challenges. In addition, there is also *Faraday rotation*, which affects only polarised light. All of these effects have one thing in common: they become more prominent at lower radio frequencies.

### Dispersion

Dispersion affects all radio pulsars without exception. In an ionized medium, radio waves travel slower than the speed of light  $c$ . Like in other types of light, this phenomenon depends on the wave frequency

<sup>14</sup> Yes, this is a literal quote from astrophysicist Donald Melrose.



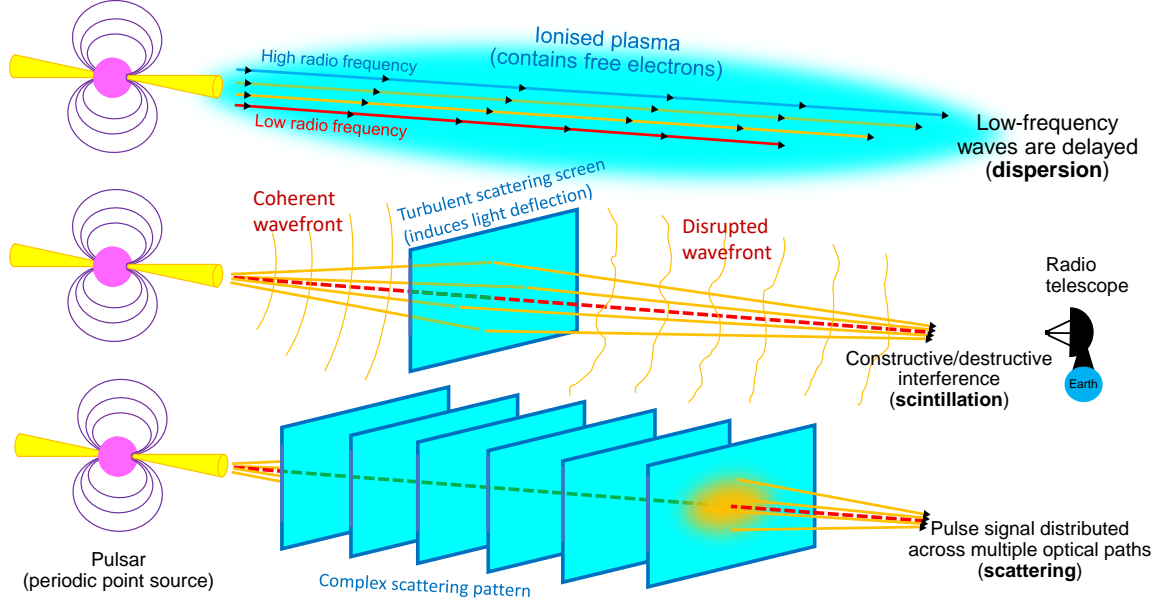


Figure 1.6: Diagram depicting how dispersion (top), scintillation (middle) and scattering (above) arise from the interaction of radio emission with the interstellar medium. These phenomena are highly correlated and often occur simultaneously, and each presents its own set of observational challenges.

and on the properties of the medium. In this case, the determining factor is the electron density  $n_e$ . The time delay of a wave with  $f$ , travelling along a LOS  $l$  for a distance  $d$  is

$$\Delta t = \frac{e^2}{2\pi m_e c} \frac{1}{f^2} \int_0^d n_e(l) dl = \mathcal{D} \times \frac{\text{DM}}{f^2}, \quad (1.8)$$

where  $m_e$  is the electron mass,  $e$  the electron charge,  $\mathcal{D} = 4.15 \times 10^3 \text{ MHz}^3 \text{ pc}^{-2} \text{ cm}^3 \text{ s}$  is the *dispersion constant*,  $f$  is expressed in MHz and DM is the *dispersion measure* in units of  $\text{cm}^{-3} \text{ pc}$ . In essence, the DM is a measure of electron column density<sup>15</sup>, and the larger it is, the longer it takes the wave takes to reach the observer. However, it is more relevant that this time also increases with decreasing  $f$  values. Therefore, pulses reach the observer at different times depending on the observed frequency.

Since we usually do not know the absolute distance to a pulsar nor the exact  $n_e$  distribution in the Milky Way, the relevant arising phenomenon from dispersion is the *differential delay* between two frequencies  $f_2 > f_1$  (in MHz)

$$\Delta t_{12} = 4.15 \times 10^6 \text{ ms} \times \left( \frac{1}{f_1^2} - \frac{1}{f_2^2} \right) \times \text{DM}. \quad (1.9)$$

This delay can smear a pulsar's signal when observed at low frequencies over wide radio bands. The

<sup>15</sup> Column density: 3-dimensional density of electrons integrated along the LOS, which results in a 2-dimensional density on the sky.






Figure 1.7: Frequency-resolved pulse profiles of **PSR J1708–4843** as observed with the MeerKAT radio telescope at different dates, in the 900 to 1,700 MHz radio band. Constructive interference produces the bright bands, while destructive interference nullifies the signal.

first pulsar discovery **PSR B1919+21** ( $P_s = 1.337$  s) has  $DM = 12.44 \text{ cm}^{-3} \text{ pc}$ , which implies only a small differential delay of 0.191 s within the narrow 81 to 82 MHz band on which it was observed (Hewish et al., 1968), allowing its serendipitous discovery. For reference, the highest-DM pulsar is the Galactic center magnetar **PSR J1745–2900** ( $P_s = 3.764$  s), with  $DM = 1778 \text{ cm}^{-3} \text{ pc}$  (Eatough et al., 2013). For pulsars with  $P_s < 100$  ms, the DM-induced differential delay quickly wraps over several rotations, leading to a complete loss of the periodic signal if dispersion is not corrected for. Therefore, when a new pulsar is discovered, the DM is quoted alongside  $P_s$  as one of the fundamental pulsar parameters.

### Scintillation

Scintillation occurs when inhomogeneities in the turbulent plasma medium induce the deflection of light rays as they move along the LOS, leading to the scattering of the radio waves in random directions. As depicted in Fig 1.6, this leads to a disruption of the originally coherent wavefront, which manifests as a drastic boost or decrease of luminosity depending on whether there is constructive or destructive interference at the observer (Cordes et al., 1986). Scintillation is also highly dependent on frequency because light deflection depends on the refractive index. Thus, scintillation may boost a pulsar's apparent brightness in one frequency band and reduce it in another. In addition, slight time-dependent changes within ISM, or even the movement of the observer or the pulsar with respect to each other can also radically change the interference pattern at the observer. Therefore, the apparent luminosity of the pulsar can change rapidly by orders of magnitude across very short time scales.

Unlike dispersion, this phenomenon is unpredictable and astronomers can either observe across large frequency bandwidths hoping to capture a good "scintil", or cross their fingers so they get lucky in their next observation. As an example, it manifests clearly in the MMGPS discovery **PSR J1708–4843** ( $P_s = 16.7$  ms,  $DM = 28.7 \text{ cm}^{-3} \text{ pc}$ )<sup>16</sup> in Fig. 1.7. For this pulsar, the S/N of the pulse after integration cross frequencies changes from non-detections to a whooping value of 40. Another exemplary case is

<sup>16</sup> This is the biggest spoiler I have dropped in this entire thesis.




Figure 1.8: Pulse profile of **PSR J1818–1422** at the frequencies of 0.94 GHz (left) and 1.62 GHz (right) as seen by the MeerKAT radio telescope (Oswald et al., 2021). The gray lines represent the radio data, while the solid black line are fits to said data. Scattering is present in both images, but is much more prominent on the left one, as seen from the significantly less sharp shape of the pulse profile. The scattering timescales of  $\tau_s(0.94 \text{ GHz}) = 73.8 \pm 0.5 \text{ ms}$  and  $\tau_s(1.62 \text{ GHz}) = 9.59 \pm 0.03 \text{ ms}$  are reported, implying  $\alpha = -3.8$  (Oswald et al., 2021).

that of **PSR J2140–2311B** ( $P_s = 12.986 \text{ ms}$ ,  $\text{DM} = 25.1 \text{ cm}^{-3} \text{ pc}$ ). It was discovered with the Green Bank Telescope in 2001 at the **globular cluster M30** during a luminosity boost thanks to scintillation (Ransom et al., 2004), and then it eluded re-detection for almost 20 years until the MeerKAT radio telescope re-observed it with its enhanced sensitivity and larger bandwidth (Balakrishnan et al., 2023).

## Scattering

Scattering is caused by the same physics as scintillation, but it manifests to the observer on a different regime. In fact, it already occurs during scintillation. It is the broadening of the pulse due to some light rays travelling along longer optical paths before reaching the observer, therefore arriving with a time delay. However, it is more prominent with more complex ISM characteristics along the LOS and therefore in pulsars with large DM values. As depicted at the bottom of Fig. 1.6, after many interactions with the ISM, the pulsar ceases to be a point-like source and instead becomes a disc on the sky, with light from the edges of the disk being slightly delayed with respect to the ones reaching from the centre. In practice, this means that the pulsar signal  $S_{\text{PSR}}(t)$  is *scattered* over a timescale  $\tau_s$ , which can be described with a convolution with an exponential:

$$\hat{S}_{\text{PSR}}(t) = \int S_{\text{PSR}}(t) \times \exp\left(-\frac{t-t'}{\tau_s}\right) dt'. \quad (1.10)$$

Therefore, it is said that scattering produces a *scattering tail* on the pulse profile.

Scattering is highly dependent on frequency because lower frequencies are more prone to suffer deflection in the ISM. As a consequence, the timescale  $\tau_s$  increases with decreasing frequency  $f$  with

the power law

$$\tau_s(f) = \tau_{\text{ref}} \left( \frac{f}{\text{GHz}} \right)^{-\alpha}, \quad (1.11)$$

where  $\tau_{\text{ref}}$  is the reference value at 1 GHz, and  $\alpha$  is the *scattering index*. In general, it is observed that most scattered pulsars are also the ones with the largest DM values (e.g. Oswald et al., 2021), which makes sense given that there are more interactions with the ISM along the LOS. The index  $\alpha$  is predicted to have a value of  $\alpha = 4$  for a Gaussian distribution of inhomogeneities in the ISM, and  $\alpha = 4.4$  for a Kolmogorov-like turbulent ISM (e.g. Rickett, 1977; Romani et al., 1986). However, typically observed values are  $\alpha = 4.0 \pm 0.6$ , depending on DM (Löhmer et al., 2004; Krishnakumar et al., 2019; Oswald et al., 2021). Fig. 1.8 depicts this effect in **PSR J1818–1422** ( $P_s = 291.5$  ms,  $\text{DM} = 619.7 \text{ cm}^{-3} \text{ pc}$ , Oswald et al. (2021)), where it is shown that the scattering tail is much more prominent at low frequencies.

Scattering is one of the main obstacles to discovering pulsars with  $\text{DM} > 1,000 \text{ cm}^{-3} \text{ pc}$ . Time scales of  $\tau_s \sim 100$  ms at 1 GHz are very typical for pulsars at  $\text{DM} \sim 500 \text{ cm}^{-3} \text{ pc}$  (e.g. Oswald et al., 2021), and since pulsars are brighter at low frequencies many surveys have historically targeted  $f < 1$  GHz. If  $\tau_s > P_s$ , the periodicity of the pulsar signal is washed out, and unlike dispersion, this effect can not be corrected. Because of this, few pulsars are known in distant regions of the Milky Way hidden behind a dense and complex ISM environment, such as the Galactic centre or the Galactic far side. The solution is often to perform observations at higher frequencies, but this is a compromise due to the steep spectral properties of pulsars described in the previous section.

## Faraday rotation

Faraday rotation is the most subtle of the ISM effects on pulsar emission: if one only records intensity information about radio light it will not be detected, as it affects the PA of linearly polarised light instead. Left-hand and right-hand circularly polarised emission present slightly different refractive indices in the presence of a magnetic field parallel to the LOS ( $B_{||}$ ), introducing a rotation on the PA

$$\Delta\text{PA} = \frac{e^3}{2\pi m_e^2 c^2 f^2} \int_0^d n_e B_{||} dl = \frac{c^2}{f^2} \times \text{RM} = \lambda^2 \times \text{RM}, \quad (1.12)$$

where  $\lambda$  is the radio wavelength and RM is the *rotation measure*, usually expressed in units of  $\text{rad m}^2$ .

Faraday rotation does not affect the S/N of a pulsar like dispersion, scintillation or scattering, but it is essential to correct for it if we want to study the polarised emission shown in Figs. 1.4 and 1.5. Without it, the addition of the pulse profile over different frequencies would be incoherent, resulting in no detection of linearly polarised emission. Measuring both the DM and the RM in a pulsar also results in an estimate of the average magnetic field strength along the LOS in the ISM, since from equations 1.8 and 1.12 it can be shown that  $\langle B_{||} \rangle = 1.23 \mu\text{G RM/DM}$ . The pulsar with the highest recorded RM is the aforementioned Galactic center magnetar **PSR J1745–2900**, with  $\text{RM} \approx -6.7 \times 10^4 \text{ rad m}^2$ , implying  $B_{||} \approx -460 \mu\text{G}$  and indicative of the highly magnetized environment at the Galactic center (Eatough et al., 2013).

### 1.2.7 The pulsar population I: properties from spin parameters

We have discussed the observational characteristics of pulsars. However, we now delve into deriving astrophysical information from their spin period  $P_s$  and its derivative  $\dot{P}_s$ <sup>17</sup>, also called *spin-down rate* owing to its negative nature. Because yes, the spin of pulsars *evolves over time*. The fastest evolving radio pulsar<sup>18</sup> is magnetar **PSR J1808–2024**, with  $P_s = 1.36$  s and  $\dot{P}_s \approx 4 \times 10^{-11}$  s s<sup>-1</sup> (Champion et al., 2020). On the other hand, the slowest evolving Galactic pulsar is **PSR J1216–6410**, with  $P_s = 3.54$  ms and  $\dot{P}_s = 1.62 \times 10^{-21}$  s s<sup>-1</sup> (Lorimer et al., 2007). So, for a pulsar to increase its  $P_s$  by 1 s, it can take as short as 1000 years or as long as 30,000 Gyrs ( $\sim 2,000$  times the age of the Universe!). And perhaps more importantly for this section: under some assumptions, physical characteristics can be extracted from the combined input of  $P_s$  and  $\dot{P}_s$ .

#### Spin-down luminosity

The *spin-down luminosity*  $L_s$  is a measure of a pulsar's energy loss over time, and the energy budget it has available for interacting with its environment. For this, we assume that the decrease of  $P_s$  is caused by the loss of rotational kinetic energy  $E_s$  over time. The rotational kinetic energy of a spinning object is a function of its momentum of inertia  $I$ <sup>19</sup> and its angular velocity  $\omega$ :

$$E_s = \frac{1}{2} I \Omega_s^2 = I \frac{2\pi^2}{P_s^2}. \quad (1.13)$$

Under the assumption of no mass loss and a constant  $I$  with a typical value of  $10^{45}$  g cm<sup>2</sup> (see Bejger and Haensel, 2002; Kramer et al., 2021a), we define:

$$L_s = |\dot{E}_s| = 4\pi^2 I \frac{\dot{P}_s}{P_s^3} \approx 13.95 \times 10^{31} \text{ erg s}^{-1} \left( \frac{\dot{P}_s}{10^{-15} \text{ s s}^{-1}} \right) \left( \frac{P_s}{\text{s}} \right)^{-3}, \quad (1.14)$$

which takes values ranging from  $10^{29}$  to  $10^{39}$  erg s<sup>-1</sup>.<sup>20</sup> For comparison, the luminosity of the Sun is  $3.8 \times 10^{33}$  erg s<sup>-1</sup>.

#### Surface magnetic field

Assuming the dipole model, we can relate  $L_s$  with the strength of the magnetic field on the surface of the neutron star  $B_s$ . In this case, we approximate a pulsar as an orthogonally rotating magnetic dipole of magnitude  $D$  and assume that  $L_s$  is equivalent to the energy loss due to the generation of

<sup>17</sup> Dot notation is equivalent to a time derivative in physics:  $\dot{x} = dx/dt$ ,  $\ddot{x} = d\dot{x}/dt$ , etc.

<sup>18</sup> The true fastest evolving pulsar would actually be magnetar **PSR J1808–2024**, with  $P_s = 7.56$  s and  $\dot{P}_s = 5.49 \times 10^{-10}$  s s<sup>-1</sup>, but it is visible only in the X-rays (Mereghetti et al., 2005).

<sup>19</sup> A moment of inertia is similar to a rotational mass. Massive objects require large forces to be accelerated, objects with large moments of inertia require large torques to spin-up.

<sup>20</sup> erg s<sup>-1</sup>: CGS units for  $10^{-7}$  W. Astronomers defined CGS units to deal with the low stellar energy fluxes on Earth, but these numbers grow very large when talking about absolute luminosities.

electromagnetic emission as predicted by classical electrodynamics (e.g. Jackson, 1998):

$$4\pi^2 I \frac{\dot{P}_s}{P_s^3} \approx \frac{2}{3c^3} D^2 \omega^4 = \frac{32\pi^4}{3c^3} \frac{B_s^2 R_{\text{NS}}^6}{P_s^4}, \quad (1.15)$$

where  $R_{\text{NS}} \approx 15$  km is the radius of the neutron star. Therefore,  $B_s$  can be estimated as

$$B_s \approx \sqrt{\frac{3c^3}{8\pi^2} \frac{I}{R_{\text{NS}}^6} P_s \dot{P}_s} \approx 10^{12} \text{ G} \left( \frac{\dot{P}_s}{10^{-15} \text{ s s}^{-1}} \right)^{1/2} \left( \frac{P_s}{\text{s}} \right)^{1/2}. \quad (1.16)$$

For the vast majority of pulsars, this estimate is the only insight into the strength of their magnetic field. Typical values range from  $B_s = 10^8$  G to  $B_s = 10^{15}$  G, and while it is only a useful ballpark estimate, it is probably within an order of magnitude from the real magnitude.

### Characteristic age

Finally, equation (1.16) also tells us about the relation between  $P_s$  and  $\dot{P}_s$ ,

$$\dot{P}_s = \frac{8\pi^2}{3c^2} \frac{B_s^2 R_{\text{NS}}^2}{I} P_s^{-1}, \quad (1.17)$$

which implies that as the  $P_s$  increases, the  $\dot{P}_s$  decreases, giving us an insight into the time evolution of pulsars. Most importantly, and under the assumption of the orthogonal dipole and a constant  $B_s$ , it allows us to estimate the amount of time it has passed since the pulsar was born with  $P_s = 0$  ms, the *characteristic age*

$$\tau_c = \frac{P_s}{2\dot{P}_s}. \quad (1.18)$$

This age can range from just 100 years to 100 Gyr ( $\sim 10$  times the age of the Universe), but it is only a very useful estimate, as we know for a fact that pulsars are born with  $P_s > 0$ . In addition, studies with unusual pulsars where a second derivative  $\ddot{P}_s$  has been measured show that relation (1.17) is not exact. Instead,  $\dot{P}_s$  evolves slower than predicted, possibly due to the decay of the  $B_s$  over time and definitely due to pulsars not being perfect orthogonally rotating dipoles, especially for young and high- $B_s$  pulsars (e.g. Kaspi and Helfand, 2002; Viganò et al., 2013). Nonetheless, the estimate of  $\tau_c$  has proven to be close to the real value in the few cases where independent age estimates exist. These independent age measurements come mostly from two places: the ages of associated supernovae, and the age of binary companions. Specific examples of this will be given later on in this introduction.

### 1.2.8 The pulsar population II: a visit to the pulsar zoo

It is cool that we have physical information about pulsars only from measuring their  $P_s$  and  $\dot{P}_s$  values. But, can we tell something about their history from any of this? The answer is yes. In fact, we can even classify pulsars into distinct populations, and place them in a single evolutionary picture. For the next section, I have prepared a tour through the broader pulsar zoo, where you will finally come to understand the difference between the first MSP discovery **PSR B1937+21** and **PSR B1919+21** at the centre of the Crab supernova remnant.

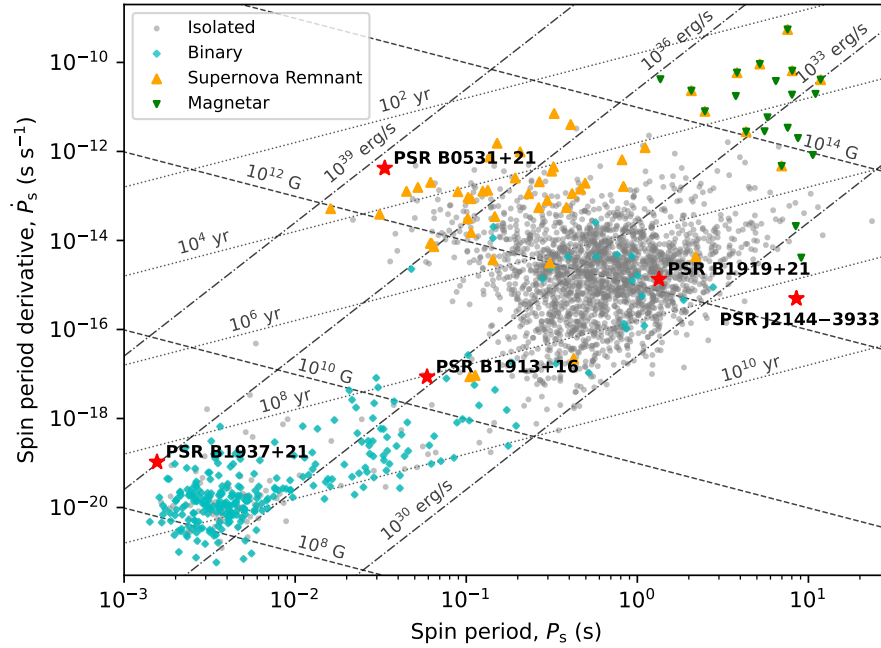


Figure 1.9: Pulsar spin vs. spin derivative diagram, colloquially known as the  $P - \dot{P}$  diagram, with all pulsars included in the ATNF Pulsar Catalogue as of February 2024. Four groups are highlighted: isolated pulsars with no binary companions (grey dots), pulsars with at least one binary companion (cyan diamonds), pulsars associated with a supernova remnant (yellow triangles) and magnetars (inverted green triangles). Lines of equal characteristic age (dots), equal surface magnetic field strength (short dashes) and equal spin-down luminosity (long dashes) are drawn. Four individual pulsars which have been mentioned in the text are highlighted: the pulsar at the center of the Crab Supernova supernova remnant (**PSR B0531+21**), the first pulsar discovery (**PSR B1919+21**), an aged pulsar from the graveyard (**PSR J2144–3933**), the pulsar in the Hulse–Taylor binary (**PSR B1913+16**) and the first millisecond pulsar discovery (**PSR B1937+21**).

### The $P - \dot{P}$ diagram

Readers and readdresses, please marvel at the specimens from Fig. 1.9! This is the  $P - \dot{P}$ <sup>21</sup> diagram, the equivalent of the Hertzsprung-Russell diagram in pulsar astronomy, where all 3,000 radio pulsars listed in the [ATNF Pulsar Catalogue](#) (Manchester et al., 2005) as of February 2024 are shown. Observe how pulsar specimens in binary systems, those associated with supernova remnants and those classified as magnetars cluster around their favourite corners of the diagram! Pay attention to the lines of equal  $L_s$ ,  $B_s$  and  $\tau_c$ , which tell you about the properties of these tribes!

Sorry, I got carried away. But it is true that in the  $P - \dot{P}$  diagram we can identify clusters of populations with common properties lying at similar evolutionary stages. Once a new pulsar is discovered, placing it in this diagram will tell us much of its story. Some of the pulsars mentioned in this text have been highlighted, so that you can already make some guesses about their nature. So, let us review what populations we have here.

<sup>21</sup> It's pronounced as "pee-pee-dot".



### Canonical and young pulsars

*Canonical pulsars* constitute the “main sequence” of pulsars. They populate the center of the  $P - \dot{P}$  diagram (bulk of grey dots in Fig. 1.9), with  $0.1 \lesssim P_s \lesssim 4$  s and  $10^{17} \lesssim \dot{P}_s \lesssim 10^{13}$  s s<sup>-1</sup>. They are characterised by moderate magnetic field strengths ( $10^{11} < B_s < 10^{13}$  G), but present a wide range of spin-down luminosity and characteristic ages. A subpopulation exists on the upper left of the main bulk, with slightly shorter spin periods ( $0.01 \lesssim P_s \lesssim 1$  s), large spin-down luminosities ( $L_s \gtrsim 10^{35}$  erg s<sup>-1</sup>), short characteristic ages ( $\tau_c \lesssim 10^5$  yr) and an abundance of supernova remnants. These are the *young pulsars*.

Young pulsars and conical pulsars form an evolutionary continuum, with all pulsars starting young and energetic in the region of the diagram populated by numerous supernova remnants. A prime example is the Crab pulsar (**PSR B0531+21**). With a fast spin period of  $P_s = 33.392$  ms, a young age of  $\tau_c = 1257$  yr and a large spin-down luminosity of  $L_s \approx 10^{39}$  erg s<sup>-1</sup>, this pulsar is constantly emitting at high energies and exciting the nebula around itself. The actual age of this pulsar is confirmed at 970 years from East-Asian historical records describing the supernova event in 1054 A.D. (e.g. Collins et al., 1999; Breen and McCarthy, 1995, and references therein), showing the accuracy of the characteristic age estimate.

With the passage of millennia, supernova remnants expand until they dissolve into the interstellar medium and pulsars become less energetic, moving rightward and downward across the  $P - \dot{P}$  diagram until they become part of the bulk of canonical pulsars. Here we find specimens such as **PSR B1919+21**, the first pulsar ever discovered. Outside of kick-starting the field of pulsar astronomy, this is a rather unremarkable middle-aged pulsar with a slow spin period of  $P_s = 1.337$  s, a low spin-down luminosity of  $L_s \approx 10^{31}$  erg s<sup>-1</sup> and an old age of  $\tau_c \approx 10$  Myr. Canonical pulsars continue to evolve in the same direction and do not get to live more than a few 100 Myr, making their lifespans rather short compared to stars like the Sun. Pulsars end their life entering the *pulsar graveyard*, the empty region at the bottom-right of the diagram where pulsar emission ceases. Some exceptional pulsars such as **J2144–3933** ( $P_s = 8.5$  s) challenge emission models by surviving deep in the pulsar graveyard with strikingly low spin-down luminosities (Young et al., 1999), but in general the name “graveyard” is quite fitting, as this is the end-point of the second life of massive stars.

### Magnetars

Magnetars are the population of highly magnetized neutron stars ( $B_s \gtrsim 10^{14}$  G) at the top-right of the diagram, which are beasts of their own (for a review, see Kaspi and Beloborodov, 2017). They are a much younger population, and half of them are associated with supernova remnants. However, only a handful of them are radio emitters, with the majority of them being detected in the X-rays instead. But those that emit in radio, they tend to have a high fraction of linearly polarised emission, sometimes almost 100% like in the case of **PSR J1745–2900** (Eatough et al., 2013). How did they achieve such ultra-strong magnetic fields, and how did their birth conditions differ from the canonical young population? Very valid questions. Unfortunately, they have no certain answer just yet. However, their evolution and ultimate fate are not much different from those of canonical pulsars, meaning they also spin down over time and enter the pulsar graveyard at the end of their lifespans.

## Recycled pulsars

Of all the pulsars in existence, 90% will follow the evolutionary path described above, be it a canonical pulsar or a magnetar. This in itself would be a pretty complete picture, but if you are paying attention, you should be telling yourself: "Wow, pulsars are cool, but what about all of these cyan dots at the bottom-left of the diagram? How do they fit with the evolutionary picture described above? And why are they all in binary systems?". Very good questions! That is why I have the second half of this introduction chapter dedicated almost entirely to them, but I will give a very exciting sneak-peak in this section.

Recycled pulsars are the detached population in the bottom-left of the  $P_s - \dot{P}_s$  diagram, with  $P_s \lesssim 100$  ms and  $\dot{P}_s \lesssim 10^{-17}$  s s<sup>-1</sup>. They are the most stable rotators of the entire pulsar population and, in many ways, they are the polar opposites of magnetars. They possess very weak magnetic fields ( $B \lesssim 1 \times 10^{10}$  G) and methuselan characteristic ages ( $10^8 < \tau_c < 10^{11}$  yr), and the vast majority of them are found in binary system, which is the key to their spin properties. These are pulsars that have been "resurrected" from the pulsar graveyard via the process of *pulsar recycling*, a stage of mass transfer from their evolving companions, resulting in the accretion of matter and, most importantly, angular momentum. This results in the *spin-up* of the pulsars and the "burying" of the magnetic field, which leads to a very small  $B_s$  and  $\dot{P}_s$  values after the recycling (Bhattacharya and van den Heuvel, 1991; Tauris, 2011).

The pulsars that undergo the most efficient recycling reach spin periods of  $P_s < 10$  ms, achieving the status of *millisecond pulsars* (MSPs), while pulsars with  $P_s > 10$  ms are usually called *partially recycled pulsars* (PRPs)<sup>22</sup>. The vast majority of these pulsars are, therefore, found in binaries. However, they can also exist in isolation, as some processes can disrupt the binary system or even destroy the binary companion. A typical PRP is the Hulse–Taylor pulsar (**PSR B1913+16**,  $P_s = 59.030$  ms), while an example of a fully recycled MSP is the first pulsar mentioned in this thesis (**PSR B1937+21**,  $P_s = 1.558$  ms). Why are these two different? This is a very intersecting question, and this is why I have the second half of this introduction dedicated to them.

## 1.3 Binary pulsars: Why are they so interesting?

Around one third of the main-sequence stars in the Milky Way are in binary system (e.g. Raghavan et al., 2010; Yuan et al., 2015), and since pulsars are born from originally main-sequence stars, they do not lag behind by much. The first discovered binary pulsar was the Hulse–Taylor one (**PSR B1913+16**,  $P_s = 59.03$  ms), found with Arecibo by astrophysicists Joseph Taylor and Russell Hulse in 1974 (Hulse and Taylor, 1975). If you have read the motivation section, you will know how enthusiastic I am about it, and with good reason. Its orbital companion is an unseen neutron star, and its study gave rise to the first detection of gravitational waves via the shirking of the orbit as predicted by Albert Einstein's theory of general relativity (Taylor et al., 1979; Taylor and Weisberg, 1982), which in its turn led to the awarding of the 1993 Nobel Prize and, much later, to the birth of gravitational-wave astronomy. This was a measurement that was only possible thanks to the clock-like rotating nature of pulsars, so it became evident that pulsars in binary systems are great laboratories for astrophysics and fundamental physics.

Many other binary pulsars have been discovered since then. Of the 3,534 pulsars listed in the ATNF

---

<sup>22</sup> This acronym is not common in the literature, but is of my own making to facilitate the reading of this thesis



database on May 2024, 404 belong to binary systems, 299 of which are Galactic (they belong to the Milky Way)<sup>23</sup>. Ironically, it turns out that Galactic double neutrons star (DNS) systems like the Hulse–Taylor binary are quite rare, with only 21 of them being known. Instead, the vast majority of Galactic binary pulsars are fully recycled MSPs with light He WD (companions ~60%), followed by sizable number (~20%) of MSPs and PRPs with CO and ONeMg companions. Furthermore, there is also a handful of exotic binaries containing non-recycled canonical pulsars. Each of these systems have brought in their own bit of extra science, both when it comes to testing astrophysics and fundamental physics.

But, how do we actually identify these systems? How certain can we be about the nature of their companions? And how can we use them to probe physics and astrophysics? In the following sections, I will show how we identify the nature of pulsar companions, and the state of the art of binary pulsar science.

### 1.3.1 Binary parameters: Keplerian orbits and the mass function

In the previous sections, I have shown we can learn much about pulsars just by knowing about  $P_s$  and  $\dot{P}_s$ . But, what about their binary companions? Can we see or measure them as well? Well, if we are lucky sometimes we can detect them with optical or infrared telescopes (e.g. Antoniadis et al., 2011; Pallanca et al., 2013), or if it is a neutron star we can even detect it as a second pulsar (Lyne et al., 2004). But in most cases, the white dwarf or neutron star orbital companions are too faint and distant to be directly detected, and the only hint we have about the companion is the orbit of the pulsar. But with this alone, we can already go a long way.

To first approximation, pulsars and their companions are point-like sources that obey Newtonian physics. In an isolated two-body system, orbits are described by **Kepler's laws of motion** (e.g. Roy, 2005). As depicted in Fig. 1.10, the pulsar and its companion trace an ellipse as they orbit around a common centre of mass. The orbit of the pulsar is described by three *Keplerian parameters*<sup>24</sup>: the *orbital period*  $P_b$ , the *semi-major axis* of the ellipse  $a_p$ , and the *eccentricity* of the ellipse  $0 < e < 1$ . However, in pulsar astronomy, often we do not know size of the orbit but its projection along the LOS, so instead of  $a_p$  we quote the *projected* semi-major axis  $x = a_p \sin i$ , where  $i$  is the *inclination angle* of the plane on which the orbital ellipses are drawn, the *orbital plane*. Appendix A provides a more detailed definition of this parameter, but in here I do share some insights: a  $i = 0$  deg orbit is completely *face-on*, meaning that we stare perpendicularly at orbital plane, while a  $i = 90$  deg orbit is completely *edge-on*, meaning that we stare at the side of the orbital plane. And finally, as shown in Fig 1.10, if  $e \approx 0$  we have a *circular orbit* of radius  $x$ , while  $e \gg 0$  describes a highly *eccentric orbit* of semi-major axis  $x$ .


Just knowing these three Keplerian orbital parameters of the pulsar already gives us very useful information about its companion. As depicted in Fig. 1.10, for an observed pulsar orbit with  $P_b$ ,  $a_p$  and  $e$ , we can guess the properties of the orbit of the companion as it is in the same orbital plane and revolves around the same *center of mass*, with the exact same orbital period  $P_b$ , eccentricity  $e$ , but with the ellipse pointing to the opposite direction and with an unknown semi-major axis  $a_c$ . How is

<sup>23</sup> The remaining ones reside in Globular Clusters, but they literally are a completely different story, as they have their own evolutionary story which is a topic for another thesis.

<sup>24</sup> Keplerian parameters, in honour of Kepler first describing the orbits of planets around the Sun with his laws of planetary motion in the early XVIIth century, in contrast to the circular orbits proposed Nicolaus Copernicus. In 1687, Isaac Newton showed that these orbits could be explained via the Universal law of gravitation.

**PSR J0034-0534:**

$$\begin{aligned} P_s &= 1.8772 \text{ ms} \\ P_b &= 1.59 \text{ days} \\ x &= 1.44 \text{ ls} \\ e &= 4.3 \times 10^{-6} \end{aligned}$$


**PSR B1913+16:**

$$\begin{aligned} P_s &= 59.030 \text{ ms} \\ P_b &= 0.323 \text{ days} \\ x &= 2.34 \text{ ls} \\ e &= 0.617 \end{aligned}$$




Figure 1.10: Depiction of the orbits of pulsars **PSR J0034-0534** and **PSR B1913+16** (Hulse–Taylor) and their unseen companions (not to scale!). The orbit of **PSR J0034-0534** is highly circular (Bailes et al., 1994; Desvignes et al., 2016), and it has a minimum companion mass of  $M_c > 0.15 M_\odot$ , making it a light He WD system where the pulsar is the gravitationally dominant body. On the other hand, the orbit of **PSR B1913+16** is highly eccentric and balanced in masses (Hulse and Taylor, 1975; Weisberg and Huang, 2016).

this relevant, you ask? Well, the position of the center of mass is determined by the distance between the masses of the orbital components and their relative positions, and therefore the semi-major axes of their orbits:

$$M_p \vec{r}_p + M_c \vec{r}_c = 0 \rightarrow M_p a_p = M_c a_c, \quad (1.19)$$

where  $\vec{r}_p$  and  $\vec{r}_c$  are the instantaneous position of the pulsar and its companion, and  $M_p$  and  $M_c$  their respective masses. In addition, as stated in Kepler's third law,  $P_b$  is a function of the orbital separation  $a_p + a_c$  and the total mass in the system,  $M_p + M_c$ :

$$P_b = \sqrt{\frac{4\pi^2 (a_p + a_c)^3}{T_\odot (M_p + M_c)}}, \quad (1.20)$$

where distances are expressed in units of light-seconds (ls), masses in  $M_\odot$ , and  $T_\odot = GM_\odot/c^3 = 4.92549094764 \mu\text{s}$  as defined in Prša et al. (2016), where  $G$  is Newton's gravitational constant<sup>25</sup>. These two equations are then combined to define the *binary mass function*

$$f(P_b, x) = \frac{4\pi^2 x^3}{T_\odot P_b^2} = \frac{(M_c \sin i)^3}{(M_p + M_c)^2}, \quad (1.21)$$

<sup>25</sup> The exact values of  $M_\odot$  and  $G$  have unavoidable measurement uncertainties, but  $M_\odot \times G$  is known to much better precision. Because of that, the IAU settled defines  $M_\odot$  and  $G$  as constants just like with  $c$ , and we can express masses in units of  $M_\odot$  free of worry as long as we do not have to translate to SI units.

which is the first hint about the the nature of the companion. Under the assumption of a pulsar mass (e.g.  $M_p = 1.35 M_\odot$ ), its value gives us an estimate of the massiveness of a system. And while the value of  $i$  is unknown a-priory, if we assume  $i = 90$  deg, the binary mass function can be solved numerically to derive *minimum companion mass* just from the measurement of  $x$  and  $P_b$ , which is already a very useful scientific probe. For instance, a mass function of  $\sim 0.2 M_\odot$  implies  $M_c > 1.0 M_\odot$ , telling us very clearly that the companion is not a light He WD, but instead a heavy CO WD, a ONeMg WD or even a neutron star. However, we should keep in mind that the assumed pulsar mass, which can be as light as  $1.1 M_\odot$  and as heavy as  $2.2 M_\odot$ , is a great source of uncertainty, and that the value of  $i$  is randomly distributed across binary systems. In fact, if the true  $i$  is close to 0 degrees, the minimum companion mass will be a gross underestimate compared to the true companion mass. Despite all, the minimum companion mass remains a good estimate for most cases: the (random) distribution of  $i$  values follows a  $\sin i$  function, with a median value of  $i = 60$  deg, so it is far more likely that it lies closer to 90 deg than to 0 deg (see Appendix A for a derivation of this fact). Therefore, it remains true that, except for the most unlucky systems, the minimum companion mass is often close to the true mass.

### 1.3.2 Binary populations: a visit to the binary zoo

Now that we know how to read spin *and* Keplerian orbital parameters to extract astrophysical parameters, can we know the nature of a pulsar binary by placing them in some kind of diagram, as we did before with the  $P - \dot{P}$  diagram? Well, you are in luck, because the answer is *yes!*

Readers and readdresses, from the makers of *the pulsar zoo*, I present to you *the binary pulsar zoo!* Make yourselves comfortable and be amazed by all the exposition in Fig. 1.11. In these four diagrams, we showcase the distribution of almost 300 specimens as listed in the ATNF database as of April 2024. Observe how different tribes appear within the  $P_b$ ,  $x$ ,  $e$  and  $P_s$  space, where each population resides in its own area. And especially, pay attention to the lines of equal minimum  $M_c$  in the  $P_b - x$  diagram, and see how the groups are clearly separated according to their minimum  $M_c$  values. Wonder at how this correlates very well with the pulsations of MSP and PRP pulsars in the  $P - \dot{P}$  diagram!

Repeated jokes aside, it is quite clear that pulsars with one type of binary companions tend to cluster around similar orbital and spin characteristics. For example, pulsars with light companions ( $M_c \lesssim 0.3 M_\odot$ ) are almost always MSPs ( $P_s < 10$  ms) and have the most circular orbits, with the orbital eccentricity going as low as  $e < 10^{-6}$ . However, pulsars with massive companions  $M_c \gtrsim 0.7 M_\odot$  tend to be PRPs ( $P_s > 10$  ms). Within this group, massive CO/ONeMg WD systems and DNS systems are distinguished by their orbital eccentricity, with double-neutron star systems (DNSs) always presenting  $e \gtrsim 0.1$ . There is even a handful of canonical pulsars in eccentric orbits with CO/ONeMg and main sequence companions!

### 1.3.3 Binary evolution I: mass transfer and pulsar recycling

How did all of these binaries form? Did Galactic pulsars download *Tinder* and find their best-matching companions? Well, that would be weird, because in the Milky Way pulsars and their companions are actually siblings, born together as main-sequence stars millions or even billions of years ago<sup>26</sup>. Their current spin and orbital characteristics are not a result of chance, but of a common (and

<sup>26</sup> However, pulsars in Globular Clusters are another story. Their stellar densities are so high that stellar encounters and the exchange of binary companions are the norm

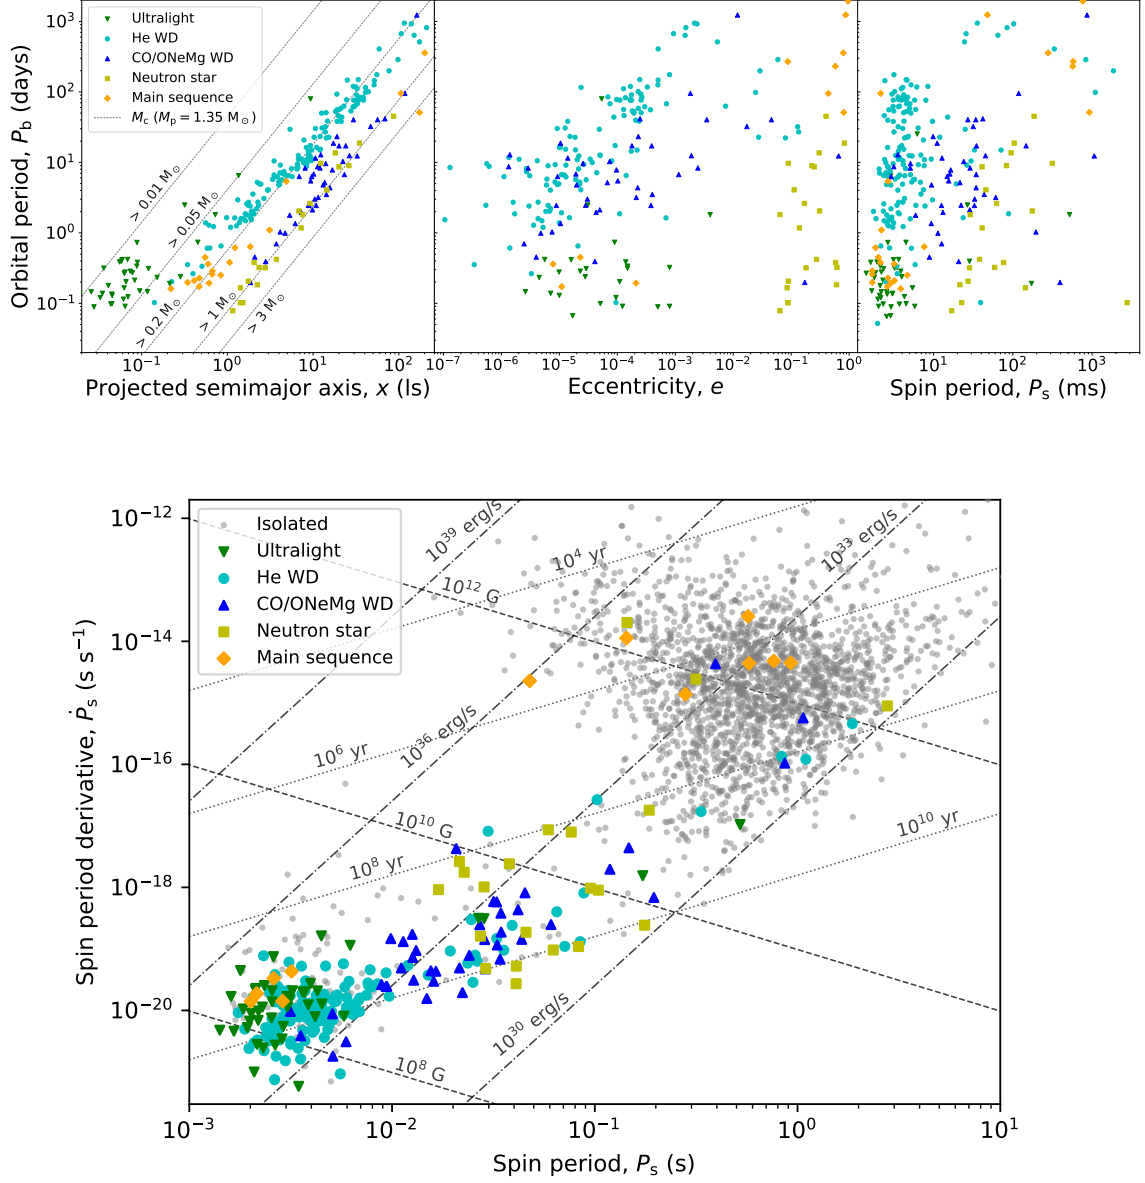


Figure 1.11: **Top:** Orbital period vs. projected semi-major axis, orbital eccentricity and spin period of the 298 pulsar binaries found in the Milky Way as listed in the ATNF Pulsar Catalogue on March 2024, highlighted by the nature of their companions. Lines of equal minimum companion mass *from the mass function* are drawn on the first diagram (dotted lines), at the values of  $M_c > 0.01, 0.05, 0.2, 1.0$  and  $3.0 M_\odot$  under the assumption of  $M_p > 1.4 M_\odot$ . **Bottom:**  $P - \dot{P}$  diagram focusing all 298 binary pulsars found in the Milky Way as listed in the <https://www.atnf.csiro.au/research/pulsar/psrcat/>

ATNF Pulsar Catalogue on March 2024, highlighted by the nature of their companions and contrasted against the isolated population (grey dots). Lines of equal characteristic age (dots), equal surface magnetic field strength (short dashes) and equal spin-down luminosity (long dashes) are drawn.

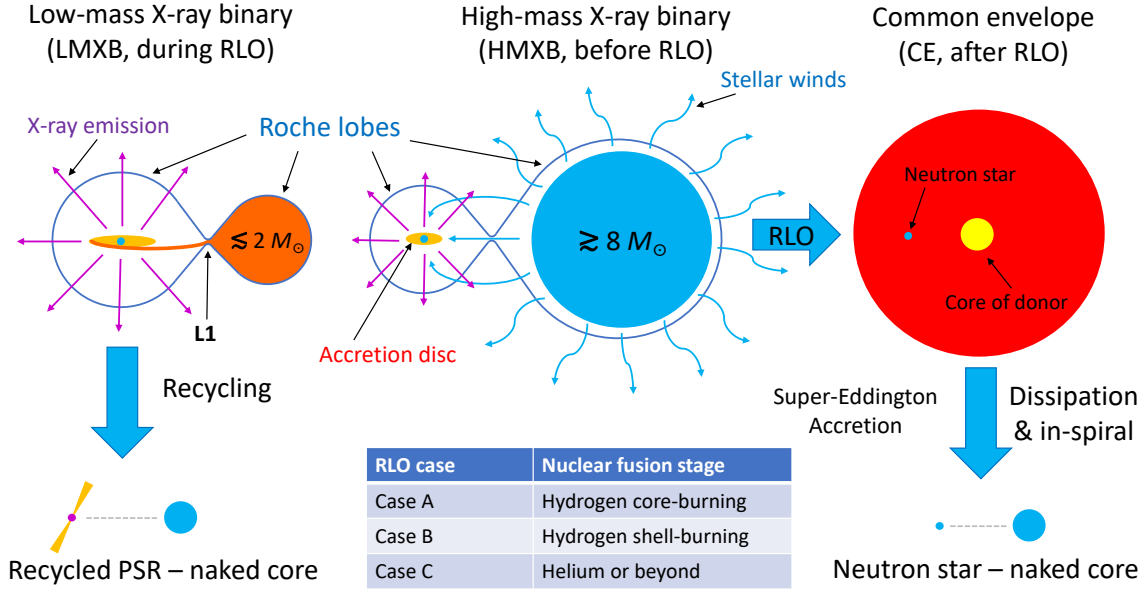


Figure 1.12: Diagrams depicting LMXB and HMXB, and their evolution. While the cases treated here involve only accreting neutron stars, many X-ray binaries host accreting stellar-mass black holes instead. As of today, 349 LMXB and 152 HMXB have been identified in the Milky Way (Avakyan et al., 2023; Fortin et al., 2023), but distinguishing whether the accreting object is a neutron star or a stellar-mass black hole is not easy. However, X-ray pulsations have been detected in some X-ray binaries, proving that some of them are indeed host accreting neutron stars, with MSPs existing only in LMXBs (Di Salvo and Sanna, 2022; Kim et al., 2023).

often violent) evolutionary history. And the key to this joint evolution is binary interaction, which occurs because stars are not point-like objects, but instead have finite physical sizes. As exposed in Sections 1.2.1 and 1.2.2, stars increase their size as they burn their nuclear fuel, leading to the transfer of mass between the stars. Now, we delve into the basics of this binary interaction, starting with *pulsar recycling*.

Pulsar recycling is the process of *spinning-up* a pulsar via the accretion of mass from its companion. During this process, the in-falling material accumulates in a spinning accretion disc. The plasma of this disc is captured by the magnetic field of the neutrons star and dragged into its surface along the magnetic field lines. This interaction induces rotational torques on the pulsar, which changes its rotation until an equilibrium between the spin and the infilling material is reached at a few ms (Radhakrishnan and Srinivasan, 1982; Ghosh and Lamb, 1979; Ghosh and Lamb, 1992). Additionally, it is generally accepted that the accretion process induces the decay of the magnetic field, even though the exact mechanisms of this process is not yet understood (Bhattacharya, 2002). Nonetheless, the final fast spin of the neutron stars means that radio emission begins anew once mass transfer is over, even if the pulsar had already entered the graveyard before mass transfer begun. Thus, our beloved recycled pulsars that inhabit the bottom-left of the  $P - \dot{P}$  diagram are born, or rather, resurrected (Bhattacharya and van den Heuvel, 1991).

Long-lasting and efficient mass transfer occurs during *Roche-lobe overflow* stages (RLO). In a binary system, the Roche lobes are the regions around each star where their gravitational field is dominant

over the other's. When a one of the binary components evolves, its expanding envelope can overflow its own Roche lobe, spilling material onto the other through the bridge of the L1 Lagrange point (Pringle and Wade, 1985; Davis et al., 2013; Tauris and van den Heuvel, 2023)<sup>27</sup>. These stages are classified according to the state of the companion when it occurs: Cases A, B and C refer to when RLO is triggered during hydrogen core-burning by a late main sequence star, hydrogen shell-burning by a giant star, or helium burning or beyond by a giant or super giant star (Kippenhahn and Weigert, 1967; Tauris and van den Heuvel, 2006; Tauris, 2011). During mass transfer, these system becomes glaringly bright in the X-rays, with total luminosity's between  $10^{34}$  and  $10^{39}$  erg s<sup>-1</sup>,<sup>28</sup> and are thus named *X-ray binaries* (Tauris and van den Heuvel, 2006).

Efficient recycling will occur only in the right conditions. Naively, one would expect that more massive donor stars can transfer more mass and thus should produce more MSPs, but it is actually the opposite. Surely at some point in your life, you have eaten too much and too fast, resulting in an unpleasant indigestion. Do not be ashamed, as it happens to neutron stars too. The accretion-induced brightness of a star can be so high that radiation pressure overcomes the gravitational force, pushing the infalling material away. Therefore, there is a maximum luminosity an accreting neutron star can achieve, the so-called *Eddington limit* (Shakura and Sunyaev, 1973; Poutanen et al., 2007)

$$L_{\text{Edd}} = \frac{GM\dot{M}_{\text{Edd}}}{2R_{\text{in}}} = \frac{1}{12}c^2\dot{M}_{\text{Edd}} = \frac{4\pi GMm_p c}{\sigma_T} = 1.5 \times 10^{38} \frac{M}{M_{\odot}} \text{erg s}^{-1}, \quad (1.22)$$

where  $G$  is Newton's constant of gravitation,  $M$  is the mass of the accreting object,  $\dot{M}_{\text{Edd}}$  is the accretion rate at the Eddington limit,  $R_{\text{in}}$  is the inner radius of an accretion disc,  $m_p$  is the mass of the proton and  $\sigma_T$  is the Thomson scattering constant. For a typical neutron star, that implies a maximum accretion-induced brightness of  $L_{\text{Edd}} \approx 2 \times 10^{38}$  erg s<sup>-1</sup> (100,000 times brighter than the Sun!) and a maximum accretion rate of  $\dot{M}_{\text{Edd}} \approx 0.04 M_{\odot} \text{Myr}^{-1}$ . Despite that limit, there is strong evidence for super-Eddington accretion onto neutron stars in the past of some pulsar systems (e.g. Tauris et al., 2011; Cognard et al., 2017), and in the present of some extragalactic ultraluminous X-ray sources<sup>29</sup> (e.g. Israel et al., 2017; Carpano et al., 2018; Bachetti et al., 2014)

What does this have to do with MSPs? As depicted in Figure 1.12, X-ray binaries are classified as *low-mass* (LMXB) or *high-mass* (HMXB) depending on the mass of the donor. LMXBs have slowly-evolving donor stars with  $\lesssim 2M_{\odot}$ , which expand until they fill their RLO, kick-starting a sub-Eddington mass transfer that can last for millions or billions of years (Bahramian and Degenaar, 2023), thus producing MSPs with  $P_s < 10$  ms. A side-effect of this long-lasting interaction is that tidal forces circularise the orbit<sup>30</sup> (Zahn, 2008), so at the end of the interaction the remaining eccentricity can be as low as  $e = 10^{-7}$ . On the other hand, HMXB have fast-evolving massive donor stars with  $\gtrsim 6 M_{\odot}$ . In these systems, mass transfer already occurs before RLO via stellar winds or as the neutron

<sup>27</sup> This phenomenon is explained via the physics of the *reduced* three-body problem (Roy, 2005), but in short, If you get too far away from your home star, your orbit will be disturbed by the presence of the other star. Liu Cixin wrote a very nice science fiction book with this exact theming.

<sup>28</sup> This occurs because the potential gravitational energy of the infilling material is turned into heat (due to friction) as the plasma accumulates on the accretion disc, making it extremely hot and bright. As a reminder, the Sun's total luminosity is  $3.8 \times 10^{33}$  erg s<sup>-1</sup>, so this is several orders of magnitude brighter!

<sup>29</sup> I wrote my master's thesis on these guys!

<sup>30</sup> When the donor star is filling or is about to fill its RLO, it takes on an elongated shape. This leads to rotation torques if the star's rotation period is not synchronized with the orbital period and if the orbit is eccentric. Therefore, tidal dissipation tends to seek tidal locking and to circularise the orbit.



star plunges into a disc of material expelled from its massive companion as it orbits it in an eccentric orbit (Tauris and van den Heuvel, 2006; van den Heuvel, 2019). When the donor in a HMXB reaches RLO, its the expanding envelope quickly overfills not only the L1 bridge but also the entire Roche lobe, spilling around the binary and triggering an unstable, super-Eddington runaway mass transfer while the envelope engulfs both the neutron star and core of the donor, putting the systems through a *common envelope* stage (CE, Tauris and van den Heuvel, 2006; van den Heuvel, 2019). In a way, the neutron star enters the outer layers of its evolving companion. The drag caused by the surrounding plasma triggers the in-spiral of the orbit of the neutron star. This same drag transfers the orbital energy of the neutron star onto the envelope, leading to its eventual dissipation if the neutrons tar does not merge with the core of the expanding star first (van den Heuvel, 2019). However, after dissipation, the neutron star will not have accreted a enough mass for recycling (Tauris et al., 2012; Lazarus et al., 2014; Tauris et al., 2017).

In both cases, the donor star ends up losing its envelope, leaving behind a neutron star in orbit with the core of the previous star. However, this is only one stage of interaction. In reality, the formation of the observed pulsars in binary systems is much more intricate and complex.

### 1.3.4 Binary evolution II: the making of the binary zoo

Now that we know how to recycle a pulsar, the question that remains is: how do all of the pulsars from Fig. 1.11 end up with their companions? Can we explain every single binary dot in Fig. 1.11? Answering this question has provided a salary of many astrophysicists for more than a few decades, and many achievements have been made in the understanding binary evolution (for a very big review on binary evolution and pulsar recycling, see e.g. Tauris, 2011; Tauris and van den Heuvel, 2023). In this section, I will try to summarise it withing what is humanly possible to write in a PhD thesis introduction. So, take good hold on the diagrams in Fig. 1.11 and 1.13, because we are in for a trip.

#### Canonical pulsars in binaries

Everyone with siblings knows that not all are treated equally. It is the same story for binary stars. In the Milky Way, they are born at the same time as a pair hydrogen core-burning main sequence stars (**MS–MS**), but with a mass asymmetry. The *primary star* is the more massive one at birth, always shown as the left-side component in Fig. 1.13. As explained in sections 1.2.1 and 1.2.2, the primary star exhausts its hydrogen reserves firs, evolving and becoming the first-born pulsar. As shown in Fig. 1.13, canonical pulsars with main sequence companions (**PSR–MS**) are an intermediate stage of binary evolution, observable after the birth of the first pulsar but before the less massive *secondary star* evolves and recycling occurs. Since the supernova event involves a mass loss and a supernova kick, these systems are eccentric. They are rare and short-lived, with the pulsar entering the graveyard and/or the secondary star evolving and turning the system in an X-ray binary within a few hundred million years. Therefore, only 7 systems of this kind are known, with **PSR B1259–63** or **PSR J0045–7319**<sup>31</sup> being prime examples (Johnston et al., 1992; McConnell et al., 1991). With characteristic ages of  $\tau_c \approx 300$  ky and  $\tau_c \approx 3$  My, they are in wide an very eccentric orbits ( $e \approx 0.8$ ) around luminous  $\sim 10 M_\odot$  main-sequence companions (Johnston et al., 1994; Bell et al., 1995).

<sup>31</sup> This pulsar is actually within the Small Magellanic Cloud, but the pulsar population there is believed to follow the same evolutionary paths as those in the Milky Way, despite not being "Galactic" in a strict sense.

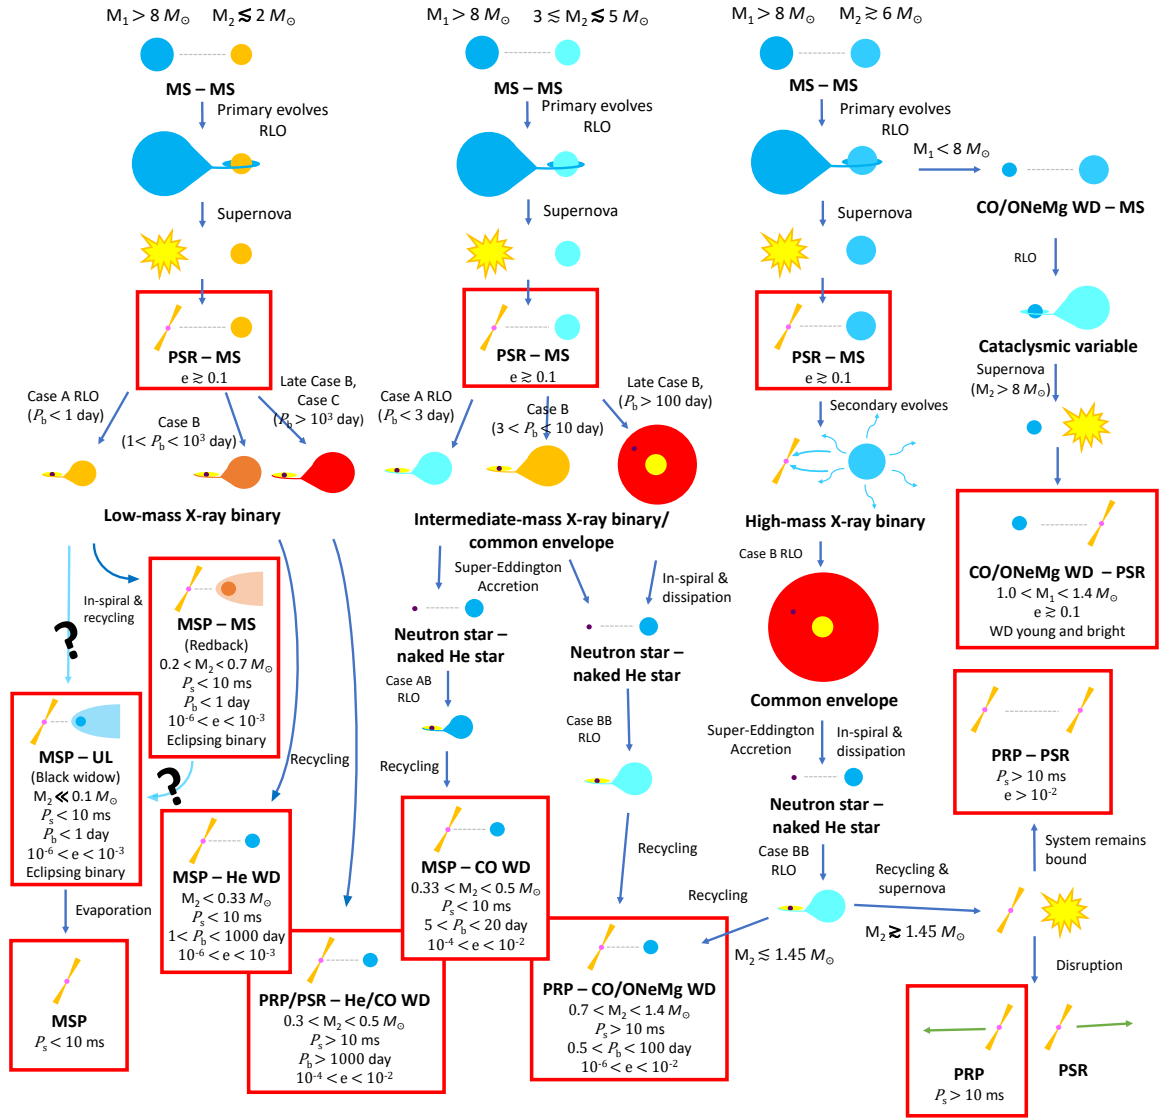


Figure 1.13: Comprehensive (but not 100% exhaustive!) evolutionary tree of pulsars in binary systems according to the masses of the secondary star and the orbital parameters. Highlighted in red squares are potentially observable pulsar systems resulting from binary evolution. The diagram has been produced thanks to the very valuable insights provided in Tauris and Savonije (1999), Tauris and Sennels (2000), Tauris (2011), Tauris et al. (2012), Chen et al. (2013), Benvenuto et al. (2014), Lazarus et al. (2014) and Tauris et al. (2017), and references therein.



However, the avid reader will have noticed in Fig. 1.13 that, before the pulsar is formed, the system goes through a first stage of mass exchange from the primary star to the secondary star, resulting in some degree of mass loss. If the primary loses too much mass and reaches  $M_1 < 8 M_\odot$ , it may become unable to trigger a supernova, becoming a CO or a ONeMg WD instead. Ironically, if the secondary has  $M_2 > 8 M_\odot$  or achieves it during the mass transfer, it may trigger a supernova and become a pulsar forming an exotic, eccentric massive white dwarf–canonical pulsar binary (**CO/ONeMg WD–PSR**, Tauris and Sennels, 2000). This is the rarest kind of binary pulsar, with only two confirmed specimens: **PSR B2303+46** and **PSR J1141–6545** (Dewey et al., 1985; Kaspi et al., 2000), both of them with ages of  $\tau_c \approx 30$  Myr and with bright optically identified  $> 1.2 M_\odot$  white dwarf companions with (Thorsett et al., 1993; Thorsett and Chakrabarty, 1999; van Kerkwijk and Kulkarni, 1999; Bhat et al., 2008; Antoniadis et al., 2011).

### Spider systems

In the closes binary systems we find stellar murderers: the spider systems, MSPs with evaporating semi-degenerate *ultralight* (UL) and main-sequence light companions in Fig. 1.11. Fully recycled MSPs with light companions hail from light systems that went through a LMXB stage. If the original orbital period was  $P_b < 1$  day, the system spirals-in during the X-ray binary stage (Verbunt and Zwaan, 1981) resulting in the shrinking of the Roche lobes and a run-away mass loss for the secondary star. This leads to the formation of *redback* (**MSP–MS**) and *black widow* systems (**MSP–UL**), where the secondary stars are in a constant state of irradiation that slowly drives them to evaporation (Kluźniak et al., 1988; Ruderman et al., 1989). In some extreme cases, the companion is entirely evaporated, leaving an isolated MSP behind as in the case of **PSR B1937+21**, the very first pulsar mentioned in his thesis. As you can see, it is not pleasant to be the secondary star in these systems<sup>32</sup>. These systems are usually observed as binary MSPs with orbital periods shorter than a day and with periodic eclipsing as the material stripped from their companions gets on the way. Several dozens of these systems are known in the Milky Way (Roberts, 2013), but whether black widows and redbacks are two distinct populations (Chen et al., 2013), or if black widows descend from redback systems (Benvenuto et al., 2014) is a matter of ongoing discussion.

### MSPs with light He WD companions

Not all LMXBs end up with stellar murders: systems with  $P_b > 1$  day before the RLO and the LMXB stage have a more pleasant evolution. Efficient recycling still occurs, but without dooming the companion into evaporation. Instead, detached and circular **MSP–He WD** systems are formed (Tauris and Savonije, 1999), where the secondary star has lost enough mass in the Case B RLO that is no longer able to fuse helium into carbon and oxygen. Thus, while for isolated stars it would take at least 60 Gyr to evolve into a He WD (see Fig. 1.1), this evolution channel has already brought them into existence within the 13.8 Gyr of the Universe. As shown in Fig. 1.11, the orbital eccentricity increases with increasing orbital period, but so does the efficiency of the recycling, with PRPs forming instead of MSPs in the most detached systems. Rappaport et al. (1995) and Tauris and Savonije (1999) also point to a direct correlation between the resulting orbital period  $P_b$  and the final mass of the white dwarf, where systems with wider orbits result in more massive He WD due to the smaller mass loss

<sup>32</sup> In fact, the names *black widow* and *redback* refer to two cannibal species of spiders where, fittingly, the female eats the male after mating encounters: the American and European *black widow spider*, and the Australian *redback spider*.

in the secondary star. In the most detached systems with  $P_b > 1000$  days, the resulting pulsar can be a PRP or even a canonical one from a lack of interaction, and a CO WD can form instead if the nuclear fusion of helium is triggered before or during the interaction, leading to the formation of the few wide, circular and light **PRP/PSR-He/CO** systems, which are shown as the widest and least recycled systems in Fig. 1.11. **MSP-He WD** are the most common type of pulsar binary in the Milky Way, with a total amount of 145 published specimens in April 2024. Since the spin evolution of MSPs is extremely slow ( $\dot{P}_s < 10^{-19} \text{ s s}^{-1}$ ), these systems are amongst the most ancient objects we see in the Milky Way, with their pulsars staying up and bright for billions and billions of years.

### MSPs with light CO WD companions

As seen in Fig. 1.13, a subset of light **MSP-CO** systems also exists, which arises from an exotic evolutionary channel involving an *intermediate* X-ray binary (IMBX, secondary star with  $3 \lesssim M_2 \lesssim 5 M_\odot$ ). In this case, the pulsar is not recycled during the first RLO after its formation, but on a later stage. After the secondary loses its hydrogen envelope during the Case A RLO in a IMBX, it leaves behind its naked, non-degenerate core-burning helium core, a *naked He star*. When helium starts to exhaust in the core, the naked He star also expands as it transitions into a helium shell-burning star, triggering a Case AB RLO, and the efficient recycling of the pulsar (Tauris et al., 2012). This class of binaries form the bridge between the sequences of pulsars with light He WD companions and heavier CO/ONeMg WD companions in the  $P_b - x$  diagram in Fig. 1.11 at the orbital period range of  $5 < P_b < 20$  days. This arising class is a relatively recent discovery, with only handful on known systems, the first one of which being **PSR J1614-2230** (Crawford et al., 2006; Tauris et al., 2011). Other well-studied cases are **PSR J1933-6211** (Crawford et al., 2006; Geyer et al., 2023) and **PSR J1125-6014** (Faulkner et al., 2004; Shamohammadi et al., 2023).

### PRPs with massive CO/ONeMg WD companions

The exhaustive visit of the binary pulsar zoo now leads us to the more massive, less common members of the white dwarf family! As shown in Fig. 1.13, massive **PRP-CO/ONeMg WD** systems follow a similar evolutionary path to that of light **MSP-CO WD** systems, but they stem from wider IMXB ( $M_2 > 3 M_\odot$ ,  $P_b > 3$  day) or HMXB ( $M_2 > 6 M_\odot$ ) systems that go through a Case B RLO (Lazarus et al., 2014). In the most massive and/or wide systems, the binary even undergoes a CE stage, which leads to an in-spiral of the system (Lazarus et al., 2014). Like in light **MSP-CO** systems, the recycling of the pulsar occurs at a Case BB RLO stage when the left-over naked He secondary moves into the helium shell-burning stage. However, unlike in the lighter systems, the rapid evolution of these more massive stars makes it so that even in this stage, the mass transfer is super-Eddington and short-lived, occurring in less than  $10^5$  years e.g. Lazarus et al., 2014; Guo et al., 2021. Therefore, as shown in Fig. 1.11, the pulsars are often only partially recycled and the resulting orbital eccentricities tend to be higher, up to  $e \approx 10^{-2}$ . As seen in the  $P_b - x$  diagram of Fig. 1.11, these systems form a well-distinguished sequence, separated from that of light **MSP-He WD** systems. Indeed, and as pointed out in McKee et al. (2020), there is a well-defined mass gap between these two types of systems, with the white dwarf in lighter systems reaching up to  $\sim 0.5 M_\odot$ , while the heavier ones lying in the range of  $0.7 \lesssim M_2 \lesssim 1.4 M_\odot$ . Within this latter group, the ones with  $M_2 \gtrsim 1.1 M_\odot$  are actually very likely to host ONeMg WDs instead of CO WDs, making them the most massive pulsar-white dwarf systems that we know of. As of April 2024, there are 44 confirmed heavy **PRP-CO/ONeMg**

**WD** systems in the Galaxy, 4 of which **PRP-ONeMg WD**, but before the work in this thesis, only two of them are confirmed to host massive ONeMg WD companions: **PSR J2222–0137** (Boyles et al., 2013; Cognard et al., 2017; Kaplan et al., 2014; Guo et al., 2021) and **PSR J1528–3146** (Jacoby et al., 2007; Berthereau et al., 2023), both of them with  $> 1.2 M_{\odot}$  white dwarfs and, in the case of **PSR J1528–3146**, an optical detection (Jacoby et al., 2006).

### Pulsars with pulsars: double neutron star systems

Double pulsar, double trouble, double excitement! Double neutron stars (DNS) systems are the direct siblings of the the most massive **PRP-ONeMg WD** systems, as seen by their large overlap in Fig. 1.11. However, they have one very significant difference: much larger orbital eccentricities, with  $e \gtrsim 0.1$ . Where does it come from, if all of them are recycled PRPs, and recycling induces a circularisation of the orbit? As shown in Fig. 1.13, DNSs system (**PRP-PSR**) follow the same evolutionary as **PRP-ONeMg WD**, but after the recycling of the primary pulsar, if the evolved naked He secondary star preserves a mass of  $\gtrsim 1.45 M_{\odot}$ , it undergoes an ultra-stripped supernova and forms another pulsar (Tauris et al., 2015). The ultra-stripped supernova entails the loss of  $0.2 - 0.4 M_{\odot}$  and a supernova kick of  $< 50 \text{ km s}^{-1}$ , either resulting in the total disruption of the system, producing the several isolated PRPs from Fig. 1.11, or introducing a significant eccentricity (Tauris et al., 2017).

Why are these systems so special? Aside from their orbital eccentricity, which leads to better mass measurements (more on this in Chapter 2) they are the only binary where we can potentially observe both the recycled primary and the canonical secondary, making them quite exciting despite compromising only 21 of known Milky Way pulsar binaries. The not-yet-mentioned-for-enough-times Hulse–Taylor binary proudly belongs to this category, with the observed pulsar being the recycled component ( $P_s = 59.03 \text{ ms}$ ,  $\dot{P}_s = 8.6183 \times 10^{-18} \text{ s s}^{-1}$ ), as it is the case the vast majority of them due to the slower evolution of PRPs with respect to canonical pulsars. However, in a handful of systems, the visible components is actually the second-born, canonical pulsar, such as in **J1906+0746** ( $P_s = 144.07 \text{ ms}$ ,  $\dot{P}_s = 2.02678 \times 10^{-14} \text{ s s}^{-1}$ , Lorimer et al., 2006). The crown jewel of Galactic DNS systems must surely be **PSR J0737–3039A/B**, the one and only true *double pulsar* system where both the recycled (A,  $P_s = 22.70 \text{ ms}$ ,  $\dot{P}_s = 1.75993 \times 10^{-18} \text{ s s}^{-1}$ ) and canonical (B,  $P_s = 2.773 \text{ s}$ ,  $\dot{P}_s = 8.92 \times 10^{-16} \text{ s s}^{-1}$ ) components are observed as radio pulsars (Burgay et al., 2003; Lyne et al., 2004; Kramer et al., 2021a).

## 1.4 Beyond just pulsars: what do we get from binary systems?

While pulsars and their binaries are very interesting by themselves, it is natural to ask what is the impact of everything I have mentioned above, other than “just pulsars”? It turns out that there is one piece of information I have hidden until now: in many binary systems, we can measure the pulsar mass,  $M_p$ , the companion mass,  $M_c$ , and the inclination angle,  $i$ , to high accuracy. And we have to give our thanks to Albert Einstein for this. The specifics of how this is done are given in Chapter 2, but for now I can say that, in the most compact systems, the theory of general relativity (GR) predicts deviations of the orbital motion of pulsars with respect to Newtonian physics, which are measurable thanks to the clock-like nature of recycled pulsars. This leads to many interesting experiments in fundamental physics.

So, we have come to understand what are the different populations of Galactic pulsar binaries and




Figure 1.14: Histogram with all confirmed neutron star mass measurements from pulsar systems as of July 2024. Most of them are pulsar mass measurements, while names with (C) indicate a companion mass measurement. Red names are confirmed neutron stars, while blue indicate that the nature of the companion star is unclear (its mostly Globular Cluster systems, where the eccentricity does not indicate whether the companion is a neutron star or a massive white dwarf). The plot is was made by Vivek V. Krishnan, and is taken from [Paulo Freire's webpage](#).

how to extract astrophysical information from them. But, what can we do with these measurements? What aspects of physics and astrophysics can we study? In this last section, I will explain which open questions are touched upon by pulsars in binary systems.

#### 1.4.1 The neutron star mass distribution

Measuring pulsar masses is the the most reliable way of probing the neutron star mass distribution. The presently sampled distribution, cointining more than 50 measurements, is shown in Fig. 1.14. For reference, the DNS **PSR J0453+1559** contains the lightest known pular, with  $M_p = 1.174 \pm 0.004 M_\odot$  (Deneva et al., 2013; Martinez et al., 2015), while the most massive pulsar from radio observations is in the black widow system **PSR J0740+6620**, with  $M_p = 2.08 \pm 0.07 M_\odot$  (Lynch et al., 2018; Fonseca et al., 2021). One of the questions that kept astronomers wondering was weather this measured distribution reflects the *birth* neutron star mass distribution, and weather there is actually a unique distribution. Indeed, it has been proposed that the neutron star mass distribution may be bimodal, with a canonical and PRP-dominated group centered at  $M_p \sim 1.4 M_\odot$  with a Gaussian spread of  $0.07 M_\odot$ , and a MSP-dominated group at  $M_p \sim 1.8 M_\odot$  with a spread of  $0.18 M_\odot$  (Antoniadis et al., 2016). This has led some to speculate that the MSP mass distribution does not reflect the neutron star birth mass distribution, and that MSP may have acquired a significant fraction of mass from their companions during the recycling process, such as in the case of **PSR J0952-0607**, a MSP with  $P_s = 1.41$  ms and  $M_p = 2.35 \pm 0.17 M_\odot$  in a black widow system (Romani et al., 2022)<sup>33</sup>. However, this postulate has been mostly refuted on the basis that massive  $M_p \gtrsim 1.8 M_\odot$  PRPs have been measured, such as in

<sup>33</sup> This pulsar's mass was measured via spectroscopical analysis of its black widow companion instead of timing, which requires some modelling assumptions. But if believed, it would technically be the most massive pulsar.

**PSR J2222–0137**, with  $P_s = 32.82$  ms and  $M_p = 1.831(10) M_\odot$  in a massive **PRP–ONeMg WD** system (Cognard et al., 2017; Guo et al., 2021).

To strengthen this point, we can look at the equation for estimating the amount of accreted mass during recycling  $\Delta M_p$  derived by Tauris et al. (2012):

$$\Delta M_p \approx 0.22 M_\odot \frac{(M_p/M_\odot)^{1/3}}{(P_s/\text{ms})^{4/3}}. \quad (1.23)$$

as a function of the original  $M_p$  and the final spin-up period. If this equation is true, recycling a  $M_p = 1.35 M_\odot$  pulsar to  $P_s = 1.5$  ms would require the accretion of  $\Delta M_p \approx 0.14 M_\odot$ , more common MSPs with  $P_s = 5$  ms would require  $\Delta M_p \approx 0.03 M_\odot$  and fast PRPs with  $P_s = 20$  ms would only require  $\Delta M_p \approx 0.005 M_\odot$ . This implies to form even the fastest-spinning MSPs, the accretion  $\Delta M_p \lesssim 0.1 M_\odot$  sufficient (Antoniadis et al., 2016). Therefore, the possible bimodality of the pulsar mass distribution is still an open question, and the only way for getting more insights are further mass measurements!

### 1.4.2 The physics of ultra-dense matter

Neutron star mass measurements are key in studying condensed matter physics. Their density is equivalent to that of atomic matter, with the possibility of quark-gluon plasma and other exotic states of matter existing inside of their cores (Ascenzi et al., 2024). But, how can we experiment with its properties without taking a spoonful of it and bringing it to Earth and making it explode in the largest fission reaction mankind has ever witnessed? The solution, once again, is to measure neutron star masses. That is because every model of ultra-dense matter predicts a *equation of state* in microscopic physics, establishing a relationship between pressure and density. This relationship is then translated into a mass–radius relationship for neutron stars, with different models predicting different maximum masses they can sustain without collapsing into a black hole, ranging from 1.5 to 2.5  $M_\odot$  (e.g. Lattimer and Prakash, 2001). Thus, measuring neutron star masses, especially the most massive ones, becomes a viable way of testing micro physics (Özel and Freire, 2016). In conjunction with neutron star radius constraints from the modelling of X-ray emission of some pulsars (e.g. Riley et al., 2019; Riley et al., 2021), this shows that pulsars are Universe-given condensed matter laboratories!

### 1.4.3 Supernova physics and kicks

The observation of young pulsars in supernova remnants give us very good insights on the magnitude of supernova kicks, which can be used to model the details of supernova physics (e.g. Janka et al., 2022). However, this comes with one limitation: we lack a measurement of the neutron star mass, which contains a significant fraction of the mass of the progenitor star. That is not the case for DNS systems, where we can both measure the mass of the second-born neutron star and get informed about the kinematics of the supernova from the remaining orbital parameters (e.g. Tauris et al., 2017; Andrews and Mandel, 2019).

Currently, there is tentative evidence from DNS systems that more massive second-born neutron stars receive larger supernovae kicks based on their larger orbital eccentricities (Andrews and Mandel, 2019), and simulations suggest that the more massive naked He-burning stars could trigger an iron




Figure 1.15: **Left:** orbital eccentricity against second-born neutron star mass in DNS systems with good mass measurements. **Right:** orbital eccentricity against orbital period for all known DNS systems. Massive systems are those that have a large second-born neutrons tar mass, while possible (unconfirmed) DNS systems are those without enough mass constraints to confirm a DNS nature of the system, but whose orbital and spin parameters are consistent with it.

core-collapse supernova instead of an electron-capture supernova after the recycling stage (Tauris et al., 2015). The numbers behind these notions are shown in Fig. 1.15, where the DNS systems according to their orbital parameters and second-born neutron star masses. In the left-side diagram, it seems apparent that there is a positive correlation between  $e$  and the mass of the second-born NS, while on the right diagram it is observed that there is a divide between the higher eccentricity/mass and the lower eccentricity/mass DNS systems, perhaps caused by a larger magnitude of the supernova kick in more massive systems. In both plots, a striking separation is observed between high and low-eccentricity systems, which has been pointed out as the distinction between iron core-collapse and electron capture supernovae (Andrews and Mandel, 2019). The right-hand side plot also shows a tentative positive correlation between  $e$  and  $P_b$  within each group, interpreted as the progenitor He star preserving more mass before the supernova event and thus resulting in a larger supernova kick. Nonetheless, the sample of DNS systems with good mass measurements is still small, so new discoveries and mass measurements are needed to confirm these notions.

#### 1.4.4 The white dwarf mass distribution

While neutron stars are extremely interesting by themselves, we should not forget about their lighter cousins: white dwarfs. Modelling of their optical emission has proven very useful in constraining their masses and other properties, such as temperature, spin, radii, magnetic fields and compositions (e.g. Hollands et al., 2020; Caiazzo et al., 2021). However, mass measurements from binary pulsar analysis is the most reliable and assumption-free way to measure their masses. Measuring white dwarf masses is necessary, for example, to explore the true physics of the Chandrasekhar limit. For example, many already-mentioned simulations of stellar evolution place it at the  $1.38 M_{\odot}$  threshold (e.g. Nomoto et al., 1984; Takahashi et al., 2013; Chanlaridis et al., 2022), but if the star is fast-rotating it could exist up to  $1.48 M_{\odot}$ , or even higher for differentially rotating white dwarfs (e.g. Yoon and Langer, 2004). Currently, very few timing mass measurements exist on this side due to the scarcity of known



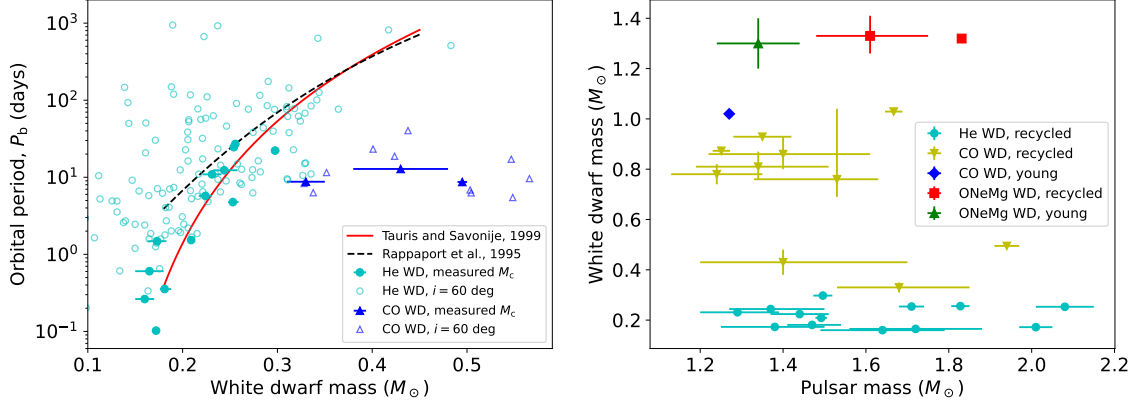


Figure 1.16: **Left:** orbital period against against white dwarf mass in light **MSP–He WD** and **MSP–CO WD** systems. Proper mass measurements are presented with error bars, while for white dwarfs without them we present the median companion mass from the mass function instead ( $i = 60^\circ$ ). The continuous and dashed lines trace the  $P_b - M_{WD}$  relationships predicted by Rappaport et al. (1995) and Tauris and Savonije (1999). The three CO WD systems with mass measurements are **PSR J1614–2230**, **PSR J1933–6211** and **PSR J1125–6014**. **Right:** white dwarf mass against pulsar mass for all systems with white dwarf companions and mass measurements, divided according to the nature of the white dwarf. The two unrecycled systems are **PSR B2303+46** and **PSR J1141–6545**.

**PRP–ONeMg WD** systems, so it is our mission to discover and measure more of them!

Other interesting questions relate to the masses of white dwarfs in binary systems. For instance, several studies predict an increasing relationship between the white dwarf mass and  $P_b$  in **MSP–He WD** binaries, independently of the original donor mass star as a result of binary evolution, which is thought to be caused directly by the delay in the onset of the RLO due to orbital separation (e.g. Rappaport et al., 1995; Tauris and Savonije, 1999). However, as shown in the left side of Fig. 1.16, this prediction has been tested only with a handful of mass measurements, while more than a hundred masses remain uncertain. It also seems that light **MSP–CO WD** binaries do not follow these prediction, as expected from their diverging evolutionary history. Finally, as shown in the right side of Fig. 1.16, the global white dwarf mass distribution in binary systems seems to be bimodal, with a significant gap separating the  $M_{WD} \lesssim 0.5 M_\odot$  systems with MSPs and the  $M_{WD} \gtrsim 0.7 M_\odot$  systems with PRPs (McKee et al., 2020; Shamohammadi et al., 2023), which is also understood from their diverging evolutionary histories. However, the  $M_c \gtrsim 1.1 M_\odot$  side of the distribution remains sparsely sampled, and it is not entirely clear that it is contiguous with the rest of the distribution. Could the white mass distribution be actually trimodal? Only more measurements will tell.

#### 1.4.5 Pulsars and gravitational-wave astronomy

We are at the advent of gravitational-wave astronomy and pulsars are more relevant than ever, directly and indirectly, even for isolated pulsars. It turns out that, if neutron stars were to be asymmetric with respect to their rotation axes, their own rotation would emit continuous gravitational waves. In the case of MSPs, their rotation frequency enter into the bands ground-based GW observatories such as *LIGO*, *Virgo* and *Kagra* ( $\sim$ kHz), making them potentially detectable sources. Despite years of searching, no detection has been made even from the fastest and closest MSPs, but efforts are still

ongoing (Wette, 2023).

On the other extreme of the gravitational-wave spectrum, the clock-like stability of pulsars turns them into a gravitational-wave detectors: by measuring correlations in the rotational variability's of pulsars at different positions of the sky, astronomers can look for gravitational-waves coming from outside our galaxy, turning the array of analysed pulsars into a Galaxy-sized gravitational-wave detector! Since pulsars accumulate years of observations, this method could potentially detect low-frequency signals ( $\sim$ nHz) coming from merging super-massive black holes in distant galaxies. First evidence for these signals have already been announced by the *pulsar timing array* collaborations (PTAs: EPTA and InPTA Collaborations, 2023; NANOGrav Collaboration, 2023; CPTA Collaboration, 2023; PPTA Collaboration, 2023). The discovery and study of more pulsars will allow for further constraints on gravitational waves.

However, more directly related with the work in this thesis is the emission of gravitational-wave radiation from coalescing pulsar binaries. Future space-based missions like the Laser Interferometer Space Antenna (*LISA*) will be able to detect gravitational emission from thousands of Galactic coalescing binaries in the  $\sim$ mHz frequencies of (e.g. Korol et al., 2017; LISA Consortium, 2023). This will include, of course, binaries containing neutron stars, making a great observing counterpart to pulsar observations. Most of these binaries will be circular, but the detection of tens of new eccentric, previously undetected DNS has already been predicted, which will allow for a targeted search for pulsated signals with the SKA (Lau et al., 2020).

But we do not have to wait all for observing gravitational waves from neutrons tar binaries: *LIGO* and *Virgo* have already observed merging neutron stars, such as GW 170817 (LIGO and Virgo Collaborations, 2017). Galactic DNS systems are the direct progenitors of these events, and thus they can be used to predict what we expect to see in the GW regime. The case of merging neutron stars is especially enticing, as they are thought to be the main seeders of heavy metals in the Milky Way (Chen et al., 2021a)<sup>34</sup>. Currently, nine pulsars are known in Galactic NS systems with a merger time shorter than the current age of the Universe<sup>35</sup>, and from this a merger rate of 30 events per million year is derived in the Milky way, or  $\sim 300$  events/Gpc<sup>3</sup> yr<sup>-1</sup> in the local Universe<sup>36</sup>, leading to a potential detection rate of  $\sim 3$  events per year with *LIGO*. Further sky surveys and discoveries of merging DNS systems will modulate the precision and accuracy of these predictions.

## Testing theories of gravity

We are now delving into a more exotic aspects of binary pulsars, but nonetheless a fundamental one. With binary pulsars, we have confirmed many of the predictions of GR. An example I have mentioned many times and I will mention as many more times as it is needed (or as I want to) is the Hulse–Taylor binary, where the orbital decay due to the emission of GW waves was observed for the first time (Taylor et al., 1979; Taylor and Weisberg, 1982). But there is much more than this. The discovery of the exceptional system in **PSR J1141–6545**, a  $P_s = 0.394$  s canonical pulsar in orbit with a fast-rotating CO WD (Fonseca et al., 2021) led to the detection of the dragging of space-time from the spin of the white dwarf, an additional prediction from GR (Antoniadis et al., 2011). Or for instance, the discovery

---

<sup>34</sup> In fact, strontium has been detected in the kilonova (explosion from the merger of two neutron stars) associated with GW 170817 (Watson et al., 2019).

<sup>35</sup> Merger time: time remaining until the two stars of the binary collide due to the orbital decay. If it is shorter than the age of the universe, these systems have the characteristics of systems that merged in the past.

<sup>36</sup> In the local Universe, there are 10,000 galaxies within 1 Gpc<sup>3</sup>



of **PSR J0337+1715**, a  $P_s = 2.733$  ms MPS with not one, but *two* He WD companions (Ransom et al., 2014) allowed for the test of the most fundamental assumption of GR: the strong equivalence principle, according to which, all objects are accelerated the same way by a gravitational field independently of their own mass, which was passed with success as well (Voisin et al., 2020). And to mention one last success of GR, we can call back to the double pulsar **PSR J0737–3039A/B**, where we can probe higher-order corrections of the propagation of light in the largest space-time curvature because pulses from pulsar A pass as close as 11,000 km from the surface of pulsar B, leading to the measurement of higher-order corrections of the Shapiro delay that are observed in no other known system, pulsar or not. The measurements, once again, agree with the prediction by GR (Kramer et al., 2021a; Hu et al., 2022).

But, is GR always right? We may also want to test this. For instance, according to GR, gravitational-wave emission occurs only due to the existence of an oscillating mass quadruple in a binary system<sup>37</sup> (Peters and Mathews, 1963). However, if one relaxes some of the assumptions of GR some funky effects arise and gravitational waves may come from an oscillating dipole, which would mean that the orbits in binary systems with a large compactness asymmetry, e.g. pulsar–white dwarf systems, would decay faster than expected. Studying pulsars in close orbits with white dwarf companions is the best way to test these theories. So far, no deviation from GR has been observed (e.g. Zhu et al., 2019; Gupta et al., 2021). Other alternative theories, such as the scalar-tensor gravity, have also been mostly discarded via pulsar experiments (Zhao et al., 2022). But who knows what new theories could be tested with newer discoveries in the future. The holy grail of tests of gravity would be the discovery of a pulsar in orbit with a black whole, where we would be able to test gravity in the most extreme conditions where all relativistic corrections would come into play or even become dominant: next to the event horizon, proving whether the no-hair theorem is right or wrong, and whether GR holds up even in its most extreme environment (Liu et al., 2014).

---

<sup>37</sup> Quadruple refers to a kind of mass distribution: a monopole is a spherically symmetric field with a single source (e.g. an electrostatic charge or a gravitational mass), a dipole is a field with a sense of direction arising from two opposite sources (e.g. a magnetic field, which has a north pole and a south pole), and a quadruple has four sources in a plane, forming two axis of opposite charge (e.g. the energy-momentum distribution of a binary system, two positive poles being the stars and the “antipoles” being the empty spaces in their orbits).



# Methods: radio data and pulsars

---

## 2.1 Before we start

We now know all about pulsars and how scientifically interesting they are. Hooray! Let us celebrate! But before we move on to the amazing science results presented in this thesis, we have another chapter ahead of us: how to observe pulsars, how to work with pulsar data, and how to extract astrophysical information from it.

Broadly speaking, there are two big disciplines in pulsar astronomy: *pulsar searching* and *pulsar timing* (Lorimer and Kramer, 2005). Pulsar searching seeks to discover new pulsars. It is about creating and implementing efficient algorithms to process the data and identify signals in it. One of the biggest challenges in this field is to find binary pulsars, whose orbital motion tampers with the periodicity of their signal. Pulsar timing is about modelling the spin of a known pulsar down to every single rotation in order to extract astrophysical information from it. From this modelling, we can measure the spin evolution, sky motion, and, in compact binaries, component masses.

In this section, I will explain the data formats and methods used in these two disciplines. However, before we delve into pulsar searches and timing, we also have to understand where the data comes from in the first place. Therefore, this chapter has three parts. First, an introduction to radio astronomy and radio telescopes. Then, a section on the methods of pulsar searching, including the search for binary pulsars. And finally, a section on pulsar timing and the measurement of relativistic effects and component masses in binary systems.

## 2.2 Radio astronomy, radio telescopes, and data acquisition

Astronomical use of radio antennas was pioneered by Karl Jansky in the 1930s, when he unexpectedly observed the emission from the supermassive black hole at the centre of the Milky Way while investigating possible sources of interference in radio communications<sup>38</sup> (Jansky, 1933). Today, astronomical radio telescopes are found all across the world. Pulsars emit from the lowest frequency radio waves all the way up to Gamma rays, but they most often observed with radio telescopes in

---

<sup>38</sup> In fact, the unit of radio flux density, the "Jansky" (Jy) is named in his honour

the decimetre wavelength range, close to  $\lambda = 21 \text{ cm}$ <sup>39</sup> or the frequency  $f = 1.4 \text{ GHz}$  (Lorimer and Kramer, 2005), usually balancing the spectral index of most pulsars with the ISM-induced effects.

In this section I will answer the following questions: How do we detect radio waves? How is radio data stored? What does a radio telescope actually see? As I go over these topics, I will introduce useful concepts such as data arrays, telescope sensitivity, beam resolution and interferometry. Finally, I will do a proper introduction of the two telescopes from which the data in this thesis comes.

### 2.2.1 From radio light to radio data: how do we record radio waves?

In Fig. 2.1, a diagram of the signal propagation from analog sampling by the antenna until digitisation is shown. Here I am going to overview the main components depicted there, but for a more in-depth understanding of radio telescopes, please check (Condon and Ransom, 2016). Radio telescopes observing around  $f = 1.4 \text{ GHz}$  are often built as reflecting parabolic antennas, where the radio waves are focused into the *receiver* by a *parabolic dish*. The *receiver* is composed of feed horns, which contain *feeds* inside. Radio waves induce voltage oscillations in the feeds, which are then propagated into the data stream. Usually, feed horns contain two orthogonal feeds (linear or circular) which sample two orthogonal polarisation modes of radio light, each of them sprouting their own data stream. The voltage signals from each data stream are amplified by the *low noise amplifiers*. They are designed to have a maximum response within the desired  $f_1 - f_2$  band. But unfortunately, they also introduce instrumental noise. Therefore, they are cryogenically cooled so that they add as little noise of their own as possible. Then, a *bandpass filter* is applied, which attenuates the signal outside the desired  $f_1 - f_2$  range. Subsequently, the signal is mixed with another *monochromatic signal* of frequency  $f_{\text{LO}}$ , inputted by a *local oscillator*, so that the whole frequency range is lowered by  $-f_{\text{LO}}$  for easier transmission with less loss. Afterwards, further amplifications and bandpass filters are applied and the data is sent to the backend for digital recording and storage.

Finally, the signal is digitised by the backend. A backend is part hardware and part software, designed to store the signal in a specific format. Our interest is in the *pulsar backends*, which aim to preserve time and frequency information in order to study the rotation of pulsars in great detail. Pulsar data is stored in a two-dimensional array of time and frequency (filterbanks), or a three dimensional array of time, frequency, and spin phase (folded archives), depending on whether we want to use it for searching or for timing. Both of these data formats have two elements in common: accurate *time stamps* measured with high-precision hydrogen maser clocks, and discrete *frequency channels* of bandwidth  $\delta f$  (channel bandwidth). Each orthogonal feed gives rise to its own data array, of the same dimensionality and size, but with different values. To store the polarization information, the complex-valued amplitudes of each array are combined into the components of the Poincaré vector

$$\begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} |X|^2 + |Y|^2 \\ |X|^2 - |Y|^2 \\ 2 \text{Re}(X^*Y) \\ 2 \text{Im}(X^*Y) \end{bmatrix}, \text{ or } \begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} |LCP|^2 + |RCP|^2 \\ -2 \text{Im}(LCP^*RCP) \\ 2 \text{Re}(LCP^*RCP) \\ 2|LCP|^2 - |RCP|^2 \end{bmatrix}, \quad (2.1)$$

for two orthogonal linear feeds ( $X$  and  $Y$ ), or for two left-handed and right-handed circular feeds ( $LCP$

<sup>39</sup> This number is not a matter of chance: 21 cm is the hydrogen line emitted by the spin flip of its single electron, a cornerstone of cosmological observations. Thus, telescope receivers covering this wavelength are often called the "21 cm receiver" or something along these lines.

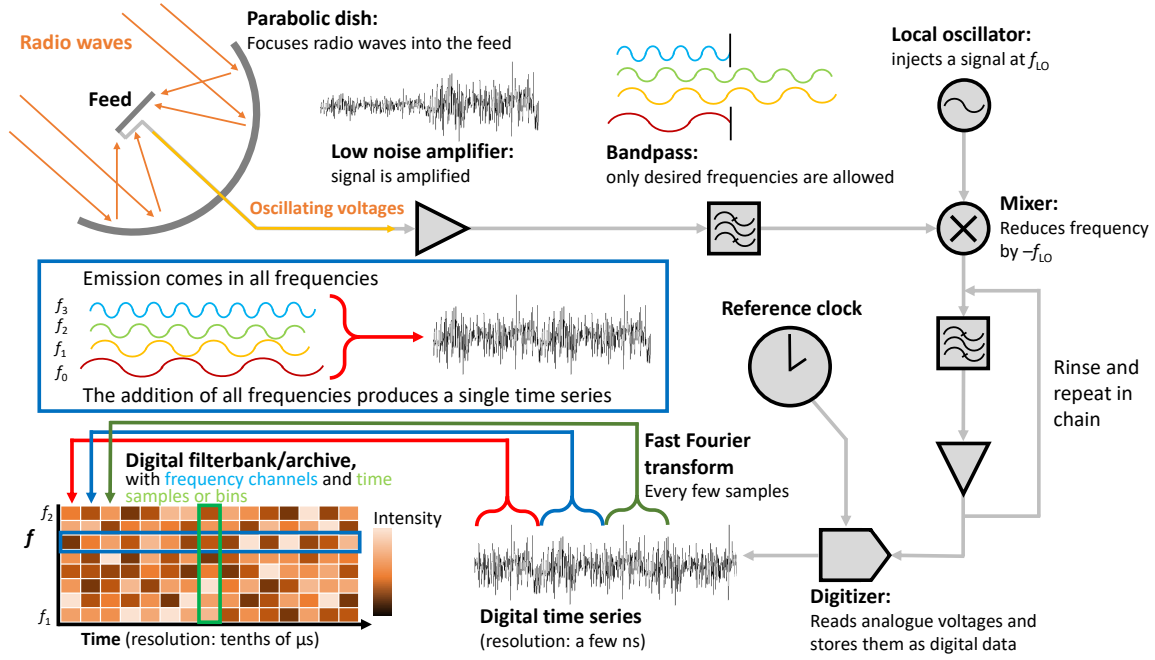


Figure 2.1: Diagram of a radio antenna with a receiver and a pulsar backend with a frequency range  $f_1 - f_2$ . The oscillating electric field from radio waves is captured by a conductive feed as oscillating voltages. The signal is then propagated, digitised and stored via several hardware devices and software tools. This diagram is drawn based on insights provided by Lorimer and Kramer (2005) and Condon and Ransom (2016), and it assumes a single feed with a single data stream. Specific data streams, the presence or absence of components and when and how digitisation occurs is highly telescope-dependent.

and  $RCP$ ), where  $X^*$  stands for the complex conjugate of  $X$  and so on. In the Poincaré vector,  $I$  is the total intensity recorded by both feeds combined, while  $Q$ ,  $U$  and  $V$  encode polarisation state of light (see Appendix 1 in Lorimer and Kramer (2005) for an in-depth explanation.)

## 2.2.2 Radioastronomy 101: what does a radio telescope see?

### Telescope sensitivity

Flux densities in radio are measured in units of Jansky, which quantify the amount of power per unit area per unit frequency ( $1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$ ). For a radio telescope, we define the *source temperature* of a signal collected by the dish as the temperature of a resistor emitting black body radiation of equal power at the same temperature (Condon and Ransom, 2016; Lorimer and Kramer, 2005). The larger is the *collecting area* of the telescope, the greater are the collected power and temperature response. This is quantified via the *telescope gain* (Condon and Ransom, 2016; Lorimer and Kramer, 2005)

$$G_T = \frac{A}{2k_B}, \quad (2.2)$$

which has units of  $\text{K Jy}^{-1}$ , where  $k_B$  is the Boltzmann constant, and the  $1/2$  factor comes from the fact that each orthogonal feed captures only half of the source flux. Multiplying a source flux  $F_{\text{src}}$

by  $G_T$  results in the *source temperature* at the telescope:  $S_{\text{src}} = F_{\text{src}} \times G_T$ , in units of K. Telescope instrumentation (e.g. the low noise amplifiers) also produce noise in the data, dominated by the *system temperature*  $S_{\text{sys}}$ , also in units of K. Therefore, for a given source with a given telescope, the *signal-to-noise ratio* of a source is then  $S/N = n_p F_{\text{src}} \times G_T / S_{\text{sys}}$ , where  $n_p$  is the number of polarisation modes (1 or 2, usually 2). From this, we can see that the factor  $G_T / S_{\text{sys}}$  is a very useful estimator of telescope sensitivity. But, what does a telescope actually see? That depends on how many antennas and horn feeds our telescope has.

### Single-dish telescopes

The region in the sky covered by an individual receiver element is called a *beam*. Beams have an angular extension which depends on the size of the dish and the observed wavelength. Usually, either the antenna or the receiver is steerable, so beams can change their location in the sky, and a receiver containing multiple feed horns can place a beam tiling on the sky, forming a low-resolution image. The beam shape is defined by the signal response to a nearby source, described by

$$S_{\text{src}}(\theta) \propto \text{sinc}^2 \pi \theta \left( \frac{D}{\lambda} \right), \quad (2.3)$$

where  $\theta$  is the angular distance between the source and the center of the beam, and  $D$  is the diameter of the dish (Condon and Ransom, 2016). The left side of Fig. 2.2 plots the recovered  $S_{\text{src}}$  as a function of  $\theta$ , showing that the main beam lobe is found within  $\theta < \lambda/D$ , and that at  $\theta = \lambda/D$  the source signal is lost. Despite that, beam sidelobes arise at  $\theta > \lambda/D$ , so the brightest sources can still be detected even if the beam is pointing away from them. Let us also keep in mind that sky positions and beam shapes are actually two-dimensional, so that  $\theta$  represents distance from the beam center in any direction. Thus, ideally, the main beam lobe is circle-shaped (or ellipsoidal if the dish is not circular) and the side-lobes are concentric rings. However, in reality, obstructions such as the focus cabin supports can modify the beam shape significantly.

The beam size is a measure of *telescope sky resolution*. It is usually measured at the full width at half maximum (FWHM), the angular diameter of the region to which the sensitivity is halved is halved:

$$\delta\theta = 2\theta_{\text{FWHM}} = 0.89 \frac{\lambda}{D}. \quad (2.4)$$

Thus, collecting area and telescope resolution are opposing parameters: increasing  $D$  increases  $G_T$  but reduces  $\delta\theta$ . In the end, this poses a problem: is it better to have a large sensitivity over a small sky area, or a small sensitivity over a large sky area? For pulsar searching, a large angular resolution is actually a disadvantage, because we want to cover as much sky area as possible, but of course, we do not want to lose sensitivity either. The obvious solution to this conundrum is to place as many beams as possible within large collecting area, but extremely large radio telescopes are difficult (and expensive!) to build and maintain. In this case, the smart solution is an interferometer.

### Interferometers

Interferometers are telescopes composed of multiple antennas distributed in a physical *antenna array* (Condon and Ransom, 2016). Each antenna has the same components of a single-dish telescope, and they can in fact operate as individual telescopes, but that would not use their full potential. To begin

with, if multiple antennas coordinate and point towards the same location in the sky, the combined  $G_T$  is the addition of all individual antenna gains, providing a significant boost in telescope sensitivity. But most importantly, the wavelength of radio light is large enough that radioastronomers can record individual waveforms. The right side of Fig. 2.2 shows a simple interferometer composed of only two antennas, separated by a *baseline* distance  $B$ . While wave crests from the same source reach each antenna at different times, we can compute the time delay between antennas,  $\Delta t$ , and then correct for it when adding the recorded signals. Such operation is analogous to the **double-slit experiment**: it results in constructive or destructive interference depending on  $\theta$ , drawn as the orange signal response in the right side of Fig. 2.2, placing telescope resolution at  $0.89\lambda/B$  instead of  $0.89\lambda/D$  (Condon and Ransom, 2016). The new beam shape arising from the coherent addition of antenna signals is called a *coherent* or *secondary beam*, which can be steered within the single-dish *primary beam* by adjusting the delay  $\Delta t$ .

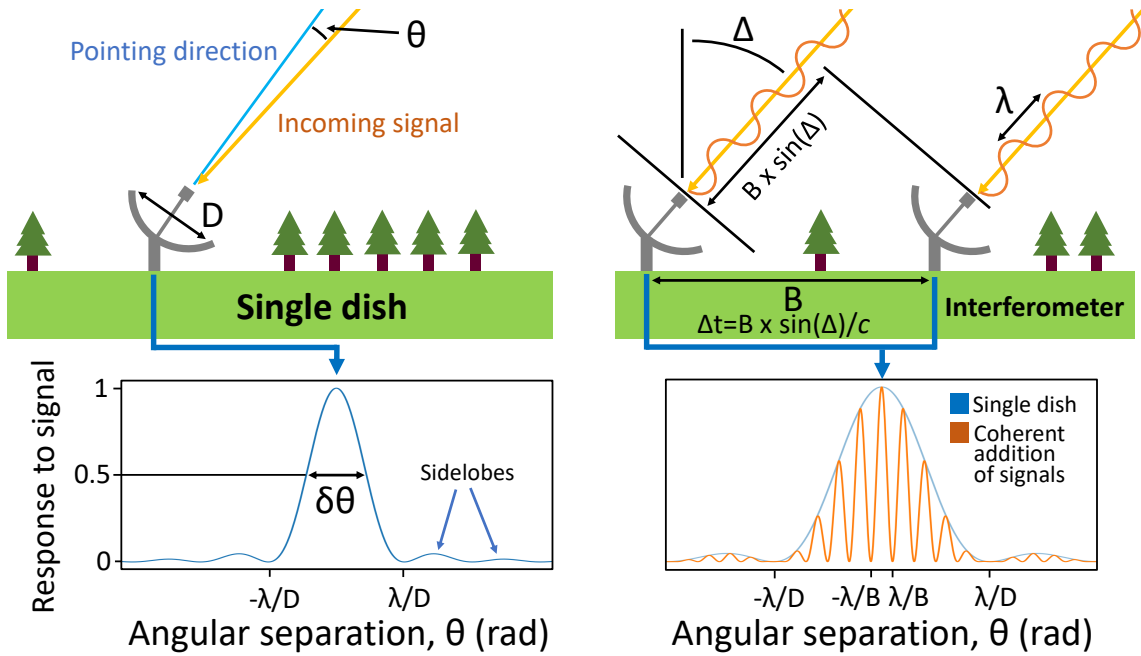


Figure 2.2: Angular resolution and signal response of single dishes and interferometers. **Left:** signal response of a single dish as a function of the pointing-to-source angular for a single dish, assuming that the beam coincides with the pointing direction. **Right:** signal response of a two-antenna interferometer as a function of the pointing-to-source angular for a single dish.

An array of more than two antennas is needed to achieve high-quality coherent beams. The coherent beam shown in the right side of Fig. 2.2 has powerful sidelobes which are limited in response only by the *primary beam* pattern, arising from the single dish beam. To solve this, we have to think once again that in reality beams are two-dimensional. A two-antenna interferometer, like the double-slit experiment, creates an interference pattern only in the direction of the baseline joining them, so the resolution is improved only in one dimension. To constrain the coherent beam in the other dimension, we need to add least a second, non-parallel baseline from a third antenna not placed along the same line as the other two. This provides constraints to the coherent beam in two dimensions and it reduces the

response of the sidelobes significantly. With the addition of further antennas with different baselines, constructive interference only occurs at the center of the beam ( $\theta \approx 0$ ) and the coherent sidelobes lose almost all of their power. At the same time, the resolution of the full array is determined by the longest baseline distance values between antennas (Condon and Ransom, 2016). In essence, with many antennas, we create an array-sized telescope without the need of building it!

For pulsar data, the technique of *beamforming* is used. It consists of assuming a sky position for a potential source within the primary beam, computing the time delay for each antenna, and then adding the voltages for the chosen sky position (van Veen and Buckley, 1988). This operation can be performed for as many possible positions as desired, the only limitation being the sensitivity drop of the primary beam (equation (2.3)) and the available computing and storage resources. The results is the formation of a digital coherent beam tiling within the primary beam, where each coherent beam acts as a pixel in an image. If a source is bright within the incoherent beam, only the coherent beams close to it will detect it. Thus, from a single observation, we can get hundreds of nested high-resolution coherent sub-observations in one go.

### 2.2.3 The telescopes of this thesis: Parkes/Murriyang and MeerKAT

In this thesis, I have worked with are the  $D = 64$  m single dish *Murriyang telescope*<sup>40</sup> at the *Parkes Observatory* in Australia, and the 64,  $D = 13.5$  m antennas of the *MeerKAT interferometer*<sup>41</sup> in the Karoo desert in South Africa (see Fig. 2.3 for pretty pictures). Both telescopes are located in the Southern hemisphere, giving them a privileged view of the Southern Galactic plane, including the Galactic centre, and reaching parts of the sky not covered by other sensitive facilities in the north.

The Parkes/Murriyang is run by the Australia Telescope National Facility (ATNF) telescope, and it has had a rich history. For instance, this very telescope *enabled the retransmission of TV signals from the Apollo 11 Moon landing*. It has been at the center of numerous pulsar projects (e.g. Manchester et al., 2001; Keith et al., 2010; PPTA Collaboration, 2023), and as stated in the very first Section of Chapter 1, at some point it had discovered more than half of all known pulsars. In my case, I have worked with data from the 21-cm *multibeam* receiver (1997–2018, Staveley-Smith et al., 1996) and the ultra wide-bandwidth low-frequency receiver (UWL, 2019–today, Hobbs et al., 2020). The former multibeam receiver was centred at  $f = 1382$  MHz, with bandwidth  $\Delta f = 400$  MHz. Its name comes from the fact that it was actually composed of 13 feed horns, producing a 13-beam tiling pattern on the sky. Nonetheless, in this thesis I have only worked on data from the central beam. The UWL receiver was centered at  $f = 2368$  MHz, with  $\Delta f = 3328$  MHz. It has  $G_T \approx 0.9 \text{ K Jy}^{-1}$  and  $S_{\text{sys}} \approx 21 \text{ K}$ , with a sky resolution resolution of  $\delta\theta = 20$  arcmin at  $\lambda = 21 \text{ cm}$ .

The MeerKAT interferometer is run by the South African Radio Astronomy Observatory (SARAO). It is a young radio telescope, with first light in 2019<sup>42</sup>, and it was built as the precursor of the square-kilometre array at middle frequencies (SKA-mid). It consists of 64 individual antennas spread through a region 10 km across (see Fig. 2.3). With a combined gain of  $G_T \approx 2.6 \text{ K Jy}^{-1}$  and a system temperature of  $S_{\text{sys}} = 20 - 25 \text{ K}$  (depending on the observing frequency and the pointing elevation,

<sup>40</sup> *Murriyang* is the Wiradjuri word for “the sky world”, home of *Biyaami*, the creator god in Australian Aboriginal mythology.

<sup>41</sup> The name *MeerKAT* is word play. KAT refers to the Karoo Array Telescope, a precursor of MeerKAT, while *meer* is the Afrikaaner word for “more”. At the same time, *meerkats* are a *species of animals* from the Karoo desert, famous for standing up and looking into the distance in groups, very much like MeerKAT’s antennas.

<sup>42</sup> That is just one year before the start of my PhD!



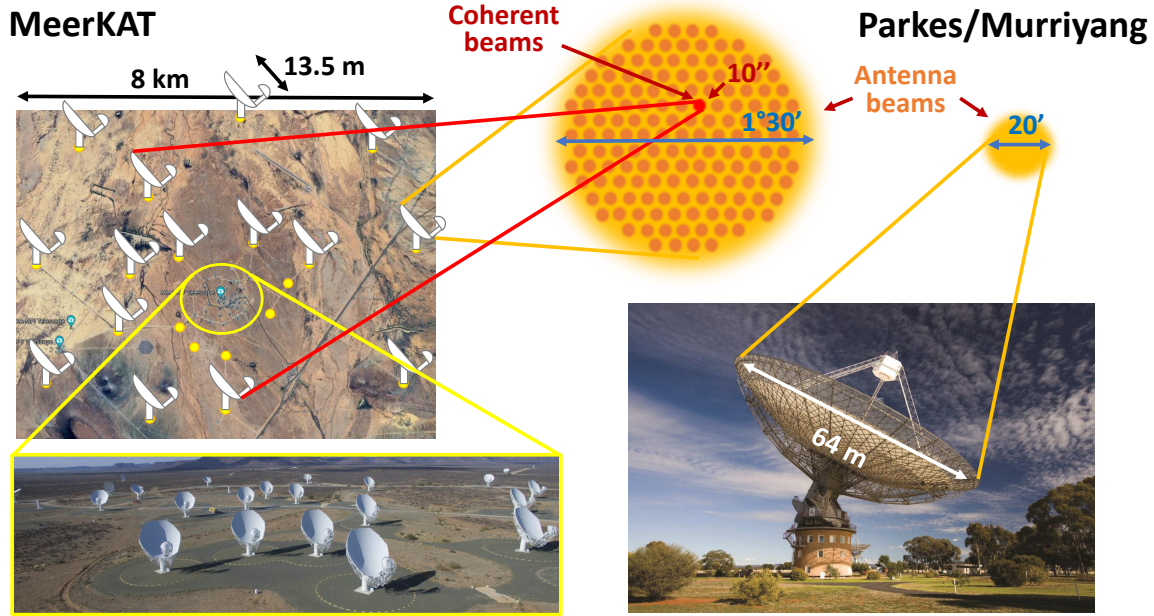


Figure 2.3: **Left:** the 64 antennas composing the MeerKAT interferometer in the Karoo desert, South Africa. The map is a screenshot from Google Maps, and yellow circles indicate the position of the antennas on the map. **Right:** the single-dish Murriyang telescope at the Parkes Observatory, Australia.

see the [MeerKAT General documentation](#)), it is the most sensitive facility in the Southern Hemisphere at decimeter wavelengths, and in particular 15 times more sensitive than Parkes/Murriyang. Because of that, it is an excellent Southern pulsar facility (Bailes et al., 2020), which has already yielded hundreds of new pulsar discoveries (e.g. Ridolfi et al., 2022; Padmanabh et al., 2023; Clark et al., 2023; Turner et al., 2024; Carli et al., 2024) and placed new constraints on highly relativistic systems (e.g. Hu et al., 2022). Currently, three receiver sets are available on each antenna: the UHF receivers ( $f = 580 - 1015$  MHz, Bailes et al., 2020), the L-band receivers ( $f = 900 - 1670$  MHz, Bailes et al., 2020) and the S-band receivers ( $f = 1750 - 3500$  MHz, Barr, 2018). With  $D = 13.5$  m and at  $\lambda = 21$  cm, it has a primary beam resolution of  $\delta\theta = 1.5$  deg. However, the filter banking beamformer user-supplied-equipment (FBUSE Barr, 2018) allows for the creation of hundreds of coherent beams within the primary beam via the technique of beamforming (Chen et al., 2021b). With the longest baseline between antennas being  $B = 8$  km, its resolution can go as small as  $\delta\theta = 10$  arcsec.

## 2.3 Pulsar searching

To have fun with pulsars, we first have to discover them! As described in Chapter 1, the first pulsar ever discovered was **PSR B1919+21**, found by Jocelyn Bell in 1968 (Hewish et al., 1968). While looking for scintillating quasars at the Mullard Radio Astronomy Observatory, she pointed the telescope at a bright radio source and was surprised to see a train of pulses drawn by the pen-chart on a sheet of paper, appearing every 1.337 s. Back in the day, “backends” worked by recording voltage powers via inc on a paper sheet with pen-charts which were then inspected by eye, just like with seismometers.

Today, processing of radio data is somewhat more sophisticated. While the brightest pulsars can still be detected by eye in a time series of voltage powers, almost all new discoveries are hidden below a layer of instrumental noise. How do we discover a new pulsar then? Do we just squint our eyes very hard? Not really, but that is what I will answer in this section.

We are on the look-out for new pulsars whose spin, DM, and orbital parameters are unknown. A modern pulsar search pipeline contains the following chain elements (Lorimer and Kramer, 2005): **observe** a point in the sky for a time range  $T$ , **store** the filterbank-format data in a cluster with available processing capacity, **de-disperse** the data at trial DM values, **de-orbit** the data at trial  $a$  values, perform a **periodicity search**, **identify signals** from this search, and perform a **visual inspection** of the frequency–phase and time–phase behaviour of these signals to decide whether they are indeed pulsars or something else. As you will see, pulsar search relies very heavily on the periodic nature of the sought-for signals. However, orbital motion may change the periodicity of the pulsar during the observation, forcing us to account for it during the search. In addition, there is also the need to distinguish between genuine pulsar signals from man-made radio-frequency interference (RFI) signals. All of these elements will be treated in this section, starting with the data format of filterbanks, and leaving RFI for the end, as it can be dealt with all across the previous steps. Let us go and find pulsars in our filterbank-mode data!

### 2.3.1 Filterbanks: your frequency- and time-resolved friends.

Search-mode pulsar data is stored in *filterbanks* files, or Flexible Image Transport System (FITS)-mode<sup>43</sup> files that can be nonetheless converted into filterbanks. Filterbanks are two-dimensional arrays with discrete time values,  $t$ , and radio frequency values,  $f$ , of length  $T$  and bandwidth  $\Delta f$ . The frequency axis is divided in *frequency channels* of bandwidth  $\delta f$ , while the time axis is divided in time samples of length,  $\delta t_s$ . During searching, we usually account for the filterbank array storing is the total intensity, the  $I$  term of the Poincaré vector equation (2.1), in order to boost the power of any astrophysical signal in the array. Nonetheless, the other terms of the Poincaré vector may be stored in the original voltages array, or in a FITS-mode file for later use. Individual filterbank files can be as large as several GB, with typical  $\delta t_s$  values of a few  $\mu s$  and time–frequency sample containing 4, 8, 16, or 32 bits depending on the dynamic range of the data. Nonetheless, such size is necessary to look for periodic signals with a large range of possible spin and DM values.

### 2.3.2 De-dispersion: make many time series

A pulsar signal, unless scintillation is present, is broad-band: it appears in every frequency channel. To search for a pulsar, we must add its signal across frequency channels to boost the  $S/N$  (Lorimer and Kramer, 2005). However, integrating across  $f$  is not as trivial as just adding the time series from each frequency channel. As explained in Chapter 1, the pulsar signal is *dispersed*, with a delay in the arrival of the pulses in the telescope between different  $f$  values. As depicted in Fig. 2.4, to correct the DM-induced delay between two frequencies expressed in equation (1.9), we shift the time axis  $t$  at each frequency channel by

$$\Delta t = 4.15 \times 10^6 \text{ ms} \times \text{DM}(f_{\text{ref}}^2 - f^2), \quad (2.5)$$

---

<sup>43</sup> FITS files are standard in astronomy, and have been endorsed by multiple institutions such as NASA or the International Astronomical Union. Designed with images in mind, they store data arrays and column headers in optimised ways.

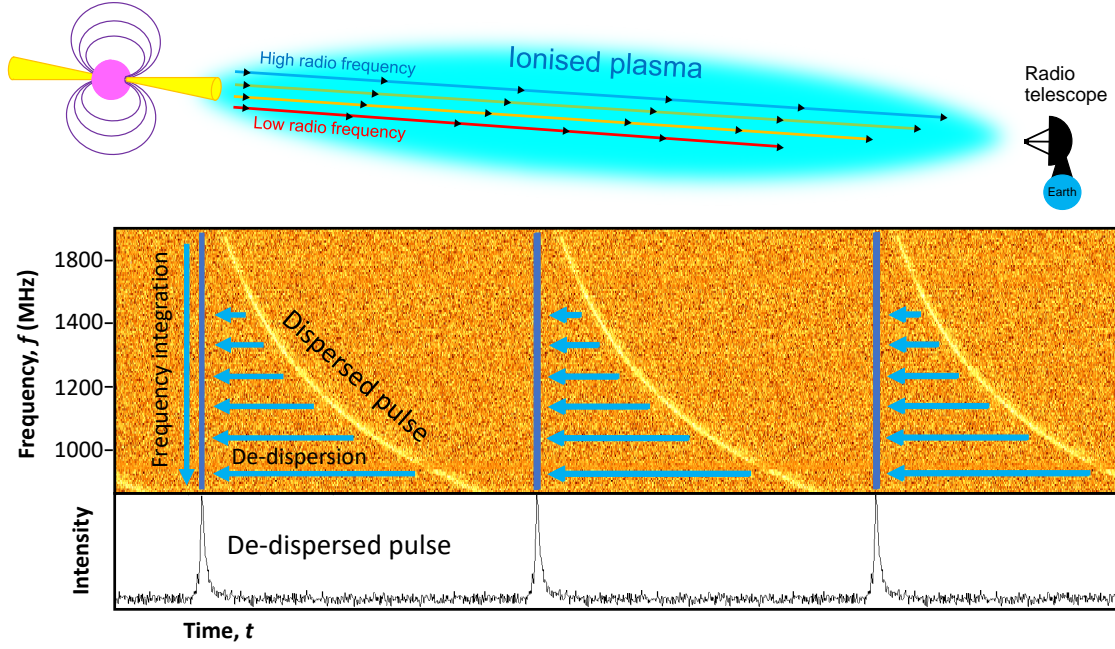


Figure 2.4: Incoherent de-dispersion of a filterbank. If the de-dispersion is performed at the correct DM, the signal is successfully recovered. Used data: MeerKAT observations of **PSR J1452–5549**, a MMGPS discovery.

where  $f_{\text{ref}}$  is a chosen reference value, before adding the time series across the frequency channels (Lorimer and Kramer, 2005). After addition, if the used DM values is correct, this results in a single time series where it is easier to recover the pulsar signal. This operation is known as *incoherent de-dispersion*.

However, there is one **big tiny** catch: we are looking for new pulsars, so we do not know their DM value! But hope is never lost. We only have to de-disperse the filterbank at many DM trial values, and search for pulsars in every resulting time series. The key to that is to come up with an efficient de-dispersion plan. The first thing we decide on is the covered *DM range*. Obviously,  $\text{DM} > 0 \text{ cm}^{-3} \text{ pc}$ , as dispersion can only be positive, but the maximum DM value is a choice. The largest known DM value is  $\text{DM} = 1778 \text{ cm}^{-3} \text{ pc}$  in **PSR J1745–2900**, located close to the Galactic centre (Eatough et al., 2013), but more distant pulsars could have even larger DM values. Nonetheless, Galactic electron density models (e.g. Cordes, 2004; Yao et al., 2017) predict that pulsars are unlikely to be found at  $\text{DM} > 3000 \text{ cm}^{-3} \text{ pc}$  without residing outside our Galaxy, so  $3000 \text{ cm}^{-3} \text{ pc}$  seems like a good upper limit.

Next step is to decide the DM step,  $\Delta\text{DM}$ , between each de-dispersion trial. This depends on the DM smearing between trials we are willing to tolerate. Following equation (1.9), we want to introduce incremental  $\Delta t_{\Delta\text{DM}}$  corrections across the band smaller than the smallest pulse width  $W_{\text{PSR}}$  we are searching for, as otherwise the pulsar signal may fall in between DM steps. For instance, if we want to find pulsars with  $W_{\text{PSR}} = 1 \text{ ms}$  in a  $f = 1000 - 1800 \text{ MHz}$  band, then we should implement  $\Delta\text{DM} < 0.35 \text{ cm}^{-3} \text{ pc}$ . In theory, the best choice is equating  $\Delta t_{\Delta\text{DM}}$  to the filter bank time resolution

( $\delta t_s$ ), so that we can search for the narrowest pulsed signals. However, we have to keep in mind that at

$$\text{DM} > 1.205 \times 10^{-7} \text{ cm}^{-3} \text{ pc} \times N_{\text{ch}} \delta t_s \times \left( \frac{1}{f_1^2} - \frac{1}{f_2^2} \right)^{-1}, \quad (2.6)$$

where  $N_{\text{ch}}$  is the number of frequency channels and  $\delta t_s$  is expressed in ms, the DM smearing within each frequency channel starts to become larger than  $\delta t_s$ , so keeping a high time resolution and a small  $\Delta \text{DM}$  becomes inefficient. After this threshold, reducing the time resolution of the filterbank and the DM steps improve the de-dispersion performance. Therefore, the optimal approach is an adaptative  $\Delta \text{DM}$  step that increases as we increase DM of our search (Lorimer and Kramer, 2005).

### 2.3.3 Single pulse searches: bright pulsars want to be found

If we have de-dispersed a filterbank with the right DM value, some pulsars are bright enough that individual pulses will appear already in the time series. We define the mean flux density of a pulsar,  $F_m$ , as the flux density received at the telescope average over  $P_s$ . Given a pulse with  $W$ , the  $S/N$  of the pulsar after de-dispersing over a bandwidth  $\Delta f$  is

$$S/N \approx F_m \times \frac{n_p G_T}{S_{\text{sys}}} \times \sqrt{\frac{\Delta f}{\delta f}} \times \sqrt{\frac{P_s - W}{W}}, \quad (2.7)$$

where the  $\sqrt{\Delta f}$  factor comes from the fact that, while during the addition of frequency channels both the noise signal and the astrophysical signal add on each other, the noise signal of Gaussian nature adds incoherently, thus increasing with  $\propto \sqrt{\Delta f}$ , while the pulsar signal adds coherently, thus increasing with  $\propto \Delta f$ . As a result,  $S/N \propto \Delta f / \sqrt{\Delta f} \propto \sqrt{\Delta f}$ . If the  $S/N \gg 1$ , then individual pulses become detectable.

### 2.3.4 Periodicity searches: time domain, or Fourier domain?

Pulsar, pulsar, are you there? Say that we have the same time series as before, but we only manage to see noise. Can we be sure there is no pulsar in it? No, it can also be that for individual pulses have  $S/N < 1$ , so we have to perform a periodicity search. There are two main algorithms that allow for it: the *fast folding algorithm* (FFA), and the *fast Fourier transform* (FFT). These are the backbone of modern pulsar searching, so let us take a good look at them!

#### The Fast Folding Algorithm

The Fast folding Algorithm relies in the operation of *folding* to coherently add individual pulses to produce a high  $S/N$  pulse profile. Folding reduces the time axis  $t$  to a single rotational *spin phase*  $0 \leq \phi_s < 1$ , defined as

$$\phi_s = \frac{t}{P_s} - \text{trunc} \left( \frac{t}{P_s} \right), \quad (2.8)$$

where the operation *trunc* stands for truncation to integer unit (assuming  $t = 0$  at the start). This operation introduces *phase bins*, a fixed number of discrete  $\phi_s$  values whose power is the addition of the power values at the nearest  $t$  discrete points. A typical folding operation results in  $N_{\text{bin}} = 1024$

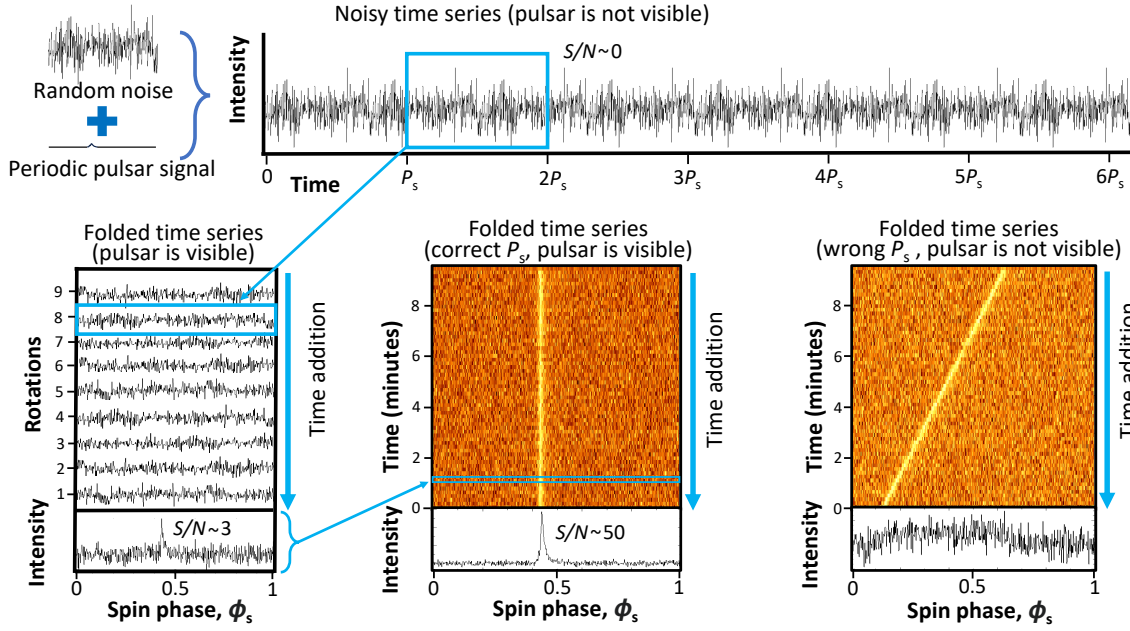


Figure 2.5: Folding of a pulsar signal. Individual pulses are much fainter than the instrumental noise, making them invisible in the time series. However, after adding several pulsar rotations, the pulsar signal starts to peak above the noise. Folding the entire observation boosts the  $S/N$  significantly, yielding a clear detection. If the folding period is incorrect, the pulsar signal does not add coherently across the observation, leading to no detection. Used data: MeerKAT observations of **PSR J1452–5549**, a MMGPS discovery.

bins of size  $\delta\phi_s = P_s/N_{\text{bin}}$ . After its implementation, the addition of several pulses may peak above the noise, with

$$S/N \approx F_m \times \frac{n_p G_T}{S_{\text{sys}}} \times \sqrt{\frac{\Delta f}{\delta f}} \times \sqrt{\frac{T}{P_s}} \times \sqrt{\frac{\delta\phi_s}{\delta t_s}} \times \sqrt{\frac{P_s - W}{W}}, \quad (2.9)$$

where  $T$  is the integration time. This works, once again, because the signal noise adds incoherently while the pulsar signal adds coherently. Thus, as shown in Fig. 2.5, where we see a depiction of the process of folding of a de-dispersed time series, after enough additions of the pulse on itself, its  $S/N$  increases significantly and a detection is made.

However, like in de-dispersion, we do not know the correct  $P_s$  value during the search. As shown in the right side of Fig. 2.5, folding the pulsar with the wrong spin period leads to a non-detection. That is where the FFA comes in: it folds the time series at many trial  $P_{\text{trial}}$  values and looks for the one that yields the highest  $S/N$ . Developed by astronomer David Staelin in 1969, this algorithm consists on taking a time series of  $N_{\text{sample}}$  samples and a first trial folding period  $P_{\text{trial},i}$  (in units of samples), and folding all the periods between  $P_{\text{trial},i}$  and  $P_{\text{trial},i+1}$  with a step

$$\Delta P = \frac{P_{\text{trial},i}}{N_{\text{samp}} - P_{\text{trial},i}} \quad (2.10)$$

in a computationally efficient way, while keeping a constant  $N_{\text{bin}}$ . Subsequently, for  $P_{\text{trial},i+1}$ , the  $\Delta P$



is increased and the process is repeated, with the FFA becoming more efficient the larger are the searched periods (Staelin, 1969; Cameron et al., 2017).

### The Fast Fourier Transform

Introducing the second-worst nightmare<sup>44</sup> of undergrad physics students, and also the most widely used periodicity search tool in pulsar astronomy (Lorimer and Kramer, 2005): the Fourier transform. A Fourier transform is an operation that takes a function and performs a spectral decomposition. If there is a repeating signal in a time series, its frequency will be highlighted by the Fourier transform. In pulsar astronomy, we implement the Fast Fourier Transform (FFT), a computationally-efficient version of the discrete Fourier transform, first described by the historical mathematician Carl Friedrich Gauss, and then re-discovered by modern mathematicians James Cooley and John Tukey (Cooley and Tukey, 1965). For a real-valued time series  $f(t)$  of length  $T$  with  $N_{\text{samp}}$  discrete time stamps  $t_n = n\delta t$ , where  $n$  runs from 0 to  $N_{\text{samp}} - 1$ , the Fourier decomposition is expressed as (Bracewell, 2000)

$$f_n = \sum_{k=0}^{N_{\text{samp}}-1} a_k e^{i\theta_k} e^{-i2\pi nk/N_{\text{samp}}}, \quad (2.11)$$

where  $i = \sqrt{-1}$  is the imaginary unit,  $a_k$  are real-valued Fourier coefficients, and  $\theta_k$  are Fourier phases. In essence, the equation above re-describes the time series as an addition of many sine functions of amplitude  $a_k$ , frequency  $\nu_k = k/T$ , and starting phase  $\theta_k$ <sup>45</sup>, where  $k$  runs from 0 to  $N_{\text{samp}} - 1$ . The frequencies  $\nu_k$  are called the Fourier frequencies of time-oscillating signals in the data, and are not to be confused with the radio frequency  $f$ . The frequency resolution of the Fourier domain is equal to  $\delta\nu = 1/T$ , while the largest covered frequency sample is  $1/\delta t$ .

The FFT is a very powerful tool for describing periodic pulsar signals. From only looking at the amplitudes  $a_k$  or, more commonly, their power  $|a_k|^2$ , we can already extract a great deal of information. Fig. 2.6 shows the  $|a_k|^2$  response of the Fourier transform for some periodic signals. For instance, the top row shows how a single sinusoidal function of period  $P_s = 200$  samples is transformed into a single spike of power at  $\nu = 0.005 \text{ samples}^{-1}$  ( $\nu = 1/P_s$ ).

However, the FFT is not a magical solution. Pulsars are not sinusoidal sources, but pulsating sources with an arbitrary pulse shape, and anything that is not a perfect sinusoidal oscillator will spread its power across higher harmonics. This is intrinsically tied to the pulse width: as shown in the second and third rows of Fig. 2.6, the narrower is a pulse profile, the more it spreads its power across high-frequency harmonics. An additional issue is that  $1/P_s$  does not correspond to any of the discrete  $\nu_k$  bins, the power is split across neighbouring bins, a phenomenon called *scalloping* (Ransom et al., 2002). Finally, real time series comes with noise, which is also picked up by the Fourier transform. *White noise* is random instrumental noise of Gaussian nature, while red noise arises from subtle long-term variations of the time series baseline. As shown in the fourth row of Fig. 2.6, in the Fourier domain white noise becomes white noise of equal power across the Fourier series, while red noise manifests as increasing noise power at low  $\nu$  values<sup>46</sup> (Lorimer and Kramer,

<sup>44</sup> I always found the Laplace transform more scary.

<sup>45</sup> Euler's formula states that  $e^{ix} = \cos x + i \sin x$ , so a time-varying complex exponential represents an oscillating function in the real-imaginary complex plane.

<sup>46</sup> In fact, their names are based on the colors detected by our eyes: red light has more power at low frequencies, while white

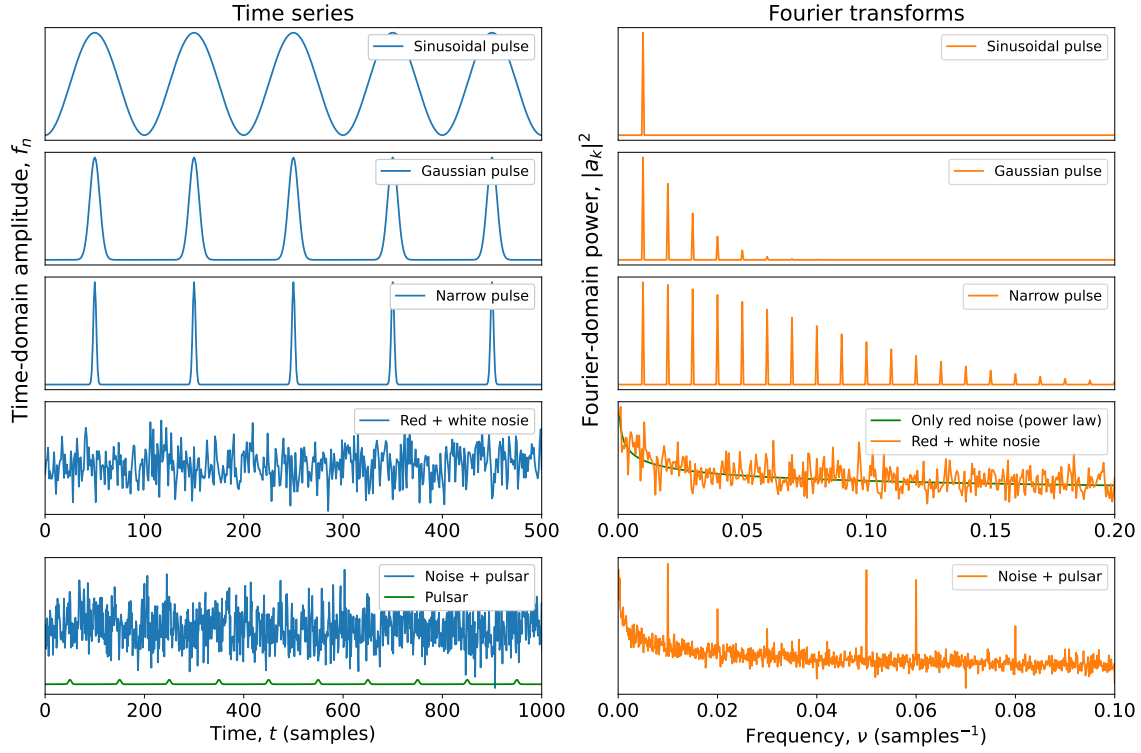


Figure 2.6: Fourier transform of four different time series. From top to bottom: a pure periodic, sinusoidal series with  $P_s = 100$  samples, two periodic pulse trains of Gaussian shapes with the same periodicity but differing pulse widths, a time series with only red and white noise, and a periodic signal with noise. I should note that the time series are much longer than shown here. In real data, a pulse can repeat hundreds or thousands of times across the time series. Also, I have used a power law in the amplitudes of red noise to generate it.

2005). As is shown in the last row of Fig. 2.6, both white and red noise can affect the power of a periodic signal in the Fourier domain by cancelling part of it, or even hiding it amongst the noise.

However, all of these problems have solutions. This power spread across harmonics can be recovered with the technique of *incoherent harmonic summing*, which consists of adding the power of Fourier bins that are multiples of each other (e.g.  $|a_k|^2 + |a_{2k}|^2 + |a_{3k}|^2 + |a_{4k}|^2$ , Taylor and Huguenin, 1969). Like in the domain, noise adds incoherently while the pulsar signal adds coherently, resulting in a boost of the  $S/N$ . Studies have shown that for pulsars with pulse duty  $\delta = W/P_s = 0.1$ , four harmonic summing operations already results in an optimal recovery of Fourier power, while for  $\delta < 0.015$ , 32 harmonic summing operations are needed (Ransom et al., 2002). Scalloping can also be corrected for by correlating the  $a_k e^{i\theta_k}$  values of neighbouring frequency beams resulting in an interpolated Fourier power value in between bins (Ransom et al., 2002), which can also be used to recover the original periodicity. And finally, red noise can be removed by running a median filter, an operation known as *de-reddening*. For instance, a running median filter with a 10-second-long window can be applied to the time series, removing red noise at  $\nu < 0.1$  Hz without erasing potential pulsar signals with pulse widths shorter than the window. On the other hand, de-reddening can be performed directly on the

---

light has equal power at all frequencies.

Fourier domain via the same operation. However, it has been shown that de-reddening in the Fourier domain can actually be more harmful to pulsar signals than in the time domain, as it removes power regardless of its nature, potentially hampering the low-frequency harmonics spikes of a pulsar signal (Singh et al., 2022).

### 2.3.5 Acceleration searches: chase the binary pulsars!

The observed spin period of a pulsar is modulated by the Doppler factor caused by the relative motion between the pulsar and the observer. Isolated pulsars have a constant, non-measurable Doppler factor due to the relative motion between them and the Solar System, but if the pulsar is in a binary, that is no longer the case and it can spell disaster for a periodicity search. Fig. 2.7 shows what it looks like when a pulsar is accelerated along the LOS: the pulsar signal is no longer periodic, so both the Fourier transform and folding with a constant  $P_s$  results in a significant loss of power (Johnston and Kulkarni, 1991). However, pulsar searchers also have come up with solutions to correct this problem and recover the periodicity of binary pulsars.

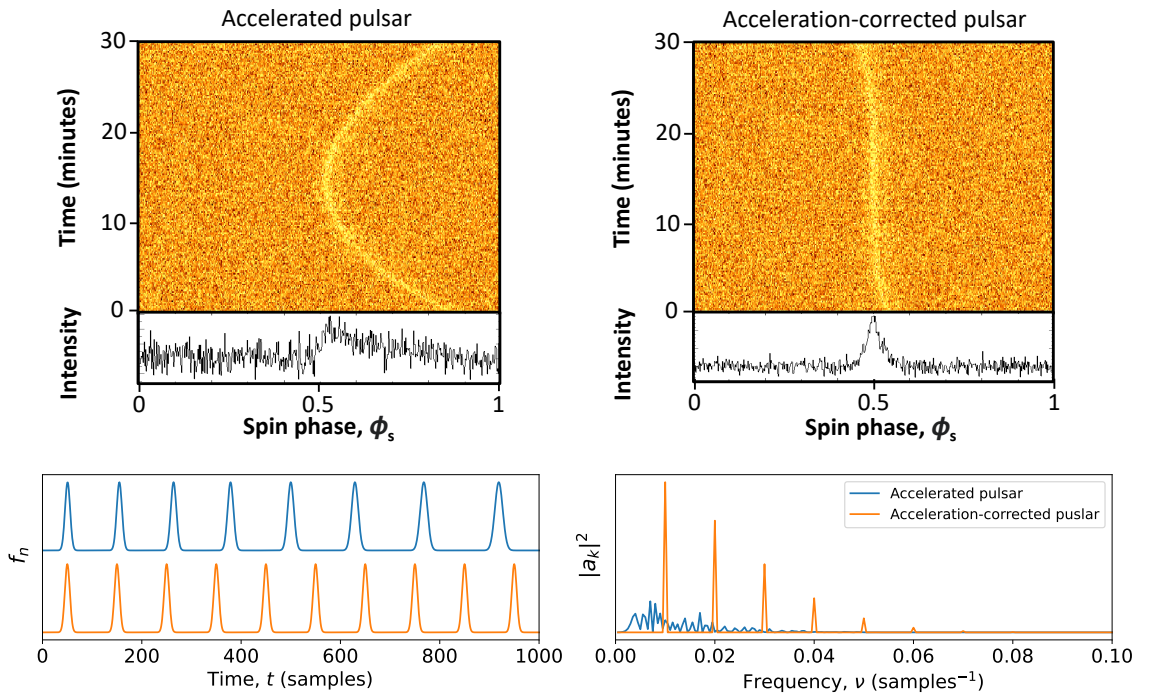


Figure 2.7: Effect of acceleration on pulsar periodicity. The top plots show the effect upon folding, where the change in the spin period leads to a non-detection. The bottom plots show the consequences of the same effect, but in the Fourier transform. In both cases, the signal is properly recovered after a time-domain re-sampling. The top row of plots have been drawn from MMGPS observations with the `pdmp` software.

An *acceleration search* assumes a constant acceleration along the LOS,  $a$ , during the observation, so that the Doppler modulation of the observed spin period,  $P_{\text{obs}}$ , can be described with a simple



linear evolution over time:

$$P_{\text{obs}} \approx P_s \left\{ 1 + \frac{v_0}{c} + \frac{a}{c}t + O(t^2) \right\} \approx P_0 + P_0 \frac{a}{c}t = P_0 + P_1 t, \quad (2.12)$$

where  $v_0$  is the velocity value at  $t = 0$ ,  $P_0$  is  $P_{\text{obs}}$  at  $t = 0$ , and  $P_1$  is a constant (apparent) spin derivative factor. This approximation is not true across the entire orbit, but it works well if  $T < 0.1P_b$  (Johnston and Kulkarni, 1991). As shown in Fig. 2.7, when folding or performing the Fourier transform after correcting for the pulsar motion with a constant acceleration, much of the pulsar signal is recovered with minimal losses.

The operation of correcting for acceleration is performed via time-domain re-sampling, which transforms the time axis from  $t$  to  $t'$  via

$$t' = t - \frac{P_1}{P_0}t^2 = t - \frac{a}{c}t^2, \quad (2.13)$$

stretching or contracting the time series. If the assumed  $a$  value is correct, the pulsar signal becomes periodic again. In essence, it is equivalent to placing the observer in the same frame of reference as the pulsar. However, since the  $a$  is not known during the search, we find ourselves with the same problem as de-dispersion: we need to perform many acceleration trials based on our data and the science objectives, adding a nested loop of  $a$  trials within each DM trial before performing the FFA or the FFT. The acceleration step,  $\Delta a$ , is decided on the observation length  $T$  and the searched pulse width  $W$ . The cumulative pulse drift given an  $\Delta a$  after a time  $T$  is

$$\Delta t_{\text{acc}} \approx \frac{\Delta a}{c}T^2, \quad (2.14)$$

and we seek to give enough  $\Delta a$  resolution so that  $\Delta t_{\text{acc}} < W$ . For instance, if we want sensitivity to pulsars with  $W = 1$  ms in  $T = 30$  min observations, we need  $\Delta a < 0.1 \text{ m s}^{-2}$ . The searched  $a$  range also depends on the systems we want to be sensitive to. From Keplerian mechanics and assuming a circular orbit and  $M_p = 1.35 M_\odot$ , the maximum acceleration for a pulsar given an orbital period  $P_b$  and companion with mass  $M_c$  is

$$|a_{\text{max}}| = c T_\odot^{1/3} \left( \frac{2\pi}{P_b} \right)^{4/3} \frac{M_c}{(1.35 + M_c)^{2/3}}, \quad (2.15)$$

which means that, for instance, if we search for PSR–black hole binaries with  $M_c = 10 M_\odot$  and  $P_b = 1$  hour, we must loop through  $-2000 < a/\text{m s}^{-2} < 2000$ .

In general, given an  $a$  range, the amount of re-sampling trials increases as  $\propto T^2$  due to the decreasing  $\Delta a$ . And given that the time series length itself increases in length, then amount of operations per trial increase with  $\propto T^3$ . At the same time, with longer observations, it is less likely that  $T < 0.1P_b$  holds true. Therefore, longer observations are not always beneficial when looking for binary pulsars, but if that is the case a cheaper alternative to time-domain re-sampling does exist: searching the spread of the signal across the Fourier series and fitting for  $P_0$  and  $P_1$  at the same time, a method called *correlation technique* (Ransom et al., 2002). While the correlation technique is more involved to implement, in long time series it is computationally more efficient than the time-domain re-sampling, and it does not require the production of a new time series for each trial, resulting in a  $P_0$  and  $P_1$

directly on the time series (Ransom et al., 2002). However, Fourier-domain can only applied to FFT searches, while FFA searches that account for acceleration have no choice but to implement the time-domain re-sampling (e.g. Wongphechauxsorn et al., 2024).

### 2.3.6 Identification and inspection of candidates

Pulsars are not the only signals that may be present in a FFA or a FFT. Sometimes, other signals masquerade as pulsars (see Section 2.3.7 right below). Sometimes, random noise piles up and is picked up by the FFT or the FFA. That is because these algorithms are agnostic to the nature of the signal. SO, how do we know what is a pulsar and what is not a pulsar? Here we look at the last steps of a pulsar search, which deal exactly with this question.

#### Noise or signal?

To identify signals in the FFT or the FFA, we have to make sure they are not noise. Random noise in the FFT can be described as Gaussian, with two degrees of freedom in each frequency bin ( $a_k$  and  $\theta_k$ ). If we combine  $a_k e^{i\theta_k}$  into a single complex value, the power of the FFT is the addition of two variables (square of the real and imaginary parts) that follow Gaussian noise, resulting in a  $\chi^2$  probability distribution of degree two for the Fourier power. If we normalise the noise variance to unity, the probability of a Fourier beam having a power  $P$  above a certain value  $P_{\min}$  is the *false-alarm probability*

$$p(P > P_{\min}) = e^{-P_{\min}}. \quad (2.16)$$

Given a time series of  $N_{\text{samp}}$ , we set  $P_{\min} = \ln N_{\text{samp}}/2$  on the basis that there should be less than one bin on which this threshold is reached by the random pile-up of noise. Therefore, signals with power  $P > P_{\min}$  are more likely to be real, becoming more significant with increasing power. This threshold is used to automatise the finding of signals in the FFT, along with peak-finding algorithms that identify the optimal frequency value of signals (Lorimer and Kramer, 2005). It should be noted that with harmonic summing, the same principle applies but with  $\chi^2$  distributions of higher degree since we are adding powers on top of powers.

For the FFA, usually matched filters of different width are run across  $\phi_s$  at every trial fold, seeking to identify a pulse in the data and quantify the  $S/N$  on every fold. Then, high  $S/N$  folds are selected as candidates (see e.g. Cameron et al., 2017).

#### Local or astronomical?

Signals detected in a periodicity search may be of astrophysical origin or of local origin (e.g. man-made, see Section 2.3.7). If our observations place multiple beams on the sky, be it from different feed horns in a single antenna or from interferometry, an astrophysical signal will be detected only in the beams closest to the sky position, with a  $S/N$  distribution across beams peaking close to the true position. On the other hand, local signals leak into all of the feed horns of an antenna, or into all of the coherent beams of an interferometer by introducing a very powerful signal into several antennas. This property can be used to distinguish pulsars from other signals, and they have been used with a large degree of success (e.g. Ng et al., 2015; Padmanabh et al., 2023).

### Fold and optimise!

If a significant signal with a DM and  $P_s$  value (and perhaps a  $a$  value) has been found in the periodicity search, and it has a reasonable distribution amongst the beams, then we should take the most significant detection and fold it, but not fully. The folded signal is to be stored in a three-dimensional array of  $f - t - \phi_s$  values with limited resolution, with folding only every few seconds, and with de-dispersion across groups of neighbouring  $f$  channels (see example Fig. 2.8). The limited resolution which allows for quick corrections of DM,  $P_s$  and  $a$  and then a subsequent complete folding and de-dispersion for a quick computation of  $S/N$  values. This is used to perform a sliding search of the optimal DM –  $P_s$  –  $a$  configuration that yields the best  $S/N$ , as shown in Fig. 2.8. Once the optimal values are found, these are reported to the astronomer.

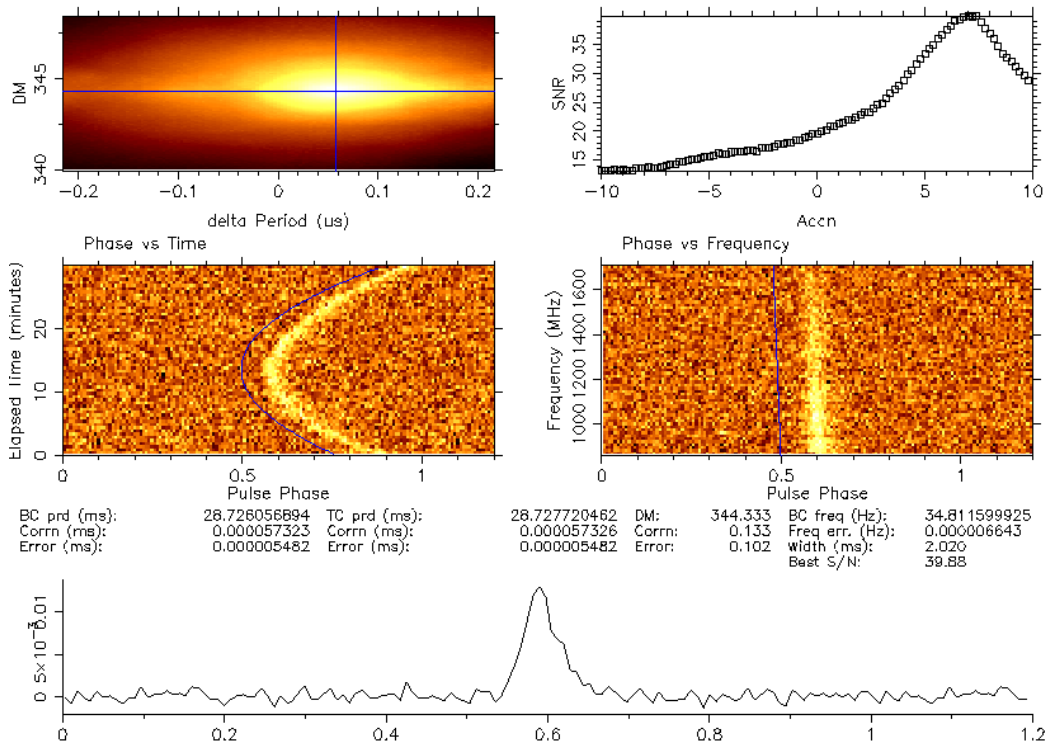


Figure 2.8: Example of a pulsar candidate with optimised DM,  $P_s$  and  $a$  parameters. The top-left plot shows the two-dimensional  $S/N$  distribution across the DM –  $P_s$  space, while the right-side plot shows the same across  $a$  trials. The middle-left plot shows the time sub-integrations as a function of  $\phi_s$ , while the middle-right plot shows a fully folded, but still frequency-resolved profile. The bottom plot shows the pulse profile, fully-folded and de-dispersed at the optimal values. Plot drawn with the `pdmp` program from the `PSRCHIVE` software

### Inspect! Pulsar or not a pulsar?

Despite all of our efforts, not all that we fold and optimise are pulsars. Take a good look at the plots in Figs. 2.8 and Fig. 2.9, which all show pulsar candidates. Can you tell which are pulsars and which are not? Which properties would you say are shared amongst those that are real pulsars? If you have guessed that they are the ones with continuous, broad-band emission cross frequency and time, then you

have guessed right! However other tell-tale signals are a non-zero DM values, and clear optimisation in the  $DM - P_s - a$  space as shown in Fig. 2.8. If that is not the case, the signal we have picked up may be something else! And on that note...

### 2.3.7 Radio-frequency interference: the nightmare of radio astronomy

If pulsars were the only periodic signals appearing in radio telescopes, the story would probably end here. Sadly, this is not the case. Optical astronomers complain about light pollution and the appearance of [satellite trails in their images](#)<sup>47</sup>. Radio astronomers complain about [WiFi, Bluetooth, radars, satellite communications, and much more](#). Many of these emissions are also periodic in nature, leading to the detection of their signals in our data analysis. RFI is bothersome, to say the least. It can be very bright, often saturating the receivers. Saturation occurs when the feeds are receiving more power than their maximum transmission capacity, leading to the out-maxing of recorded power in certain bands and leakage onto others, often eclipsing pulsar signals. As shown in Fig. 2.9, RFI is even capable of producing very pulsar-like signals, prompting search algorithms to reporting false pulsar candidates.

Radio telescopes are strategically built in locations secluded from RFI interference. For instance, the MeerKAT telescope in South Africa is located in the Great Karoo desert, with the closest town being Carnarvon, located 90 km away. To visit that telescope, one has to put all electronic devices on flight mode<sup>48</sup> and drive into the site with RFI-proof diesel cars (gasoline cars ignite their fuel with sparks, which interfere with telescope observations). However, no matter what, telecommunications are everywhere and RFI always comes in. Sometimes it is also the fault of astronomers. In 2011, astrophysical-looking signals of suspected terrestrial origins were detected with Parkes/Murriyang (Burke-Spolaor et al., 2011). In 2015, they finally discovered the source: the opening of doors of microwave ovens at the observatory (Petroff et al., 2015). Not even warming up your lunch is safe!

So, despite all of these efforts, we have to deal with RFI, which only becomes more abundant over time (see footnote 47). Luckily, for pulsar-format data, we have a few tricks in our selves, many of which are already common practice in pulsar astronomy (see the pulsar handbook Lorimer and Kramer, 2005). During this section I have ignored this aspect of data analysis, but since we already know all that comes into play when searching for pulsars, it is much more straightforward to understand how to deal with it at different steps of the pipeline.

#### Zero-DM filtering

RFI is of local origin or, at the very most, of low Earth orbit origin. That means that periodic RFI signals have  $DM = 0 \text{ cm}^{-3} \text{ pc}$ , making their identification rather easy. One common strategy is to build a *zero-DM birdie list*: perform FFT analysis at  $DM = 0 \text{ cm}^{-3} \text{ pc}$ , identify powerful signals, and then ban them by zero-weighting the corresponding Fourier frequency bins when searching at higher DM values. This works for both broad-band signals appearing across all frequency channels (e.g.,

---

<sup>47</sup> My least sincere regards to owners of private corporations throwing out hundredths of highly-reflecting satellites into low-Earth orbit.

<sup>48</sup> Remember to warn your relatives about flight mode before spending the night on a radio observatory. My mother once thought that I had crashed the van on the way to the Effelsberg telescope in Germany because I stopped answering to her messages. In a panic, she told my friends back home I was probably in a hospital or worse, and attempted to organise an on-line "Miquel search" with them. It was dramatic to learn about it the next morning.

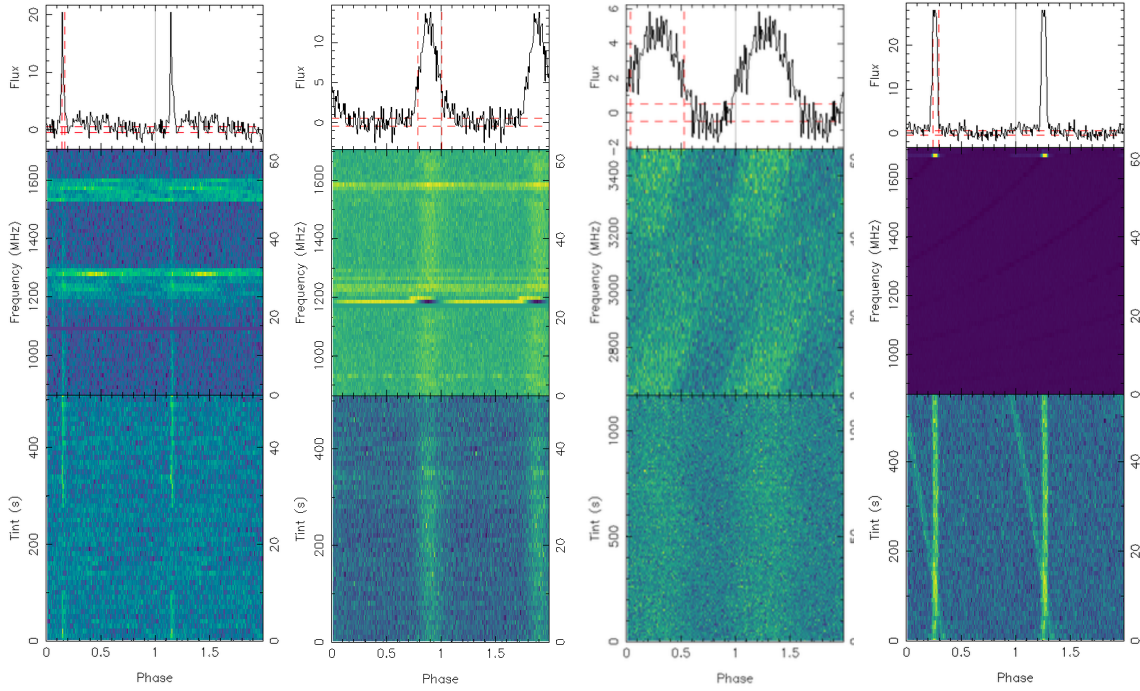


Figure 2.9: Examples of RFI affecting a few MeerKAT observations from the MMGPS. The top row shows the signal profiles, while the second and third rows show the  $f - \phi_s$  and  $t - \phi_s$  behaviour of several pulsar candidates. The signals are already de-dispersed. The first one is a legitimate pulsar detection despite RFI activity on some frequency bands. The second one has  $DM = 0$  and it is RFI activity leaking into other frequencies due to receiver saturation, resulting in a spurious pulsar candidate. In the third one, broad-band RFI creates another spurious candidate, in this case masquerading as a  $DM > 0 \text{ cm}^{-1} \text{ pc}$  candidate thanks to an (un)lucky jump in the emission phase at two different bands. In the fourth one, a satellite enters the field of view, shining in a very narrow band with dual periodicity. The data is from the MMGPS survey, and the plots were drawn with [PulsarX](#).

second and third plots in Fig. 2.9), and *narrow-band* signals appearing only in a few or even a single frequency channels (e.g. fourth plot in Fig. 2.9).

Another option is *zero-DM time-domain removal*: subtracting the frequency-averaged value at every time in the dispersed filterbank ( $DM = 0 \text{ cm}^{-3}$ ). This operation will remove any broad-band RFI signal that has equal power across the entire observed radio band. However, it will not remove narrow-band signals, or signals with changing power across the band, as the one seen on the second and fourth plots in Fig. 2.9.

### Masking unwanted signals

Sometimes RFI is persistent but only affects a limited set of frequency channels. This is seen in the first and second plots of Fig. 2.9, where the regions of  $1.2 < f < 1.3 \text{ GHz}$  and  $1.5 < f < 1.6 \text{ GHz}$  are heavily contaminated, and in the fourth plot of Fig. 2.9, where it shows up as an extremely narrow-band periodic signal. One solution is to zero-weight the affected frequency channels, at the expense of also losing any potential pulsar signal contained in them. However, in practice, the pulsar  $S/N$  may actually increase due to the reduced noise baseline. The same idea is applicable in the time domain:



some observations may be affected only during limited time ranges. In such case, time-domain masks are also a viable option.

### Smart filtering algorithms

The previous filtering methods have the advantage of being effective and somewhat simple to implement. However, they can also be a bit blunt, as they remove significant chunks of data and potentially sacrifice astrophysical signals. More sophisticated measure statistical properties of the entire data array (e.g. mean power, standard deviation of power, kurtosis), and zero-weight those samples that significantly deviate from the expected values. This method can be applied to search-mode filterbanks or fold-mode archives, and it is also the least biased one because it targets individual data points instead of time ranges or frequency ranges, often resulting in clean observations with a high degree of success (e.g. Men et al., 2023; Morello et al., 2019; Lazarus et al., 2020).

## 2.4 Pulsar timing

Imagine that after implementing a search pipeline, we have discovered many interesting pulsars in binary systems. Hooray! Let us celebrate again! But, how do we go about measuring this pulsar? The key is, once again, their periodicity. The rotation of a pulsar is like a ticking clock, and by measuring variations from the expected periodicity we can infer unmodelled processes happening in the pulsar system. For this, we must build a model of the pulsar’s rotation that predicts every single rotation, and then compare it against the measured *times of arrival* (ToAs) of pulses in our telescope. That technique is called *pulsar timing*. I can not understate the precision of pulsar timing compared to everything that came before in this chapter: for some pulsars, the ToA uncertainty can be as low as a few hundred ns ( $\sim 10^{-7}$  s, e.g. Desvignes et al., 2016). This means that we can measure a pulsar’s rotation rate down to less than one thousandth of the fraction of the spin phase!

The extremely powerful technique of pulsar timing has yielded many milestone tests of gravity and discoveries. For instance, the first indirect detection of gravitational-wave radiation in the Hulse–Taylor binary was made thanks to the continued timing of its pulsar, **PSR B1913+16** (Taylor et al., 1979). In fact, almost all of the binary measurements quoted in Chapter 1 were done via pulsar timing. Here I will share one discovery I have not mentioned yet: one of the lesser known facts of astronomy is that the first discovered exoplanet was not **51 Pegasi b** in 1995, but instead **PSR B1257+12 A** and **PSR B1257+12 B** in 1992<sup>49</sup> was (Wolszczan and Frail, 1992). And, you guessed, it, that discovery was made via pulsar timing.

In this section, I will go over the techniques that made these measurements possible. We will begin with getting the simplest orbital solution out of a newly discovered pulsar. From there, I will move on how to describe pulsars via timing, and what information can we extract with this technique. In the end, I will conclude with how to measure component masses from pulsar timing, as this is one of the pillars holding the science results of my thesis. But first, let us understand the data format we use in pulsar timing: folded pulsar archives.

---

<sup>49</sup> To be fair, 51 Pegasi b was the first exoplanet discovered orbiting a *Sun-like* star (Mayor and Queloz, 1995), which is much more media-friendly than planets orbiting pulsars. Nonetheless, learning about **PSR B1257+12** sparked my first interest in pulsars. An look at me now, writing a thesis about them. Because of that it will always have a special place in my heart.

### 2.4.1 Pulsar archives: the backbone of pulsar timing

In searching we worked with high-resolution filterbank files. In pulsar timing we work with fold-mode *pulsar archives*. Pulsar archives were already introduced in Section 2.3.6 as low-resolution  $f - t - \phi_s$  three-dimensional arrays. However, pulsar archives do not always have low resolution. In fact, for pulsar timing, we want keep the original frequency channel resolution and a large number of phase bins, typically 1,024, in order to ensure more precise measurements during timing, while also keeping the four components of the Poincaré vector. However, pulsar archives have indeed low resolution in the time axis, preserving partial folds every few seconds that we call *subintegrations*. Additionally, a pulsar archives store a timing model and a DM value for further folding and frequency integration when needed. The advantage is that new DM values and spin evolution models can be installed at any time before the resolution is reduced for measurements, making them versatile data objects.

Pulsar archives can be achieved in two ways. Search-mode filterbanks can already be folded into pulsar archives. However, this has one big disadvantage: intra-channel DM smearing and time-smearing due to uncorrected orbital evolution from the recorded filterbank. For this reason, once the spin, DM, and orbital parameters of a pulsar are well-known, we record and store observational data directly in high-resolution folded archives. Key to this is the process of storing *coherent de-dispersion*, which used the DM value to change the phase of voltage oscillations via equation (1.9) before producing the frequency channels (Hankins and Rickett, 1975). Then, when pulsar archives are created, the pulsar signal is already de-dispersed within individual frequency channels, allowing for high-precision timing of even the pulsars with the largest DM values.

### 2.4.2 Orbital solutions: the low-precision cousins of timing solution

Once a pulsar is discovered in a binary system, we only have a single observation with an observed spin period,  $P_{\text{obs}}$ , and perhaps a measured LOS acceleration,  $a$ , or apparent period derivative  $\dot{P}_{\text{obs}}$ . That is definitely not enough information to extract the Keplerian parameters, and much less to build a timing model. However, more measurements from further observations give us more hints about the orbital behaviour of a pulsar. In this first section, I will show how to achieve a first orbital solution just from measuring  $P_{\text{obs}}$  and  $a$  at different epochs.

#### Measuring orbital periods at different epochs

The key to solving a pulsar binary is to observe the evolution across of  $P_{\text{obs}}$  across time. That is done by re-observing the pulsar at a regular cadence over several days, weeks or even months. In every observation, an optimised value of  $P_{\text{obs}}$ , and also hopefully  $a$  or  $\dot{P}_{\text{obs}}$ , are obtained. If the  $P_{\text{obs}}$  variability range is not very large, just folding the filter bank at every observation and re-optimising for the best  $P_{\text{obs}}$  and  $a$  values like it is done in Section 2.3.6 is enough, but if the involved accelerations are high, we may need to perform an acceleration search on the de-dispersed time series to re-detect the pulsar. Also,  $a$  and  $\dot{P}_{\text{obs}}$  are completely exchangeable via  $a = c \times \dot{P}_{\text{obs}} / P_{\text{obs}}$ .

To achieve a reliable  $\dot{P}_{\text{obs}}$  measurement, it is important to correct for the motion of the Earth and the telescope with respect to the pulsar. The main effect is the yearly change in the direction of Earth's orbital velocity, which shifts the Doppler effect of the pulsar across the year. This leads to a difference between the spin period observed on site, the *topocentric*  $P_{\text{topo}}$ , and the intrinsic period that would be



observed at the *Solar System Barycenter* (SSB)<sup>50</sup>, the barycentric  $P_{\text{SSB}}$ :

$$P_{\text{topo}} = P_{\text{SSB}} \sqrt{\frac{c - \vec{v}_E \vec{K}}{c + \vec{v}_E \vec{K}}} \approx P_{\text{SSB}} \left\{ 1 - \frac{28.9 \text{ km s}^{-1}}{c} \cos b_E \sin \left( 2\pi \frac{t}{\text{yr}} - l_E \right) \right\}, \quad (2.17)$$

where  $\vec{v}_E$  is the Earth's velocity vector (of magnitude  $28.9 \text{ km s}^{-1}$ ),  $\vec{K}$  is the Solar System-to-pulsar unit vector, and  $b_E$  and  $l_E$  are the ecliptic latitude and longitude of the pulsar (see Appendix A for an explanation of sky coordinate systems). The time axis,  $t$ , now stands for the time of the year on which the pulsar is observed (from 0 to 1 years, with 0 at March 20th, see Appendix A).

After several  $P_{\text{BBS}}$  measurements, we obtain a  $P_{\text{BBS},i}(t_i)$  time series, and maybe even an extra  $a_i(t_i)$  time series, where  $i$  labels each observation. And with that, we can think about solving the binary.

### Orbital modulation of the spin

The observed spin period of a pulsar is modulated by the Doppler effect within its own system. The intrinsic spin period of the pulsar  $P_s$  is shifted into the observed  $P_{\text{SSB}}$  according to

$$P_{\text{SSB}} = P_s \sqrt{\frac{c + \vec{v}_{\text{PSR}} \vec{K}}{c - \vec{v}_{\text{PSR}} \vec{K}}} \approx P_s \left\{ 1 - \frac{x}{c} \dot{E} \sin \omega \sin E + \frac{x}{c} \dot{E} \sqrt{1 - e^2} \cos \omega \cos E \right\}, \quad (2.18)$$

where  $\vec{v}_{\text{PSR}}$  is the pulsar velocity vector,  $x$ ,  $e$  and  $\omega$  are orbital parameters introduced in Chapter 1,  $\omega$  is the *angle of periastron* and  $E$  is the *eccentric anomaly* (for a derivation, see e.g. Section 2B of Dhurandhar and Vecchio, 2001).

The angle of periastron,  $\omega$ , is easy to understand, as it only defines the orientation of the orbit within the orbital plane. In Section 1.3.1, I introduced the orbital eccentricity  $e$ , which defines the shape of the orbital ellipse. The angle  $\omega$  defines the pointing direction of the periastron of the ellipse (see Fig. 1.10), running from 0 to 360 deg. If  $\omega = 0$  deg or  $\omega = 180$  deg, the periastron is pointing perpendicularly away the LOS. If  $\omega = 90$  deg, it is pointing away from the observer. And if  $\omega = 270$  deg it is pointing towards the observer. To see a more detailed depiction of the angles describing the orbital orientation of the system, consult Appendix A.2.2.

The eccentric anomaly  $E$  is a bit more complex. I give a precise definition of it in Appendix A.2.1, but for our purposes here, it is enough to understand it as a measure of the orbital phase based on the periodicity of the orbit  $P_b$  and its eccentricity  $e$ :


$$E - e \sin E = \frac{2\pi}{P_b} (t - T_0) \quad \text{and} \quad \dot{E} = \frac{2\pi}{P_b} \frac{1}{1 - e \cos E}, \quad (2.19)$$

where  $T_0$  the time of passage of periastron. Taking  $k$  as any integer, the time derivative  $\dot{E}$  is the largest at  $t = T_0 + kP_b$  and  $E = 2\pi k$  (periastron) and the lowest at  $t = T_0 + (k + 0.5)P_b$  and  $E = 2\pi(k + 0.5)$  (apastron).

<sup>50</sup> Solar System Barycenter: the center of mass of the Solar System, which is the point around which the Earth orbits. It is close to, but not always within the Sun, as the Sun orbits around it too and Jupiter makes the Sun move a lot.

**PSR B1913+16:**

$P_s = 59.030$  ms  
 $P_b = 0.323$  days  
 $\chi = 2.34$  ls  
 $e = 0.617$   
 $\omega = 292.545$  deg

**PSR J0034-0534:**

$P_s = 1.8772$  ms  
 $P_b = 1.59$  days  
 $\chi = 1.44$  ls  
 $e = 4.3 \times 10^{-6}$   
 $\omega = 314$  deg




Figure 2.10: Barycentric spin and acceleration evolution given the orbits depicted in Fig. 1.10, and drawn according to equation (2.18). On the left, the evolution of the observed  $P_{\text{SSB}}$  is shown, while on the right the correlation between  $a$  and  $P_{\text{SSB}}$  is shown.

Fig. 2.10 depicts the time evolution of the observed  $P_{\text{SSB}}$  and  $a$  values for a couple pulsars, according to equation (2.18). The top row depicts the highly eccentric Hulse–Taylor pulsar (**PSR B1913+16**) within a single orbit, resulting in an asymmetric spin evolution over time. On the other hand, the bottom row depicts the highly circular **PSR J0034–0534**, resulting in a symmetric sinusoidal evolution. Fig. 2.10 also shows the value of  $a$  as a function of  $P_{\text{SSB}}$ , showing that the  $a - P_{\text{SSB}}$  evolution draws a closed curve which is ellipse-shaped for circular systems, and gains complexity for eccentric systems.

**Searching for  $P_b$** 

As you may have guessed, the key to finding an orbital solution is to fit the time series  $P_{\text{BBS},i}(t_i)$  to equation (2.18). But that is not always a trivial task. Fitting for equation (2.18) is almost impossible without a reasonable guess of  $P_b$  as a prior. That is because, unlike the other Keplerian parameters,  $P_b$  is not constrained within a limited range of possible values and can go from zero to infinity. If we take  $\Delta t_{\text{obs}}$  as the typical time between  $P_{\text{BBS},i}(t_i)$  measurements, we can identify two regimes. In the  $\Delta t_{\text{obs}} < P_b$  regime, individual orbital cycles are well-sampled and finding  $P_b$  becomes trivial. However, if  $\Delta t_{\text{obs}} > P_b$  then we may be skipping several orbits between each observation and the series  $P_{\text{BBS},i}(t_i)$  appears to be random. For the alter case, several algorithms exist to estimate the value of  $P_b$ . Here are some of them:

**The acceleration ellipse.** A circular orbit draws an ellipse in the  $a - P_{\text{SSB}}$  space, as is shown for **PSR B1913+16** in Fig. 2.10. This ellipse is described by

$$a^2 = b_2 P_{\text{SSB}}^2 + b_1 P_{\text{SSB}} + b_0, \quad (2.20)$$

where

$$P_s = -\frac{b_1}{2b_2}, P_b = \frac{2\pi c}{P_s \sqrt{-b_2}}, \text{ and } x = \frac{P_b}{2\pi P_s} \sqrt{P_s^2 - \frac{b_0}{b_2}}, \quad (2.21)$$

with  $x$  expressed in units of ls (Freire et al., 2001). If at least three high-quality  $a - P_{\text{SSB}}$  points have been measured, fitting the ellipse results in a direct constraint on  $P_s$ ,  $P_b$ , independently of  $T_0$ ,  $e$  and  $\omega$ . Fig. 2.11 shows two examples of the acceleration ellipse fit with real data, taken from Freire et al. (2001), showing the effectiveness of this method.




Figure 2.11: Example fits to the acceleration ellipse, from only a handful of observations from pulsars **Tucanae 47 W** and **Tucanae 47 S** (the first globular cluster pulsars shown in this thesis!), taken from Freire et al. (2001). The resulting  $P_b$  and  $x$  constraints are 0.1322(16) days and 0.240(3) ls (left), and 1.24(14) days and 0.79(14) ls (right), consistent with their timing parameters.

**The Lomb-Scargle periodogram.** This algorithm is similar to the FFA algorithm used during pulsar searching, but with the difference that it is designed for unevenly sampled time series (Lomb, 1976; Scargle, 1982). It implements a periodicity search in the  $P_{\text{BBS},i}(t_i)$  time series via folding at different trial  $P_{\text{b,trial}}$  values, with each fold changing the observation times  $t_i$  into phase values:

$$\phi_i = \frac{t_i}{P_{\text{b,trial}}} - \text{trunc} \left( \frac{t_i}{P_{\text{b,trial}}} \right). \quad (2.22)$$

Then, on the folded time series,  $P_{\text{BBS},i}(\phi_i)$ , a sinusoidal fit of frequency  $\Omega_{\text{trial}} = 2\pi/P_{\text{b,trial}}$  and starting phase  $\theta_{\text{trial}}$  is performed. Finally, the fit quality

$$\chi^2 = \sum_i \left[ \frac{P_{\text{BBS},i}(\phi_i) - \sin(\Omega_{\text{trial}}\phi - \theta_{\text{trial}})}{\Delta P_{\text{BBS},i}} \right]^2, \quad (2.23)$$

where  $\Delta P_{\text{BBS},i}$  is the measurement uncertainty, is reported as a function of the folding  $P_{\text{b,trial}}$ , and

the correct  $P_b$  should stand above the others (van der Plas, 2018). Because of that, it is particularly sensitive to circular orbits as long as there is enough orbital coverage. This algorithm is a good alternative to fitting the acceleration ellipse if the  $a$  measurements are not very significant.

**The roughness periodogram.** The Lomb-Scargle periodogram loses sensitivity if the orbit is highly eccentric. The roughness periodogram circumvents that by using a different metric to compute the quality of each folding trial. Instead of performing a sinusoidal fit, it measures the *roughness* of the fold by adding the difference between consecutive  $P_{\text{BBS},i}(\phi_i)$  points:

$$R = \sum_{i=0}^{n-1} [P_{\text{BBS},i}(\phi_i) - P_{\text{BBS},i+1}(\phi_{i+1})]^2, \quad (2.24)$$

where  $n$  is the number of measurements (for  $i = n$ , it loops back to  $i = 0$ ). This way, and if there is good orbital coverage, at the correct  $\Delta P_b$  consecutive points in the  $P_{\text{BBS},i}(\phi_i)$  should draw a smooth curve, resulting in a minimum  $R$  value (Bhattacharyya and Nityananda, 2008). Fig. 2.12 demonstrates the effectiveness of this method by showing  $R$  as a function of  $P_{b,\text{trial}}$ , and the  $P_{\text{BBS},i}(\phi_i)$  series at the best  $P_b = 18.791$  days for the highly eccentric **PSR J0514–4002** ( $e = 0.888$ ), taken from (Bhattacharyya and Nityananda, 2008).

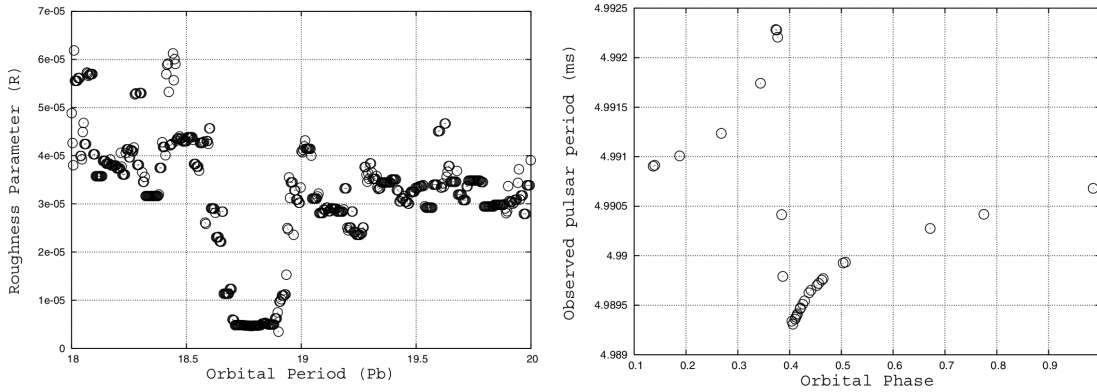


Figure 2.12: Example of a roughness fit to **PSR J0514–4002** observations, from (Bhattacharyya and Nityananda, 2008). On the left,  $R$  is plotted against trial folding periods. On the right, the  $P_{\text{BBS},i}(\phi_i)$  is folded at the correct orbital period of  $P_b = 18.791$  days.

### Fitting for the full-Keplerian description

Once a good value for  $P_b$  has been found, fitting for the remaining orbital parameters becomes an easier job. That is done by minimising the function

$$\chi^2 = \sum_i \left[ \frac{P_{\text{BBS},i}(t_i) - P_{\text{BBS,model}}(t)}{\Delta P_{\text{BBS},i}} \right]^2, \quad (2.25)$$

where  $P_{\text{BBS,model}}(t)$  follows equation (2.18), and  $\Delta P_{\text{BBS},i}$  is the measurement uncertainty. During modelling, we need to transform the  $t_i$  coordinates to  $E_i$  for every trial model via equation (2.19). Since there is no analytic solution to this transformation, it has to be done numerically via algorithms

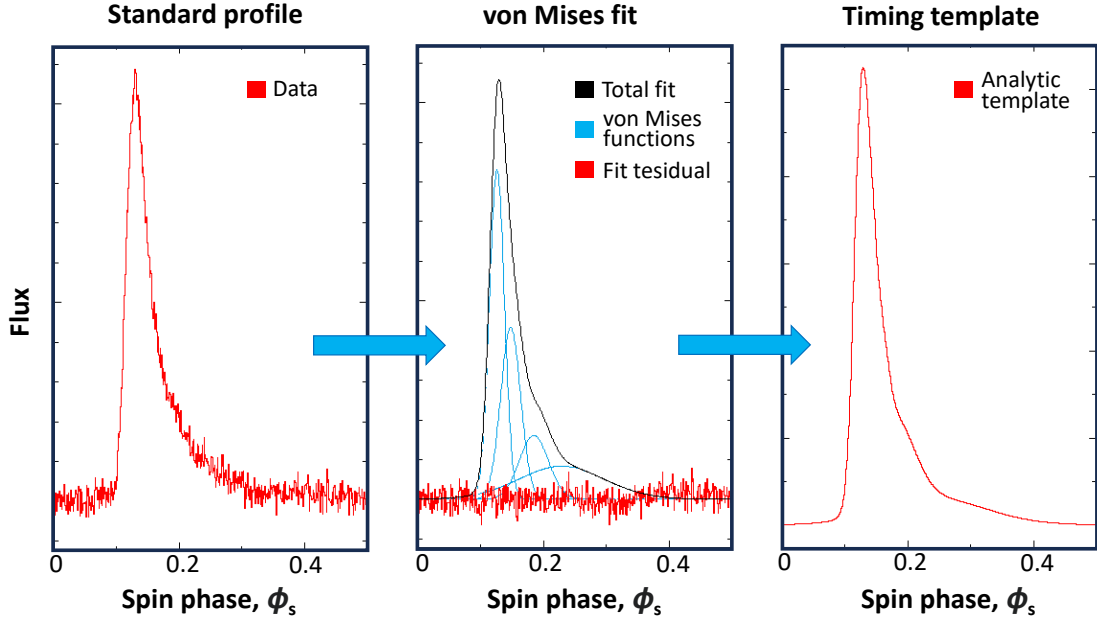


Figure 2.13: Depiction of the making of a timing template from 72 hours of **PSR J12227–6208** Parkes observations at  $f = 968$  MHz, with  $\Delta f = 416$  MHz. In the left plot, we see the standard profile of a scattered pulsar from real data. In the middle plot, four von Mises functions have been fit to the standard profile. In the right plot, the addition of the four von Misses functions results in an analytic timing template.

like the [Newton-Raphson method](#) or the [bisection method](#) (to learn how to solve transcendental equations, see Gautschi, 2011). Nonetheless, by giving adequate priors to the Keplerian parameters, it is easy to make  $\chi^2$  converge into a minimum value.

The best test of the quality of a solution is to see if it recovers the pulsar signal in future observations via direct folding. If so, then it is a successful *orbital solution*. From orbital solutions we can already extract insightful scientific information, such as the minimum companion mass. However, while an orbital solution tracks the  $P_{\text{SSB}}$  evolution over time successfully, but it does not track every single rotation of the pulsar. We are not happy with just that. In pulsar timing, we aim for much higher precision via the measurement of pulse ToAs.

### 2.4.3 The principles of pulsar timing I: measuring ToAs

How do we measure a ToA in the first place? Can we actually detect every single rotation in a pulsar? In reality, we can not do that except with the brightest of pulsars. As stated several times already, to recover a pulsar signal we have to fold and integrate the pulse several times over. It is not different for the measurement of ToAs: we measure the time of arrivals of integrated pulses. To do that, we compare them against a *timing template*, an idealised version of the pulse profile, and the first tool we need to build before we do any timing measurements.

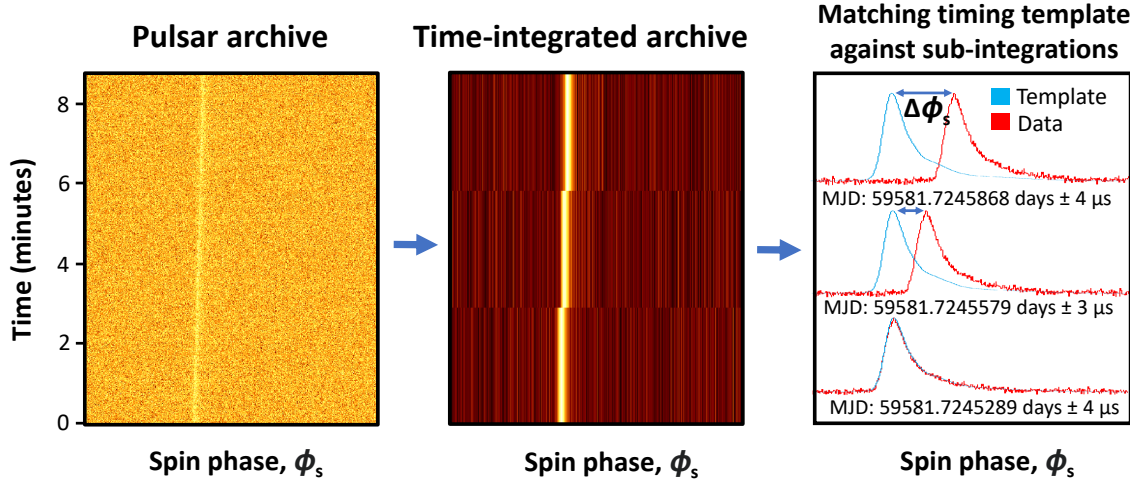


Figure 2.14: Measuring ToAs from a pulsar archive. On the left, we have a pulsar archive with a high time resolution, showing the pulsed signal as a function of the spin phase. In the middle, the archive has been integrated (scrunched) in time to only three sub-integrations. On the right, ToAs are being measured from those sub-integrations with a timing template. Data from **J1452–J1452-5549**, a MMGPS discovery.

### The timing template

A timing template is an idealised, noise-less version of a pulse profile (Lorimer and Kramer, 2005). To build it, we need a high  $S/N$  pulse profile to use it as a reference, which can be achieved by folding and integrating the pulsar signal many times over. For instance, the high  $S/N$  profile shown in the left side of Fig. 2.13 was achieved with the addition of 72 hours of Parkes observations of **PSR J1227–6208**. Of course, how many hours we can integrate for depends on the quality of your preliminary orbital solution. So, if all we have is an orbital solution, we can only get this profile from within a single observation.

Once we have a high  $S/N$ , we proceed to trace a noise-less version of it. The most straightforward way of doing it is to smoothen the Gaussian profile if the  $S/N$  is high enough. However, standard profiles with significant noise will result in spurious structures on the timing template. In those cases, it is best to perform an analytical fit. A common choice is to fit an addition of von Mises functions (Gaussian functions which are cyclical in  $\phi_s$ ). These simplest pulse profiles are well-described by the addition of two or three von Mises functions, while more complex profiles need more components. The final addition of von Mises functions can be used as a timing template. An example of this is shown in Fig. 2.13, where the addition of four von Mises successfully describes the highly scattered pulse profile of **PSR J1227–6208**.

### Measuring ToAs

Fig. 2.14 shows the process of measuring ToAs from folded pulsar archives. The fundamental idea is that we compare the timing template against high  $S/N$  sub-integrations of pulsar archives. A sub-integration is achieved by folding the pulse profile over a  $\Delta T$  interval, resulting in a low-resolution pulsar archive in the time axis. If we are confident in the folding solution, there should be minimal smearing of the pulsar signal during this operation. Then, the timing template and the pulse profile of

every sub-integration are compared to measure the pulse drift in the spin phase,  $\Delta\phi_s$ . Since the phase bins of individual sub-integrations preserve absolute time information, the combination of this time information and  $\Delta\phi_s$  results in an absolute ToA value, often reported in units of [Modified Julian Dates](#) (MJDs), as is common in astronomy. For reference, MJD is the number of days that have transpired since midnight on November 17, 1858 at the Greenwich meridian. Furthermore, ToAs are measured in the reference frame of the telescope, and are therefore topocentric.

The ToA measurement uncertainty  $\sigma_{\text{ToA}}$  can be extremely small, down to a few  $\mu\text{s}$  or even hundreds of ns. In fact, timing precision can have better resolution than that of  $\phi_s$  bins in the pulsar archive, thanks to sophisticated interpolation methods. For example, the Fourier phase gradient scheme (PGS) performs a Fourier transform of both the template and the sub-integration profile to measure  $\Delta\phi_s$  as a global shift of the complex phases of Fourier coefficients (Taylor, 1992). A variation of the PGS is the Fourier-domain Markov-Chain Monte Carlo (FDM), which computes more accurate ToA uncertainties from one-dimensional Markov-Chain Monte Carlo chains (Wang et al., 2022a). Finally, a commonly used alternative is the Gaussian interpolation shift (GIS), which measures the template–profile correlation function for every  $\Delta\phi_s$  with bin resolution, and then interpolates the results with a Gaussian function to find the best  $\Delta\phi_s$  at sub-bin resolution (Hotan et al., 2005). All of these methods are effective, albeit some are better than others depending on the individual pulse shape and  $S/N$  conditions of pulsars (Wang et al., 2022a). Nonetheless, ToA uncertainty always ends up being inversely proportional to the  $S/N$  profile in every sub-integration:

$$\sigma_{\text{ToA}} \propto \frac{S_{\text{sys}}}{n_p G_T F_m} \times \frac{1}{\sqrt{\Delta T \Delta f}} \times \sqrt{\frac{W}{P_s - W}}. \quad (2.26)$$

Therefore, it is evident that a choice is at play here: increased ToA precision implies a sparser sampling of the ToAs. It is a choice of the astronomer whether to prioritise precision or sampling cadence, but too little of either of them can hamper the experiment.

ToA measurements are to be performed across all observations that we want to include in our experiment. Generally, we use the same timing template for all observations from the same receiver to have a global origin for  $\Delta\phi_s$ . However, for data from different telescopes or receivers we have no choice but to implement different timing templates and then account for a possible global phase shift between different data sets. Nonetheless, the final result is a simple list of topocentric ToA values measured as measured by each telescope/receiver. That I will now write down as a series  $\{t_{\text{topo},i}, \text{tel}_i\}$ , where  $t_{\text{topo},i}$  gives the topocentric value of the  $i$ -th ToA, and  $\text{tel}_i$  the telescope where it was measured. This series can include thousands of measurements, and span several years or even decades from first to last.

### Multi-frequency timing

The process explained above is only part of the story. In reality, we often perform *multi-frequency timing*, where we measure several ToAs at different  $f$  values. For this, instead of de-dispersing the entire pulsar archive, we de-disperse only within separated contiguous bands, essentially reducing the resolution of the frequency channels. This presents numerous advantages. For instance, due to scattering or intrinsic profile evolution, the pulse width,  $W$ , may be narrower in some frequencies than in others, improving the ToA quality in some sub-bands according to equation (2.26). Multi-frequency timing also allows us to include frequency information in our timing models, such as DM and its time



derivatives. To perform this kind of timing, two-dimensional timing templates are needed, which is just a fancy way of saying that we need to build a timing template for each frequency sub-band. In the end, this means that ToA measurements also include frequency information, so the time series becomes  $\{t_{\text{topo},i}, f_i, \text{tel}_i\}$ , where  $f_i$  is the central frequency of the band where the  $i$ -th ToA has been measured.

#### 2.4.4 The principles of pulsar timing II: the timing formula

With  $\{t_{\text{topo},i}, f_i, \text{tel}_i\}$ , we finally have the information to build a model of every single rotation of the pulsar. For that, we take  $\phi_s$  as a global rotation count that increases with time, which can be expressed as in terms of the *spin frequency*,  $\nu_s = 1/P_s$ , and its time evolution (Lorimer and Kramer, 2005):

$$\phi_s(t) = \phi_0 + \nu_0(t - t_0) + \frac{\nu_1}{2}(t - t_0)^2 + \frac{\nu_2}{6}(t - t_0)^3 + \dots, \quad (2.27)$$

where the time axis,  $t$ , can run across years or even decades,  $\phi_0$  is  $\phi_s$  at  $t = t_0$ ,  $\nu_0 = \nu_s(t_0)$ ,  $\nu_1 = \dot{f}_s(t_0)$ ,  $\nu_2 = \ddot{f}_s(t_0)$ , and so on (that is right, we are jumping straight into the spin evolution introduced in Section 1.2.7). The idea is simple: every time that  $\phi_s = \phi_0 + k$ , where  $k$  is any integer, we expect to measure a ToA. Now, we only have to compare our measurements  $\{t_{\text{topo},i}\}$  with the prediction.

However, the picture above only works if the radio telescope and the pulsar remain still relative to each other in an inertial frame of reference. In reality, many effects come into play due to the ISM, the way we track time, and the orbital and gravitation influences of the Solar System. To correct these effects, we need to correct from the frequency-dependant topocentric ToA measurements made at the telescope,  $\{t_{\text{topo},i}, f_i, \text{tel}_i\}$ , to barycentric measurements that would be performed in an inertial frame of reference at the SSB,  $\{t_{\text{SSB},i}\}$  via

$$t_{\text{SSB}} = t_{\text{topo}} + t_{\text{corr}} - \mathcal{D}/f^2 + \Delta_{\text{R}\odot} + \Delta_{\text{S}\odot} + \Delta_{\text{E}\odot}, \quad (2.28)$$

which is known as the *timing formula*, (Lorimer and Kramer, 2005).

#### Clock corrections

The first term in equation (2.28),  $t_{\text{corr}}$ , stands for the correction from Coordinate Universal Time (UTC) to Terrestrial Time (TT). ToAs are measured using UTC as the reference at the observatory, but UTC is not a smooth timescale: leap seconds are added by the International Earth Rotation Service (IERS) to account for changes in the rotation of Earth. In addition, observatory clocks may suffer time jumps due to technical issues, which are reported to astronomers. We remove these corrections to have the ToAs running on a smooth timescale, changing from UTC to the International Atomic Time (TAI), which in turn is transformed into TT by a constant offset. The TT is considered to be time as tracked by an ideal atomic clock at sea level<sup>51</sup>, and therefore a reliable timescale.

<sup>51</sup> “Sea level” deserves a more accurate definition. Here I should say ‘Earth’s geoid at sea level’, which is the altitude of equal gravitational potential that coincides (on average) with the sea level. This altitude, however, is not constant, as it is affected by the presence of continental masses or lack thereof in different parts of Earth, the shifting ocean masses due to changing tides, and the Earth’s centrifugal force from its rotation. Entire space-missions have been launched to measure this geoid.

### Dispersion measure corrections and DM time evolution

The second term in equation (2.28),  $-\mathcal{D}/f^2$ , is where the combination of  $\{t_{\text{topo},i}, f_i\}$  values comes into play. It is none other than the correction for measuring the ToA at the measured radio frequency due to the DM-induced delay, where  $D = 4.15 \times 10^6 \text{ ms} \times \text{DM}$ . Using  $f \rightarrow \infty$  as the reference frequency, we apply equation (1.9) to remove this delay. At the same time, fitting this parameter to match the predicted ToAs with the observed ones provides an accurate measurement of DM. That is relevant not just for its absolute value (which is fundamentally biased by the construction of a multi-frequency timing template) but especially for the measurement of its time evolution. This way, timing models can incorporate the DM value and its time derivatives, caused by changes in the cut of the ISM by the LOS due to the relative motion between the observer and the pulsar:

$$\text{DM}(t) = \text{DM}_0 + \text{DM}_1(t - t_0) + \frac{\text{DM}_2}{2}(t - t_0)^2 + \frac{\text{DM}_3}{6}(t - t_0)^3 + \dots, \quad (2.29)$$

which are all measured at the reference epoch  $t_0$ . It is here that the importance of multi-frequency timing becomes apparent, as without it DM derivatives would be absorbed into apparent  $f_s$  derivatives, biasing our measurement of the spin evolution.

### The Earth–Sun Rømer delay and the pulsar sky position

The third term in equation (2.28),  $\Delta_{\text{R}\odot}$ , models the Rømer delay due to the orbital motion and rotation of the Earth. Radio waves propagate at the speed of light, and Earth’s orbit has a radius of eight light-minutes. Because of that, and depending on the position of the pulsar in the sky, we may have up to a 16-minute propagation delay between the expected ToA one day and six months later. This correction implemented by computing the delay between the arrival of a wave-front at the telescope and at the SSB, almost coinciding with the Sun:

$$\Delta_{\text{R}\odot} = \frac{1}{c} \vec{R}_{\text{SSB} \rightarrow \text{tel}} \vec{K}_{\text{SSB} \rightarrow \text{PSR}} = \frac{1}{c} \left( \vec{R}_{\text{SSB} \rightarrow \text{EB}} + \vec{R}_{\text{EB} \rightarrow \text{tel}} \right) \vec{K}_{\text{SSB} \rightarrow \text{PSR}}, \quad (2.30)$$

where  $\vec{R}$  is the distance vector from the SSB to the telescope, which is the addition of the distance vectors from the SSB to Earth’s Barycenter<sup>52</sup> (EB) and from the EB to the telescope, and  $\vec{K}_{\text{SSB} \rightarrow \text{PSR}}$  is the unit vector from the SSB to the pulsar system. I must stress the importance of accounting for the changing position of the telescope with respect to the EB: a photon would take 21.3 ms to reach the center of the Earth from its surface, so failing to correct for the telescope’s changing position due to Earth’s rotation will lead to unaccounted delays of this magnitude, which are very significant compared to  $\mu\text{s}$  resolution of ToAs. Therefore, the  $\{\text{tel}_i\}$  labels are needed to locate the telescope and apply the adequate correction to the ToAs.

The motion of the telescope must be input as an external model. Earth’s orbit is actually slightly eccentric and its trajectory is affected by other Solar System bodies (mostly the Moon and Jupiter), so an analytic description would not be accurate. Instead, Earth’s orbital motion is described by Solar System Ephemeris such as [DE430](#), which computed via numerical simulations and published by NASA’s [Jet Propulsion Laboratory](#). On the other hand, Earth’s rotation is described by [IERS](#), and is

<sup>52</sup> Earth barycenter is the center of mass of our planet, located pretty much in its center. The technical name for that is the *geocenter*, but the abbreviation “EB” is more convenient for equations.

also inputted as an external model. Thus, the only variable in equation (2.30) is  $\vec{K}_{\text{SSB} \rightarrow \text{PSR}}$ , which actually describes the pulsar position in the sky. Therefore, fitting the  $\Delta_{\text{R}\odot}$  actually provides a very accurate fit of the pulsar position and its evolution over time due to the motion between the SSB and the pulsar system. If we simplify the orbit of the Earth to a circle, the evolution of  $\Delta_{\text{R}\odot}$  would be

$$\Delta_{\text{R}\odot} \approx 500 \text{ s} \times \cos b_{\text{E}} \cos \left( 2\pi \times \frac{t}{\text{yr}} - l_{\text{E}} \right), \quad (2.31)$$

where the pulsar's ecliptic latitude,  $b_{\text{E}}$ , and longitude  $l_{\text{E}}$  come explicitly into play. For intuition's sake, I will point out that the time derivative of this simplified equation is none other than equation (2.17). Thus, the Rømer delay can be seen as the cumulative effect of Doppler factor due to Earth's orbital motion. Or alternatively, that the Doppler factor arises from the apparent changes in the rate of arrival of pulses due to Earth's orbital motion.

The position of the pulsar is recorded in right ascension (RA, from 0 to 24 hours<sup>53</sup>) and declination (DEC, from 90 to -90 degrees), which are longitude and latitude coordinates on the equatorial system of sky coordinates, with the origin (RA, DEC) = (0, 0) being at the equator and marking the position of the Sun at the spring equinox (to see a deeper explanation of this coordinate system, please check Appendix A). In addition, the pulsar position actually drifts on the sky over time due to the relative motion between the Solar system and the pulsar system. This results in time derivatives of RA and DEC,  $d\text{RA}/dt$  and  $d\text{DEC}/dt$  which are measurable as long-term variations of Rømer delay. That information is encoded in the *proper motion vector*,  $\vec{\mu} = (\mu_{\text{RA}}, \mu_{\text{DEC}}) = (\cos \text{DEC} \times d\text{RA}/dt, d\text{DEC}/dt)$ , which gives the true angular velocity vector of the pulsar on the sky, often measured in units of milli-arcseconds per year ( $\text{mas yr}^{-1}$ ).

### The Shapiro delay

The fourth term in equation (2.28),  $\Delta_{\text{E}\odot}$ , models the Shapiro delay. According to GR, gravity is not a force but a phenomenon arising from space-time curvature. A consequence of this is that even massless particles such as photons, and therefore radio pulses, are affected by gravitational fields. One of these effects is the apparent slow-down of light as it propagates through curved space-time<sup>54</sup>. In particular, massive bodies in the Solar System delay the arrival of radio pulses at the telescope with their gravitational fields if are close to the LOS (Shapiro, 1964), especially the Sun and Jupiter, which can induce delays of up to 120  $\mu\text{s}$  and 200 ns respectively. These delays must be corrected according to equation

$$\Delta_{\text{S}\odot, \text{body}} = -2 \frac{GM_{\text{body}}}{c^3} \ln \left[ R_{\text{tel} \rightarrow \text{body}} (1 + \cos \theta_{\text{PSR} \leftrightarrow \text{body}}) \right], \quad (2.32)$$

where  $M_{\text{body}}$  is the mass of the Solar System body,  $R_{\text{tel} \rightarrow \text{body}}$  is the time-dependant distance between the telescope and said body, and  $\theta_{\text{PSR} \leftrightarrow \text{body}}$  is the angular distance between the pulsar and the body in the sky (see for example Taylor and Weisberg, 1989). In the case of the Sun,  $R_{\text{body}}$  is almost constant, but in the case of Jupiter and the other planets it changes significantly through the year. Like Earth's trajectory, the motion of planets (and the Sun) is given by the Solar System Ephemeris, so this is an inputted effect and not a measurement.

<sup>53</sup> In astronomy, hours are also used as an angular unit. Each hour is 60 minutes, and each minute 60 seconds.

<sup>54</sup> Although related, this phenomenon must not be confused with gravitational lensing, a phenomenon which also induces a time delay due to the deflection of light rays but that is not measurable in the case of pulsars.

### The Einstein delay

The fifth term in equation (2.28),  $\Delta_{\text{EO}}$ , is the Einstein delay. According to special relativity, when measured by an external observer, the internal clock of an object moving at a relative speed slows down as said speed increases. This phenomenon is called *time dilation*. At the same time, time also ticks slower inside of gravitational fields, a phenomenon called *gravitational redshift*. The Einstein delay corrects for these effects occurring in the telescope as seen from the inertial frame at the SSB. However, unlike the previous two delays, it does not describe an absolute delay at a given time but rather the differential of time passage. The absolute delay is computed by integrating

$$\frac{d\Delta_{\text{EO}}}{dt} = \frac{W_{\text{geoid}}}{c^2} + \sum_{\text{bodies}} \left[ \frac{GM_{\text{body}}}{c^2 R_{\text{tel} \rightarrow \text{body}}} \right] + \frac{(\vec{v}_{\text{EB}} + \vec{v}_{\text{tel}})^2}{2c^2} \quad (2.33)$$

from MJD = 43144.0003725 days to each ToA, where  $W_{\text{geoid}} = 6.969290134 \times 10^{-10} c^2$  is the Earth's gravitational and potential at the sea-level geoid,  $\vec{v}_{\text{EB}}$  is the vector of Earth's velocity relative to the SSB, and  $\vec{v}_{\text{tel}}$  the vector of the telescope's velocity with respect to the EB due to Earth's rotation (see for example Taylor and Weisberg, 1989; Edwards et al., 2006). The first term in the equation removes the contribution of Earth's gravitational redshift by correcting from the time passage at the TT frame to the time passage in a rest-frame at the EB, and it has a magnitude of  $22 \text{ ms yr}^{-1}$ . The second term removes the gravitational redshift from other Solar System bodies other than Earth. The largest contributions are that of the Sun and Jupiter, which changes over time but is always close to  $0.35 \text{ s yr}^{-1}$ . The third term corrects for the time delay caused by the telescope's motion, which also changes due to the Earth's changing orbital speed, but has an average rate of  $0.15 \text{ s yr}^{-1}$  from the orbital motion contribution, plus up to  $38 \text{ } \mu\text{s yr}^{-1}$  from the rotation contribution if the telescope is located at the equator (if the telescope is in any of the poles, rotation brings no contribution<sup>55</sup>). Once again, these effects are inputted as an external model and are therefore not fitted for. For the last term, the location label  $\{\text{tel}_i\}$  of each ToA becomes relevant once again, as it is necessary to compute the correct time-varying value of  $\vec{v}_{\text{tel}}$ .

### Putting it all together

If all of these corrections are applied correctly, then there are three fundamental phenomena that can be modelled from this scheme: spin evolution from equation (2.27), DM evolution from the  $-D/f^2$  term in equation (2.28) and equation (2.29), and the sky position and the proper motion vector from the Rømer delay in equation (2.28). These are modelled by the timing parameters  $\{p_j\} = \{f_0, f_1, f_2, \text{DM}_0, \text{DM}_1, \text{DM}_2, \text{RA}, \text{DEC}, \text{and } \vec{\mu}\}$ , which give a timing model,  $\phi_s(t_{\text{SSB}}, \{p_j\})$ . Given the timing model, we can compute the timing residuals

$$\Delta t_{\text{ToA},j} = P_s \times [\phi_s(t_{\text{SSB},i}) - \phi_0 - k_i], \quad (2.34)$$

where  $k_i$  is the closest integer rotation count to a given  $t_{\text{SSB},i}$ . If our timing model is correct, these timing residuals should distribute themselves as a Gaussian distribution around 0 with their ToA

<sup>55</sup> Yes, this means that we age more slowly on the equator than on the poles (at sea level at least). So, you now know where to move if you want to live a slightly longer life. However, it is debatable whether the improvement would come from relativity or from the much more enjoyable climate compared to the poles. Fun fact: moving to the South Pole is a viable life choice, there are [hundreds of scientists already doing so](#).

uncertainty for standard deviation, both across time and frequency, as is shown in top plots of Fig. 2.15, which depict ToAs residuals from a good fit as a function of time, and as a function of frequency. In both cases, the residuals cluster around 0, meaning that the spin, position and DM of the pulsar are well-modelled. However, if something is missing or is incorrect in the model, it will result in structured residuals, as is seen in the bottom row of Fig. 2.15. The left plot shows that an incorrect  $f_1$  parameter results in a quadratic drift of their residuals, while an incorrect  $\vec{\mu}$  vector results in an improper correction of the Rømer delay, leading to yearly periodic residuals of evolving amplitude. It is precisely these structure residuals which are measured, leading to constraints on the timing parameters.

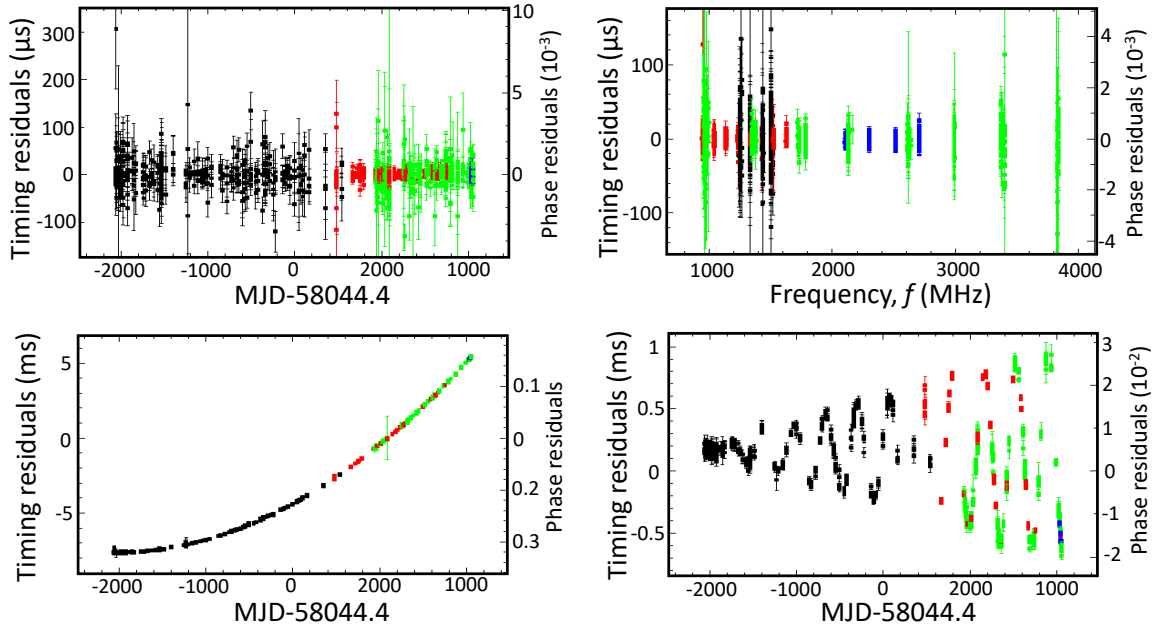


Figure 2.15: Examples of ToA residuals from more than 10 years of timing from pulsar **J1227–6208**, from four different receivers at two different telescopes. The top plots shows the timing residuals when the timing model is correct. In the left plot, they are shown against time (MJD), while in the right plot against frequency. The bottom plots have an incorrect  $f_1$  value (left), and an incorrect  $\vec{\mu}$  (right), resulting in structured residuals across time.

Speaking of which, the timing model is found by minimizing the chi-squared statistic

$$\chi^2 = \sum_i \left[ \frac{\phi_s(t_{\text{SSB},i}) - \phi_0 - k_i}{\sigma_{\text{ToA},i}/P_s} \right]^2, \quad (2.35)$$

(see the [Nobel Prize lecture from Joseph H. Taylor in 1993](#) for the first case of implementation of this formula). Ideally, the  $\chi^2$  function follows a multi-dimensional Gaussian probability distribution in the space of timing parameters,  $\{p_j\}$ , where  $j$  stands for the  $j$ -th timing parameter, described by

$$\rho(\{p_j\}) \propto \exp \left[ -\frac{\chi^2(\{p_j\}) - \chi_{\min}^2}{2} \right], \quad (2.36)$$

where  $\chi^2_{\min}$  is the best-fit model. In this distribution, the mean values  $\{p_j\}$  and their standard deviations  $\{\sigma_j\}$  are considered to be their measurement and their uncertainty. Fitting softwares often quote the multi-dimensional co-variance matrix of this distribution (for a good introduction on numerical analyses check out see Press et al., 1986), making it quite convenient for us to read the uncertainties. Then, the significance of a measurement is quoted as  $p_j/\sigma_j$ , and we usually we consider a detection is made if  $p_j/\sigma_j > 3$ . However, as we shall see, parameters can be degenerate with each other, often resulting in a more complex behaviour of  $\rho(\{p_j\})$ , prompting us to develop more sophisticated methods of exploring of the  $\chi^2$  function to achieve more reliable uncertainties.

### 2.4.5 Advanced pulsar timing I: the binary timing formula

So far, we have only considered the timing models for isolated pulsars. But in this thesis, we are interested in binary pulsars. Therefore, something is missing here. I'll tell you what it is: the same corrections that have been applied because of Earth's orbital motion are to be applied for the pulsar's binary motion in its binary system too! That is, we need to move to the frame of reference of the rotating pulsar, which implies an extra conversion to pulsar time (Lorimer and Kramer, 2005)

$$t_{\text{PSR}} = t_{\text{SSB}} - \Delta_{\text{R,PSR}} - \Delta_{\text{S,PSR}} - \Delta_{\text{E,PSR}}, \quad (2.37)$$

where the correction terms have the same physical meaning as in equation (2.28), but are of opposite sign, because we are moving from the binary system barycenter (the same as the SSB except for a constant, non-measurable Doppler factor caused by the radial motion between the Solar System and the pulsar system) to the frame of reference of the pulsar. However, unlike for Earth, where we input the orbital motion as an external model, now the orbital motion of the pulsar is the object of measurement. As we will see, the corrections in equation (2.37) go beyond simple Keplerian orbits we have described pulsars with. That is because, for the most compact binaries, the we need to account for relativistic corrections to Newtonian physics, leading to a first-order *post-Keplerian* (PK) description. This discription is thanks to the introduction of PK parameters, which can be fitted for (and therefore measured) independently, leading to a complete and accurate description of the ToAs.

In this section, I will present the binary timing models used in this thesis, and the parameters that they use for the PK description. All of these models aim to describe the same physics, but as we will see, some are more convenient than others depending on the binary system that we are studying.

#### The Damour-Deurelle model for eccentric systems

For eccentric systems, we model the three corrective terms in the right side of equation (2.37) with the Damour-Deruelle relativistic model (DD, Damour and Deruelle, 1986). This model implements the Rømer delay,  $\Delta_{\text{R,PSR}}$ , as a function of the pulsar's eccentric anomaly:

$$\Delta_{\text{R,PSR}} = x(\cos E - e) \sin \omega + x \sin E \sqrt{1 - e^2} \cos \omega. \quad (2.38)$$

In here, I should note that I am still describing a Keplerian orbit, ignoring the relativistic deformation of the orbital ellipse occurring in the most accelerated binaries, but that has not been measured in any of the systems studied in this thesis. Given that, I also like to point out that the equation above is the time integration of equation (2.18), just as it happened with Earth's Rømer delay.

Despite providing a Keplerian description of the orbit, this model already allows first-order relativistic corrections with the implementation of first-order PK parameters, starting with the time derivatives of Keplerian parameters such as  $\dot{\omega}$ ,  $\dot{x}$  and  $\dot{P}_b$ . The first two, are included by allowing a time variation of the parameters in equation (2.38), while the latter introduced as a modification of the time-dependence of  $E$ , expanding equation (2.19) into

$$E - e \sin E = \frac{2\pi}{P_b} \left[ (t - T_0) - \frac{1}{2} \frac{\dot{P}_b}{P_b} (t - T_0)^2 \right]. \quad (2.39)$$

Therefore, fitting for  $\Delta_{R,PSR}$  results not only in the measurement of five Keplerian parameters, but potentially three PK parameters as well.

In the DD model, the Shapiro and Einsein delays,  $\Delta_{S,PSR}$  and  $\Delta_{E,PSR}$ , are also modelled as a function of  $E$ . Here, the Shapiro delay is now caused by the propagation of radio pulses through the gravitational field of the companion, and its amplitude is

$$\Delta_{S,PSR} = -2r \ln \left[ 1 - e \cos E - s \left( \sin \omega (\cos E - e) + \sqrt{1 - e^2} \cos \omega \sin E \right) \right], \quad (2.40)$$

where  $r = T_\odot M_c$  and  $s = \sin i$  are the *range* and *shape* parameters of the Shapiro delay, which are determined by companion mass and the orbital inclination angle. As is shown in the left columns of Fig. 2.16, which plots  $\Delta_{S,PSR}$  as a function of  $E$  for different  $s$  values, the periodic modulation of the Shapiro delay has a clear peak. This peak occurs when the true anomaly

$$A_T = 2 \arctan \left[ \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2} \right], \quad (2.41)$$

another angle of orbital phase that is measured from the pulsar system barycenter (see Appendix A.2.1 for a detailed introduction), takes a value of  $A_T = \pi/2 - \omega$ . That is known as *superior conjunction*, the point where the pulsar is behind the companion in its orbit, and therefore radio pulses propagate deeper into its gravitational field. The left columns of Fig. 2.16 also show how the inclination angle can drastically change the Shape of the delay. With inclination angles  $i$  close to 90 deg ( $s$  close to 1, edge-on orbits), the global amplitude of the periodic delay is very large and peaks sharply at superior conjunction, but with lower values (face-on orbits) the amplitude is drastically reduced and the peak at superior conjunction becomes less sharp, making it harder to measure.

Finally, the cumulative Einstein delay is, for once, rather easy to describe, following a simple sinusoidal in  $E$ ,

$$\Delta_{E,PSR} = \gamma_E \sin E, \quad (2.42)$$

where  $\gamma_E$  is the amplitude of the Einstein delay, which can be as large as several milliseconds in some systems. The parameters  $r$ ,  $s$  and  $\gamma_E$  constitute another set of measurable PK parameters. Along with the previous Keplerian and PK parameters, they are also new timing parameters to be included in timing models.

### The Laplace-Lagrange model for circular systems

In Section 1.3.2, we showed that some binaries have ridiculously low eccentricities, with the most circular ones reaching  $e \sim 10^{-7}$ . While in theory the DD model still applies, in many cases the ToA






Figure 2.16: Full timing residual modulation by Shapiro delay (left) and residual modulation by Shapiro delay after the full amplitude is absorbed into the Rømer delay (right), for different  $r$  and  $s$  values and as a function of the eccentric anomaly. In all cases, the Shapiro delay is not fit, leaving structured timing residuals behind. The top row shows the resulting delay from  $r = 6.9 \mu\text{s}$  ( $M_c \approx 1.4 M_\odot$ ) and  $s = 0.995$  ( $i \approx 84 \text{ deg}$ ). In the bottom row,  $s = 0.6$  ( $i \approx 37 \text{ deg}$ ). If Shapiro delay is fitted for, it would result in flat residuals. The simulation assumes a circular system for the sake of simplicity.

uncertainty is too high to measure  $e$  and  $\omega$  with a large significance, leading to fits with negative  $e$  values or random  $\omega$  and  $T_0$  values. The Laplace-Lagrange model (ELL1, Lange et al., 2001) fixes this by replacing the  $e$  and  $\omega$  parameters with the Laplace-Lagrange parameters

$$\epsilon_1 = e \sin \omega \quad \text{and} \quad \epsilon_2 = e \cos \omega, \quad (2.43)$$

and changing time of passage of periastron,  $T_0$ , to the time of passage of the ascending node,  $T_a = T_0 - \omega \times P_b / 2\pi$ , the point in the orbit where the pulsar crosses the sky plane moving away from the observer. With these parameters, the Rømer delay term is rewritten as

$$\Delta_{\text{R,PSR}} = x \left( \sin \Phi + \frac{\epsilon_2}{2} \sin 2\Phi - \frac{\epsilon_1}{2} \cos 2\Phi \right), \quad (2.44)$$

where  $\Phi$  is analogous to the Eccentric anomaly, but is now measured from  $T_a$ :

$$\Phi = \frac{2\pi}{P_b} \left[ (t - T_a) - \frac{1}{2} \frac{\dot{P}_b}{P_b} (t - T_a)^2 \right]. \quad (2.45)$$

This model also allows for the measurement of  $\dot{x}$ ,  $\dot{P}_b$  and  $\dot{\omega}$ , the later being inferred from  $\dot{\epsilon}_1$  and  $\dot{\epsilon}_2$ , although for non-eccentric systems they are expected to be extremely small. The Shapiro delay also

becomes significantly simpler, being described by the same parameters:

$$\Delta_{\text{S,PSR}} = -2r \ln [1 - s \sin \Phi]. \quad (2.46)$$

However, due to the extreme circularity of these orbits, the Einstein delay is not included in this model. That is mainly because in circular orbits the distance to the companion and the orbital speed of the pulsar does not change, turning the Einstein delay into a constant, non-measurable factor. In addition to that, the Einstein delay is degenerate with  $\dot{x}$  in systems with very low  $\dot{\omega}$  values, so its inclusion can lead to a bias in the measurement of both parameters (for a discussion of this phenomenon, see Ridolfi et al., 2019).

### The orthometric parametrisation of the Shapiro delay

The  $r$  and  $s$  parameters describe the full amplitude of the Shapiro delay, but they come with a few problems. To begin with, they can be highly correlated at low  $i$  values, meaning that a large combination of  $r$  and  $s$  values can describe the same ToAs, and thus making their uncertainties grow. In addition, the Rømer delay and the Shapiro have the same periodicity in  $E$ , making them hard to distinguish during measurement. Since the Rømer delay has an amplitude of  $x/c$  (seconds), the  $\mu\text{s}$  amplitude of the Shapiro delay gets the short end of the stick and is absorbed into the Rømer delay during fitting at low inclination angles. The right column of Fig. 2.16 demonstrates this well. In the  $s = 0.995$  case, after fitting for *only the Rømer delay*, the residual amplitude of Shapiro delay is halved. And for the  $s = 0.6$  case, almost all of its amplitude is absorbed into the Rømer delay description. The general implications are that for low  $s$  values, it becomes difficult to measure of  $r$  and  $s$ .

To circumvent these correlation, Freire and Wex (2010) came up with a new description of the Shapiro delay. For circular systems, we can take equation (2.46) and decompose it into harmonics

$$\Delta_{\text{S,PSR}} = -2r(a_0/2 + a_1 \sin \Phi + a_2 \cos 2\Phi + a_3 \sin 3\Phi + a_4 \cos 4\Phi + \dots). \quad (2.47)$$

Comparing to equation (2.44), we can see that the  $-2ra_1$  component is absorbed into  $x$  and  $-2ra_2$  is absorbed into  $x\epsilon_1/2$ , showing that these harmonics are indeed not measurable. Then, the Shapiro delay is only measurable if the third and higher harmonics are still significant. In this context, the *orthometric ratio* between the amplitude of consecutive harmonics

$$\varsigma = \left( \frac{1 - \cos i}{1 + \cos i} \right)^{1/2}, \quad (2.48)$$

with  $|a_k| = 2\varsigma^k/k$  for  $k > 1$ , is included as a timing parameter included as a timing parameter. Also, the third harmonic  $h_3 = r\varsigma^3$  becomes the measurable *orthometric amplitude* of the Shapiro delay. Thus,  $h_3$  and  $\varsigma$  substitute  $r$  and  $s$  in the timing model. This formalism can also be extended to eccentric systems, and it has been included in modified versions of the DD and ELL1 models: the DDH and ELL1H models. Finally, Freire and Wex (2010) show that unlike  $r$  and  $s$ ,  $h_3$  and  $\varsigma$  are not degenerate with each other, and that  $h_3$  is in general measured with higher significance than  $r$  even in the same data sets.

### 2.4.6 Advanced pulsar timing II: biting timing solutions with *dracula*

We have already seen the chi-squared statistic

$$\chi^2 = \sum_i \left[ \frac{\phi_s(t_{\text{SSB},i}) - \phi_0 - k_i}{\sigma_{\text{ToA},i}/P_s} \right]^2, \quad (2.49)$$

which we seek to minimise to find our timing solution. But, how many free parameters (variables) do we have to fit for? With spin evolution, DM evolution and proper motion, isolated pulsar parameters can already require eight parameters. To that, we have to add five Keplerian parameters to describe the Rømer delay. And if we include PK parameters we may even go beyond 15 free parameters in total. That means that  $\chi^2$  is actually a function with more than 15 variables!

Therefore, the process of finding a timing solution for a set of ToAs is not as trivial as finding an orbital solution like we did in Section 2.4.2, as these fits only required six parameters. Even if we find a solution, how do we know we have found an absolute minimum in  $\chi^2$ ? It turns out that the  $\chi^2$  may have multiple *local minima* in the  $\{p_j\}$  space. That is because a timing model that predicts the rotation counts  $\phi_s(t_{\text{SSB},i})$  is also used to predict the closest integer rotation count to each ToA,  $k_i$ , and it can indeed happen that different timing solutions are consistent with the same set of ToAs but with different  $k_i$  values. But we know that only one real, astrophysical solution can be true at the same time. Such solution is the one that tracks every single rotation of the pulsar, the one that we call the *phase-connected* or *coherent* solution.

The key to finding the only true phase-connected solution is to turn the problem on its head: instead of explicitly searching for the best point in the  $\{p_j\}$  space, we search for the best  $\{k_j\}$  combination and take the  $\{p_j\}$  of parameters that comes with it. Here I explain how to do that with *dracula* algorithm (Freire and Ridolfi, 2018), starting from the orbital solutions presented in Section 2.4.2.

#### The limits of orbital solutions

Before I went onto describe timing solutions, we already saw orbital solutions as good descriptions the binary system in Section 2.4.2. But the precision of orbital solutions is extremely limited. Their problem is that, while they track the rotations of a pulsar within individual observations thanks to  $P_{\text{obs}}$  and  $\dot{P}_{\text{obs}}$ , they can not track them *between observations*. That is illustrated in Fig. 2.17, where two pairs of consecutive observations of an hypothetical binary pulsar are shown, A–B and C–D. Within individual observations,  $P_{\text{obs}}$  is measured by tracking the drift of the pulse within the archive, which in this case is depicted to change between observations due to the binary nature. From these measurements we then achieve an orbital solution, but this solution has no information of the pulse drift *in between observation*. As shown in the diagram, after the end of observations A and C, the pulse may drift and even wrap around the spin phase an unknown number of times before being captured again in observations B and D. Thus, we can say that orbital solutions are phase-connected *within individual observations* that last a few minutes, but not across the hours or days that separate said observations.

#### From orbital solutions to a jumped timing solutions

Before searching for the one true timing solution, we have to translate the orbital solution into a solution describing the ToAs. For that, we ignore multi-frequency timing and the PK-description of

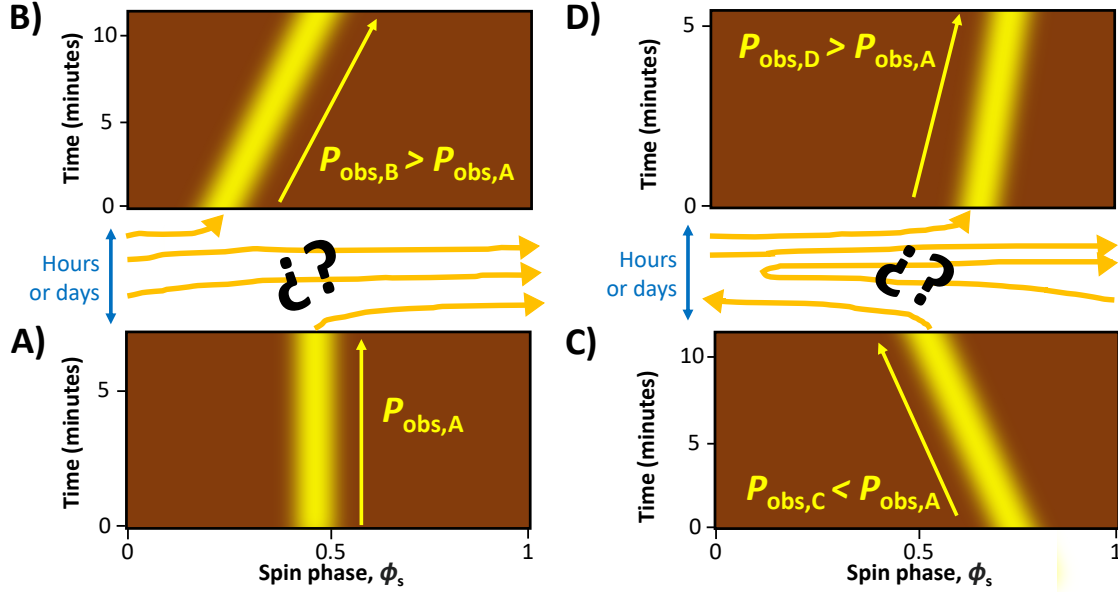


Figure 2.17: Lost rotation counts due to the innacuracy of an orbital solution. Plots A, B, C, and D represent time–phase plots of a folded observations of the same binary pulsar, with B coming days or hours after A, and D after C (the time axis is vertical and increases upwards). The yellow pulse is tracked within each observation, allowing for a good measurement of  $P_{\text{obs}}$ , but in the hours between the observations between the drift is not observed and we loose track of the number of rotations. To illustrate this, the pulse (orange arrow between plots) is drawn to wrap around several times between A and B, and C and D. In theory this wrapping could happen tenths or hundreds of times.

the pulsar. We only seek to achieve a first phase-connected timing solution, not a high-precision timing solution, and full-band ToAs and the DD or ELL1 models without PK parameters will be able to provide that. That is because we are happy to let DM and spin evolution to be confused as long as we have a preliminary model that tracks every rotation, and for instance the  $\mu\text{s}$  amplitude of the Shapiro delay has no effect on the ms-wide cycles of each rotation, so we can let it be absorbed into the Rømer delay. However, we may need to include  $f_1$  or RA and DEC in the description if the system requires it.

To seek the first preliminary timing solution, we change the  $P_{\text{obs}}$  and  $\dot{P}_{\text{obs}}$  measurements for a set of at least three full-band topocentric ToAs from each observation (in order to track spin evolution due to binary effects within each observation). As is shown by the green dots in the top plot of Fig. 2.18, if we use the orbital solution as a timing model, we will find that the timing residuals are flat in time within individual observations, but that entire observations are away from 0 by any random value between  $-1/2 < \Delta t_{\text{ToA}}/P_s < 1/2$  due to the plotting program doing its best to guess the closest rotation counts  $\{k_i\}$  to each ToA from a solution that does not account for them to begin with.

Fitting for these ToAs will result in a crash when whichever software we are using tries to find a solution that minimises the residuals without a viable  $\{k_i\}$  set. Instead, we do nice trick: we introduce variable time jumps between the observations and fit for the timing parameters. This will result in a “cheated” timing solution that flattens all residuals to 0, as shown with the red dots in the top plot of Fig. 2.18. I say “cheated” because despite it looking good, we have allowed the solution to ignore the

rotation counts between the last ToA and the first ToA of consecutive observations,  $\Delta k_i = k_{i+1} - k_i$ , and therefore to ignore the global rotation count. This comes at the price of introducing of as many variables as jumps there are to the timing model: the jumps themselves. That is a *jumped timing solution*, which does not have more information than a simple timing solution, yet.

### The *dracula* algorithm and the solution tree.

The *dracula* algorithm explores the set of  $\Delta k_i$  values between observations in a systematic way to identify the best solution that results in an absolute minimum in  $\chi^2$  (Freire and Ridolfi, 2018). It works by expanding the phasse-connection from within withing observations by connecting discrete parts, until all parts are joined and the global rotation count is accounted for. Its starts wit the set of ToAs, the jumped timing solution, and the finite set of variable jumps between observations. It also relies on the evolution of the reduced *reduced* chi-squared statistic

$$\chi_{\text{red}}^2 = \frac{\chi^2}{N_{\text{free}}} = \frac{\chi^2}{N_{\text{ToAs}} - N_{\text{par}}}, \quad (2.50)$$

where  $N_{\text{free}} = N_{\text{ToAs}} - N_{\text{par}}$  is the number of degrees of freedom, with  $N_{\text{ToAs}}$  is the number of ToAs and  $N_{\text{par}}$  is the number of timing parameters. A good fit results in  $\chi_{\text{red}}^2 \approx 1$  (model within ToA uncertainties), with  $\chi_{\text{red}}^2 < 1$  implying some over-fitting (model too precise given the ToA uncertainties), and  $\chi_{\text{red}}^2 \gg 1$  implying that the model does not describe the data at all (ToAs are not consistent with the model).

The dracula algorithm starts by removing a *single jump* between two slected, consecutive observations, and then fitting again for a global solution. This time, whichever fitting software we are using will likely converge onto a solution with  $\chi_{\text{red}}^2 \approx 1$  that flattens the residuals, as it only needs to adapt the timing parameters  $\{p_j\}$  to a set of six consecutive ToAs while it is allowed to adjust the remaining jump values between other observations. This solution has already found its  $\Delta k_i$  there where the jump has been removed. But here comes the core of the algorithm: the  $\Delta k_i$  value could be a wrong one, arising from convenience by a small number of ToAs. Because of that, the algorithm forces the fitting program to also report the timing solutions if we instead account for  $\Delta k_i + 1$ ,  $\Delta k_i + 2$ ,  $\Delta k_i - 1$ ,  $\Delta k_i - 2$ , etc., between the two observations. And it will turn out that many of these solutions are also viable! As shown in the bottom-left plot of Fig. 2.18, several solutions may exist, with their  $\chi_{\text{red}}^2$  values following a parabola around the best  $\Delta k_i + n$ , where  $n$  is not necessarily 0.

And now here comes the beauty of the dracula algorithm: it assumes that all solutions under a  $\chi_{\text{red}}^2$  threshold (e.g.  $\chi_{\text{red}}^2 < 2$  or  $\chi_{\text{red}}^2 < 3$ ) are potentially the correct one, with each of them starting a *solution branch*. Then, for each of the accepted solutions, it proceeds to remove a second jump, re-fit with different  $\Delta k_i + n$  values, and once again pick all solutions under the  $\chi_{\text{red}}^2$  for the next jump removal, creating further branches, and so on until the last jump is removed<sup>56</sup>. Along the way, some branches will die, producing no “offspring” solutions under the  $\chi_{\text{red}}^2$  threshold, and this will be indicative that  $\Delta k_i + n$  in their original branching point was an incorrect one. If we are lucky, in the end only a single branch will survive: the one containing the true timing solution that tracks the global rotations counts.

<sup>56</sup> I like to imagine the process of exploring the  $\Delta k_i + n$  as seeing how does a clock react if you change the rotation phase of some of its gears while one of its sides is disconnected, and then it is reconnected reconnected, and where every jump represents a different gear. I also fantasize with them as giant wooden gears that make *satisfyingly deep clock-turning noises*, but that’s just me.

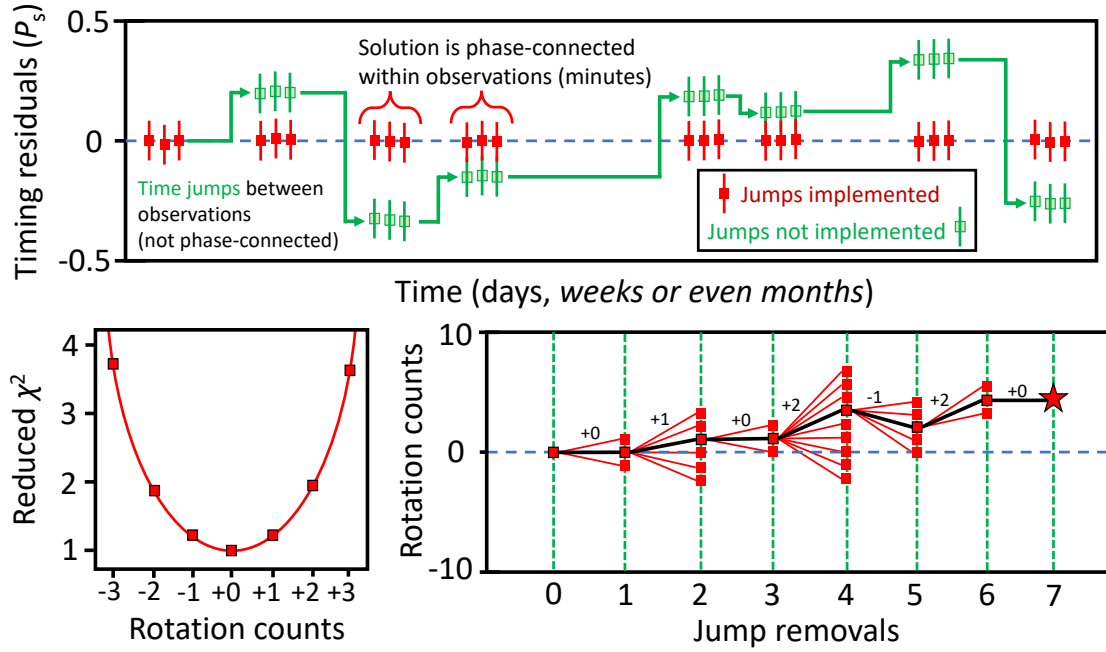


Figure 2.18: Depiction of the *dracula* algorithm based on Freire and Ridolfi (2018), the description in the text, and my experience implementing it. At the top, I show how the timing residuals of an orbital solution looks like (green ToAs), and how the timing residuals are improved if we allow the solution to ignore the rotation counts between observations (red ToAs). At the bottom-left, I show the behaviour of  $\chi^2_{\text{red}}$  as a function of the rotation counts  $\Delta k_i + n$  (or just  $n$  in the plot), during the removal of a jump between two consecutive observations. In the bottom-right, the path of that the true solution follows across the solution tree is depicted, with other dead branches drawn after the first iteration (in reality, dead branches may survive up to several jumps, but that would be too complicated to draw). This is an idealised case with only eight observations and where  $n$  never goes beyond  $\pm 2$ , but for the worst pulsars we may have dozens of observations and hundredths of viable solutions in every jump removal, each of them spawning hundreds more in the next one.

This one does not need to stem always from the lowest  $\chi^2_{\text{red}}$  solution in the branching points: there could be dead branches that started with a lower  $\chi^2_{\text{red}}$  than the final surviving solution. Therefore, it is important to explore the entire  $\Delta k_i + n$  space at every jump removal.

However, if more than one solution survives at the end, it means that we do not have enough data to count all the rotations counts just yet. Then, only new ToAs derived from new observations will be able to solve the pulsar.

### 2.4.7 Mass measurements I: physics beyond Isaac Newton

Now we have everything to model every single rotation of a pulsar thanks to our beloved timing solutions! Everything has been leading to this point: we are now ready to do some hardcore astrophysical measurements. As we learned in Sections 1.3.2 and 1.3.4, there are three keys to knowing the story of a pulsar in a binary: its spin ( $P_s$  and  $\dot{P}_s$ ), its Keplerian orbital parameters, and the stellar masses of the binary components. The first two have already been achieved through the previous sections, and now our main focus will now shift on to measuring the stellar masses in a pulsar binary. The key to those are the PK parameters.




Figure 2.19: Visual depiction of the post-Keplerian orbital phenomena depicted by GR: the rate of advance of periastron  $\dot{\omega}$ , depicted as the gradual rotation of the orbital ellipse, the Einstein delay, which causes the internal clock of the pulsar to run at different rates at different parts of the orbit, the Shapiro delay, which delays the arrival of radio pulsars of the telescope when they have to propagate through the gravitational field of the orbital companion, and the gradual orbital decay due to the emission of gravitational waves. All of them are modelled by one PK parameter, except the Shapiro delay, which is modelled by two different PK parameters.

So far, we have treated the PK as relativistic phenomena that can be modelled independently from each other. However, if our understanding of gravity is correct this is not the case, and we can use that to our advantage. In 1915, Albert Einstein published the general theory of relativity, or general relativity (GR, Einstein, 1915a), which significantly reformulated how gravity works. Instead of a force, the main agent of gravity was now the curvature of space-time, and this description led to the prediction of several new phenomena. These include time dilation due to the presence of a gravitational field, the perturbation of light ray trajectories due to the presence of a gravitating body, and deviations from Keplerian orbits in the presence of strong gravity and high speeds, amongst others (to learn more about GR, check e.g. Misner et al., 1973). Many of these effects have already been introduced in Sections 2.4.4 and 2.4.5, namely: the secular time evolution of several Keplerian orbital parameters ( $\dot{\omega}$ ,  $\dot{x}$  and  $\dot{P}_b$ ), the Shapiro delay Parameters, and the Einstein delay parameters, whose phenomenology are illustrated in Fig. 2.19. According to GR, these corrections can be modelled from the pulsar mass, the companion mass, and the orbital inclination angle. Let us see how.

### The rate of advance of periastron

Let us remember: in a Keplerian orbit, the *periastron* is the point of closest approach to the center of mass, the “inner tip” of the eccentric orbit in the diagram of Fig. 1.10. The position of the periastron with respect to the observer is described by the angle  $\omega$ , which is measured via the Rømer delay in the



pulsar system. However, it turns out that  $\omega$  is not fixed in GR. In an isolated two-body system with point masses, GR predicts a secular evolution of  $\omega$  over time, a precession of the orbital ellipse around the axis perpendicular to the orbital plane, a phenomenon measured via the *the rate of advance of periastron*<sup>57</sup>

$$\dot{\omega} = 3 \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{(T_\odot M_t)^{2/3}}{1 - e^2}, \quad (2.51)$$

where  $M_t = M_p + M_c$  is the total mass of the system, expressed in units of  $M_\odot$ . Therefore, a precise measurement of the relativistic  $\dot{\omega}$  value translates in a precise measurement of the total system mass. This is best measured in compact systems with low  $P_b$  and high  $e$  values, and in some cases it is so significant that it can already be measured via equation (2.18) even before a timing solution is achieved. For instance, in the Hulse–Taylor binary<sup>58</sup>, with  $P_b = 0.323$  days and  $e = 0.617$ , it has a value of  $\dot{\omega} = 4.226585(4) \text{ deg yr}^{-1}$ , which implies  $M_t = 2.82737(4) M_\odot$  (Weisberg and Huang, 2016). For the double pulsar, with  $P_b = 0.102$  days and  $e = 0.088$ , the advance of periastron is even larger, with  $\dot{\omega} = 16.899323(13) \text{ deg yr}^{-1}$ . Such incredible value means that the orbit precesses over itself every 19 years<sup>59</sup>, and it leads to an incredibly precise system mass measurement of  $M_t = 2.586145(3) M_\odot$  (Kramer et al., 2021a).

### The orbital decay

In recent years, ground-based gravitational-wave detectors such as *LIGO* have become very famous due to the detection of gravitational wave merger events. The first detection in 2015 (LIGO and Virgo Collaborations, 2016) sparked the start of the brand-new field of gravitational-wave (GW) astronomy. However, the often-forgotten precursor of this revolution was the observation of the *orbital decay* due to the emission of GWs in the Hulse-Taylor binary (Taylor et al., 1979; Taylor and Weisberg, 1982), which was worth a Nobel Prize<sup>60</sup> and convinced the scientific community of their existence. According to GR, the loss of energy via the emission of gravitational radiation translates into a shrinking of the orbit, which for wide orbits (much wider than the radii of the involved stars, usually the case with neutron stars) leads to a secular decrease of the orbital period over time with a predicted value of

$$\dot{P}_b = -\frac{192}{5} T_\odot^{5/3} \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{1 + (73/24)e^2 + (37/96)e^4}{(1 - e^2)^{7/2}} \frac{M_p M_c}{M_c^{1/3}}, \quad (2.52)$$

according to Peters (1964). Once again, a precise measurement of  $\dot{P}_b$  results in a constraint of a combination of the mass components,  $M_p$  and  $M_c$ . As expected, massive systems with short  $P_b$  and large  $e$  values favour a significant measurement. For the Hulse-Taylor binary, it has been measured at  $\dot{P}_b = -2.398(4)10^{-12} \text{ s/s}$  (Weisberg and Huang, 2016), implying a shortening of 76  $\mu\text{s}$  per year. Impressively subtle, is it not? What is not subtle is the *cumulative delay* of every orbit, measured as an advance of the time of passage of periastron over time,  $T_0$ . The left side of Fig. 2.20

<sup>57</sup> In fact, this correction was one of the first experimental tests of GR: it finally solved the until-then unexplained anomalous precession of the Mercury's orbit around the Sun (Einstein, 1915b).

<sup>58</sup> Did you really think I would not mention this guy again?

<sup>59</sup> In fact, in August 18th the orbit of the double pulsar completed one full precession (360 degrees) since the first ToA measurement

<sup>60</sup> To be specific, the prize was awarded for the discovery of a lab allowing to measure GR effects, including the orbital decay.

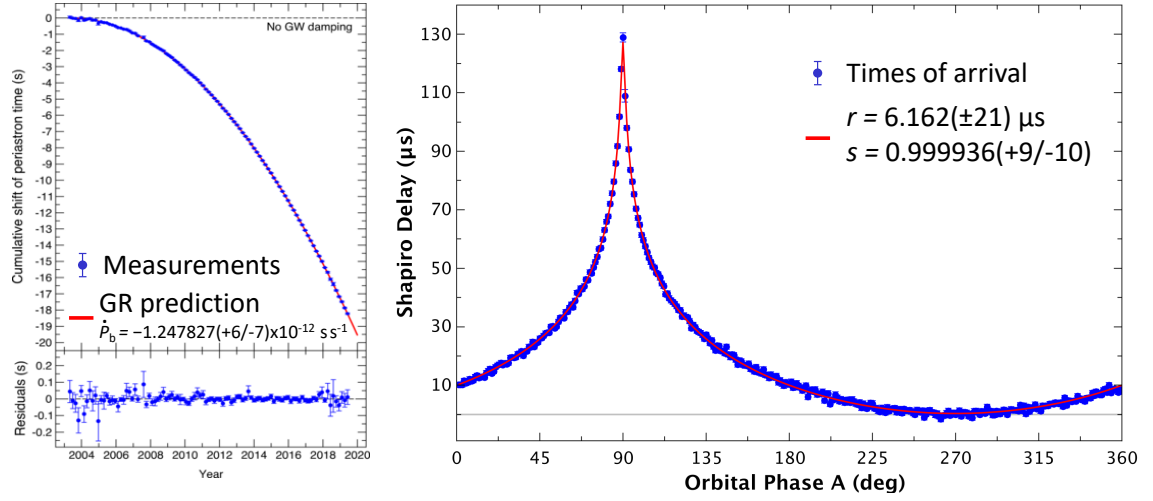


Figure 2.20: Measurement of the orbital decay and the Shapiro delay in the double pulsar system. **Left:** cumulative shift of the time of passage of periastron due to the progressive shortening of the orbit. The blue dots are Keplerian fits for  $T_0$  every 60 days, and the red line is the predicted evolution based on the mass measurements on this system. **Right:** amplitude of the Shapiro delay as a function of the orbital phase. The blue dots are the residuals drawn by the ToAs, while the red line is the Shapiro delay model provided by the  $r$  and  $s$  parameters. Plots adapted from Kramer et al. (2021a), where they were first presented.

shows this shift for the double pulsar system. After almost 20 years of timing, the orbital decay of  $\dot{P}_b = -1.247827(+6/-7) \times 10^{-12} \text{ s s}^{-1}$  manifests as an advance of  $T_0$  by 20 whole seconds.

Another consequence is the direct shrinking of the orbit, measurable via  $\dot{x}$ . That is entirely correlated with  $\dot{P}_b$  via

$$\frac{\dot{x}}{x} = \frac{2}{3} \frac{\dot{P}_b}{P_b}. \quad (2.53)$$

For instance, in the Hulse–Taylor binary it implies a physical shrinkage of 3.5 meters per year. However, such an effect is often too small to be measured with significance. In practice, measurements of  $\dot{x}$  are dominated by geometric effects induced by the proper motion of the system in the sky, and thus only  $\dot{P}_b$  is a probe of the orbital decay. And even in the case of  $P_b$ , it is often hard to disentangle between the contribution of GR and other effects, as it is explained at the end of this section.

### The Shapiro delay

I have already introduced the Shapiro delay parameters in Section 2.37. A measurement of the range and shape,  $s = T_\odot M_c$  and  $s = \sin i$  provides a direct constraint on the companion mass and the inclination angle. But we have to remember that, thanks to the mass function from equation (1.21),  $M_p$ ,  $M_c$  and  $i$  are not independent, so a constraint on  $i$  also implies a constrain in the  $M_p$ – $M_c$  space and *vice-versa*. However, I have yet to explicitly introduce the dependence of the orthometric Shapiro delay parameters  $h_3$  and  $\varsigma$  on  $M_c$  and  $i$ , so here they are according to the definitions from Freire and Wex (2010):

$$h_3 = r \varsigma^3 = T_\odot M_c \left( \frac{1 - \cos i}{1 + \cos i} \right)^{3/2} \quad \text{and} \quad \varsigma = \left( \frac{1 - \cos i}{1 + \cos i} \right)^{1/2}. \quad (2.54)$$

Once again, they provide constraints on  $M_c$  and  $i$ , so the same story as the case of  $r$  and  $s$  applies, but with slightly more complicated mathematics. Also, it is quite straightforward to change from one description to the other.

The strongest Shapiro delay ever measured comes from the Double pulsar system, which has an impressively lucky  $i = 89.19$  deg, meaning that we see the system almost entirely edge-on. With this orbital orientation, the radio pulses of pulsar A go as close as 11,000 km from the surface of pulsar B (a distance shorter than Earth's diameter!). With a range  $r = 6.162$  s, this means that the full amplitude of the Shapiro delay is of 130  $\mu$  s at superior conjunction (Kramer et al., 2021a), as is shown on the right side of Fig. 2.20.

### The Einstein delay

The Einstein delay is measurable in eccentric pulsar systems as a sinusoidal modulation of the timing residuals in  $E$ , as expressed in equation (2.42). The amplitude of this delay is our PK parameter of interest, which according to Blandford and Teukolsky (1976) has a GR-predicted value of

$$\gamma_E = T_\odot^{2/3} \left( \frac{P_b}{2\pi} \right)^{1/3} \frac{M_c (M_c + M_t)}{M_t^{4/3}} e. \quad (2.55)$$

As seen here, only if the system has some eccentricity does the Einstein delay become measurable. Ironically, and according to the equation above, this is the only parameter that becomes easier to measure in systems with larger  $P_b$  values. However, in practice, it suffers from the same problem as the Shapiro delay: it is usually absorbed into the Rømer delay due to its modulation by  $\sin E$ . But not all hope is lost, because the degeneracy is broken for systems with large  $\dot{\omega}$  values, as the shape of the Rømer delay is affected by the precession of the orbit while the Einstein delay remains unchanged. Therefore, in reality compact systems with large  $e$  values are the best candidates for a measurement of  $\gamma_E$ . In the Hulse–Taylor binary, this delay has an amplitude of  $\gamma_E = 4.407(4)$  ms, while in the double pulsar it is  $\gamma_E = 0.384$  ms, which compared to their respective spin period of 59 ms and 23 ms, is very measurable (Weisberg and Huang, 2016; Kramer et al., 2021a).

### Contamination from the Galactic field and the proper motion

The Solar System and the pulsar system do not exist in the vacuum, but instead they orbit around the center of the Milky Way and have even have their own peculiar motion due to, for example, supernova kicks in the pulsar system. This results in two effects: a gradual change in the apparent orbital orientation of the system, and a changing Doppler effect due to apparent and real accelerations.

We have already seen several times that the Doppler effect modulates the observed spin period of a pulsar. However, outside of the Doppler correction caused by the orbital motions of the pulsar and the Earth, there is also a *global Doppler factor* due to the relative velocities between the Solar System and the pulsar system,  $D$ . This global Doppler correction is not measurable, and therefore we do not worry much about it, but *its time derivative*,  $\dot{D}$ , does indeed cause some effects measurable with pulsar timing, changing any length or periodicity,  $L$ , according to

$$\frac{\dot{L}}{L} = - \left( \frac{\dot{D}}{D} \right)_{\text{Shkl}} - \left( \frac{\dot{D}}{D} \right)_{\text{Gal}} = \frac{1}{c} \left[ |\vec{\mu}|^2 d + \vec{K} (\vec{a}_{\text{PSR}} - \vec{a}_{\text{SSB}}) \right], \quad (2.56)$$

where  $|\vec{\mu}|$  is the magnitude of the proper motion vector,  $d$  is the system distance,  $\vec{K}$  is the SSB-to-pulsar system unit vector, and  $\vec{a}_{\text{PSR}}$  and  $\vec{a}_{\text{SSB}}$  are the Galactic acceleration field measured at the pulsar system and the SSB (or a nearby mass, such as a third stellar component in a very wide orbit like in Grunthal et al. (2024)), respectively. The first term describes the Shklovskii effect (Shklovskii, 1970), which accounts for the fact that while pulsar proper motion is measured on the surface of the celestial sphere, its real velocity lies on a three-dimensional space. This means that as the system moves away from the LOS, the velocity components projected on the celestial sphere gradually spills on to the radial velocity component, increasing the Doppler factor over time. This effect is most significant in nearby pulsars with large  $|\vec{\mu}|^2 d$  factors (given a transverse velocity on the sky  $v_T$ ,  $|\vec{\mu}| = v_T/d$ , so the factor decreases with distance). The second term describes the effect the relative acceleration between the SSB and the pulsar system due to the gravity field in the Milky way, which also leads to a derivative of the Doppler factor (Damour and Taylor, 1991).

These two effects contribute to the measurement of  $\dot{P}_s$ ,  $\dot{P}_b$  and  $\dot{x}$ . In fact, the Doppler derivative contributions to  $\dot{P}_b$  may even dominate over the GR contribution. Therefore, it must be corrected for if we want to know their intrinsic values. Unfortunately, the system distance  $d$  and the Galactic acceleration field can only be estimated and usually are not known with certainty. The distance may be known from parallax measurements, but for the Galactic acceleration field we always have to rely on models of matter distribution on the Milky Way, drawn from large surveys of stellar proper motions in the Galaxy (e.g. McMillan, 2017).

In the case of  $\dot{x}$ , there is another contribution that even dominates over the Doppler effect derivative. That contribution also exists for  $\dot{\omega}$ , it being the geometric proper motion effect described by Arzoumanian et al. (1996) and Kopeikin (1996):

$$\frac{\dot{x}}{x} = 1.54 \times 10^{-16} \text{ s}^{-1} \cot i \left( -\frac{\mu_{\text{RA}}}{\text{mas yr}^{-1}} \sin \Omega_a + \frac{\mu_{\text{DEC}}}{\text{mas yr}^{-1}} \cos \Omega_a \right) \text{ and} \quad (2.57)$$

$$\dot{\omega} = 2.78 \times 10^{-7} \text{ deg yr}^{-1} \csc i \left( \frac{\mu_{\text{RA}}}{\text{mas yr}^{-1}} \cos \Omega_a + \frac{\mu_{\text{DEC}}}{\text{mas yr}^{-1}} \sin \Omega_a \right), \quad (2.58)$$

where  $\Omega_a$  is the longitude of the ascending node, a newly introduced angle that defines the orientation of the system. If  $i$  defines the inclination of the orbital plane and  $\omega$  the orientation of the orbit within the orbital plane,  $\Omega_a$  defines the global orientation as the angle between the East direction in the sky and the point of the ascending node (see Appendix A for a visual definition). These effects account for the changing perspective with which we see the orbital plane due to the system moving away from the LOS, and of course it can be positive or negative depending on its orientation. In the case of  $\dot{\omega}$ , in the vast majority of systems where it is measurable the GR contribution is orders of magnitude larger than the proper motion contribution. However,  $\dot{x}$  is dominated by the proper motion contribution, and its measurement, along with that of  $|\vec{\mu}|$  and a constraint on  $i$ , allows for a significant constraint on the value of  $\Omega_a$ .

#### 2.4.8 Mass measurements II: getting proper mass constraints

Now that we know how to relate the measured PK to the component masses in the system, it is time to actually constrain the masses. This is what we all have been waiting for, and the last section of this Chapter. Obtaining good measurements of  $\dot{\omega}$ ,  $\gamma_E$ , any combination of Shapiro delay parameters and

perhaps even a corrected and significant  $\dot{P}_b$  values results in independent constraints on the  $M_c$ – $M_p$  space, but we want to put it all together. As is shown in Fig. 2.21, if we project everything on to the  $M_p$ – $M_c$  we realise that each measured PK parameters draws a curve based on the expressions presented in Section 2.4.7. If GR remains correct and we have ensured that our measurements are performed properly, all of these curves must converge into a single point, which tells us the values of the mass components in the system. That is illustrated in Fig. 2.21, which shows the  $M_c$ – $M_p$  diagrams for the Hulse–Taylor binary and double pulsar system, where the PK parameters converge into the points of  $M_p = 1.438 \pm 0.001 M_\odot$  and  $M_c = 1.390 \pm 0.001 M_\odot$  (Weisberg and Huang, 2016), and  $M_A = 1.338185(+12/-14) M_\odot$  and  $M_B = 1.248868(+13/-11) M_\odot$  (Kramer et al., 2021a).




Figure 2.21: Example  $M_p$ – $M_c$  diagrams from PK parameters from the Hulse–Taylor pulsar binary **PSR B1913+16** (left) and the double pulsar system **PSR J0737–3039A/B** (right) and from the measurement of the GR-derived PK parameters. For the double pulsar system, the PK parameters have been measured with pulsar A, except for the mass ratio which has been measured from both pulsar A and pulsar B ( $R = x_A/x_B = m_B/m_A$ ) and the geodetic recession of pulsar B,  $\Omega_B^{\text{geod}}$ , which has been measured from eclipses of of pulsar A from B’s magnetosphere (Breton et al., 2008; Lower et al., 2024). Figures adapted from Freire and Wex (2024) by Norbert Wex.

But, how do we achieve the uncertainties in the mass measurements? Well, they are derived from the uncertainties of the PK parameters. The PK parameters do not draw a single curve, but a range of possible curves within the PK value ranges given by their uncertainties, so the crossing of PK parameters actually creates a region of compatibility with observations within the  $M_c$ – $M_p$  space. In more precise terms, each PK parameter  $p_j$  provides a Gaussian probability density distribution in the  $M_c$ – $M_p$  space

$$\rho_j(M_c, M_p) \propto \exp \left[ -\frac{\left( p_{j,\text{best}} - p_j(M_c, M_p) \right)^2}{2\sigma_j^2} \right], \quad (2.59)$$

where  $j$  labels the PK parameters,  $p_{j,\text{best}}$  is the best value at  $\chi_{\min}^2$  and the functions  $p_j(M_c, M_p)$  are the ones listed in Section 2.4.7. From these constraints, the final probability distribution in  $M_c$ – $M_p$  is

the product of the probability distributions<sup>61</sup> from all measurements

$$\rho(M_c, M_p) = \prod_j \rho_j(M_c, M_p), \quad (2.60)$$

which we can integrate across  $M_c$  or  $M_p$  to achieve individual one-dimensional probability distributions for the values of  $M_p$ ,  $M_c$  and  $i$ . Therefore, at least two PK measurements are needed to achieve component mass measurements, and with three or more the system becomes over-determined and the constraints will come from the combination of the two best measurements.

However, as I have mentioned earlier, it can happen that the timing parameters are degenerate amongst each other and that  $\chi^2$  is not symmetric. That is especially true when some PK parameters are measured to low significance, which often results in uneven distributions of  $\chi^2$  and thus in non-Gaussian uncertainties. To circumvent this, instead of measuring PK parameters as independent timing parameters, we assume that GR is true from the start and model all the terms in equation (2.37) as direct functions of the Keplerian orbital parameters,  $M_c$ , and  $M_p$ , essentially giving the PK parameters their GR-predicted values. That is implemented by the DDGR timing model (Taylor and Weisberg, 1989), a modified version of the DD timing model that does exactly that with the exception of  $\dot{x}$  and  $\dot{P}_b$ , which are also allowed to take independent excess values to account for the Doppler derivative and proper motion contributions. From here, we derive probability density distribution in the  $M_c$ – $M_p$  space directly from the distribution of  $\chi^2$  values,

$$\rho(M_c, M_p) \propto \exp \left[ -\frac{\chi^2(M_c, M_p) - \chi_{\min}^2}{2} \right], \quad (2.61)$$

accounting for any possible asymmetries in the  $\chi^2$  distribution in the PK effects. That results in a more accurate description of the allowed regions in the  $M_c$ – $M_p$  space and more accurate probability distributions for  $M_c$ ,  $M_p$  and  $i$ , even in cases where individual PK parameters are not measured with a large significance (see for example Splaver et al., 2002).

And with that, we have all that is needed to see the amazing science results presented in this thesis. Now we move onto the direct application of all of these methods in three different science projects: the MPIfR-MeerKAT Galactic Plane Survey (MMGPS), the follow-up, solving and timing of the MMGPS DNS discovery **PSR J1208–5936**, and the mass measurements in the massive **PRP–ONeMg WD** system **PSR J1227–6208**.

---

<sup>61</sup> I am a bit sneaky and I have introduced some notions of Bayesian statistics here. Frequentist statistics explore the parameter space to find the best-fit solution, while Bayesian computes the probability of a parameter combination being the correct one, and then compute the most likely values from there. For mass measurements, we often use the later.

---

## New Galactic discoveries with MeerKAT: searching, finding, and solving.

---

### 3.1 Before we start

The MeerKAT telescope, introduced in Section 2.2.3, is the most sensitive pulsar search facility in the Southern Hemisphere. Having observed the sky since 2019, it has already yielded **hundreds of new discoveries**. Many of them come from targeted searches in Globular Clusters (e.g. Ridolfi et al., 2022), unidentified Galactic Gamma-ray sources (e.g. Clark et al., 2023), Galactic supernova remnants (e.g. Turner et al., 2024) and from neighbouring dwarf galaxies like the Magellanic Clouds (e.g. Carli et al., 2024; Prayag et al., 2024). All of these projects are part of the **Transients and Pulsars with MeerKAT (TRAPUM)** Stappers and Kramer, 2016), which encompasses many science topics within pulsar astronomy. However, the TRAPUM projects I have been the most involved with, and the one this chapter is mostly about, are the **Max-Planck-Institut für Radioastronomie (MPIfR)-MeerKAT Galactic Plane surveys (MMGPS)**, Padmanabh et al., 2023). This project is a collaborative effort that had been in planning a decade before MeerKAT was even built, and I was lucky to join when the MMGPS search pipeline had reached and almost definitive state. In a way, I truly am standing on the shoulders of giants<sup>62</sup>

The main focus of this chapter is to review the science case of the MMGPS, its implementation, and its science results. I start with a section describing the MMGPS pipelines, which includes my contributions to said pipeline and our *modus operandi* during the surveys. It uses many of the pulsar searching methods introduced in Section 2.3, but here they are applied in the specific conditions of the survey. The section after that one is all about pulsar discoveries in the MMGPS surveys and a significant taste of science results. Both isolated and binary pulsars discoveries I have worked with are described in great detail, borrowing from many of the timing methods described in Section. 2.4. Finally, a small side-project I led in during the testing of the MeerKAT S-band receivers is described: the search for a pulsar in the supernova remnant **SNR 340.6+03**, and in fact the first pulsar search with the S-band receivers.

---

<sup>62</sup> In science, it is all giants standing on the shoulders of giants all the way down.



### 3.2 The MPIfR-MeerKAT Galactic Plane survey

Table 3.1: The several surveys of the MMGPS. Every survey is performed with the dedicated MeerKAT receivers: UHF, L-band, and S-band.

Survey	Duration (hours)	Latitude range (deg)	Longitude range (deg)	Time/Observation	Frequency (MHz)
MMGPS-L	800	$ b_G  < 5.2$	$-100 < l_G/\text{deg} < -10$	10.6 mins	856 – 1712
MMGPS-S	1380	$ b_G  < 1.5$	$-80 < l_G/\text{deg} < 15$	21.2 mins	1968 – 2843
MMGPS-UHF	400	$ b_G  < 11$	$-62 < l_G/\text{deg} < 15$	8.4 mins	544 – 1088
MMGPS-Sgr A*	200	$ b_G  = -0.05$	$l_G = -0.04$	22 hrs	1968 – 3500




Figure 3.1: Sky coverage of the several MMGPS surveys, whose observing regions are highlighted as dashed areas. The sky map is presented in Galactic coordinates, with the Galactic center at  $(b_G, l_G) \approx (0, 0)$  deg and the Galactic plane following  $b_G \approx 0$  deg. The regions not visible from the FAST and GBT telescopes are highlighted by the green shaded closed curves. The area covered by the FAST Galactic Plane Survey Snapshot (GPPS, Han et al., 2021) is also drawn for reference. Plot originally shown in Padmanabh et al. (2023).

#### 3.2.1 Science objectives: on the hunt for faint binary pulsars

The MMGPS are a family of comensal surveys that encompass imaging, line emission detection, and pulsar searching from the same set of telescope observations (Padmanabh et al., 2023). Of these, I have only been involved with the pulsar search aspect. In this regard, the aim of the MMGPS is to discover faint and/or distant binary pulsars in the southern Galactic plane that may have been overlooked by previous Murriyang surveys in the same sky region, helping us investigate the open questions of

binary evolution discussed in Chapter 1. In addition, the discoveries will improve population models, enhance the sensitivity of PTA arrays with the discovery of new highly stable MSPs, and increase the sample of intermittently emitting/nulling pulsars (Padmanabh et al., 2023).

To achieve these aims, four different surveys were proposed, all of which are listed in Table 3.1 (Padmanabh et al., 2023). The first of these surveys, the MMGPS at L-band (MMGPS-L) started in late 2020 and was concluded in late 2022. The second one, the MMGPS-S at S-band, started in early 2022 and it is currently ongoing. The UHF-band survey (MMGPS-UHF) started in September 2024, and the Galactic center survey (MMGPS-Sgr A\*) has also begun observations. With 10-min long observations, the MMGPS-L survey was four times more sensitive than the southern High Time Resolution Universe survey at low latitude survey performed with Parkes/Murriyang at the same frequency and sky region (HTRU-low Keith et al., 2010). The currently on-going MMGPS-S covers a narrower latitude range and implements a longer integration time of 20 minutes, seeking to find distant high-DM pulsars in the deep Galactic plane, where pulsars are too dispersed or scattered to be detected at lower frequencies. In this regard, the MMGPS-S is the deepest search ever performed at  $f > 2000$  MHz in this sky region, untapping an unexplored side of the pulsar population. The MMGPS-UHF will do the opposite: it will search for nearby and bright low-DM pulsars in a wider Galactic latitude range, a potential potential pulsar treasure trove. And finally, the only targeted survey is the MMGPS-Sgr A\*, which will accumulate 200 hours of observations of the Galactic center at  $f > 2000$  MHz, aiming to find the elusive pulsars in the closest vicinity of the supermassive black hole at the center of our Galaxy. As is depicted in Fig. 3.1, each of these surveys cover different but overlapping regions of the Galactic plane. Most of these regions are not observable by powerful radio telescopes in the Northern Hemisphere, such as the Green Bank Telescope (GBT) in West Virginia, USA, and the Five-hundred-meter Aperture Spherical Telescope (FAST) in Guizhou, China. Therefore, they aimed at exploring new territory from the start.

### 3.2.2 The MMGPS pipeline: from observations to discoveries

The MMGPS pipeline, developed by the MMGPS members before my arrival, implements a FFT-based pulsar search with a time-domain acceleration search on *quasi-real-time* processing, with emphasis on “quasi”. Real-time processing would imply that no data is stored and that the search is done during the observation, but we can not reach that level just yet (it would be cool if we did). In more traditional surveys (Manchester et al., 2001; Keith et al., 2010, e.g. the PMPS or the HTRU), the data is kept for long-term storage, meaning that processing could take long time spans, and that the data could even be archived and re-processed with improved search techniques years after the observations took place to scrape the bottom of the  $S/N$  distribution. But unlike in more traditional surveys, there is no long-term data storage in the MMGPS. Instead, after every observing batch, beam-formed data is temporarily stored in filterbank format at the Accelerated Pulsar Search User-Supplied Equipment (APSUSE) in the MeerKAT site, searched, inspected for pulsar candidates, and then *deleted*. Given a dynamic range of 8 bits, a time sampling resolution of 153  $\mu$ s, 1024 frequency channels, 10-min long observations with 480 coherent beams, 20 pointings per session and two observing sessions per week, this meant that at the peak of MMGPS-L processing we were observing, storing, processing, and burning **350 TB of data per week**.

We did not do it this way out of enjoyment, but out of memory space limitations in APSUSE. This also means that the data had to be processed *fast*. With coherent 480 beams per observation, searching is highly highly parallelised. A side effect of that was that we had to deal with candidates from  $480 \times 20$

beams per session, where each beam can give potentially thousands of candidates. To say it was quite an effort would be an understatement. So, in this section I review how we discovered 80+ new pulsars step by step. We start with the observing plan: how were observations performed and how were beams distributed in the sky. But we do not linger much on it. Instead, we move onto the search pipeline, RFI treatment, and finally on to the managing of pulsar candidates. So, without further ado, let us begin.

### An observing plan with many, many beams




Figure 3.2: Tiling of primary and coherent beams in the MMGPS-L for seven neighbouring pointings. The black dashed circles represent the  $\theta_{\text{FWHM}}/\sqrt{5}$  radii of primary beams (88% power drop with respect to maximum), each of them corresponding to a different observation. Blue ellipses represent the coherent beams computed by FBUSE, with an overlap factor of 0.5 on average. Plot originally shown in Padmanabh et al. (2023).

Every survey begins with an observing plan. That means asking: “How many observations are required to cover an area and for how long?”. For the MMGPS, the choice of spacing between the central positions of primary beams of  $2\theta_{\text{FWHM}}/\sqrt{5}$  (overlap at 88% power), equivalent to 0.4 deg for the MMGPS-L, 0.2 deg for the MMGPS-S and 0.6 deg for the MMGPS-UHF according to equation (2.3). As is shown in Fig. 3.2, the primary beams are distributed with an hexagonal tiling, ensuring an efficient overlap where there is minimal variation in telescope gain on the overlap points at the edges of beam<sup>63</sup>. For convenience, we call an observation of a spot in the sky for a prolonged amount of

<sup>63</sup> Nature has known long before us that hexagonal pieces are the most efficient at covering an area. That is why *honeycombs* or *basalt pillars* have hexagonal tiling schemes, and now telescope pointings as well. Therefore, we should also stack toilet paper rolls in hexagonal tilings.

time a *pointing*. In general, several pointings with neighbouring sky positions are performed in a row during an observing session. In the case of MMGPS-L, a total of 4,140 pointings were performed, while for the MMGPS-S 3,620 pointings are planned, of which 21% are complete as of the writing of this sentence<sup>64</sup>.

These observations are processed by the Filterbanking Beamformer User-Supplied Equipment (FBUSE) which, as its name suggests, uses the interferometric nature of the telescope to perform multiple beamforming operations in real-time. This, the survey beam is always filled with a tiling of 480 coherent beams. In the first 20% of pointings of the MMGPS-L, the coherent beams were distributed within the primary beams with an overlap factor of 0.5. However, the size and shape of coherent beams is highly dependent on the used antenna array and its orientation with respect to the source on the sky, meaning that the covered area within the survey beam was inconsistent across pointing. This was corrected by implementing a fixed hexagonal coherent beam tiling as is shown in Fig. 3.2, ensuring no gaps in coverage within and between primary beams (Padmanabh et al., 2023). This makes the overlap factor variable across pointings, but in general it averages at 0.5. The tiling is computed by the fittingly-named *Mosaic* software (Chen et al., 2021b) for every observation, and then each beam is recorded on APSUSE in a filterbank format with a sampling time of 153  $\mu$ s and 1,024 frequency channels. For each pointing, 481 beams are recorded: 480 coherent ones from interferometry, and an incoherent beam from the incoherent addition of all antenna signals.

### The search pipeline

The MMGPS implements an FFT-based, time-domain re-sampling acceleration search on the filterbanks stored in APSUSE, which has several available GPU processing nodes, ideal for processing multiple beams in parallel<sup>65</sup>. The searched DM ranges goes from 0 to 3000 to  $\text{pc cm}^{-3}$ , with a variable step that increases with increasing DM, while the searched accelerations range from  $-50$  to  $50 \text{ m s}^{-2}$ . The DM step plan is generated by *DDplan.py* from *PRESTO* (Ransom, 2011), but the search itself is performed with the *peasoup* GPU-based pipeline (Morello et al., 2019; Barr, 2020). The incoherent de-dispersion aspect of the search is handled by the *dedisp* GPU-based library (Barsdell et al., 2012).

The *peasoup* pipeline implements an FFT periodicity search at every DM value, acceleration acceleration value and every beam, including the incoherent beam. De-reddening is performed in the Fourier on the Fourier domain, and eight harmonic summing operations are performed before searching for signals in the FFT. The end results of the pipeline are *candidate files*, which contain the topocentric spin period, DM, and acceleration values of identified signals in the FFT, as well as the  $S/N$  of the signal the Fourier after eight iterations of harmonic summing. For further inspection, a maximum of 1,000 candidates with  $S/N > 8.5$  are chosen from every beam.

### RFI mitigation

Ah, the nightmare of radio astronomy, here we go again. In the MMGPS, we perform RFI mitigation on the filterbank before the search, and during the filtering of *peasoup* candidates. In fact, RFI

<sup>64</sup> August 27, 2024. I am so sad to be missing out on les Festes de Gràcia right now, but I will enjoy les Festes de Santa Tecla de Tarragona on September 23.

<sup>65</sup> GPU processors extremely efficient at running codes in parallel, which is convenient when you have 480 coherent beams to process per observation. As a fun fact, the first GPU processors were developed to run video games with high graphical demands

mitigation in the filterbanks was one of my early contributions to the project. However, my RFI scheme turned out to be sub-optimal (see Section 3.2.3 to read about it) and most of its components were displaced in favour of the `filtool` program from `PulsarX`, a MMGPS custom-made software (Men et al., 2023). `filtool` zero-weights outlying pixels in the time–frequency space based on several consecutive filters, and produces a new, clean filterbanks ready to be searched. The first filter is a Skewness-kurtosis<sup>66</sup> analysis applied on every frequency channel. Frequency channels which are clear outliers within a given time window are removed. Then, a zero-DM filter is applied, mitigating broad-band signals present at  $DM = 0 \text{ pc cm}^{-3}$ . Finally, a Kadane filter<sup>67</sup> is applied to identify long-duration signals in the time series and remove them. This way, many RFI signals are not picked up by `peasoup` in the first place.

The second stage of RFI filtering occurs during the selection of `peasoup` candidates from the same pointing with the `candidate_filter` package. This package identifies candidates which may be related to one another by harmonic, DM or acceleration, thus coming from the same signal. Once a single common origin of these signals has been identified, the spatial distribution of Fourier  $S/N$  across coherent beam positions is modelled with an exponential drop from the central position, both to identify the best coherent beam for each candidate and to distinguish astrophysical signals from RFI. This step ends up reducing the number of candidates by a factor of three or four, leaving only a few thousands per pointing, which are sent to the next step.

### Folding and automatic classification

Once the DM, acceleration and fundamental spin period of each candidate has been identified, the filterbank from the beam with the best Fourier  $S/N$  is folded with the best parameters with the Yunpeng Men’s custom-made `PulsarX` software. This software performs the optimisation procedure described in Section 2.3.6 and returns optimized values for the spin, DM and acceleration of the pulsar. The optimised pulsar archive is then plotted with the `PSRCHIVE` software (Hotan et al., 2004; van Straten et al., 2012), resulting in plots of frequency–spin phase, time–spin phase, intensity–spin phase, and DM optimisation curves ready for inspection by observers. Many of these candidates will be RFI or random noise that has piled up into a signal by chance, but some will be pulsars, even if they are known ones. And amongst the flood of candidates, we may get lucky and discover a new pulsar!

But here we face an issue. This pipeline results in thousands of candidates for a single pointing. Given that at the peak of MMGPS-L processing we performed 20 pointings per observing session, with two observing sessions per week, this meant that we would have to inspect tens of thousands of pulsar candidates every week. Do you see yourself looking at tenths of thousands of images every week? Yeah, me neither. Maybe it is fine for a week or two, but after a couple months we would all be wanting to quit our careers and become poetry writers. So, instead, we make an AI look at them for us. We use the Pulsar Image-based Classification System (`PICS`), a machine learning algorithm that implements a neural network to mimic the behaviour of expert human observers (Zhu et al., 2014)<sup>68</sup>.

<sup>66</sup> Skewness and kurtosis are a measure of deviation from Gaussianity in the data. High skewness means that the noise distribution is asymmetric, and a value kurtosis deviating from 0 means that the data is distributed with a fat or sharp-peaked curve.

<sup>67</sup> The Kadane algorithm finds the largest contiguous summation within a one-dimensional array (here the time series within each frequency channels).

<sup>68</sup> Neural networks can be trained to imitate humans on many fields. For instance, they can be taught to complete a `Super Mario World` level.



The PICS classifier was trained with two sets of pulsar candidates plots already classified by humans: first pulsar candidates from the now-defunct [Arecibo telescope](#) (the PALFA survey Cordes et al., 2006), and then with a set of TRAPUM candidates from earlier MeerKAT surveys. From this training, upon inspecting every new pulsar candidate, the model now gives a score between 0 and 1, where 1 stands for “100% a pulsar” and 0 for “100% not a pulsar”.

The filtered and folded [peasoup](#) candidates are thus filtered once again by inspection of their folded plots with PICS. Subsequently, if the PICS score goes above 0.1 for a particular candidate, it is let through and shown to the human observers. This step filters out 90% of remaining candidates.

### Inspect and delete!

Now comes the most fun part, *human inspection, classification and the deletion of data*. For that, batches of candidate plots from each observing session are distributed amongst the many MMGPS members. We then classify the thousands of pulsar candidates from each pointing into five categories: RFI, noise, tier-two candidates (T2), tier-one candidates (T1), known pulsars, and non-boresight pulsars. Previously discovered pulsars are classified as known pulsars, and we are bound to find all them in the sky<sup>69</sup>. Non-boresight pulsars are known bright pulsars, but their positron is outside of the primary beam, meaning that they are being picked up by side-lobes and thus should be ignored. The T2 classification is given to “maybe this is a new pulsar, but it is low  $S/N$  and it could also noise, but we should probably keep it just in case” kind of candidates, while the T1 classification is given to “I found a new pulsar!” kind of candidates.

By the end of the inspection, the objective is to delete the filterbanks of coherent beams not containing any known pulsars, T2 or T1 candidates, so that space on APSUSE is cleared and the next batch of observations can be stored in. The most repeated sentence by the survey coordinator Ewan Barr was “Inspect these candidates and send back the classifications by the end of the week so we can make new observations!”. After that, the sky positions of T2 and T1 candidates are noted down and scheduled re-observation to confirm the existence of the pulsars.

Here are some details on how this classification worked. The candidates selected by PICS were folded with [PulsarX](#), and the classification process was automatised via the implementation of [CandyJar](#), a user interface for candidate viewing, classification and saving of the candidate classification developed by Vivek Venkatraman Krishnan. Fig. 3.3 shows what this interface looks like for in my tiny laptop screen, along with an RFI and noise candidate. After the initial classification of candidates by all the observers, I was in charge of collecting all T2 candidates and to organize a group inspection of to either move them to the T1 category, or to demote them to RFI or noise. These were basically 20-min long sessions after the weekly MMGPS meeting on Friday, and the main reason why we almost skipped lunch every Friday. The decision was made via voting: if at least three inspectors agreed for it to be upgrade a candidate to T1, it was scheduled for re-observation. This effort was not in vain, because a handful of T2 candidates were indeed re-observed and confirmed as discoveries.

### 3.2.3 My contribution to RFI excision during search

As stated in Section 3.2.2, [PulsarX](#) is the current-use software for RFI-excision before running [peasoup](#). But on the earlier times of the survey, we experimented with different approaches of RFI

<sup>69</sup> My eternal gratitude to David Kaplan for the [Pulsar Survey Scraper](#), which searches pulsars compatible with input DM and sky position across several data bases. It made our life much easier.

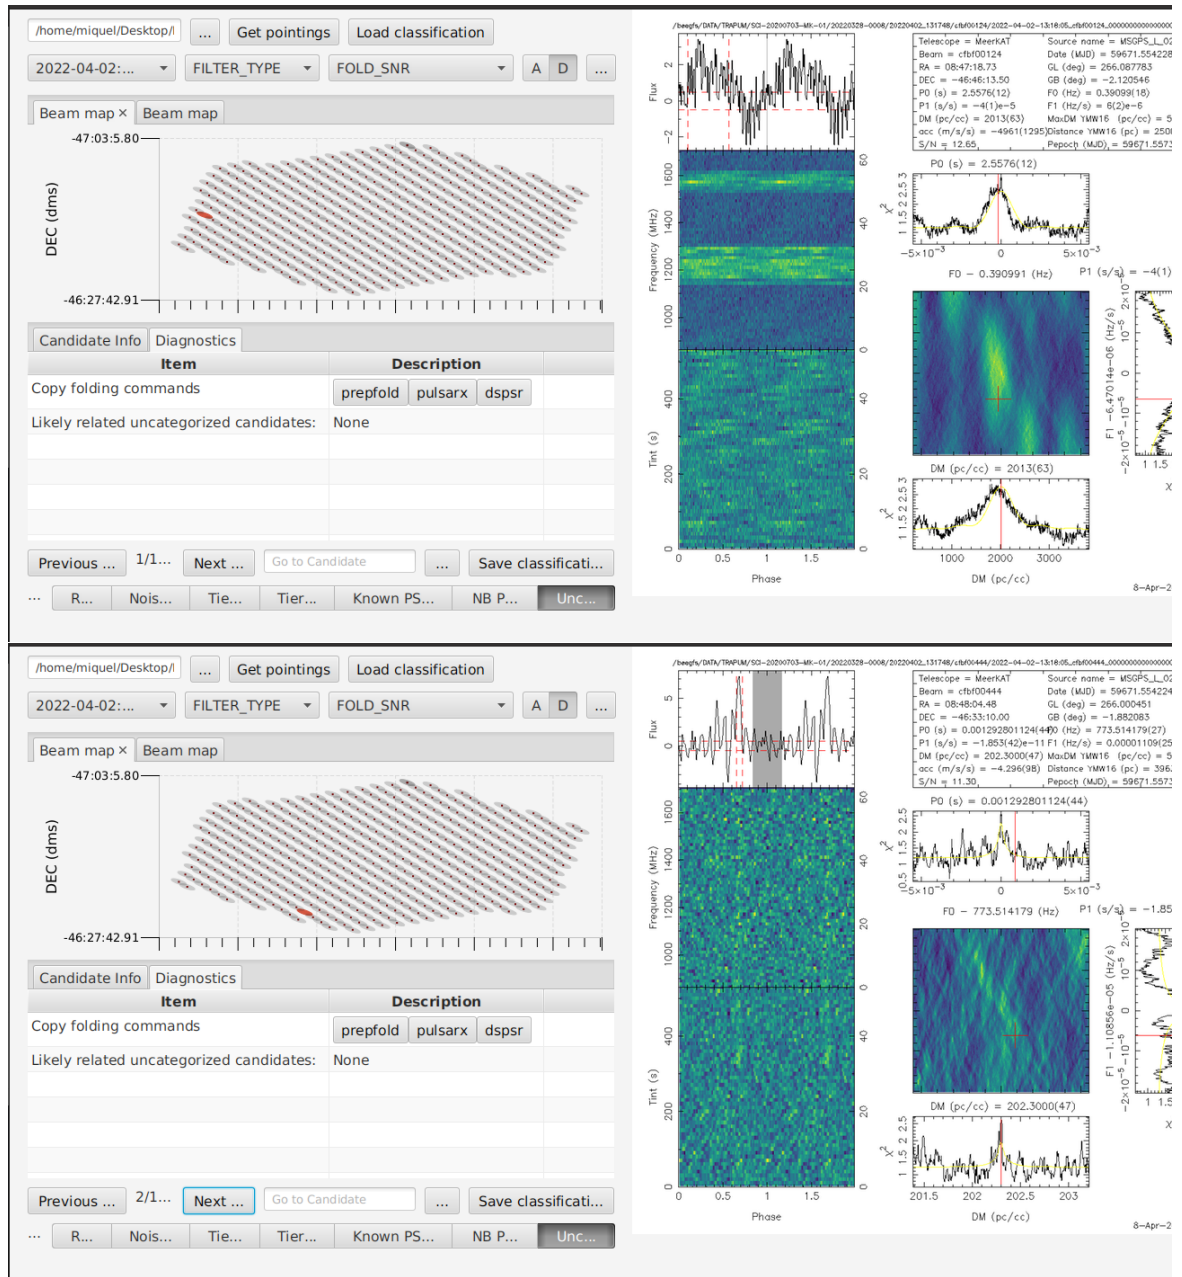


Figure 3.3: MMGPS-L pulsar candidates folded with **PulsarX** in the **CandyJar** user interface. The top plot shows a typical RFI candidate, while the bottom plot shows a typical noise candidate. The image is not trimmed, it just happens that my laptop screen is too small to show the entire candidate, but it does its job nonetheless. The top-left of each candidate shows the location of the corresponding coherent beam, the bottom-left gives us information about the candidate, including a comparison of its spin and DM parameters with those of nearby known pulsars, and the right shows the candidate plot. We have seen similar screens thousands upon thousands of times.



excision. The key factor in our experimental approaches was the incorporation of the interferometric and multibeam nature of the survey, which was a first in pulsar astronomy. Therefore, we had to try what had not been tried before. In this context, and under the guidance of the MMGPS team, I developed a multibeam approach to derive frequency masks and zero-DM Fourier frequency lists. The frequency masks were derived with `rfifind` from `PRESTO`, a dedicated pulsar search software (Ransom, 2011), while the Fourier frequency lists were achieved with custom software. The program is called `multiTRAPUM`, and it is coded with a mixture of `shell/bash` and `python` scripts. In here, I describe the masking approach, show why it was sub-optimal, and share important lessons learned from the experience.

### Mask your RFI with `rfifind`

The program `rfifind` generates frequency masks, which are used by a pulsar search software to zero-weight frequency channels with an abundance of RFI signals. That is implemented by ignoring the selected channels during de-dispersion, so that the signals never make it to the FFT. The frequency masks are derived directly from the filterbank: every frequency channel is divided into time ranges a few seconds long, and then the maximum power, power spread and mean power are computed for every interval. Then, for every frequency channel, the fraction of time intervals in which any of these three quantities go over a significance threshold is computed. If this fraction is larger than a user-set value, the channels is masked.

### The multibeam approach with `rfifind`

Implementing `rfifind` mask in single-dish surveys is straightforward, but in the MMGPS we have 481 beams per observation. The MMGPS pipeline is highly parallelised, meaning that a single run call is made for the 481 beams of each pointing, so producing a mask from every beam is not a viable strategy. Therefore, the aim of my project was to make a single frequency mask for all of the beams in the pointing. The naive solution is to make a mask from just one coherent beam. After all, because coherent beams are formed by adding signals from all of the antennas, RFI is brought into all of the beams even if it affects only one antenna. However, that is not true for localised RFI sources such as satellites or other artificial emitters located a distance from the antenna array. The blue-marked channel masks displayed in Fig. 3.4 correspond masks derived from real filterbanks from coherent beams at different corners of the beam tiling, and it shows subtle differences across them, with a small fraction of signals flagged for masking in only some of the beams.

Therefore, our aim became to produce a mask that could account for information from all of the beams, so we implemented the scheme shown the bottom of Fig. 3.4: select a handful of coherent beams across the tiling, splice their filterbanks in segments of equal time duration, select contiguous slices and join them into a single *joined filterbank*, and then perform a single `rfifind` run on it. Key emphasis on the *single run*: `rfifind` is not exactly fast, so having to run it only once per pointing saves both computing resources and time. The resulting mask is shown as the green-marked channels in Fig. 3.4, and it is made evident that it is the most comprehensive mask that picks up signals from all beams. The splitting and joining of filterbank pieces was implemented with Ewan Barr's `sigpyproc`, a `python`-based pulsar search data manipulation package. In addition, around the same time we started experimenting with `filtool`, we included it in the pipeline before slicing the filterbanks, so that we would not over-mask any frequency channels that were already cleaned by it.




Figure 3.4: Production of a frequency channel mask for a pointing with multiTRAPUM. Usually seven beams are chosen from the coherent beam tiling: the central one and a corona of six from the edge. These beams are then sliced every few seconds with sigpyproc, and then selected slices are joined into a single beam. Then rfifind is run on the joined beam to produce a comprehensive channel mask. The masks marked in blue are derived from each selected beam, and they are shown for comparison. The only mask produced by multiTRAPUM is the one derived from the joined filterbank, marked in green. The shown masks are derived from actual coherent beams from the same MMGPS-L observation.

### Fourier-frequency birdies

In addition to frequency masks, we also included a search for signals in the FFT at zero-DM from the same set of beams. These signals are called *birdies*, and they are given to peasoup as a range of Fourier frequencies to ignore when identifying potential pulsar signals in the FFT. When ignored, we say that these Fourier bins are *zapped*.

Our approach was similar to the one used in multi-beam single-dish surveys, where individual beams correspond to a different feed at the focal point of the telescope, such as in HTRU surveys with Parkes/Murriyang (e.g. Ng et al., 2015). From the selected set of beams, we de-dispers each filterbank at  $DM = 0 \text{ pc cm}^{-3}$  with prepdata, perform a FFT with realfft, and de-redden the FFT series with rednoise. All of these programmes are available in PRESTO. Then we open the de-reddened FFT series and search for significant signals in them with the custom-made zapFourier.py script.

The identification of signals in the FFT is rather simple. First, we convolve the FFT power with a series of sinc functions,

$$\hat{a}_k^2 = \max_i \left[ \sum_{k'} a_k^2 \frac{\text{sinc}\left(\frac{\pi}{i}(k - k')\right)}{\sqrt{2i - 1}} \right], \quad (3.1)$$

where  $\text{sinc}(x)$  is set to 0 outside of  $-\pi < x < \pi$ , and where  $i$  runs from 1 to 7. In the case of  $i = 1$ , the Fourier power is not convolved, while for  $i > 1$  individual bins of Fourier power are averaged with

weighed neighbouring bins. The  $i$  value resulting in the largest power is chosen for the power value of every individual bin. That is done because the FFT power series is actually very noisy. This smoothing function does not detract power from the Fourier bins, while at the same time it joins Fourier bins affected by the same signal so that we can actually catch whole birdies instead of fragments of them. And finally, all bins with smoothed power

$$\hat{a}_k^2 > \ln \frac{N_{\text{FFT}}}{2} \quad (3.2)$$

are flagged as potentially containing an RFI signal following the reasoning from equation (2.16), where  $N_{\text{FFT}}$  is the number of bins in the FFT. For 10-min-long observations with a sampling of 153  $\mu\text{s}$ , this usually means  $\hat{a}_k^2 > 15$ .

The story does not end here, of course, as we are talking about a multibeam approach. The final decision on whether to flag certain Fourier bins based on how many times does the same signal appear across all the studied beams. If the same bin is flagged for zapping in more than half of the beams, then it is zapped. Otherwise, the signal is ignored, as it could be either be either localised RFI or a low-DM astrophysical signal. Fig. 3.5 draws the FFT power as a function of Fourier frequency after such zapping is done on a real MMGPS-L observation, showing that many zero-DM signals have been identified and zapped away. However, some signals going above the threshold are not removed, as they may not have appeared in more than half of the beams.




Figure 3.5: Fourier frequency birdies highlighted in the zero-DM, de-reddened FFT of the central beam of a real MMGPS-L pointing. On the left, the power of FFT up until 1000 Hz (signals of 1 ms) is shown. On the right, we zoom in into a 433.3 Hz signal (2.3 ms). The 433 Hz frequency is of common use for signal transmissions between low-power devices, for example headphones or remote door openers. How the FFT would look like before and after zapping the birdies is shown in red and black, respectively. Some signals above the zapping threshold survive because particular spikes of power may not appear in more than half of the beams.

### Implementation and deprecation

The multiTRAPUM pipeline was used not only in MMGPS, but also for other TRAPUM projects such as the Fermi-sources and nearby galaxies targeted searches (Clark et al., 2023; Carli et al., 2024). In the case of MMGPS-L, it was used in approximately the first 30% of pointings. When the filterbanks

arrived in APSUSE from an observing session, a call to run it on every pointing was made, producing the frequency masks and the Fourier frequency birdies list before calling `peasoup` to start the search pipeline. The masks were successful to a degree, with every pointing having between 20% and 30% of the band masked. In fact, `multiTRAPUM` became a useful diagnostic of the RFI environment during each observation. Occasionally, the fraction of zapped channels went above 50% and that was indicative of some heavy RFI interference, and that would set the pointing for re-observation.

However, the approach was flawed in a couple of fundamental aspects. First, strong RFI environments lead to baseline changes in the filterbanks from different coherent beams. This means that when running `rfi find` on the joined beam, sudden jumps between the signal power can be picked up as RFI, leading to over-masking of a small fraction of pointings. That problem could be dealt with by running `filtool` before slicing the filterbanks, as it also normalises the data during cleaning. But the biggest, most important flaw was that `rfi find` targets the same frequency channels as `filtool`. `filtool` cleans the RFI-affected channels, meaning that a RFI signal may be removed but that those channels can still contribute towards  $DM > 0 \text{ pc cm}^{-3}$  signal during de-dispersion. On the other hand, the masking of channels removes all signals, affecting  $S/N$  or pulsars in the FFT.

This led to the deprecation of the use of `multiTRAPUM`. A red flag was raised when some known pulsars were not detected in the expected MMGPS-L pointings. Upon inspection of the corresponding pointings, it was later found that these pulsars were far from the center of any coherent beams, meaning that their  $S/N$  was bound to be low in the FFT. Nonetheless, while they had not been reported by the standard `peasoup` run, a quick test `peasoup` run on the best beam without masking reported the signals with  $S/N > 8.5$  by a small margin. That was attributed to the reduced band on which the signal was de-dispersed: according to equation (2.9), the masking of 25% of frequency channels leads to a 14% loss in  $S/N$ . Therefore, it was deemed that the masking approach was “overkill”, and that it could potentially make us miss some discoveries at the bottom of the  $S/N$  distribution. In the light of this, `multiTRAPUM` was deprecated in favour of only running `filtool` on every filterbank before running `peasoup`. Regarding the Fourier birdies, they worked well but it was considered that the improvement they brought was minimal. Instead a list of fixed birdies from typical RFI signals seen across multiple observations was implemented.

Nonetheless, the development, implementation and later deprecation of `multiTRAPUM` was a very valuable learning experience about the behaviour of interferometric multibeam data not only for me, but also for the MMGPS collaboration as a whole. It helped us understand our RFI environment, how it evolved from observation to observation, and what was to be expected in general. And finally, we learned that the noise baseline can actually change across filterbanks from coherent beams from the same observation, a fact which we would have not noticed otherwise, and that will likely be taken into consideration in future multibeam RFI excision approaches.

### 3.2.4 Managing discoveries: confirmation, localisation, and follow-up

The first discovery came quite late into the survey. The survey had been running for months and all we got as thousands upon thousands of RFI, noise and known pulsars. We had a few promising T2 candidates, but upon re-observation the signal were nowhere to be found, indicating that they were good-looking noise, but noise after all. Morale was starting to drop and we even started wondering if we were doing something wrong and, especially, whether our survey was as sensitive as predicted. And then, on May of 2021 we struck gold. I was the lucky one to have it in the batch of candidate plots to inspect. After looking at too many noise plots, the high- $S/N$  gem shown in Fig. 3.6 appeared

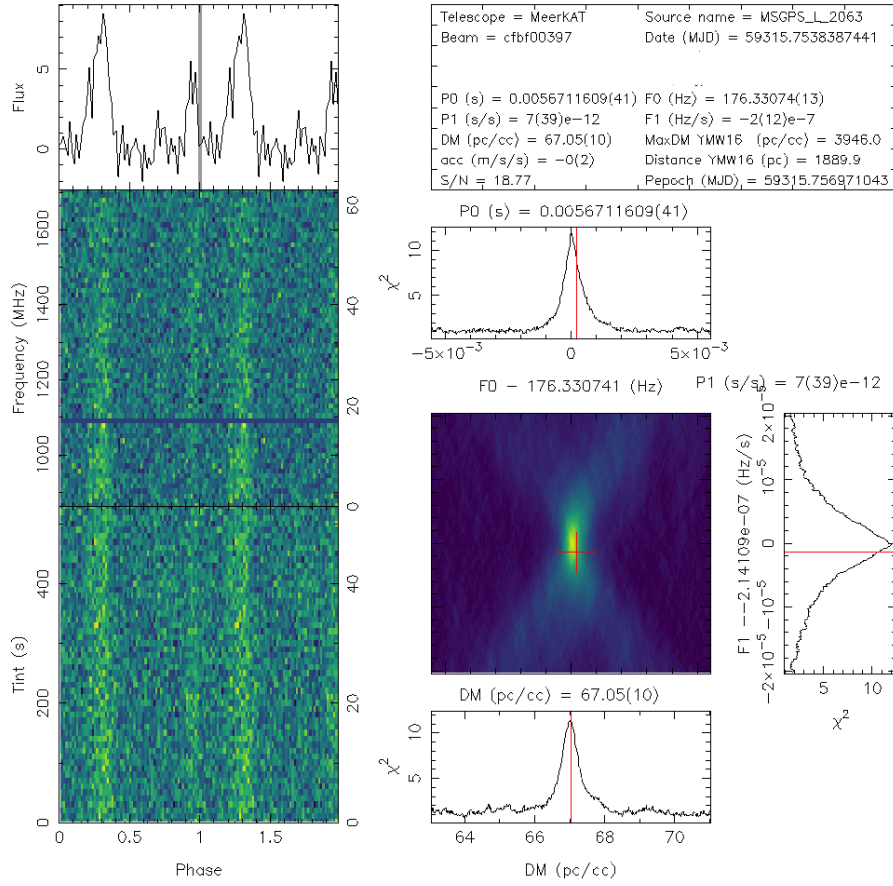


Figure 3.6: PulsarX discovery plot of **PSR J1306–6043**. The tell-tale pulsar characteristics of the realness of this candidate are: a clean vertical signal in the frequency–spin phase and time–spin phase plots, a clear pulse profile, and very clean optimisation curves in the  $f_0$ ,  $f_1$ , and DM spaces.

before my eyes in CandyJar. But it was not possible, I thought. It could not be so obvious. I thought it so much that I did not consider it a discovery but as a trickery of the incredibly bright Vela pulsar, **PSR B0833–45**<sup>70</sup>, which has the tendency to appear in too many of our pointing as a non-boresight pulsar. So, I slept it off and forgot about it across the weekend, until on Monday I confirmed that no, that was not Vela. Instead, it was our first discovery, **PSR J1306–6043**.

Our mattermost chat exploded with celebratory messages. People had been waiting for this event for way too long, so there was no restraint. But for that, I had made sure before sharing it with the collaboration that it really was a new pulsar. And after that, we had to manage the discovery. After this pulsar, we established protocol to manage every new discovery. Here is how it goes.

<sup>70</sup> The Vela pulsar is another of the very early discoveries (Large et al., 1968). With  $P_s = 89$  ms and  $DM = 67.771$  pc cm<sup>-3</sup>, it is a young pulsar associated with a 11,00-year old supernova remnant nebula, and one of the brightest in the sky. For more info, see Sushch et al. (2011) and references therein.

### Known or not known? Find the fundamental period!

The first part of confirming a discovery is to make sure it is not a known pulsar in disguise. Already in the FFT search, we implement harmonic summing. But sometimes harmonic summing does not work properly, and `peasoup` reports an harmonic of a known pulsar, for instance half or a quarter of its true period. The multibeam filtering step of the MMGPS pipeline searches for these multiplicities and corrects for them, but some pulsars are so bright that ridiculous harmonics such 1/32th of the fundamental period are still reported and they end up appearing in our candidates or, most annoyingly, it reappears multiple times with different DM values. Vela and other bright pulsars are specially good at this, and **PSR J1306–6043** even shames a similar DM value with it.

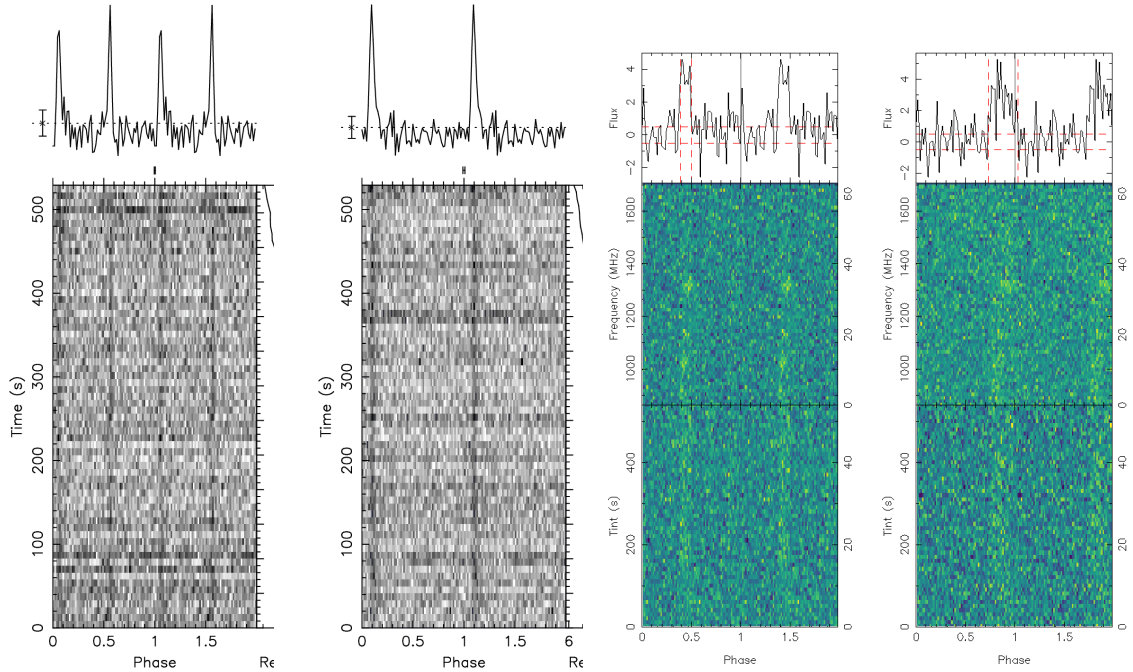


Figure 3.7: Folding at consecutive harmonics of **PSR J1208–5935** (left) and **PSR J1015–5368** (right). In both cases, spin phase runs from 0 to 2, so the profile repeats at  $\phi_s > 1$ . For **PSR J1208–5935**, the folds at 57.41 ms and 29.71 ms are shown. For **PSR J1015–5368**, the folds at 20.8 ms and 10.4 ms are shown. The **PSR J1208–5935** plots were made by me with me with PRESTO, while the **PSR J1015–5368** plots were both presented found as independent candidates in CandyJar. In both cases, the fold is from the highest  $S/N$  coherent beam in the discovery observation.

So, the first exercise when finding a potential discovery is to compare its spin period, DM and position are compatible with that of other pulsars. If that is not the case, we must still refold the filter bank at harmonics of the reported period to find the fundamental period,  $P_{\text{fund}}$ . That is identifiable by the following characteristics: when folding at  $n \times P_{\text{fund}}$ , where  $n$  is an integer, the unique pulse is repeated multiple times across  $0 < \phi_a < 1$  if  $n > 1$ . On the other hand, when folding at  $P_{\text{fund}}/n$ , the pulse profile becomes wider with increasing  $n$ . This is shown clearly in Fig. 3.7 for two other MMGPS-L discoveries, **PSR J1208–5935** and **PSR J1015–5368**. In the case of **PSR J1208–5935**, folding at twice the true period results in the pulse profile appearing twice. In the case of **PSR J1015–5368**, both the fundamental period and the half-period were found in CandyJar, with the



half-period actually having a better  $S/N$  than the fundamental one, showing that it is indeed important to make sure we are seeing the true spin period. For that pulsar, the hint is that at the half-period the pulse gets noticeably wider. Only when the fundamental period has been identified and it is not the same as that of any known pulsar, we can claim a discovery.

### Re-observe and confirm!

Any experiment must be reproducible for its measurements and conclusions to be considered. In a pulsar search, the pulsar has to be detected more than once to show it is not just good-looking RFI or noise. For that, we can try two things. First, we may fold the filterbanks of neighbouring coherent beams and see if a detection is repeated. If the pulsar is bright enough and close enough to another beam, this will already yield another detection, albeit with a lower  $S/N$ . That is not a confirmation by itself, as all filterbanks are derived from the same photons, but if it behaves as predicted it confirms at least an astrophysical-like behaviour of the signal regarding beam response.

The true confirmation must come from a follow-up observation. If the candidate presents no acceleration in the discovery, a simple fold should recover the signal. However, if the pulsar is in a binary, re-detection may actually require to perform another search due to the changes in observed spin period. In this regard, we have no standard way of performing this search. On many occasions, Ewan Barr re-runs a `peasoup` search with a restricted DM search range. For cases when that is available, I installed `PULSAR_MINER` (Ridolfi et al., 2021) on APSUSE, a PRESTO-based search pipeline. PRESTO itself is also a FFT-based search software which implements the Fourier-domain correlation technique to find accelerated pulsar signals (Ransom, 2011). Once a pulsar has been detected a second time, we can finally claim a discovery with no shadow of doubt.

### Localisation

As mentioned a few times already, the interferometric nature of MeerKAT allows for multiple detection of the same pulsar in several neighbouring coherent beams. These detections are expected to have varying  $S/N$  depending on the closeness of the beam to the true pulsar position. We use this to our advantage to *triangulate* the pulsar position when there are multiple detections.

Localisation is done in three steps. First, we measure the  $S/N$  of the pulsar across several neighbouring beams via folding of the pulsar signal. Ideally, we want at least three detections with  $S/N > 8$  for reliable results. I like to use the `prepfold` program from PRESTO to perform a quick fold and get a fast  $S/N$  of the pulse profile, but `PulsarX` or `PSRCHIVE` are legitimate choices too. What matters is that the software stays consistent across all folds.

Afterwards, we need to understand the response of the coherent beams to a signal in a sky position relative to them. For that, we simulate the two-dimensional coherent beam shape given the antenna array used during the observation, the sky direction from the point of view of the telescope, and the time at which the observation occurred. Given the interferometric nature of MeerKAT, the coherent beam shape is highly sensitive to which antennas are included into the observation (typically 40–50 out of 64) and how they are oriented with respect to the sky position (remember the discussion on the shape of coherent beams in Section 2.2.2). Therefore, this step is crucial if we are to guess the true position of the pulsar based on its relative position from the coherent beams. Thankfully, the necessary information is available for every observation. For that, we use `Mosaic` to achieve realistic beam shapes at the moment of the observation.



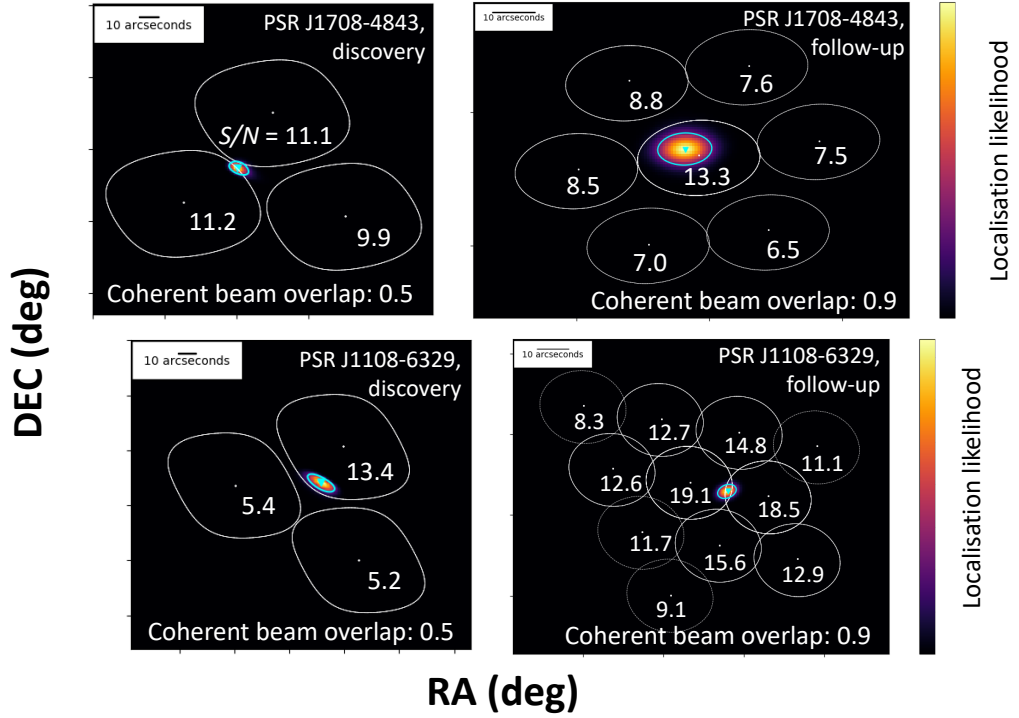


Figure 3.8: Localisation maps of MMGPS pulsars based on the distribution of detection  $S/N$  across neighbouring coherent beams from a single observation. The white contours are sensitivity curves of the coherent beams, set at 0.5 of the maximum power on the left, and at 0.9 of the maximum power on the right. The tighter is the tiling, and the more detections are performed, the more precise is the localisation.

Finally, we use the [SeeKAT](#) software (Bezuidenhout et al., 2023) to find the sky position that is most consistent with the distribution of  $S/N$  values across beams. That is done by minimizing<sup>71</sup>

$$\chi^2(\text{RA}, \text{DEC}) = \sum_{\text{beams}} [S/N_{\text{measured}} - S/N_{\text{predicted}}(\text{RA}, \text{DEC})]^2, \quad (3.3)$$

where  $S/N_{\text{measured}}$  is our set of measured  $S/N$  values, and  $S/N_{\text{predicted}}$  is the expected  $S/N$  of value each beam given their shapes as simulated by [Mosaic](#), and their relative position with respect to the source position (RA and DEC). As an example, Fig. 3.8 illustrates the localisation of two different MMGPS discoveries process by mapping the probability distribution density

$$\rho(\text{RA}, \text{DEC}) \propto \exp \left[ -\frac{\chi^2(\text{RA}, \text{DEC}) - \chi_{\min}^2}{2} \right]. \quad (3.4)$$

If performed correctly, this exercise can give us the pulsar position down to an uncertainty of just a few arcseconds within a single observation. That is highly exceptional and pioneering. Surveys with single-dish radio telescopes need to perform multiple observations at slightly different positions to

<sup>71</sup> You may noticed that, unlike other cases we have seen, this  $\chi$  statistic is not normalised by the measurement uncertainties. That is because the uncertainty of a  $S/N$  measurement is always 1 by definition.

triangulate a pulsar location, but the interferometric nature of MeerKAT can give us this information already at discovery. However, it can happen that a pulsar is too faint or too close to the center of one of the coherent beams to yield multiple detections in the original pointing. In such case, the localisation can be performed during a follow-up re-observation using a tighter coherent beam overlap of 0.9, ensuring that several beams detect the pulsar this time as is shown in the right side of Fig. 3.8. After localisation has been confirmed, future observations of the same pulsar are performed by pointing at the true position and performing.

#### Follow-up

If the pulsar is in a binary, that is immediately picked up either by the non-zero acceleration value at discovery, or by a change in  $P_{\text{obs}}$  in the confirmation observation. If that is the case, a follow-up program is established for each discovery, first to find an orbital solution, and then for timing with continued follow-up.

Before an orbital solution is found, we schedule short re-observations every few days, ideally twice per week. In the vast majority of cases, these observations need to be searched for due to the evolving nature of the observed spin period. The observations are made in the APSUE search-mode, producing filterbanks of the same characteristics as the search survey, lasting between five to ten minutes. These observations need are searched for in a similar way as the localisation follow-up observations: with a **peasoup** search for a quick detection, and in the case that is not achieved, with further searches. For the pulsars I was in charge of, I relied on **PULSAR\_MINER** or on a custom-made **PRESTO** pipeline. The objective is, of course, to measure  $P_{\text{obs}}$  and  $\dot{P}_{\text{obs}}$  in every observation.

After an orbital solution has been found from  $P_{\text{obs}}$  and  $\dot{P}_{\text{obs}}$  measurements, observations are moved to a monthly cadence as it is no longer crucial to have a high observing cadence. The story here gets a bit blurry, as every pulsar is a different case, but in general we try to move onto timing after the orbital solution has been achieved. Previous filterbanks are then folded into high resolution pulsar archives with the orbital solution, and new observations are performed on the best pulsar position with the Pulsar Timing User-Supplied Equipment (PTUSE, Bailes et al., 2020). PTUSE is another cluster at the MeerKAT site which is capable of coherent de-dispersion and real-time folding of observed pulsars. This way, high-quality folded pulsar archives are produced directly instead of filterbanks. From here, the science that we get out of the experiment is highly pulsar dependant. Speaking of which, it is time to move onto the discoveries.

### 3.3 MMGPS discoveries: an isolated and binary pulsar jamboree

You like methods, I like methods, we all like methods. But do you know what else do we all like? Science results. Do not worry, I will not be shy anymore: this section will be all about MMGPS science results. The MMGPS has made **83 new pulsar discoveries so far**<sup>72</sup>: 78 in the MMGPS-L (Padmanabh et al., 2023), three in the MMGPS-S, and one in the MMGPS-UHF. Of the 78 MMGPS-L pulsars, 17 are recycled pulsars in a binary orbit. All of these 17 pulsars have *at least* an orbital solution, so we are quite certain about their nature.

<sup>72</sup> All discoveries can be consulted in the **TRAPUM webpage**, so you can look up their spin and DM parameters to your hearth's content.

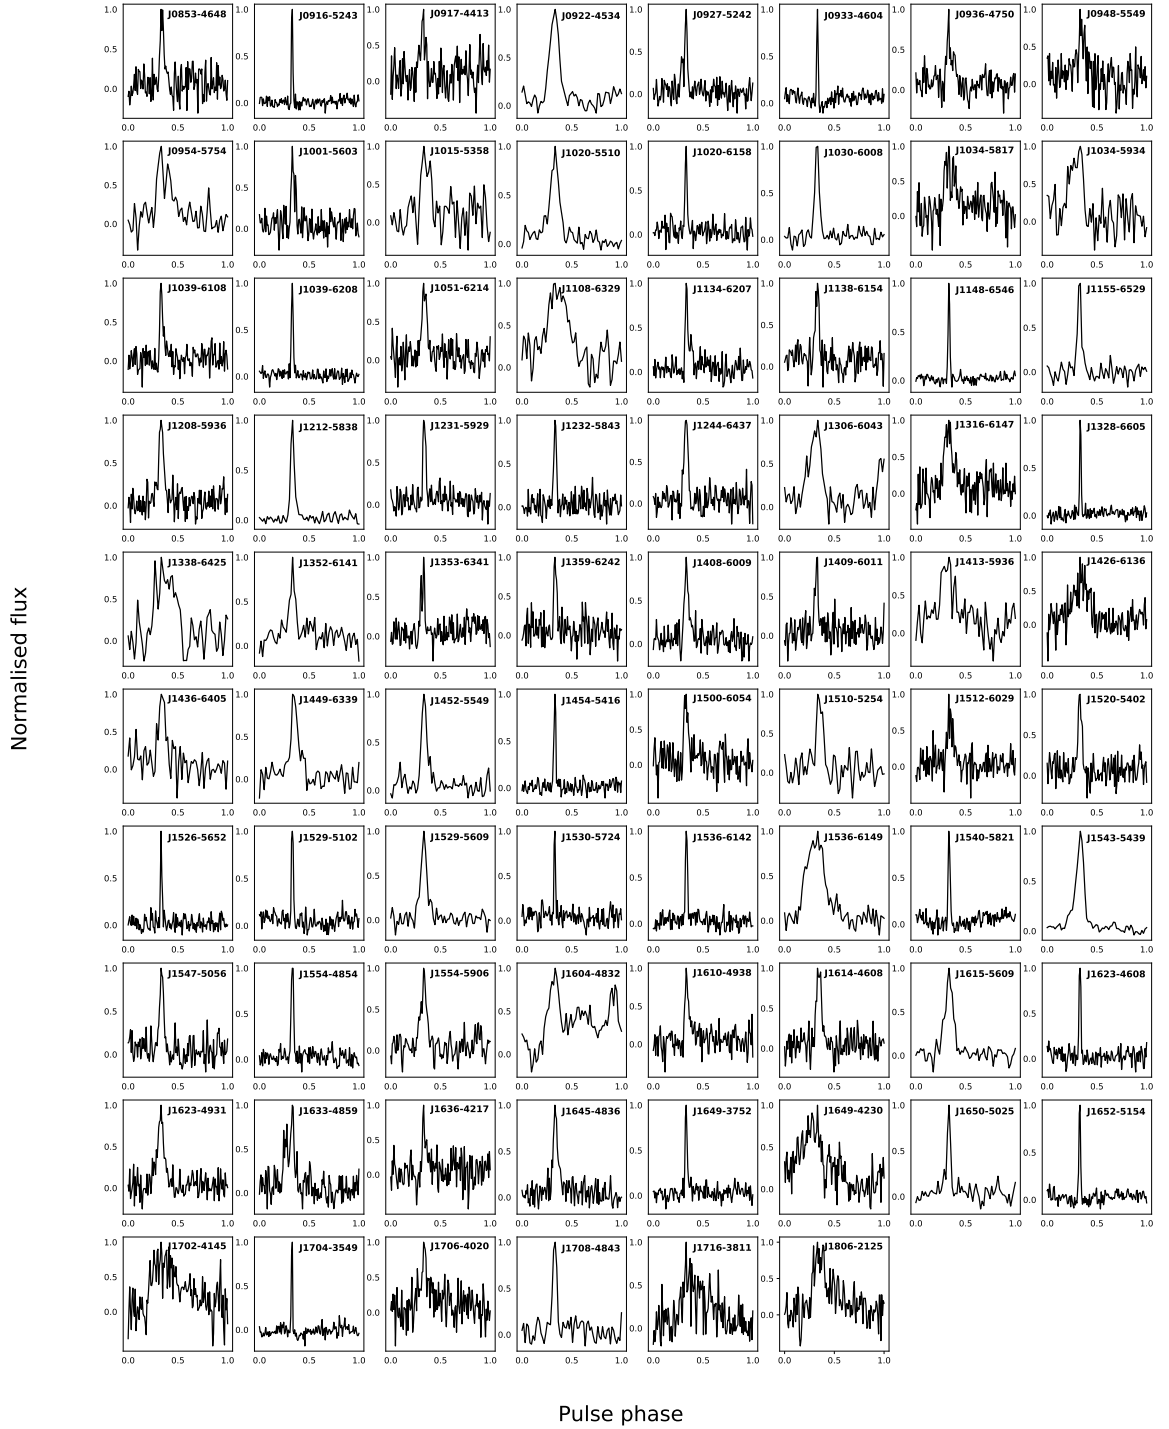


Figure 3.9: Folded pulse profiles of all 78 MMGPS-L discoveries from the discovery observation. Plot originally presented in Padmanabh et al. (2023)

The MMGPS-L discoveries are very diverse, as shown already by the folded pulse profiles displayed in Fig. 3.9. The fastest-spinning one is the binary MSP **PSR J1615–5609**, with  $P_s = 3.36$  ms, while the slowest one is the aged **PSR J0933–4604**, with  $P_s = 3.67$  s. At the same time, the highest-DM discovery is the highly scattered **PSR J1716–3811**, with  $DM = 1219$  pc cm $^{-3}$ . All three MMGPS-S discoveries consist of high-DM isolated canonical pulsars: **PSR J1011–5703**, **PSR J1239–6307**, and **PSR J1614–5218**, with  $P_s = 0.282$ ,  $0.234$ , and  $0.510$  s, and  $DM = 1037$ ,  $717$ , and  $673$  pc cm $^{-3}$ , with **PSR J1239–6307** being affected by heavy scattering. Their profiles are also shown in Fig. 3.10. Their high-DM nature is not surprising: this is the distant and faint population the MMGPS-S was designed to probe. Unfortunately, most of our discoveries do not have a  $\dot{P}_s$  measurement from timing yet, so a  $P_s$ – $\dot{P}_s$  depicting their evolutionary placement is not possible.




Figure 3.10: Folded pulse profiles of the three current MMGPS-S discoveries: **PSR J1011–5703**, **PSR J1239–6307**, and **PSR J1614–5218**. The horizontal axis is the spin phase, running from 0 to 1. Plots made with **PulsarX** from the discovery observations.

Of all of these discoveries, I happened to be the first one to see 14 of them in the CandyJar tool<sup>73</sup>. Usually, first-seers are also assigned follow-up responsibilities, meaning that I have worked in detail on a subset of these 14 pulsars. In this section I share the observational and scientific details of them. I start with an exposition of the isolated pulsars I have worked with. Then, I follow with an overview of all binary discoveries in the MMGPS-L. And after that, I dedicate a section to every MMGPS binary pulsar I have worked with.

### 3.3.1 Details about isolated MMGPS-L discoveries

The isolated pulsars I have been involved with are: **PSR J1353–6341**, **PSR J1449–6339**, **PSR J1134–6207**, **PSR J1231–5929**, and **PSR J1652–5154**, all of them MMGPS-L discoveries. In addition, I also worked on the localisation and follow-up of **PSR J1452–5549**. The pulse profile of these pulsars are shown in Fig. 3.11. Here is a section dedicated to each of them.

#### PSR J1353–6341

This is a canonical pulsar with  $P_s = 2.076$  s and  $DM = 439$  pc cm $^{-3}$ . Its pulse profile is rather pretty, with two distinct components consistent with a conical beam shape. It was the second MMGPS

<sup>73</sup> In Appendix D I share the music I was listening to during each of these discoveries, and my personal thoughts about it. Please do check it if you are interested.




Figure 3.11: Folded pulse profiles of the pulsars discussed in this section. From left to right: **PSR J1353–6341**, **PSR J1449–6339**, **PSR J1452–5549**, **PSR J1134–6207**, **PSR J1231–5929**, and **PSR J1652–5154**. Plots made with **PSRCHIVE**, **PRESTO** and **PulsarX**.

discovery, so it brought a lot of excitement with it. It was observed on 30/04/2021 and seen in CandyJar with a folded  $S/N = 13$  on 22/05/2021. A follow-up observation at 0.9 overlap on 28/05/2021 quickly re-detected it in 10 different beams, with the best beam having  $S/N = 20$ . This enabled a SeeKAT localisation within less than three arcsec. The follow-up also confirmed its isolated nature, so it was not scheduled for regular follow-up. The next observation was only in 16/08/2021 in PTUSE-mode. From it, polarised emission was detected and a measurement of  $RM = -159 \text{ rad m}^2$  was achieved.

However, the most interesting result from **PSR J1353–6341** is that its pulsed emission is intermittent, with on/off intervals lasting for more than one minute. That is seen in all observations, including the discovery observation, as is shown in Fig. 3.12. In those time–phase plots, the pulsar signal turns off for entire minutes, spending 20 to 40% of the observations in the *off* state. Pulsars of this kind are quite rare, with only a few handful of so-called *nulling* or *nulling pulsars* having been observed before (e.g. Kramer et al., 2006; Wang et al., 2007; Lyne et al., 2017). These pulsars typically have  $P_s \gtrsim 1 \text{ s}$  and present peculiar behaviours. For example, the spin-down rate of **PSR B1931+24** is larger during the *on* state than in the *off* state, pointing towards changes in the spin-down torque in the different states (Kramer et al., 2006). The on/off states of known nulling pulsars can last from a short as a few minutes to as long as entire weeks, with a large diversity in the ratio of time spent in the on/off states

(again, Kramer et al., 2006; Wang et al., 2007; Lyne et al., 2017). In the most extreme cases, we have *rotating radio-transients* (RRATs), a pulsar-like neutron star population with characteristically long periods of  $P_s \gtrsim 1$  s and variability at the pulse-to-pulse level, with only a small fraction of the pulses being detectable (e.g. McLaughlin et al., 2006; Zhang et al., 2024).




Figure 3.12: Time-phase plots of **PSR J1353–6341** from 30/04/2021 (left, plotted with PulsarX), and from the 16/30/2021 observation (right, plotted with PSRCHIVE).

**PSR J1353–6341** fits quite well with these populations, making a notable addition to it due to the large clarity in the transition of its states. For this reason, it is a very interesting discovery that is being follow-up in dedicated MeerKAT observing program along with other isolated MMGPS discoveries led by my fellow PhD student Denisha Pillay. The main aim of this program is to track any further profile changes and to measure its spin evolution.

### PSR J1449–6339

This is an isolated PRP with  $P_s = 29.47$  ms and  $DM = 75$  pc cm $^{-3}$ . It was discovered on 01/07/2021 with  $S/N = 19$  in an observation from 16/06/2021. It was localised with SeeKAT to within less than three arcsecs thanks to multi-beam re-detections in the the 0.9-overlap follow-up observation of 03/07/2021, with the best beam having  $S/N = 25.6$ . No acceleration detected in any of the observations, but its short spin period made it a tempting PRP candidate in a potentially wide orbit. While I was the first one to see it and had performed localisation, my fellow PhD student Shalini Sengupta took charge of the follow-up effort. With the help of the MMGPS member Vivek Venkatraman Krishnan, a phase-connected timing solution was determined on 16/10/2024 without the need of implementing *dracula*: the pulsar was confirmed to be isolated. However, a significant  $\dot{P}_s = -3.24(14) \times 10^{-18}$  s s $^{-1}$  was measured from the timing solution. The negative value is explained via an inaccurate sky position or an unmodelled proper motion vector. While these issues would produce year-long sinusoidal modulation of the ToA residuals (as is shown in Fig. 2.15), if the timing solution spans less than a year, it only produces a slope that is measured a spurious  $\dot{P}_s$  value.

Nonetheless, if it had been a canonical pulsar, a much larger spin-down rate of  $\dot{P}_s \approx 10^{-15} \text{ s s}^{-1}$  would have been observed. Since that is not the case, confirms that it is indeed a recycled PRP. The lack of a binary companion suggests that this pulsar started as the first-born neutron star in the DNS formation channel, but that the supernova kick of the second-born neutron star disrupted the system, breaking the binary after the recycling phase and sending each neutron star into their own orbits across the Galaxy (end-point at the bottom-right of Fig. 1.13). Given the recycled nature of the pulsar, it could be hundreds of millions of years old, and its former orbital companion is probably lost to the pulsar graveyard in a far-away corner of the Milky Way.

### PSR J1452–5549

This is another fast-spinning isolated pulsar, with  $P_s = 75.25 \text{ ms}$  and  $\text{DM} = 184 \text{ pc cm}^{-3}$ , discovered on 11/06/2021 from a 20/05/2021 observation with an amazing  $S/N = 28$ . This pulsar was found by Shalini Sengupta, but it was handed to me for localisation and follow-up<sup>74</sup>. Given its spin period, it was possible for it to be a PRP, so to follow-up race was on. I started working on it from the localisation observation of 15/06/2024. A solid localisation smaller than two arcsec was achieved with SeeKAT thanks to 10 multibeam detections at 0.9 overlap, with a whooping  $S/N = 52$  on the best beam. However, it also became clear that it isolated, so the follow-up efforts were drastically reduced. Given the lack of a timing solution, no constraints on  $\dot{P}_s$  are available and thus it is unclear whether **PSR J1452–5549** is a relatively young canonical pulsar, or an isolated PRP like **PSR J1449–6339**. A phase-connected timing solution should be relatively straight-forward to achieve, but my efforts shifted towards the timing of binary pulsars instead. Nonetheless, thanks to its high  $S/N$ , it is ideal for making plots about how to treat pulsar data in a thesis about MeerKAT. Just so you know it, you have already seen this pulsar in Figs. 1.3, 2.4, 2.5 and 2.14. High  $S/N$  pulsars are really handy for plotting.

### PSR J1134–6207, PSR J1231–5929, and PSR J1652–5154

Unfortunately, not many measurements have been performed by on these pulsars other than reporting their spin periods and DM values from the discovery observations. These would be  $P_s = 0.689, 0.410$ , and  $0.600 \text{ s}$ , and  $\text{DM} = 662, 356$ , and  $266 \text{ pc cm}^{-3}$ . They are all classical middle-aged canonical pulsars that happen to be high-DM, found with  $S/N = 17, 17$ , and  $27$ . They were observed on 17/08/2021, 08/11/2021, and 06/02/2022, and discovered in 30/08/2021, 29/11/2021, and 20/02/2022. Given their high  $S/N$  and the over-abundance of classical canonical pulsar in the survey, they were slated for re-observation sometime at the end of all the MMGPS-L observations. They were part of a successful 2023 MeerKAT open-time proposal written by my fellow PhD student Denisha Pillay, and they are currently under follow-up in a joint observing program along with many other isolated MMGPS-L (the same one as **PSR J1353–6341**). They were re-detected and localised in this project, and they are currently under monitoring to see if any changes occur in their folded profiles, and to measure  $\dot{P}_s$  for an adequate placement in the  $P_s - \dot{P}_s$  diagram.

<sup>74</sup> I fail to remember why we swapped **PSR J1449–6339** and **PSR J1452–5549**, especially since they are very similar. Perhaps it was because **PSR J1452–5549** was deemed more likely to be binary and I was already in charge of a couple of them.



### 3.3.2 The MMGPS-L binary jamboree

Readers and readdresses! Please marvel at our 17 *brand-new* binary specimens! See all of them falling in their respective sequences in the  $P_b - x$  diagram presented in Fig. 3.13! We have a large variety in exposition: 10 MSPs and 7 PRPs, 15 circular and two eccentric, all with a wide range of masses! There is something for everyone to like!




Figure 3.13:  $P_b - x$  and  $P_b - P_s$  diagrams with the MMGPS-L binaries compared against binaries in the Galactic field as listed at the ATNF database in early 2024. Lines of equal minimum companion mass from the mass function in equation (1.21) are shown in the  $P_b - x$  diagram. Circular discoveries are highlighted with red stars, while eccentric discoveries are highlighted with purple stars.

All of the systems plotted in Fig. 3.13 are listed in Table 3.2 along with a binary type classification for a more detailed description. Almost all evolutionary end-points in Fig. 1.13 are represented. However, please note that all the results presented in this section are from the orbital solutions, so no precise  $M_c$  measurements are reported in Table 3.2. Instead, the minimum and median  $M_c$  values are derived from the mass function in equation (1.21) under the assumption of  $i = 90$  and  $60$  deg. Remember, the true values are likely to cluster between the minimum  $M_c$  and slightly above the median  $M_c$ , with very few true values being much larger than the median  $M_c$  (see Appendix A.2.2 for more details about the distribution of  $i$  values).

For starters, we are not short of recycled pulsars with white dwarf companions. We have **PSR J0954–5754**, **PSR J1108–6329**, and **PSR J1615–5609**, three most classical MSPs with  $P_s \approx 4$  ms and  $P_b = 1$ –10 day circular orbits with  $M_c = 0.1$ – $0.3 M_\odot$  companions, likely **MSP–He WD** systems descending descending from the Case B RLO – LMXB evolutionary channel.

**PSR J1306–6043**, **PSR J1352–6141**, **PSR J1536–6149**, **PSR J1543–5439**, and **PSR J1554–5906** are close cousins to the first three systems. They are all slightly slower-spinning MSPs with  $P_s = 4$ –9 ms and have somewhat heavier companions with  $M_c = 0.15$ – $0.4 M_\odot$ . The striking difference is their long orbits with  $P_b \sim 100$  days. From the slightly larger  $M_c$  values, it seems that the positive  $P_b - M_{WD}$  correlation discussed in Section 1.4.4 holds true, as is shown in the  $P_b - M_c$  diagram for white dwarf systems in the left side of Fig. 3.14. However, this result has to be taken with a grain of salt until a precise mass measurements from pulsar timing are available.

Table 3.2: Binary systems discovered by the MMGPS-L, grouped by their evolutionary history. Pulsars with  $e \sim 0.0$  mean that their eccentricity value is not constrained due lack of a phase-connected timing solution, but that the orbit is circular nonetheless. The median and minimum companion mass is estimated from the mass function in equation (1.21) assuming  $i = 90$  and  $60$  deg. The binary type is deduced from the parameters at the evolutionary end-points in Fig. 1.13.

PSR	$P_s$ (ms)	$P_b$ (days)	$x$ (ls)	$e$	min/med $M_c$ ( $M_\odot$ )	Binary type
<b>J0954–5754</b>	4.835	3.839	2.227	$\sim 0.0$	0.12/0.15	<b>MSP–He WD</b>
<b>J1108–6329</b>	4.228	0.907	1.073	$\sim 0.0$	0.15/0.18	<b>MSP–He WD</b>
<b>J1615–5609</b>	3.359	8.033	6.017	$\sim 0.0$	0.20/0.25	<b>MSP–He WD</b>
<b>J1306–6043</b>	5.671	85.93	40.42	$1.3 \times 10^{-4}$	0.29/0.36	<b>MSP–He/CO WD</b>
<b>J1352–6141</b>	4.738	86.88	35.42	$\sim 0.0$	0.25/0.31	<b>MSP–He WD</b>
<b>J1536–6149</b>	6.875	111.88	29.69	$\sim 0.0$	0.17/0.21	<b>MSP–He WD</b>
<b>J1543–5439</b>	4.314	95.80	35.35	$\sim 0.0$	0.23/0.28	<b>MSP–He WD</b>
<b>J1554–5906</b>	8.701	105.13	39.44	$\sim 0.0$	0.24/0.29	<b>MSP–He WD</b>
<b>J1338–6425</b>	4.088	18.80	16.12	$\sim 0.0$	0.39/0.48	<b>MSP–CO WD</b>
<b>J1510–5254</b>	4.780	12.78	17.96	$\sim 0.0$	0.50/0.61	<b>MSP–CO WD</b>
<b>J1015–5358</b>	20.81	0.333	2.991	$2.8 \times 10^{-4}$	1.18/1.45	<b>PRP–ONeMg WD</b>
<b>J1034–5934</b>	34.47	1.988	2.947	$\sim 0.0$	0.26/0.32	<b>¿PRP–CO WD?</b>
<b>J1413–5936</b>	21.67	1.247	5.291	$\sim 0.0$	0.77/0.94	<b>PRP–CO WD</b>
<b>J1529–5609</b>	36.05	0.805	4.927	$8.4 \times 10^{-5}$	1.04/1.28	<b>PRP–ONeMg WD</b>
<b>J1708–4843</b>	16.66	0.544	2.176	$< 10^{-5}$	0.50/0.61	<b>PRP–CO WD</b>
<b>J1155–6529</b>	78.88	3.668	15.34	0.267	1.24/1.52	<b>double NS</b>
<b>J1208–5936</b>	28.71	0.632	4.256	0.348	1.06/1.30	<b>double NS</b>

We transition into heavier systems as we move down in Table 3.2. **PSR J1338–6425**, and **PSR J1510–5254** belong to the arising class of fully recycled **MSP–CO WD** systems that bridge the two main sequences in the  $P_b - x$  diagram in Fig. 3.13. As is seen in Fig. 3.14, these systems deviate significantly from the theoretical  $P_b - M_c$  relationships discussed in Section 1.4.4. With  $P_s = 4\text{--}5$  ms,  $P_b = 10\text{--}20$  days, and  $M_c = 0.4\text{--}0.6 M_\odot$ , they are likely to descend from the Case A RLO – IMXB evolutionary channel, with recycling occurring in the Case AB RLO stage during the helium shell-burning stage of the companion, as is depicted in Fig. 1.13. If that were to be confirmed, they would be the fourth and fifth addition to this class, right after **PSR J1614–2230** (Crawford et al., 2006; Tauris et al., 2011), **PSR J1933–6211** (Crawford et al., 2006; Geyer et al., 2023), and **PSR J1125–6014** (Faulkner et al., 2004; Shamohammadi et al., 2023).

**PSR J1015–5358**, **PSR J1034–5934**, **PSR J1413–5936**, **PSR J1529–5609**, and **PSR J1708–4843** bring us to the truly heavy **PRP–CO/ONeMg WD** systems. This group is more diverse, with

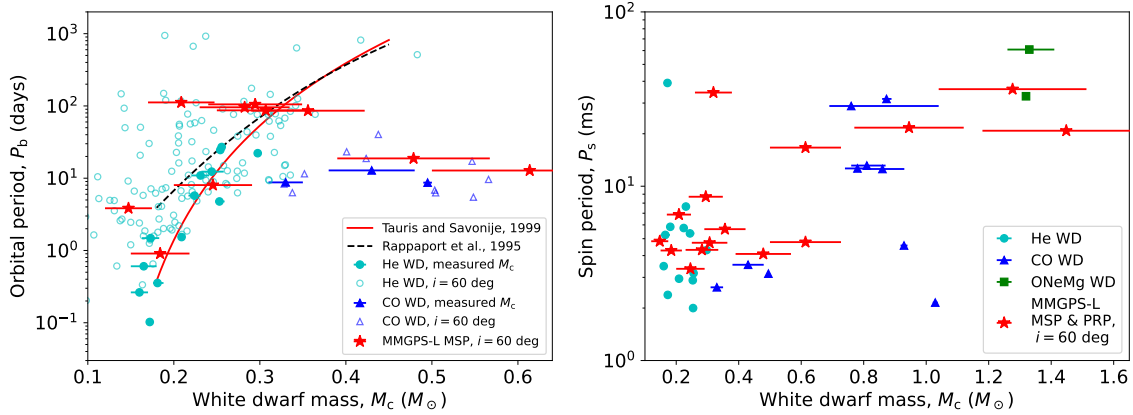


Figure 3.14: Galactic recycled pulsar–white dwarf  $P_b - M_c$  and  $P_s - M_c$  diagrams updated with the MMGPS-L discoveries. For MMGPS-L systems, only median  $M_c$  values derived from the mass function (equation 1.21) are plotted. The  $M_c$  uncertainty is derived from the difference between the minimum and median  $M_c$  values. Only MSPs are included in the left diagram, as the predictions from Rappaport et al. (1995) and Tauris and Savonije (1999) apply only to them. The same sources for known systems are the same as in Fig. 1.16.

$P_s = 15\text{--}40$  ms and  $M_c = 0.5\text{--}1.2$ , but they all share orbital periods of  $P_b \lesssim 2$  days, indicative of IMBX and HMBX evolution from Case B RLO, with recycling occurring in a posterior Case BB RLO stage. As is shown in the  $P_s - M_{\text{WD}}$  diagram in Fig. 3.14 **PSR J1015–5358** and **PSR J1529–5609**, with  $M_c > 1 M_\odot$ , are likely to be exceptional **PRP–ONeMg WD** systems that have gone through a HMXB CE stage, like **PSR J2222–0137** (Guo et al., 2021) and **PSR J1528–3146** (Jacoby et al., 2007; Berthereau et al., 2023). However, unlike those systems which have  $P_b > 2$  days, the ones discovered by the MMGPS-L have both  $P_b < 2$  day, with **PSR J1015–5358** going as low as eight hours. This means that they will be treasure troves of gravity tests, with significant post-Keplerian parameters to measure via pulsar timing.

**PSR J1413–5936** and **PSR J1708–4843** likely descend from an IMXB system given their PRP nature, with  $P_s > 10$  ms, but they are quite exceptional for their own reasons. **PSR J1708–4843**, with  $P_s = 16.66$  ms, seem to be the lighter of the heavy PRP systems, if the  $M_c = 0.5\text{--}0.6 M_\odot$  value is to be trusted. Potentially, it could lie in the  $M_c$  gap between **MSP–CO WD** systems and **PRP–CO WD** systems, as is seen in the  $P_s - M_{\text{WD}}$  diagram in Fig. 3.14. However, it could also just be unlucky and have a low orbital inclination angle that biases its minimum  $M_c$  derived from equation 1.21 towards a low value. Out of 17 systems, at least some should be unlucky enough for this to happen. Only a proper mass measurement of the Shapiro delay via pulsar timing (or lack thereof) will tell. A potentially worse victim of *low-inclination-angle-itis* is **PSR J1034–5934**. Its  $P_s = 34.47$  ms and  $P_b = 1.99$  days would be indicative of a PRP evolution with a close but massive companion, but its  $M_c = 0.26\text{--}0.32 M_\odot$  says otherwise. It could be an extremely exceptional system formed via an unknown evolutionary channel, but my bet is on a very low inclination angle playing tricks on equation 1.21. If that is true, then the parameters of the Shapiro delay will be hard to measure and its mass measurements will never be very precise. On the other hand, if we actually manage to measure a Shapiro delay implying a large inclination angle and a low companion mass, we will give theorists a very bad headache (or maybe they will become really excited about it).

Finally, **PSR J1155–6529** and **PSR J1208–5936** are our two beloved DNS systems. They stick out from the other systems via their large eccentricities of  $e = 0.267$  and  $e = 0.348$ , but otherwise they followed the same evolutionary channel as **PRP–ONeMg WD** systems until the very end, when the companion underwent a supernova instead of becoming a white dwarf after the recycling phase. These two systems are the only MMGPS discoveries with proper mass measurements thanks pulsar timing, confirming their DNS nature. Their total masses from an  $\dot{\omega}$  measurement are  $M_t = 2.61 \pm 0.02$  and  $2.586(5) M_\odot$ , highly consistent with the addition of two neutron star masses. Regarding **PSR J1208–5936**, Section 3.3.4 and even Chapter 4 are coming up, so more details will follow soon. However, **PSR J1155–6529** was analyzed by my fellow MMGPS member Marina Berezina, so you will have to wait until she makes a publication about it to know more!

From all of the systems listed above, only the ones with proper constraints on their  $e$  values are the ones with timing solutions. Achieving these timing solutions was no easy feat. In the following sections, I proceed to explain the observational and modelling challenges, and further science results derived from intensive follow-up of the binaries whose follow-up I have been involved with.

### 3.3.3 PSR J1306–6043: the tutorial pulsar

Not only was this the first discovery of the MMGPS-L survey, but it also happened to be in a binary. I was the lucky one to find it while listening to the blissful melody of *Stickerbush Symphony*<sup>75</sup> on 14/05/2021. It was reported by CandyJar with  $S/N \approx 19$  from a 11/04/2021 observation, and in fact its discovery plot has already been shown in Fig. 3.6. With  $P_s = 5.67$  ms and  $DM = 67$  cm pc<sup>-3</sup>, it is a fully recycled MSP with two pulse components in the profile: a bright one and a fainter one, separated from each other by  $\sim 0.3$  phase turns. This means that it is probably showing us both of its magnetic poles, but a more secure statement on this will depend on an analysis of its polarised emission.

The binary nature of this pulsar became evident only upon re-observation. The discovery presented no detectable acceleration, but the 19/05/2021 re-observation revealed a changing  $P_{SSB}$ . On 19/05/2021 it presented  $P_{SSB} = 5.67098$  ms, while on 11/04/2021 it had presented  $P_{SSB} = 5.67132$  ms. Subtle, but significant. This pulsar’s orbital modulation is so subtle that we did not even need to perform a search: just folding with PRESTO/prepfold and optimizing for the spin yielded a detection of the best  $P_{SSB}$  value detection, with no measurable acceleration or  $\dot{P}_{SSB}$  within the five minutes of the observation. An example of this is shown on the folded time–phase plot in the left of Fig. 3.15, where the drifting pulse is seen upon simple refolding with PSRCHIVE/dspsr. By correcting this drift, a new  $P_{SSB}$  measurement was achieved.

The follow-up was facilitated by the bright nature of the pulsar, so in a way this was the perfect training ground for more challenging pulsars to come in the future. In the first re-observation of 19/05/2021, the pulsar was detected on 13 different beams at 0.9 overlap, with the bast beam yielding an integrated  $S/N = 30$ . This allowed for a tutorial-like, first successful sky localisation with SeeKAT, and a re-observation schedule with MeerKAT was set on a cadence of once every few days, with each observation yielding a new  $P_{SSB}$  value but no acceleration or  $\dot{P}_{SSB}$ .

The observations without an acceleration or  $\dot{P}_{SSB}$  were highly indicative of either a very light companion, or a long orbital period, according to equation 2.15. It turned out to be the later: consecutive  $P_{SSB}$  measurements across the days tracked a slow sinusoidal modulation across the months, as is shown in the right side of Fig. 3.15. With that information, it became clear that this was

<sup>75</sup> Please check Appendix D to see my comments on the music I have been listening to when every discovery happened.




Figure 3.15: Orbital spin modulation of **PSR J1306–6043**. On the left, a time–phase fold from 25/05/2021 is shown, along with the corresponding  $P_{\text{SSB}}$  measurement (blue line parallel to the pulse). On the right, consecutive  $P_{\text{SSB}}$  measurements across  $\sim 90$  days are shown, starting at 11/04/2021 and ending at 07/07/2021.

a circular binary with a  $P_b \sim 100$  day orbit.

This pulsar is now solved, but the last observation I worked on was from 07/07/2021. After that point, follow-up responsibility for this pulsar was moved from me to my back-then fellow PhD student Prajwal Voraganti Padmanabh. He went on to find the orbital solution as soon as a single orbital cycle was complete, and on 26/08/2022 he achieved a timing solution via the implementation of Paulo Freire’s *dracula* on both MeerKAT and Parkes/Murriyang observations. The final orbital period is  $P_b = 85.93$  days and the minimum companion mass is  $M_c = 0.29 M_\odot$ . It is rather heavy for a He WD, so if the inclination angle is not too high this could actually be a CO WD instead. Unfortunately, the wide and circular nature of the system means that any PK effects in the system are too small to measure via pulsar timing, and so far no Shapiro delay has been detected either. Therefore, it is unlikely that we ever know the stellar masses of this system in high precision.

However, this pulsar had one last surprise left for us. A significant Gamma-ray counterpart had been detected by the **Fermi telescope** (Padmanabh et al., 2023) at the same sky position. Some radio MSPs are also Gamma-ray MSPs, so there was huge opportunity here. Gamma-ray timing is rather different than radio timing, as Gamma-ray telescopes perform individual photon detections instead of recording continuous radio waves. I do not intend to delve into time-domain Gamma-ray astronomy now, but I can tease a cool result. The Fermi telescope has been recording Gamma photon detections from all of the sky since 2008. With the radio timing solution at hand, Prajwal folded the photon detections from the Gamma-ray source position and indeed detected Gamma-ray pulsations. This extended the timing baseline from one year to 14 years, yielding very precise  $P_s$  and  $\dot{P}_s$  measurements. As exciting as this is, more details about this pulsar are to be found only in a future publication by Prajwal Voraganti Padmanabh!

### 3.3.4 PSR J1208–5936: the pulsar that built CandyCracker and dracula2

This pulsar was our third discovery. It showed up in *CandyJar* on 30/05/2021 from a 06/05/2021 observation. It was rather bright, with a  $S/N = 19$ , but it was found at a coherent beam at the very

edge of the tiling. This means, indeed, that it was very close to being missed. It has  $P_s = 28.71$  ms and  $DM = 344.4$  pc cm $^{-3}$ . From the first observation it was a very good binary candidate: the LOS acceleration of  $a = -17 \pm 11$  m s $^{-2}$  was a tell-tale sign. Indeed, upon re-observation on 04/06/2021, it was re-detected with  $P_{SSB} = 28.6969$  ms instead of the  $P_{SSB} = 28.7058$  ms from the discovery. That made it a very good PRP with a massive binary companion candidate.

Following-up this pulsar was not as trivial as with **PSR J1306–6043**. As the  $P_{SSB}$  evolved rapidly and over a wide range, re-detections were not as simple as folding and re-optimising. Instead, we had to run PEASOUP acceleration searches on the incoming beams and look for its signal in the list of FFT candidates. After that, we folded it with PSRCHIVE to achieve more accurate  $P_{SSB}$  and  $a$  measurements, as is shown in the time–phase plots at the top of Fig. 3.16. Usually, this worked quite well despite the not-too-impressive  $S/N$ . We achieved multiple beam detections on 04/06/2024, with the best one presenting  $S/N = 9$  due to the short observing window of five minutes. Nonetheless, we performed a good localisation with SeeKAT, and regular re-observations were scheduled every few days, with two observations a few hours apart within the same day.

Solving its orbit was the first challenge we faced in the learning curve of MMGPS-L binary pulsars. It became evident after a few observations that the orbital period was shorter than the observing cadence, meaning that that  $P_{SSB}(t)$  series looked more like a random mess rather than the elegant and smooth curve provided by orbital motion, as is shown in the  $P_{SSB}(t)$  plot in Fig. 3.16. Our first attempts at solving this pulsar was to estimate  $P_b$  by fitting an ellipse to the  $P_{SSB} - a$  measurements as described in Section 2.4.2, but that failed quite miserably as the measurements were not at all described by an ellipse, as is shown in the  $a - P_s$  plot in Fig. 3.16. Nonetheless, that was a crucial hint: this is not a circular binary, but an eccentric one.

To solve a short and eccentric orbit, we had to resort to more sophisticated methods. The one that proved effective was computing the *roughness periodogram* introduced in Section 2.4.2,

$$R = \sum_{i=0}^{n-1} \left[ \frac{P_{BBS,i}(\phi_i) - P_{BBS,i+1}(\phi_{i+1})}{\phi_i - \phi_{i+1}} \right]^2, \quad (3.5)$$

but with an extra normalisation factor  $\phi_i - \phi_{i+1}$  to account for possible gaps in orbital coverage, so that points that are far apart in orbital phase during the fold do not bias the  $R$  value upwards. This search was implemented in the custom-made python script `estimateOrbit.py`, which was coded specifically to solve this pulsar but is now available at the [CandyCracker](#) repository for any other pulsar binary waiting to be solved. In addition, `estimateOrbit.py` also includes an implementations of the Lomb-Scargle periodogram of a  $P_{SSB}(t)$  series and the ellipse fit of the  $a - P_{SSB}$  measurements.

To go to the specifics, in **PSR J1208–5936**, long observations ( $T > 20$  min) or for observations with very large  $a$  values, we could actually split the observation into multiple parts and achieve several independent, close-by  $P_{SSB}$  measurements. That was done with the observations from 05/05/2021, 11/06/2021, and 16/08/2021. Finally, the roughness search was successful. The bottom of Fig. 3.16 shows a refined plot of the reciprocal roughness,  $1/R$ , against trial  $P_b$  values from measurements until 31/08/2021, with a clear peak standing out at  $P_b = 0.632$  days, well below the observing cadence. With the  $P_b$  estimate as a prior, we ran `pyfitorbit`, a python code written by the MMGPS member Gregory Desvignes that fits equation 2.18 to a  $P_{SSB}(t)$  series, measuring the Keplerian parameters from the orbital solution. Indeed, `pyfitorbit` converged to a Keplerian description with  $P_b = 0.632$  days,  $x = 4.257$  ls, and  $e = 0.348$  with a minimum  $M_c = 1.06 M_\odot$ , confirming all of previous suspicions of the system being an eccentric DNS. The orbital solution is plotted in the bottom



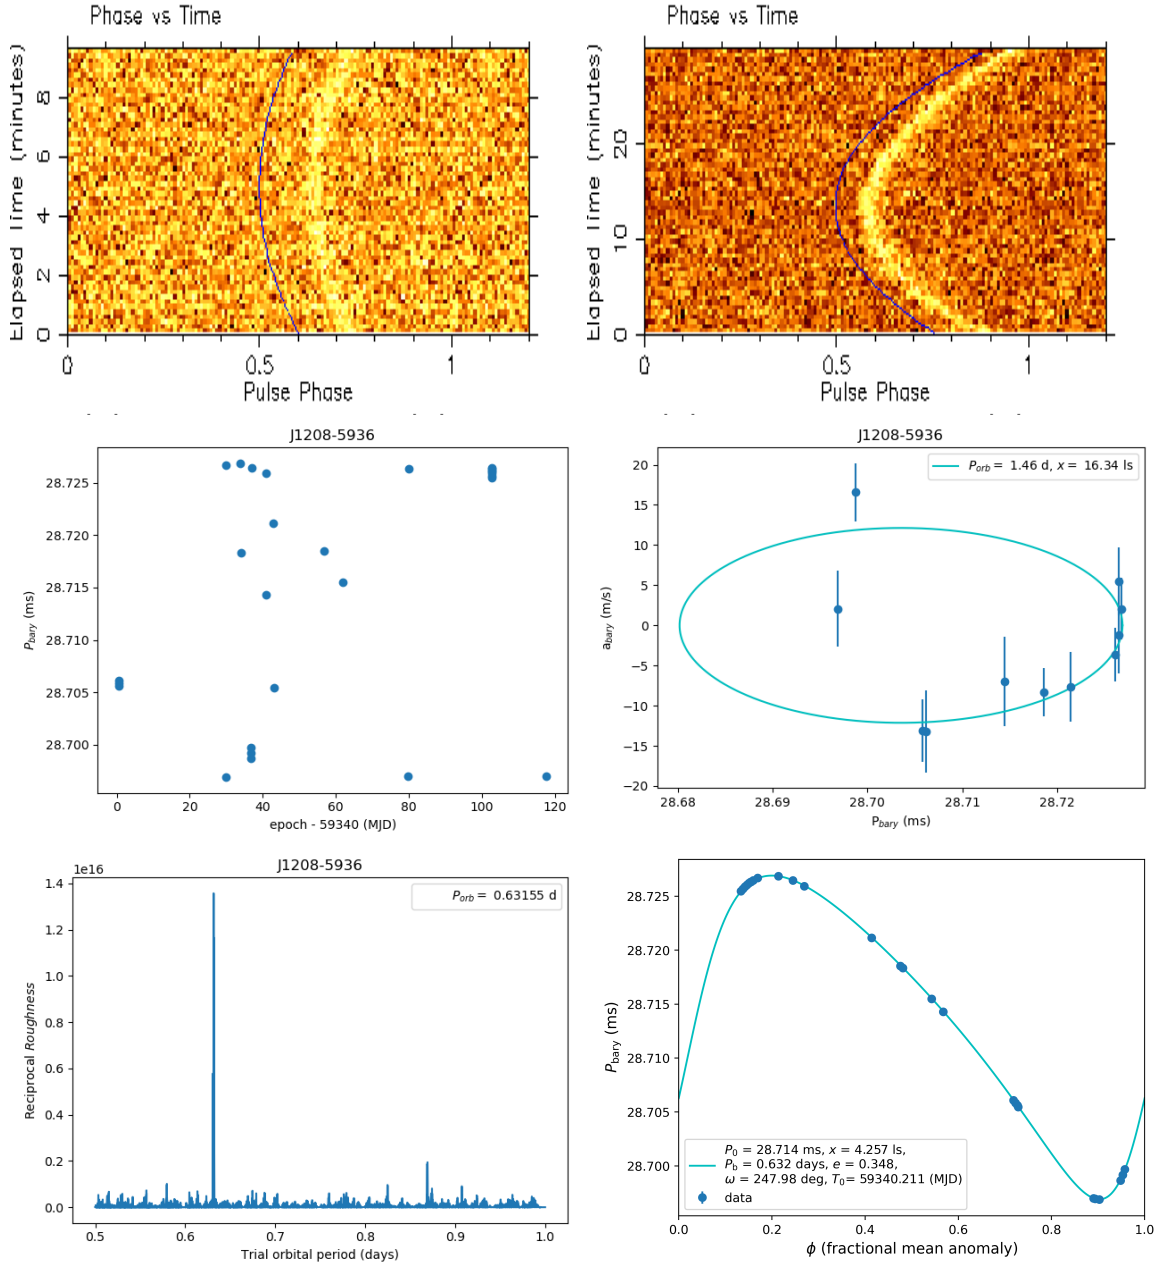


Figure 3.16: Orbital spin modulation of **PSR J1208–5936**. At the top, two folded phase–time plots from observations from 11/06/2021 and 16/08/2021 showing the  $P_{\text{SSB}}$  and  $a$  measurements (blue curve tracing the pulse). In the middle row, the  $P_{\text{SSB}}(t)$  series and  $P_{\text{SSB}} - a$  measurements from 06/05/2021 to 31/08/2021. In the  $P_{\text{SSB}} - a$  plot, the best-yet-wrong-fit ellipse is drawn with a cyan curve, which describes an erroneous circular orbit with  $P_b = 1.46$  days and  $x = 16.34$  ls. The bottom row shows the reciprocal roughness ( $1/R$ ) against trial  $P_b$  values, and the  $P_{\text{SSB}}$  measurements folded at the correct  $P_b$  values with a cyan curve tracing the predicted spin modulation from the successful orbital solution.



row of Fig. 3.16, where the the observations from 05/05/2021, 11/06/2021, and 16/08/2021 are shown as multiple close-by points in a row.

However, our job was far from done, as the next step was to achieve a timing solution and to measure PK parameters. Given the eccentric nature of the system, we were eager to measure the relativistic periastron advance  $\dot{\omega}$  and achieve a first mass constraint of the total system mass, but we had yet to achieve phase-connection for that. For that, I implemented the *dracula* into *dracula2.py*, another *python* script available in *CandyCracker*. This script is a wrapper of *TEMPO2* (Hobbs et al., 2006; Edwards et al., 2006), a popular pulsar timing software that fits timing models onto ToAs. *dracula2.py* implements time jumps between the observations and proceeds to remove them and explore the  $\chi^2$  of solutions as it introduces or removes rotation counts, as is explained in Section 2.4.6. In my implementation, I chose to remove the jumps in order of increasing time between observations, starting with the shortest.

For **PSR J1208–5936**, I derived three full-band ToAs per observation, and fed them to *dracula2.py* using the orbital solution as a starting guess. It turns out that the orbital solution was close as it quickly found a unique solution that described all of the ToAs with only data from 06/05/2021 to 18/10/2021. Mind that this baseline is less than a year long. Usually, timing solutions take longer than one year to achieve due to the yearly periodicity of the Earth’s Rømer delay, which is measured to constrain the pulsar position on the sky. Without at least a year of timing baseline, this position is very hard to constrain, but the precise localisation provided by SeeKAT thanks to the interferometric nature of MeerKAT allowed us to bypass this limitation by having RA and DEC as fixed parameters. This phase-connected timing solution already gave us an  $\dot{\omega} \approx 1 \text{ deg yr}^{-1}$  measurement, which implied a total system mass of  $M_t \approx 2.6 M_\odot$ .

From that point onward, with the pulsar already solved, timing observations were moved to a regular monthly cadence. Many more exciting measurements and science came from this pulsar, but writing them here is not the right place. For that, there is the entirety of Chapter 4 dedicated to it right after this one, and I refer you to it. But here is a teaser: we improved the mass constraints in this system and constrained the neutron star merger rate in the Milky Way. Please go check it out!

### 3.3.5 PSR J1708–4843: the vanishing pulsar

This pulsar was found quite deep in the survey. CandyJar showed it on 20/08/2021 from an observation from 30/07/2021, with a decent  $S/N = 17$ . It came with  $P_s = 16.66 \text{ ms}$ ,  $DM = 29 \text{ pc cm}^{-3}$ , and, most importantly,  $a = 11.07 \pm 0.89 \text{ m s}^{-2}$ . Therefore, it looked to be a case similar case to that of **PSR J1208–5936**, but with the advantage that it came after it. This meant that all the binary solving tools were ready and available, so this pulsar was destined to be a walk in the park.

It was not, even though at first it seemed like it would. As with **PSR J1208–5936**, we started with PEASOUP searches to re-detect the pulsar in the FFT. The first follow-up observation was made on 24/08/2021, and we achieved multiple beam detections, with  $S/N = 23$  in the best beam. The spin period had changed from  $P_{\text{SSB}} = 16.6561 \text{ ms}$  to  $P_{\text{SSB}} = 16.6624 \text{ ms}$ , and a LOS acceleration of  $a \approx -2 \text{ m s}^{-2}$  was detected. Everything was going great! Then, we re-observed it on 31/08/2021 and the pulsar was not detected. Gone, PEASOUP did not find its signal at all. We had a drought across September, where no MeerKAT observations were performed, but when the time to it observe the pulsar again on 19/09/2021, we also failed to detect it. But before panic ran amoc, we re-detected it on 28/09/2023. Thankfully, it was back! Only to disappear again during the two observations in 03/10/2021.

It turns out that **PSR J1708–4843** is a heavily scintillating pulsar. As explained in Section. 1.2.6, scintillating pulsars can be a big headache for astronomers, as changing refraction patterns in the ISM can significantly change the received flux from a source within a radio band. Sometimes, these changes are enough to almost make **PSR J1708–4843** vanish. Some of these observations are shown in Fig. 3.17, where it can be seen that the intensity of the signal changes across the radio bands and observing dates. And that was a big issue, because if we were to solve this binary we could not afford to detect it in only half of our observations.

Therefore, we set ourselves on a mission to recover all of the missing detections. After all, the pulsar had not stopped emitting, it was just too faint. PEASUP is a search software designed for large surveys, but in the way it is implemented it does not allow us to go into the deepest levels of  $S/N$  within an observation. And that is for good reason, as otherwise the MMGPS-L would have never been completed. But that meant that we had to try another angle to find vanishing pulsar.

With that in mind, we shifted to **PULSAR\_MINER**, a pulsar search software designed by TRAPUM member Alessandro Ridolfi for globular cluster surveys<sup>76</sup> (Ridolfi et al., 2021). **PULSAR\_MINER** implements an FFT-based, Fourier-domain acceleration search, which is more expensive than the time-resampling acceleration search implemented in the MMGPS-L. However, since we knew the DM value of the pulsar, we only had to perform a single de-dispersion trial, facilitating the search. In itself, **PULSAR\_MINER** is also a wrapper of PRESTO, the pulsar-search software on which it is based. One aspect that **PULSAR\_MINER** adds is a smarter candidate shifting and selection from FFT signal detection before folding pulsar candidates with PRESTO/prepfold. Unfortunately, **PULSAR\_MINER** did not fold the missing **PSR J1708–4843** detections right away, but while inspecting the list of pre-folding candidates we found its spin period amongst the many discarded signals. Upon manual folding, the missing detections were recovered with a shy  $S/N \approx 6$ . Some of these detections are shown in Fig. 3.17, where one really needs to squint the eyes to see the signal in the frequency–phase plots. All of that makes me wonder about how close we were to miss this pulsar if scintillation had been more unfortunate in the discovery observation, and about how many pulsars like that have we actually failed to discover.

That being said, with the recovered observations we achieved eight detections between 30/07/2021 and 03/10/2021, five of which had good acceleration measurements. Therefore, it was only a matter of time that we solved this orbit! As with **PSR J1208–5936**, the orbital period seemed to be shorter than the observing cadence, so we began with a  $P_b$  search. The first attempt was to fit an ellipse in the  $a - P_{\text{SSB}}$  space, and this time we were successful, as is shown in Fig. 3.18. The fit implied an orbital period of just over 12 hours, making it the most compact system we had found at that point. It also indicated a circular nature. Given the circularity, instead of a roughness search like in **PSR J1208–5936**, we implemented the Lomb-Scargle periodogram search to refine the  $P_b$  value. The periodogram presented a prominent peak at  $P_b = 0.544$  days, consistent with the ellipse fit. Upon fitting with `pyfitorbit` with this  $P_b$  as a guess, we easily found the remaining orbital parameters:  $x = 2.178$  ls and  $e \sim 0$ . With a  $M_c > 0.5 M_\odot$ , this system was placed amongst the several other known compact **PRP–CO WD** systems. Finally, deriving three full-band ToAs from each of these eight observations, we ran `dracula2.py` and easily found the phase-connected solution with barely over three months of data, all thanks to the precise localisation provided by the interferometric nature of

<sup>76</sup> Globular cluster searches are rather different from blind sky searches like the MMGPS surveys. These searches accumulate many hours in a single target (a globular cluster) and then the data is stored, allowing pulsar searchers to scrape the bottom of the  $S/N$  barrel.




Figure 3.17: Folded and acceleration-corrected frequency–phase plots from **PSR J1708–4843** observations, showing the changes in  $S/N$  across bands due to scintillation. The observations are 30/06/2021 (discovery), 24/08/2021, 31/08/2021, 19/09/2021, 28/09/2021, and 31/05/2022. The last one is added to show that scintillation can not only decrease the signal, but also increase it. All observations are between five and ten minutes long. The 31/08/2021 and 19/09/2021, 2021 observations were folded after detection with PULSAR\_MINER.

MeerKAT.

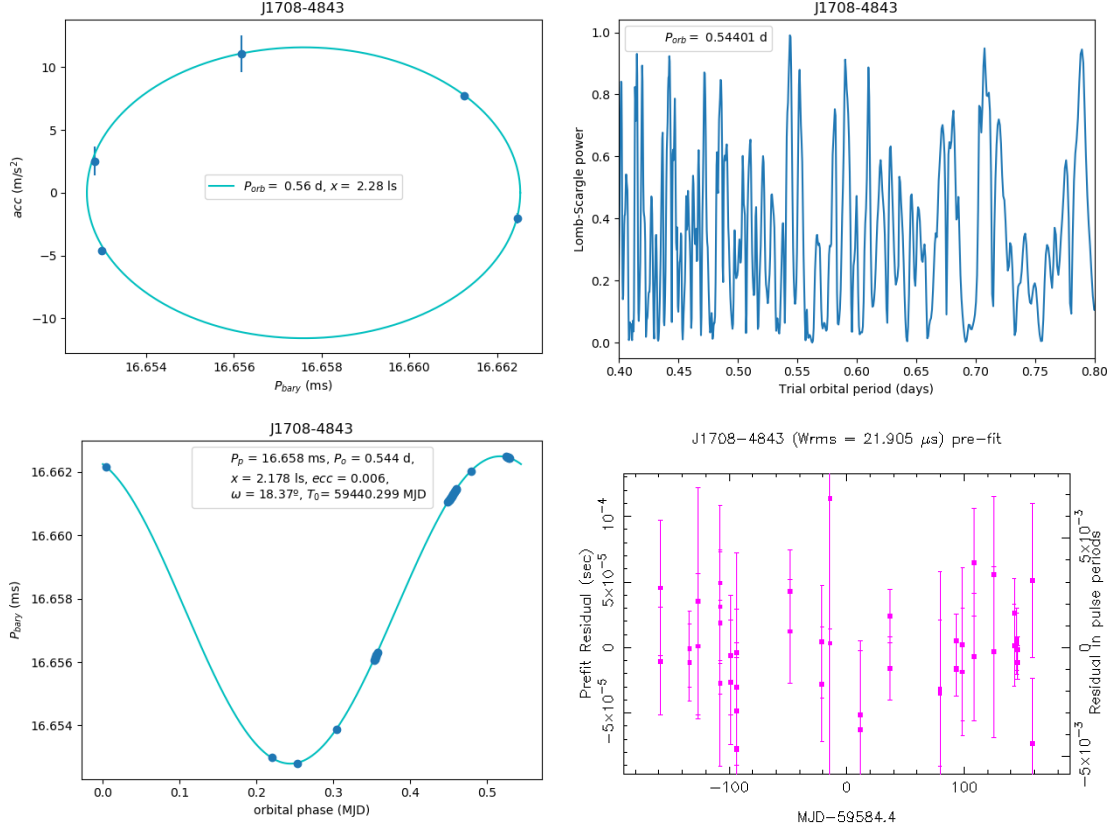


Figure 3.18: Orbital fit of **PSR J1708–4843** with the ellipse on the  $a - P_{SSB}$  space (top-left), Lomb-Scargle periodogram between  $P_b = 0.4 - 0.6$  days (top-right), the full-Keplerian fit to the  $P_{SSS}(t)$  series including data until 03/10/2021 (bottom-left), and timing residuals as a function of time including data until 09/09/2021 (bottom-right).

After that, the follow-up of this pulsar was moved to regular monthly observations. With the phase-connected timing solution, re-detection became a triviality. The most recent timing solution is from 09/09/2022, including data until that date. The currently implemented timing model is ELL1 (Section 2.4.5), and its eccentricity is so small that we can not get a proper measurement yet. Instead, we have the following constraints on the Laplace-Lagrange parameters from equation (2.43)  $\epsilon_1 = 3(8) \times 10^{-6}$  and  $\epsilon_2 = -4(5) \times 10^{-6}$ . That is in part because the timing resolution is not too great, with an average ToA uncertainty of  $\sim 20 \mu$ s due to the faint nature of the source. Given the compact nature of the system there is hope for a future measurement of orbital decay: we expect  $\dot{P}_b > -10^{-14} \text{ s s}^{-1}$  from GR alone. Additionally, a Shapiro delay measurement is potentially measurable with a dense orbital campaign. Eventually, we should be able to achieve precise mass measurements from this system.

### 3.3.6 PSR J1108–5936 and PSR J1543–5439: soft-looking but hard-to-bite pulsars

I have decided to lump these two pulsars together because they share a similar story, although the pulsars themselves are quite different from each other.

**PSR J1108–5936** was discovered on 23/09/2021 from an observation in 09/09/2021. It presented a folded  $S/N = 15$ ,  $P_s = 4.28$  ms,  $DM = 233$  pc cm $^{-3}$  and  $a = -2 \pm 1$  m s $^{-2}$ . It did not present many noteworthy characteristics, other than a rather wide pulse profile with a pulse duty of  $\sim 30\%$  and a rather steep spectrum, as is shown in Fig. 3.19. It was re-observed and localised on September 28, and its binary nature was confirmed with the change in orbital period. Instead of searching with PEASOUP, for this pulsar we proceeded with acceleration searches with PULSAR\_MINER directly, recovering its signal and measuring  $P_{SSB}$  in every observation without a problem. Subsequent observations showed that, once again, its orbital period was below the observing cadence, but by that point the expertise in this are made solving it a breeze. We identified its orbital period with the roughness periodogram with only eight observations in 21/02/2022, and then successfully ran `pyfitorbit` to find the remaining orbital parameters, finding  $P_b = 0.907$  days,  $x = 1.073$  ls,  $e \sim 0$  and  $M_c > 0.15 M_\odot$ . The orbit is drawn at the bottom of Fig. 3.19. A most-classical, compact **MSP–He WD** system. One pulsar done.

**PSR J1543–5439** was discovered in 01/02/2022 from a 15/01/2022 observation. It was one of the brightest discoveries, with  $S/N = 50$  in CandyJar. With  $P_s = 4.31$  ms and  $DM = 102$  pc cm $^{-3}$ , but no detectable LOS acceleration and a rather uninteresting profile (see Fig. 3.19), it looked like an isolated MSP, or an MSP with a very light orbital companion or in a very wide orbit, like **PSR J1306–6043**. Indeed, it turned out to be a **PSR J1306–6043** clone. As in that pulsar, its spin modulation was so subtle that we did not need to perform searches to detect it: just refolding and optimizing yielded a good re-detection and  $P_{SSB}$  measurement, always without any measurable acceleration. This was so much the case that we did not even bother to do an intense follow-up like for the other binary discoveries. The confirmation and localisation observation was performed in 11/02/2022, and it was not observed again until 08/03/2024, and then again in 04/04/2024 and 06/04/2024. At that point, I looked at the time series and said, “Yes, I see the sinusoid”. Then, I tampered with the Keplerian parameters manually until I got a spin modulation curve that coincided with the five points in the  $P_{SSB}(t)$  series, and then ran `pyfitorbit` with them as a prior to find the solution while keeping  $e = 0$  as a fixed value. Needless to say, as is shown in Fig. 3.19, it was a success and I felt very proud about it. Quick and elegant, as all pulsars should be. Indeed, it is a wide-orbit pulsar with  $P_b = 105.13$  days,  $x = 95.8$  ls,  $e \sim 0$  and  $M_c > 0.23$ .

And then, we made three ToAs per observation for both pulsars, ran `dracula2.py`, and found beautiful timing solutions with the ELL1 models. That is what I would like to say, but can not, because it is definitely not true. These two pulsars are playing hard to solve, and have resisted all of my attempts at achieving a phase-connected timing solution so far. Every time `dracula2.py` runs, solutions start piling up in numbers of thousands upon thousands at every jump removal until the program runs out of computing time. I have attempted, for example, to restrict the number of solutions that can arise from every branch at every jump removal to 20 (this means 20 times the amount of solutions that have reached at the jump), or to be extremely restrictive and allow only the lowest of  $\chi_{red}^2$  values. In the best case, it manages to reach the end with dozens of viable solutions, only for them to break when new ToAs from new observations are added to the fit.

Obviously, some fundamental issue is at play here. Apparently, this is an issue that Prajwal Voraganti Padmanabh also faced with **PSR J1306–6043**, so there is a lesson to learn here: do not be fooled by apparently easy pulsar binaries. I have a couple theories of why are we facing these issues. In the case



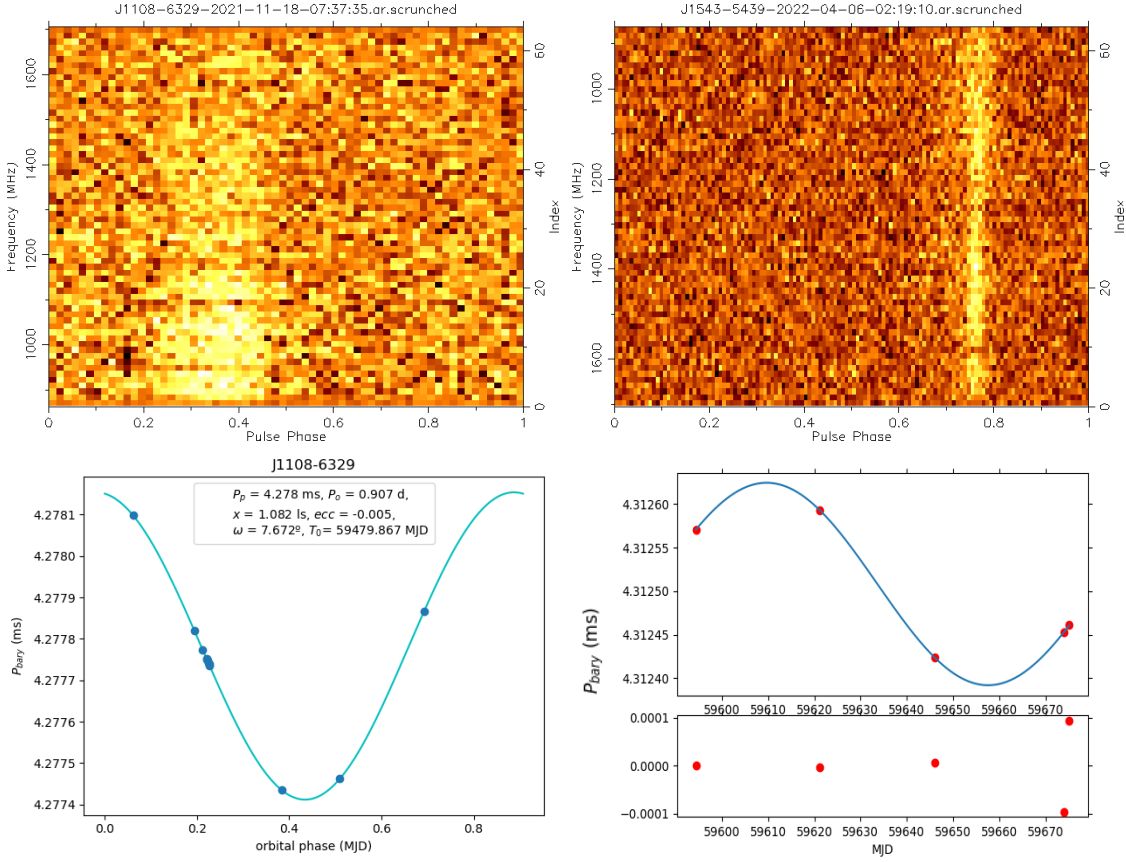


Figure 3.19: Frequency-resolved profiles and orbital solutions of **PSR J1108–5936** (left) and **PSR J1543–5439** (right). The **PSR J1543–5439** orbital plot is captured from the `pyfitorbit` user interface, showing the spin evolution in time rather than in orbital phase. The bottom is this same plot shows the *spin residuals*, the difference between the measured and predicted  $P_{\text{SSB}}$  values.

of **PSR J1108–5936**, the wide pulse profile also widens the ToA uncertainty. Currently, its average ToA uncertainty lies at  $\sim 15$   $\mu\text{s}$ , which for a  $P_s = 4.28$  ms pulsar implies a spin-phase fractional uncertainty of 3%. Compared to, for example, the 1% of **PSR J1208–5936** and **PSR J1708–4843** during the search for phase-connection, it is a large number. It means that its ToAs are consistent with potentially more phase turns between observations, allowing them to accommodate more potentially viable solutions at each jump removal. That, of course, makes it trickier to solve. In the case of **PSR J1543–5439**, the average ToA uncertainty is  $\sim 8$   $\mu\text{s}$  thanks to its narrower profile and larger brightness, but its wide orbit means that we are probing few orbital cycles with our timing baseline, implying a larger uncertainty in  $P_b$ . That means that we will have to wait until more observations covers several orbital cycles to achieve a phase-connected solution, which would take more than a year of observing.

However, while these issues do not help, the main problem comes from the spin nature of pulsars. Previous pulsars solved by `dracula2.py` were all PRPs with  $P_s > 15$  ms, while both **PSR J1108–5936** and **PSR J1543–5439** are MSPs with  $P_s \approx 4$  ms. For the PRPs, we chose to set  $f_1 = 0$  and the SeeKAT-derived sky position as a fixed parameters. But a true  $f_1$  value of  $\sim 10^{-20}$  or an off-set on the position by a few milliarcseconds can easily lead to timing residuals of an amplitude of a few

ms. For PRPs with  $P_s > 15$  ms with timing baselines shorter than one year, that is not an issue as these residuals may be absorbed as a deviation in  $f_0$ . However, for a  $P_s \approx 4$  ms MSP, this can spell disaster. Therefore, we have to account for these parameters when running `dracula.py`, resulting in an exponential increase of solutions due to the large array of  $f_1$ , RA, and DEC values taken by intermediate solutions. To work around this issue, on Paulo Freire's suggestion, I am currently experimenting with a  $\chi^2$  penalisation factor to equation (2.35),

$$\chi_{\text{new}}^2 = \chi^2 + \left( \frac{\text{RA} - \text{RA}_{\text{SK}}}{\Delta \text{RA}} \right)^2 + \left( \frac{\text{DEC} - \text{DEC}_{\text{SK}}}{\Delta \text{DEC}} \right)^2, \quad (3.6)$$

where RA, DEC,  $\Delta \text{RA}$ , and  $\Delta \text{DEC}$  are the positions of pulsars according to the TEMPO2 solutions and their uncertainties, and  $\text{RA}_{\text{SK}}$ ,  $\text{DEC}_{\text{SK}}$  are the SeeKAT pulsar positions, so that `dracula2.py` gives priority to solutions that are close to the localised position when looking at  $\chi_{\text{red}}^2 = \chi_{\text{new}}^2 / N_{\text{free}}$  from equation (2.50). This way, the interferometric nature of MeerKAT should also give us a prior when solving for pulsars.

Nonetheless, these two systems still show promise for when phase-connected timing solutions are available. **PSR J1108–5936** in particular could have a measurable relativistic orbital decay of  $\dot{P}_b > 5 \times 10^{-15} \text{ s s}^{-1}$ , which if combined with a potential Shapiro delay detection could lead to precise mass measurements. **PSR J1543–5439**, on the other hand, would have to rely on a potential Shapiro delay detection from mass constraints. However, its better use would be to be added as a low-frequency gravitational-wave probe in the PTA experiments due to its stable rotation and bright nature.

### 3.3.7 PSR J1015–5358: the evil twin of PSR J1708–4843

We had a headache with **PSR J1708–4843** due to its scintillating nature, which made it disappear on half of its observations. That was nothing compared to **PSR J1015–5358**, which decided to disappear before we could localise it with multibeam detections. Hints of its fidgety nature showed up already at discovery: it was found with a measly folded  $S/N = 11$  in the candidates. I considered it a T2 candidate until it showed up a second time in CandyJar. Apparently, its  $S/N$  was so low that the filtering step of the MMGPS pipeline failed to relate two harmonics of the same signal, so it showed up twice, first at  $P_s = 10.4$  ms with  $S/N = 11$ , and then at  $P_s = 20.8$  ms with  $S/N = 10$ . This was on 04/04/2022, from a 17/02/2022 observation, and both of these candidates were already shown in Fig. 3.7. It turned out that the fundamental one is the  $P_s = 20.8$  ms signal. In addition this pulsar is at  $\text{DM} = 31 \text{ pc cm}^{-3}$  and it was detected with a whooping LOS acceleration of  $a = -28 \pm 3 \text{ m s}^{-2}$ .

After this first detection, excitement ran rampant. Could it be another DNS? A very compact system? Or perhaps a very heavy companion? The pulsar was so faint that it only yielded a single-beam detection during discovery, so we had to confirm it. It was re-observed with a 0.9 coherent beam overlap for localisation on 23/04/2022, but the signal was not found by PEASOUP on any of the beams. The same happened on the two observations on 10/M05/2022, so we started getting nervous. At that point, we had already gone through something similar with **PSR J1708–4843**, but the lack of proper localisation or any detection besides the discovery added extra layers of doubtfulness. Where we observing at the wrong position? Had the original signal been just a noise fluke? Follow-up observations all came in with twenty coherent beams around the discovery position, so running PULSAR\_MINER was to be too time consuming. Instead, we proceeded with a simpler strategy this time around: we went directly to PRESTO, de-dispersed every incoming filterbank with the DM value



from the discovery with `prepdata`, performed an FFT with `realfft`, de-reddened with `rednosie`, and applied `accelsearch` to perform Fourier-domain acceleration search on the individual FFTs. Then we identified all signals close to the fundamental spin period in the candidate files and folded them with `prepfold`, no exception. All of that was done by a code snippet I wrote, so in a way it was my own tiny search pipeline.

This implementation was successful. The first observation where we achieved a detection was on 12/06/2022, with eight different multi-beam detection and a fold  $S/N = 26.5$  at the best beam. The ISM gods had smiled upon us and, for once, boosted this pulsar’s brightness to our favour. Immediately, we proceeded to perform localisation with `SeeKAT` and retroactively went on to search for the pulsar signal at the best beam of previous follow-up observations. We successfully recovered the signal with very shy  $S/N = 10^{77}$ . Despite the low  $S/N$ , we achieved five detections with  $P_{\text{SSB}}$  and  $a$  measurements, with  $a$  values reaching up to  $40 \text{ m s}^{-2}$ , making it the most accelerated pulsar discovered in the MMGPS-L. Fig. 3.20 shows the frequency-resolved profile of this pulsar in some of these defections, both faint and bright, and as a curiosity it seems that a second, very faint pulse component exists  $\sim 0.5$  turns away from the main pulse, barely visible in the frequency-resolved profiles of Fig. 3.20. This points towards another orthogonal rotator, where the magnetic-to-spin axis angle is  $\sim 90$  deg, that is showing us both of its magnetic poles.

We solved the orbit with only five detections. That is mostly because the  $a$  measurements were very significant, as is shown in the  $a - P_b$  ellipse diagram in Fig. 3.20. The ellipse fit showed it has a circular orbit, which made us a bit sad at first because we really wanted this one to be another eccentric DNS with many significant PK effects to measure. However, we quickly realised that despite its circularity this system would still be a scientific trove. The ellipse fit indicated an orbital period of just eight hours, the most compact one in the MMGPS-L, and that was already very promising. These highly accelerated observations allowed for several independent measurements of  $P_{\text{SSB}}$  within single observations, which enabled a smooth fit of the Keplerian parameters with `pyfitorbit`, shown in Fig. 3.20. With  $P_b = 0.333$  days and  $x = 2.989$  ls, the minimum companion mass is an impressive  $M_c = 1.18 M_\odot$ , potentially making it the most compact, massive **PRP–ONeMg WD** system known to date. That is not to be understated: the previously known conformed and most compact **PRP–ONeMg WD** system is **PSR J2222–3146**, with  $P_b = 2.45$  day and  $M_c = 1.33^{+0.14}_{-0.12} M_\odot$ . This means that this pulsar is a great opportunity for testing the Chandrasekhar white dwarf mass limit and experimenting with binary evolution models and testing alternative theories of gravity. In particular, to test multi-polar gravitational-wave emission in systems with a large asymmetry in the compactness of the components (a neutron star vs. a white dwarf) via the measurement of PK effects via pulsar timing, as is discussed in Section 1.4.5. From this system, we expect a very measurable, GR-predicted  $\dot{P}_b > 10^{-13} \text{ s s}^{-1}$  and even a significant Shapiro delay if the inclination angle is a favorable one. Measuring these parameters will be key in testing whether GR holds its ground or not.

But for that, we need first a timing solution. A phase-connection solution was already achieved with `dracula2.py` using the ELL1 model on July 2022, whose timing residuals are shown in Fig. 3.20. The latest solution available to me being from August 2022, which presents an eccentricity constraints of  $\epsilon_1 = -2.5 \pm 0.8 \times 10^{-4}$  and  $\epsilon_2 = -2.2 \pm 0.6 \times 10^{-4}$ . After that point, the pulsar was handed onto Vivek Venkatraman Krishnan for further follow-up, timing, and mass measurements.

<sup>77</sup> Please note that the  $S/N$  estimates are software-dependant. `PulsarX` and `PRESTO/prepfold` tend to give low  $S/N$  estimates, while `PSRCHIVE/pdmp` like in this case tends to give larger  $S/N$  values. That is not an issue unless we are trying to perform localisation, in which case we just have to use the same software to compute all  $S/N$  consistently within a single observation.

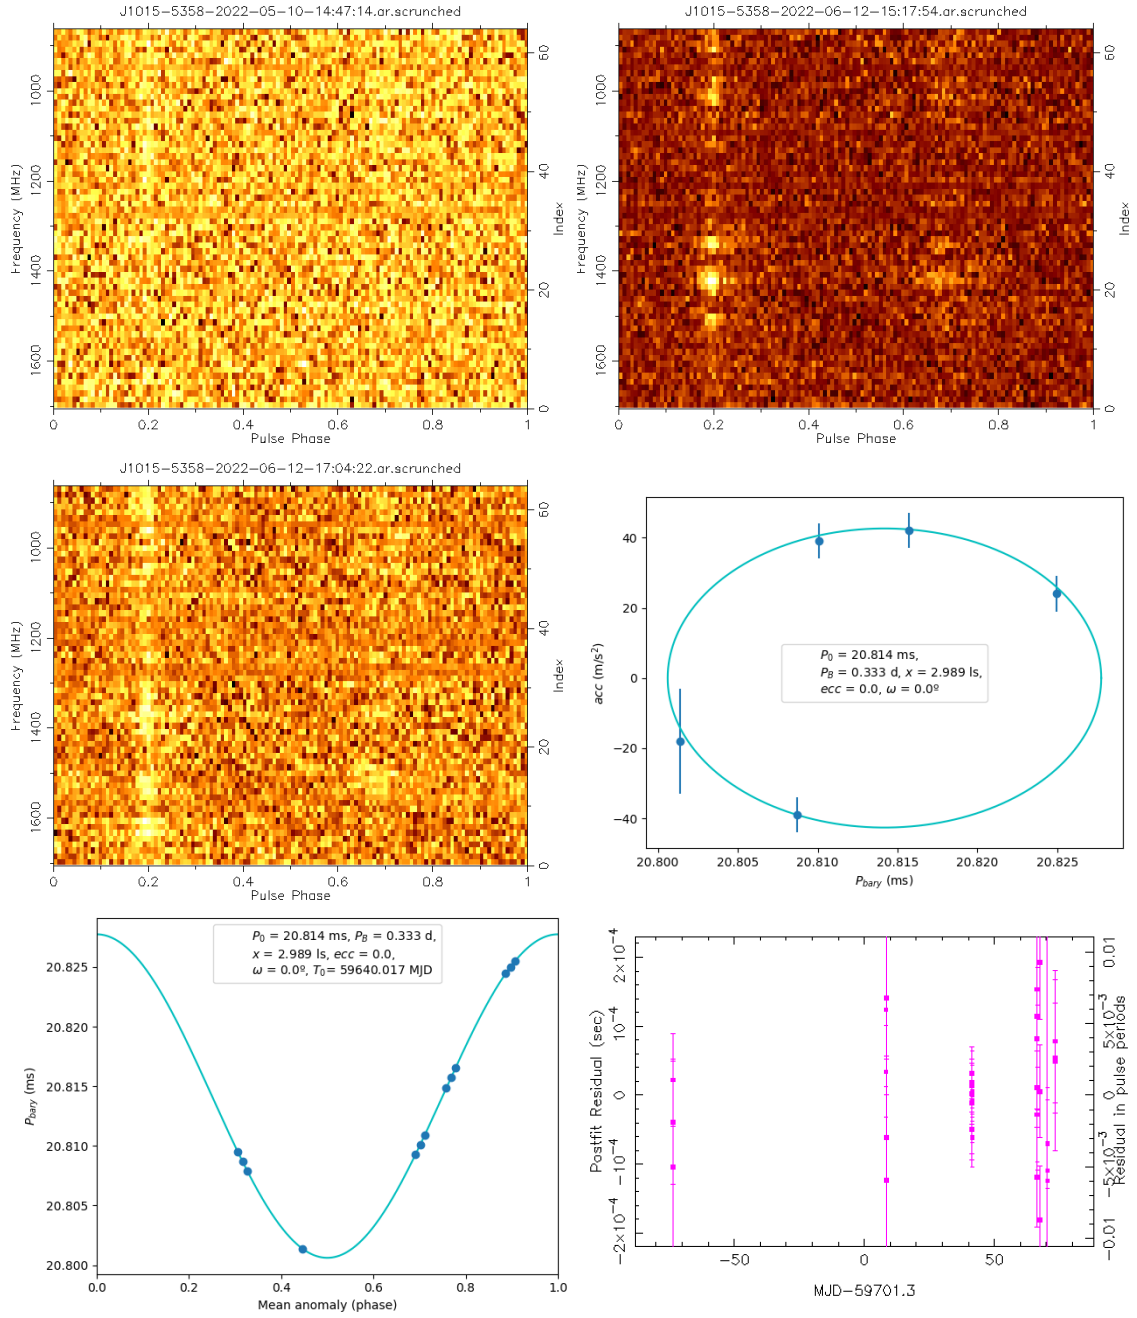


Figure 3.20: Folded and acceleration-corrected frequency–phase profiles of **PSR J1015–5358** from 10/05/2022 and 12/06/2022, showcasing the large variability in observed flux from this pulsar. That is followed by the ellipse fit in the  $a - P_{\text{SSB}}$  space, an orbital fit to the  $P_{\text{SSB}}(t)$  series, and the timing residuals from the first phase-connected timing solution.

### 3.4 SNR G340.6+0.3 and the MeerKAT S-band receivers: the search for the missing pulsar

If you wanted to read more about search-based science projects with MeerKAT in this thesis, you are in for a tiny bitter-sweet treat. **SNR G340.6+0.3** is a dusty, optically-faint, radio-bright supernova remnant located in the Galactic plane, with J2000.0 RA = 16:47:41 and DEC = -44:34:00, or  $l_G = 340.6$  deg and  $b_E = 0.34$  deg (Chawner et al., 2020). It was discovered with the [now-decommissioned Molongo telescope](#)<sup>78</sup> in 1973 (Kluzniak et al., 1988) and it has been studied through the years (e.g. Dubner et al., 1996), but without any evidence for pulsar activity being found in its interior (Chawner et al., 2020). It has a shell-like structure approximately 6 arcmins across, and it is estimated to be at a distance of 15 kpc (Kothes and Dougherty, 2007), but all images so far were quite blurry owing to the low-resolution of radio telescopes.

Then MeerKAT came along, and it promised and unprecedented high-resolution view of the morphology of radio bright supernova remnants thanks to its interferometric nature, as demonstrated by [the spectacular image of the Galactic center at 1.4 GHz image published in 2022](#) (Heywood et al., 2022). The highest resolution prospects came from the high-frequency S-band receivers, so **SNR G340.6+0.3** was included in a test imaging field at  $f = 2.6$  GHz. The image, shown in Fig. 3.21, was a success that showed the nebulous features of the supernova shell with unprecedented detail. However, it brought a promising surprise with it: a very tempting, previously unseen point source at its center, highlighted in Fig. 3.21 by the black markers.

When pulsars are observed in the imaging domain, which implies integrating all of flux over time while ignoring the pulsed nature, they appear as rather unremarkable point sources due to their small size, as if they were any other radio-bright star in the field. Unfortunately, MeerKAT the S-band receivers had not yet been calibrated, so it was not possible to measure if any fraction of the received light was polarised. Therefore, the only way to know if there was a pulsar in that faint dot was to search for it in the time domain. And thus, this search became a little side-project of mine. Now, it is time for me to tell how I performed this search that, sadly, ended without success. Nonetheless, it became the first pulsar search performed on data from the MeerKAT S-band receivers, making it a worth-while experiment with the potential for a significant scientific outcome.

The search observation was performed on 12/06/2022 and it was 20 minutes long. The chosen S-band receiver configuration was S3, which provides radio coverage in the  $f = 2.4 - 3.3$  GHz band. Like MMGPS observations, it came with a coherent beam tiling of 0.5 overlap, all of which were stored in APSUSE. A first, preliminary search was actually performed by Ewan Barr on all coherent beams with `peasoup`, without success. After that, I searched the filterbanks from seven coherent beams: one at the central position of the point source, and a belt of six more around it<sup>79</sup>. Of course, the first thing I did was to produce a frequency mask from the central beam. Rather than cleaning RFI, this was done to probe the RFI environment in the MeerKAT S3 band, an exercise that had not been yet performed on filterbanks. The masks was computed with `PRESTO/rfifind` and is shown on the right side of the of Fig. 3.21. It it selected less than 10% of the frequency channels. This serves to show how much cleaner is the S-band at S3 compared to the L-band, where we used to flag more than 20% channels consistently when implementing frequency masks in the MMGPS-L. These masks

<sup>78</sup> The Molongo telescope, aside from discovering supernova remnants and being a classic pulsar discovery machine, has also performed masterpiece experiments such as [broadcasting the sound of a LEXUS car engine to the Galaxy](#).

<sup>79</sup> It's hexagons all the way down.




Figure 3.21: Left: image of **SNR G340.6+0.3** at 2.6 GHz, taken by MeerKAT and provided to me by Sarrvesh Sridhar. It presents an unprecedented resolution on this object. Within the supernova remnant, the inconspicuous point-source is highlighted by the two black markers. Right: RFI signals picked up by **rfifind** on the filter bank from the central beam from the 06/12/2022 observation. The identified signals are marked in red.

were not implemented, as like in the MMGPS, it was deemed that **Pulsar/filtool** would perform a better job at cleaning the RFI without masking.

After this initial test, I performed two small-scale searches on the central beam and the six around it: one FFT-based, Fourier-domain accelerated search with **PRESTO**, and one FFA-based search with **riptide** (Morello et al., 2020; Morello et al., 2023). The **PRESTO**-based search was implemented to account for any possible LOS acceleration in the source, and it essentially was a re-use of the search scripts I implemented to re-detect **PSR J1015–5358** in Section 3.3.7, but with the inclusion of multiple DM steps. During the search, I considered signals with a pulse widths down to 0.35 ms, which given the high explored frequencies implied a DM step of  $\Delta\text{DM} \approx 1 \text{ pc cm}^{-3}$ . The explored DM ranged from 0 to 3000  $\text{pc cm}^{-3}$ , which essentially allowed us to detect signals up to a distance of 50 kpc in that sky position. Finally, although it is not expected from a young pulsar born in a supernova, I added acceleration to the search to account for any binary motion. **PRESTO/accelsearch** does not work with acceleration ranges, but rather with Fourier frequency derivative ranges,  $z = \dot{\nu}T^2$ , where  $\nu$  is the Fourier frequency and  $T$  the observation length, which implies a scaling of the equivalent LOS acceleration of the pulsar signal

$$a = c \frac{\dot{\nu}}{\nu} = c \frac{z}{\nu T^2}, \quad (3.7)$$

(see e.g. Andersen and Ransom, 2018). Therefore, choosing a range of  $z = \pm 250$  implies a max acceleration range of  $a = \pm 53 \text{ m s}^{-2}$  at  $P_s = 1 \text{ ms}$  and  $a = \pm 5300 \text{ m s}^{-2}$  at  $P_s = 100 \text{ ms}$ , making it

more than enough to cover any possible acceleration, which is not expected for in the first place given that any pulsar in **SNR G340.6+0.3** should be young and isolated. Finally, a small code snippet of my own making clusters candidates of similar  $P_s$  and DM values and chooses the highest  $S/N$  candidate amongst them, and the remaining candidates are folded and plotted with PRESTO/prepfold.

On the other hand, the FFA-based search with `riptide` did not include any acceleration. Accounting that the FFT-search should find any signal with  $P_s < 100$  ms, the FFA search only searched for signals with  $P_s > 100$  ms in the same DM range, using 256 phase bins for folding during the  $P_s = 0.1 - 0.5$  s search, 512 phase bins during the  $P_s = 0.5 - 2.0$  s search, and 1024 bins during the  $P_s = 2.0 - 120$  s search. It should be noted that, as explained in Section 2.3.4, the FFA algorithm becomes more efficient at searching through ranges of increasingly-valued large periods, so outside of the de-dispersion steps this was a very cheap search in computational terms. Additionally, `riptide` already has de-dispersion, folding, and plotting routines which makes our job very easy. For plotting, I scrunched the candidates from the  $P_s = 0.1 - 0.5$  s to 128 spin phase bins, and the  $P_s = 0.5 - 120$  s candidates to 256 spin phase bins.

With that, hundreds of candidate plots were produced after running both of these pipelines. Unfortunately, the vast majority of them were either noise with some RFI candidates that had survived the wrath of `filtool`. Some good-looking candidates are shown in Fig. 3.22. The most promising one is shown at the very bottom of Fig. 3.22. With  $P_s = 136.616$  ms and  $DM = 1331$  pc cm<sup>-3</sup>, it is one of the least noise-like looking candidates and its DM is compatible with the distance to the supernova remnant. However, the same source was re-observed on 16/09/2023, in the S1 configuration<sup>80</sup> and none of the most promising candidates were re-detected, indicating that indeed, they were just good-looking noise. From the radiometer equation in equation (2.9), we estimate that this search was sensitive down to signals with with a mean flux  $F_m \approx 30$  mJy and a pulse duty of  $\sim 10\%$  within the searched band. That means that any existing pulsed signal is to be below this threshold. At a distance of 15 kpc (Caswell et al., 1983), this implies a maximum pseudo-luminosity of  $L_m \approx 27000$  mJy kpc<sup>2</sup> at 2.4 GHz, much higher than that of any known pulsar, especially at 2.4 GHz.

The lack of a pulsar discovery is not entirely surprising, as there are several reasons to believe that this point source does not contain an active pulsar. The first red flag is that **SNR G340.6+0.3** has been imaged at 0.8 – 1.4 GHz on several occasions before before (Caswell et al., 1983; Dubner et al., 1996; Whiteoak and Green, 1996; Kothes and Dougherty, 2007) without any detection of a point source prior to MeerKAT. If the point-source identified in the MeerKAT test image at  $\sim 2.6$  GHz had really been a pulsar, given the steep spectrum of pulsars, it should have been even brighter at lower frequencies. Even today there are published L-band images of this field without any detection of this source (Goedhart et al., 2024). Additionally, if there was an active pulsar inside this nebula, it would be accompanied by a pulsar wind nebula of excited gas around the pulsar, a *plerion*, which has been seen in many other sources<sup>81</sup> However, that does not mean the point-source in in Fig. 3.21 is spurious. But instead of a pulsar, it probably is a background source far behind the supernova that happens to coincide visually with the center of **SNR G340.6+0.3**. Perhaps it is a background star within the Milky Way, or perhaps the active nucleus of a distant galaxy hundredths of millions of light years away.

<sup>80</sup> In fact, it was re-observed along with the three MMGPS-S discoveries as part of a big confirmation observing batch.

<sup>81</sup> You can see one of such plerions within **SNR G0.9+0.1** in the [Galactic center image](#) by Heywood et al. (2022).



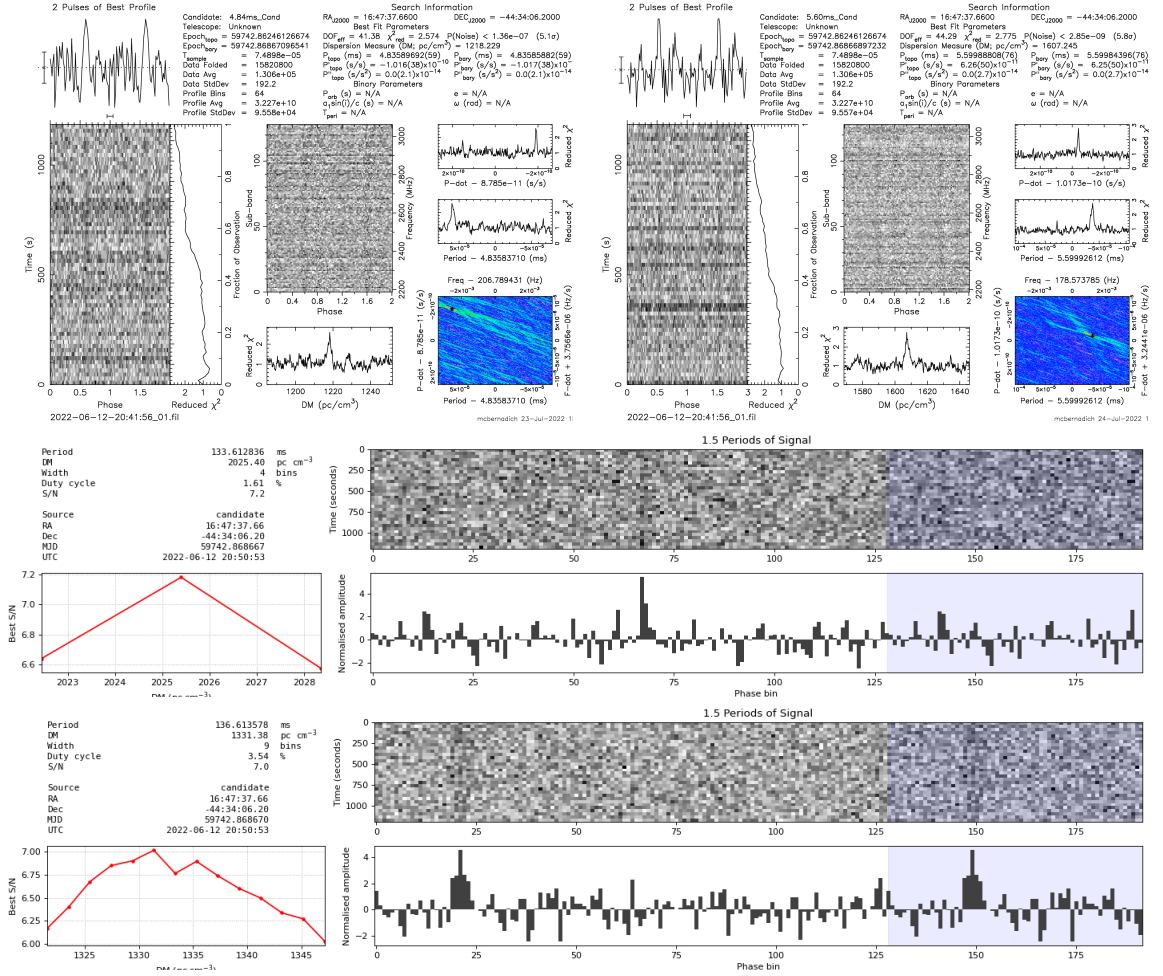


Figure 3.22: Promising pulsar candidates from the SNR G340.6+0.3 observation from 6/12/2022. At the top, two candidate plots from the PRESTO-based search are shown. At the bottom, two candidate plots from the riptide-based search are shown. None of these candidates, and several others, were ever confirmed.



---

## The eccentric DNS PSR J1208–5935 and a neutron star merger rate update

---

### 4.1 Context and contributions

In Section 3.3.4 I teased amazing science results from the MMGPS-L discovery **PSR J1208–5936**, the most compact of our DNS discoveries. If you are looking for them, you have reached the right chapter. This is a summary of Colomi Bernadich et al. (2023), *A&A*, **678**, A187, with the addition of further explanations from some of the methods that were not explained in Chapter 2, and with some plots that were not included in the original paper. The publication is included in original form in Appendix B, so please do check it if you want to see the original work.

As the first author, I was responsible for writing all of the text in the publication. I also performed all of the data analysis described in it. That is: localising **PSR J1208–5936**, solving its orbit, performing the timing analysis, constraining the component masses, modelling the scattering of the pulsed emission, searching for the RM, searching for pulsations from the putative neutron star companion, constraining the neutron star merger rate, and simulating the evolution of PK parameter measurements in the future. I also performed remote operations of the Parkes/Murriyang telescope to collect data for **PSR J1208–5936** and other TRAPUM discoveries under the observing project P1054, led by Marta Burgay. Co-authors that follow my name in the list are: Vishnu Balakrishnan, Ewan B. Barr, Marina Berezina, Marta Burgay, Sarah Buchner, David J. Champion, Weiwei Chen, Gregory Desvignes, Paulo C. C. Freire, Kathrin Grunthal, Michael Kramer, Yungpeng Men, Prajwal V. Padmanabh, Aditya Parthasarathy, Denisha Pillay, Isabella Rammala, Shallini Sengupta and Vivek Venkatraman Krishnan. Each of them have made significant contributions regarding data collection, the building of the MMGPS pipelines that enabled the discovery of **PSR J1208–5936**, participating in the candidate inspection, providing codes and programs which I used for the science analysis or helping me build them, or providing technical and/or scientific guidance during the analysis. And with all of that being said, let us move forward towards the scientific summary of the paper.

## 4.2 Summary of the publication

### 4.2.1 Orbital campaign

This continues from the point we left it in Section 3.3.4, when a phase-coherent timing solution had been achieved with `dracula2.py`. After that, observations were moved to a monthly cadence with MeerKAT, and a dense, semi-coherent orbital campaign accumulating 11 hours in total was performed on March 2022 to attempt a detection of the Shapiro delay.

An orbital campaign is a cycle of consecutive observations in a short span of time with a specific aim, usually a Shapiro delay measurement. In Section 2.4.5, and particularly in Fig. 2.16, it was shown that much of the Shapiro delay amplitude is not actually measurable and is instead absorbed into the Rømer delay. However, Fig. 2.16 also shows that there are selected orbital phases where the *residual amplitude* described by  $h_3$  and  $\varsigma$  has maximum deviation from the flat residuals, both positive and negative: three peaks and three troughs, with the highest peak remaining at superior conjunction. To achieve a good measurement of the Shapiro delay, we want to target these maximal points in time with our telescopes and measure ToAs specifically there. In addition, it is also favourable to perform an observation at the time of passage of periastron, which is when the pulsar moves at maximal speed, for optimal constraints on  $\omega$  and  $\dot{\omega}$ .

Here are some details on how the orbital campaign went which are not specified in the publication. To decide the observing times, we simulated ToAs with the fake plug-in of the timing software TEMPO2, while including a Shapiro delay signal under the assumption of  $i = 60$  deg. We then fitted for the timing model while ignoring the Shapiro delay, resulting in the residuals plot shown in Fig. 4.1 in this chapter. From the residuals, we extracted selected observing times from simulated ToAs at the desired orbital phases. We had 14 observations accumulating a total 11 hours. The respective orbital phases of each of these observations are highlighted by red circles in Fig. 4.1.




Figure 4.1: Simulated residual Shapiro delay amplitude in **PSR J1208–5936** assuming  $i = 60$  deg. The red circles mark the orbital phases covered by observations.

### 4.2.2 Timing and mass measurements

After the orbital campaign, the science experiment continued as is described in the paper. The observations from the orbital campaign were joined together to form a high- $S/N$  timing templates, and we derived ToAs from the discovery observation in 04/06/2021 up until 10/05/2022. We fitted these ToAs to the DDH model (DD with  $h_3$  and  $\varsigma$  instead of  $r$  and  $s$ ), leading to an extremely significant measurement of the rate of periastron advance,  $\dot{\omega} = 0.916(1) \text{ deg yr}^{-1}$ , which, under the assumption of GR, implies a total system mass of  $M_t = 2.586(5) M_\odot$ . However, the other three measured parameters were not as significant: the Einstein delay,  $\gamma_E = 2.93(98) \text{ ms}$ , the orthometric amplitude of the Shapiro delay,  $h_3 = 1.10 \pm (0.97) \mu\text{s}$ , and the orthometric ratio of the Shapiro delay,  $\varsigma = 0.41 \pm 0.18$ , where all constrained only to  $3\sigma$  significance or less. It could be said that the Shapiro delay is not detected at all, despite the efforts from the orbital campaign.

Nonetheless, a low-significance detection of the Einstein delay and no Shapiro delay detection are still very useful. Upon implementation of equation (2.59), the non-detection of Shapiro delay constraints and their uncertainty the inclination angle to  $i < 65 \text{ deg}$ , while the mild Einstein delay detection also puts a maximum cap to the companion mass of  $M_c < 1.7 M_\odot$ . These two constraints work against each other by putting limits on the opposite side of the  $M_c - M_p$  and  $M_c - i$  parameter spaces. At the same time, the very significant  $\dot{\omega}$  measurement also constraints the parameter space to the very narrow line defined by  $M_c = 2.586(5) - M_p$ . As is shown in the mass diagrams in Fig. 2 of the publication, that results in significant global constraints of the parameter space to a strip following the  $M_c = 2.586(5) - M_p$  line and within  $73 < i < 41 \text{ deg}$ .

For individual constraints on  $M_p$ ,  $M_c$  and  $i$ , we marginalised the probability distributions and quoted their median value as the measurement, and their 68% percentiles as their  $1\sigma$  uncertainties. From the product of individual PK parameters probabilities measured with the DDH model using equations (2.59) and (2.60), we measured  $M_p^{\text{DDH}} = 1.24_{-0.18}^{+0.12} M_\odot$ ,  $M_c^{\text{DDH}} = 1.36_{-0.12}^{+0.18} M_\odot$ , and  $i^{\text{DDH}} = 59 \pm 9 \text{ deg}$ , while from the  $\chi^2$  mapping of DDGR solutions using equation (2.61) we measured  $M_p^{\text{DDGR}} = 1.26_{-0.25}^{+0.13} M_\odot$ ,  $M_c^{\text{DDGR}} = 1.32_{-0.13}^{+0.25} M_\odot$ , and  $i^{\text{DDGR}} = 57 \pm 12 \text{ deg}$ . Both of these measurements are in consistency with each other, and with a neutron star nature for the companion star, confirming the DNS nature of the system.

### 4.2.3 Science from PSR J1208–5936

Aside from performing mass measurements, we also discussed the astrophysical properties of **PSR J1208–5936**. The first discussed topic is the scattering present in its profile. We fitted the parameters from equation (1.10) to the profile at several frequencies to measure the evolution of the scattering timescale,  $\tau_s$ , over the observed frequency,  $f$ . A fit of equation (1.11) on the  $\tau_s(f)$  values resulted in a significant measurement of  $\tau_{\text{ref}} = 693 \pm 12 \mu\text{s}$  at the reference frequency of  $f = 1 \text{ GHz}$  and a scattering index of  $\alpha = -2.8 \pm 0.2$ , which is lower than the typically expected value of  $\alpha = 4$  (e.g. Oswald et al., 2021). No polarised emission from this pulsar was detected in the PTUSE observations, leading to no RM measurement or any fit of the rotating vector model.

Another interesting aspect of this paper comes from comparing the spin and orbital parameters of **PSR J1208–5936** with that of other DNS systems. Its spin period of  $P_s = 28.71 \text{ ms}$  and orbital period of  $P_b = 0.632 \text{ days}$  is consistent with the  $P_b - P_s$  relation described in Tauris et al. (2017), which approximately states that

$$P_s \approx 44 \text{ ms } (P_b/\text{days})^{0.26}. \quad (4.1)$$

This is an empirical relation, but one that makes theoretical sense: a large  $P_b$  also indicates a likely large orbital period in the progenitor **PSR–naked He star** system (see Fig. 1.13). Large orbital periods mean a large orbital separation as well, and thus a delayed onset of the Case BB RLO and a shorter recycling stage, leading to larger  $P_s$  value for the remaining PRP. As is shown in Fig. 8 of the publication, and within a reasonable spread of values, **PSR J1208–5936** does not deviate from this norm, further supporting its DNS nature.

Something that does deviate a bit from the norm is its orbital eccentricity. A similar correlation has been observed in the  $e - P_b$  space, which once again makes sense: naked He stars with a delayed Case BB RLO onset retain more mass before exhausting their nuclear fuel, leading to larger mass loss during the second supernova event in the system and perhaps to a larger supernova kick (Tauris et al., 2015; Tauris et al., 2017). That means that the widest systems are also expected to be the most eccentric ones, an hypothesis that has remained true so far. However, as is shown in Fig. 9 of the publication, the eccentricity of **PSR J1208–5936** stands above that of its DNS peers at similar orbital periods.

This makes **PSR J1208–5936** an interesting case from the point of view of binary evolution. The excess of orbital eccentricity could be a matter of chance. After all, the orbital configuration after the second supernova event has a degree of randomness due to the distribution of possible magnitudes and directions of the supernovae kicks (Tauris et al., 2017). But it could also be a hint that the companion of **PSR J1208–5936** had somehow retained more mass before the supernova, leading to a larger mass loss and supernova kick magnitude during the event. If that were to be true, we should also expect the neutron star companion to be a massive one. Indeed, there is a similar, positive  $M_c - e$  correlation for PRP with neutron star companions has been predicted exactly with the same arguments (Tauris et al., 2015; Andrews and Mandel, 2019; Sengar et al., 2022), and even tentatively observed (Andrews and Mandel, 2019; Sengar et al., 2022). However, as is shown in Fig. 11 of the publication, the value of  $M_c$  is not precise enough to make a statement in this regard. Nonetheless, the inevitable increase in timing precision in the future will shed light on this discussion. If it were to be confirmed that the companion of **PSR J1208–5936** is indeed massive, it would speak in favour of the massive naked He star progenitor hypothesis and, perhaps, it could even tell us if the supernova event was an electron-capture supernova, or an iron core-collapse supernova.

Finally, we performed a search for pulsed signals from the companion. After all, if it is a neutron star it could be visible as a canonical pulsar, and having both detected would allow for even more precise mass measurements. For the search, we used the APSUSE-mode data from the orbital campaign observations. Assuming a series of inclination angles in the range of  $i = 30\text{--}90$  deg, we de-modulated the orbital motion of the companion with software **pysolator** (Ridolfi, 2020), a program built by our fellow TRAPUM member Alessandro Ridolfi which de-accelerates a de-dispersed time series with the assumption of Keplerian orbital parameters. Given that it is the companion of **PSR J1208–5936**, we assumed  $\omega' = \omega + 180$  deg and varying  $x$  as a function of  $\sin i$  and  $M_c$  in consistency with equation (1.21). On the de-dispersed and isolated time series, we performed both FFT and FFA searches with PRESTO/accelsearch and riptide, respectively. The codes used in these searches are available in the **demodulate-search** GitHub repository. It essentially another modified version of the scripts used to re-detect the MMGPS discovery **PSR J1015–5358**, so these scripts went a long way. Sadly, these searches ended without success, meaning that either the companion are not beaming towards us or that it entered the pulsar graveyard long before discovery, perhaps millions of years ago.

#### 4.2.4 Updating the neutron star merger rate

After all the strictly radio pulsar-related science summarised above, one of the main draws of the publication is the update of the neutron star merger rate provided by the discovery of **PSR J1208–5936** and the performance of the MMGPS-L. It turns out that this system will merge within a time shorter than the age of the Universe, meaning that systems like it are progenitors to the neutron star merger events observed by ground-based gravitational wave observatories such as *LIGO* and *Virgo*. The best example of such event is **GW170817** (LIGO and Virgo Collaborations, 2017), a neutron star merger that occurred in a galaxy at a distance of 140 million light years and that revolutionised multi-messenger astronomy thanks to its simultaneous detection in Gamma-rays, and later detection in optical and radio (LIGO Scientific Collaboration et al., 2017). It is thought that in this kind of events neutron-rich matter is released into the environment as the neutron stars tear each other apart through tidal forces right before colliding and collapsing into a black hole, seeding Galaxies with many of the metals we use in our daily life (Chen et al., 2021a).

Therefore, estimating the abundance of merging neutron star binaries in the Milky Way is key to understanding how much this process contributes to the chemical enrichment of our Galaxy, and thus the discovery of **PSR J1208–5936** motivates us to perform an update of this estimate. However, the methods used for this computation have not been discussed in this thesis yet, making this section hard to read without an expert eye. Therefore, it is time to finally explain the neutron star merger rate to its fullest, starting by a computation of the merger time of **PSR J1208–5936**.

#### The merger time of PSR J1208–5936

In Section 2.4.7, it was already stated that the emission of gravitational-wave radiation leads to a decay of the orbit over time. Equation (2.52) in particular describes the shortening of  $P_b$  as a function of the Keplerian binary parameters. However, to compute the merger time, it is more relevant to look at the evolution of the orbital separation between the two components of the binary,  $a = a_p + a_c = (x_p + x_c)/\sin i$ . As described in Peters (1964), it evolves with

$$\dot{a} = -\frac{64}{5} \frac{G^3}{c^5} \frac{1 + (73/24)e^2 + (37/96)e^4}{a(1 - e^2)^{7/2}} M_p M_c (M_p + M_c). \quad (4.2)$$

Our objective is to compute how long it takes until we reach  $a = 0$ . However, the equation above is not enough by itself. Gravitational-wave emission not only shrinks the orbit, but it also *circularises* it<sup>82</sup>, so we also need to account for the time evolution of the orbital eccentricity

$$\dot{e} = -\frac{304}{15} \frac{G^3}{c^5} \frac{e + (121/304)e^3}{(1 - e^2)^{5/2}} M_p M_c (M_p + M_c). \quad (4.3)$$

Now we have all that we need to compute the merger time of a binary system: we just have to integrate the two equations above jointly over time until  $a = 0$  is reached. I did that numerically via an

<sup>82</sup> More radiation is emitted during the periastron passage, when the stars have maximal acceleration, so in the next orbit the stars do not reach as “high” as before. It is similar to atmospheric drag for satellites in eccentric, low-Earth orbits: they suffer more friction at closest approach to Earth, circularising their orbits as well. For this reason, satellites in eccentric low-Earth orbits need thrusters to correct their trajectories and remain stable.

implementation of the Runge-Kutta method for multiple equations (to learn more about numerical methods in physics, see Gautschi, 2011).




Figure 4.2: Evolution of the orbital separation and eccentricity in the binary system **PSR J1208–5936** due to its gravitational-wave-induced orbital decay, assuming  $i = 60$  deg. The birth point has been achieved assuming the system to have the same age as the pulsar characteristic age of **PSR J1208–5936**, which is probably a huge over-estimate, especially given that it comes out larger than the age of the universe.

With that, and assuming  $i \approx 60 \pm 10$  deg from the timing constraints, the merger time for **PSR J1208–5936** is  $\tau_m = 7.2 \pm 0.2$  Gyr, as is shown in Fig. 4.2 in this chapter. The important point is that it is shorter than the age of the Universe, dated at 13.8 Gyr, implying that the neutron star merger events observed by *LIGO* and *Virgo* could have origin in systems like this one. Therefore, **PSR J1208–5936** is a suitable candidate to update the neutron star merger rate. For reference, Table 1 of the publication lists the merger time for many known DNS systems in 2023, highlighting those merging within the age of the Universe. In that list, **PSR J1208–5936** has the longest merging time amongst them.

### Estimating the neutron star merger rate in the Milky Way

Now comes the part where I explain how to estimate the population merging neutron star systems in the Milky Way from Galactic DNS observations. That is pretty much glossed over in Section 6 of the publication but here I go into slightly more detail for the non-expert eye. We use the same method already established in Kim et al. (2003), and used in recent estimates of the same quantity (e.g. Pol et al., 2020; Grunthal et al., 2021), so you can also check those reference for an in-depth explanation.

This estimate is a game of statistics: we guess the size of the population of neutron stars in the Galaxy from what we think we do *not* see with our telescopes. Let us say that there is number  $N_{\text{tot},i}$  of merging DNS systems with orbital characteristics  $i$  within our Galaxy, distributed across different luminosities, sky positions and distances from Earth. From this population, with all radio surveys so far, we were sensitive only to a fraction of pulsars

$$\lambda_i = \alpha_i N_{\text{tot},i}, \quad (4.4)$$

where  $\alpha_i$  is the dimensionless fraction and  $\lambda_i$  the actual amount of systems we can detect. However,  $\lambda_i$  is not the exact number of discoveries we will achieve, as it is also a matter of chance. Was there a good RFI environment on the observation day? Did a potential discovery have its flux at



the telescope scintillated up or down? Did we catch it in a lucky or unlucky part of its orbit that facilitated or hampered its detection? Was the astronomer paying good attention to the candidate plots in CandyJar? Because of that, we describe probability of performing  $N_{\text{disc},i}$  discoveries given  $\lambda_i$  detectable systems with a Poisson distribution,

$$p(N_{\text{disc},i}|\lambda_i) = \frac{\lambda_i^{N_{\text{disc},i}}}{N_{\text{disc},i}!} \exp(-\lambda_i). \quad (4.5)$$

Now, in reality, from all the merging DNS systems we have only have  $N_{\text{disc},i} = 1$  (for instance, we only have one **PSR J1208–5936**). That means that the probability  $p(N_{\text{disc},i}|\lambda_i)$  can be marginalised to a likelihood

$$p(N_{\text{disc},i} = 1, \lambda_i) = \mathcal{L}(\lambda_i) = \lambda_i \exp(-\lambda_i). \quad (4.6)$$

This likelihood  $\mathcal{L}(\lambda_i)$  in itself behaves like a probability distribution for  $\lambda_i$  given our single discovery. However, for each pulsar, we want to have a likely contribution of pulsars of this kind onto the Galactic neutron star merger rate. For that, we first transform the probability distribution for  $\lambda_i$  into a probability for the total number of pulsars of this population,  $N_{\text{tot},i}$ , with

$$p(N_{\text{tot},i}) = \mathcal{L}(\lambda_i(N_{\text{tot},i})) \left| \frac{d\lambda_i}{dN_{\text{tot},i}} \right| = \alpha_i^2 N_{\text{tot},i} \exp(-\alpha_i N_{\text{tot},i}). \quad (4.7)$$

From here, we note that the average contribution of a population  $i$  to the neutron star merger rate is the number of its components divided by the amount of time each individual takes to merge,

$$\mathcal{R}_i = \frac{N_{\text{tot},i}}{\tau_{\text{life},i}} f_{\text{b},i}, \quad (4.8)$$

where  $\tau_{\text{life},i} = \tau_{\text{m}} + \tau_{\text{age}}$  is the addition of the merger time plus the age of the system, and  $f_{\text{b},i}$  is the *beaming factor*, the inverse of the fraction of pulsars that beam towards us. So, if we now convert from  $p(N_{\text{disc},i})$  to

$$p(\mathcal{R}_i) = p(N_{\text{disc},i}) \left| \frac{d\mathcal{R}_i}{dN_{\text{tot},i}} \right| = \left( \frac{\alpha_i \tau_{\text{life},i}}{f_{\text{b},i}} \right)^2 \mathcal{R}_i \exp\left(-\frac{\alpha_i \tau_{\text{life},i}}{f_{\text{b},i}} \mathcal{R}_i\right), \quad (4.9)$$

we have obtained a probability distribution for the merger rate contribution of the  $i$ -th population given the fraction of systems we can recover with our telescopes  $\alpha_i$  ( $\tau_{\text{life},i}$  and  $f_{\text{b},i}$  are pre-set parameters characteristic for each pulsar). Finally, to determine the global contribution of all systems in the Milky, we have to find the probability distribution  $p(\mathcal{R} = \sum_i \mathcal{R}_i)$ , which is none other than a convolution of the individual distributions presented in equation 4.9. Thus far, these were most of the missing details in Section 6 of the publication. The following steps are also explained in the same section of the publication.

From here, all we have to do is find the value of  $\alpha_i$  for each merging DNS in the sky. That is done by, simulating the pulsar many times on the sky with varying distances, luminosities, and sky positions, and counting how many times do we detect it to derive the factor  $\alpha_i$ . So, in essence,  $\alpha_i$  is a sensitivity factor of our radio surveys towards pulsars like that one. Key in that is that to accurately estimate of the not-yet detected systems is that we compute  $\alpha_i$  accounting for *all major radio surveys* up to this

point, not just the MMGPS. So, this simulation experiment is done for seven different surveys with Arecibo, Parkes and the GBT, all of which are listed in Section 6 of the publication or the footnote of Fig. 4.3 in this chapter, plus the MMGPS.




Figure 4.3: Sky coverage of radio pulsar surveys used in the neutron star merger update experiment, in Galactic coordinates of longitude and latitude. These include: the Pulsar Arecibo L-band Feed Array survey (PALFA, Cordes et al., 2006), the Low-latitude High Time-Resolution Universe pulsar survey (HTRU-Low, Keith et al., 2010), the Parkes High-latitude pulsar survey (Burgay et al., 2006), the Parkes Multibeam Pulsar Survey (Manchester et al., 2001, PMPS, ), the Arecibo 1990 survey (Wolszczan, 1991), and the Green Bank North Celestial Cap survey (GBNCC, Stovall et al., 2014), in addition to the MMGPS-L.

The details of the simulation are specified in Section 6 of the publication, and the Appendices of the publication<sup>83</sup>. However, here is a summary: we used the pulsar population simulation software **PsrPopPy2** (Bates et al., 2014) to spread the pulsar across the sky many times over, and then simulate the surveys and achieve a  $S/N$  for the pulsar within the closest pointing from this survey. If we achieved  $S/N$  larger than a survey-dependent threshold, we considered it detected. Aside from the sensitivity of the telescope used in each survey and the integration times, the sky area covered by the survey is also important, as if the pulsar falls outside the observations it yields no detection. A particular challenge of the inclusion of MMGPS-L was accounting for the interferometric nature of the survey. Instead of one single beam like in the previous surveys, we had to implement the coherent beam tiling of pointing. So, for the MMGPS the  $S/N$  computation had two layers: first a factor from the survey beam following the same principle as the previous surveys, and then another factor from the nearest coherent beam in the tiling, as is detailed in the Appendices of the paper. The codes for the merger rate update are found in the **DNS-merger-rate-2022** GitHub repository, and they are a modification to those from Pol et al. (2020). Building these scripts was possible thanks to the extremely providential advice from my fellow PhD student Kathrin Grunthal, who had performed an earlier tampering of these codes in her own merger rate update (Grunthal et al., 2021).

<sup>83</sup> Appendices within appendices: *appendixception*.

The result of the experiment is that the increased sensitivity of the MMGPS-L increased the  $\alpha_i$  values of all known DNSs. For **PSR J1208–5936** we obtained  $\alpha = 0.01157$  and  $\mathcal{R} = 0.04^{+0.19}_{-0.03}$  mergers/Myr, the smallest contributions amongst all merging DNS systems (see Table 4 in the publication). The global merger rate in the Milky Way turned out to be  $\mathcal{R}_{\text{MW}} = 25^{+19}_{-9}$  mergers/Myr within 90% credible intervals, a slight reduction with respect to the previous estimates (e.g. Pol et al., 2020; Grunthal et al., 2021), as is shown in Fig. 13 of the publication. The change is not extremely significant with respect to the previous estimates, but as shown in Fig. 13, as this experiment is repeated in recent years, it goes down consistently: we are exploring the sky with our radio surveys deeper and faster than we are discovering new systems, meaning that the Milky Way contains fewer of these merging systems than previously expected.

### Making a prediction for *LIGO*

Now that we have estimated the neutron star merger rate in the Milky Way, we can proceed and gift a nice candy to the gravitational-wave community by predicting of how many of these events will they detect. Since in our Galaxy it takes on average 20,000 to 60,000 years for any of these events to occur, we are not likely to see one coming from the Milky Way within our lifetimes. However, the events seen by *LIGO* are of extragalactic nature, so our aim is to extrapolate the neutron star merger rate to a volume of  $1 \text{ Gpc}^3$  around our Galaxy.

That is done at the very end of Section 6 of the publication. For that, we assume that the abundance of merging DNS systems in a galaxy correlates with the rate of star formation in it. At the same time, we assume that the rate of star formation is traced by the amount of blue light emitted by the galaxy<sup>84</sup>. Kopparapu et al. (2008) computed the density of blue light emanating from the local universe from an optical catalogue of galaxies, and we used their relation to scale the merger rate in the Milky Way to  $\mathcal{R}_{\text{local}} = 293^{+222}_{-103}$  mergers  $\text{Gpc}^{-3} \text{ yr}^{-1}$  in the local universe. In turn, we know that in its **current configuration** *LIGO* is the most sensitive ground-based detector so far, being able to detect neutron star merges up to distance of  $\sim 175 \text{ Mpc}$  (or 570 million light years). Integrating this volume over the rate density given above, we achieve  $\mathcal{R}_{\text{LIGO}} = 6.73^{+5.12}_{-2.37}$  events  $\text{yr}^{-1}$ .

### 4.2.5 Final remarks and future prospects

**PSR J1208–5936** is a very exciting system to experiment with. Timing measurements improve slowly over time as we add more and more observations. After a few years of timing, we will detect the proper motion of this pulsar in the sky, helping us to understand the proper motion contributions to a measurement of  $\dot{x}$  or  $\dot{P}_b$  discussed in Section 2.4.7. Additionally, constraining the Shklovskii and galactic acceleration contributions will also improve our constraint on  $\dot{P}_s$ , which is probably heavily contaminated by these effects and has not yielded a significant measurement from timing yet. The same goes for the PK parameters. Fig. 14 of the paper shows how we expect that the relative uncertainties of the  $\dot{P}_b$ ,  $\gamma_E$ , and  $\dot{\omega}$  measurements will evolve over time. We expect that after three years since submission of the publication (two years from now) we will achieve a  $10\sigma$  detection of the Einstein delay, bringing down the mass uncertainties to  $\pm 0.1 M_\odot$ , and the inclination angle uncertainty down to  $\pm 5 \text{ deg}$ . After a decade, when we may achieve our first  $\dot{P}_b$  measurement, which when measured with precision will allow for a self-consistency test of GR.

<sup>84</sup> It is a well-known fact of astronomy that young stars are bluer than old stars, so a larger fraction of blue light emanating from a galaxy indicates a larger abundance of young stars and therefore of star formation.



## PSR J1227-6208: the most massive known PRP–ONeMg WD system?

### 5.1 Context and contributions

So far I have presented a mixture of search and timing results. Now, we leave searching behind and focus on an entirely timing project. Let me introduce **PSR J1227–6208**: a recycled  $P_s = 34.53$  ms pulsar at  $DM \approx 363 \text{ cm}^{-3} \text{ pc}$  in a circular  $P_b = 6.72$  day orbit with projected semi-major axis  $x = 23.2$  ls. What makes it special, you say? Glad you asked. Assuming  $M_p = 1.35 M_\odot$ , if we apply equation (1.21), we obtain a minimum companion mass of  $M_c = 1.27 M_\odot$ . Indeed, this may be the most massive **PRP–ONeMg WD** system we know so far. However, unlike the MMGPS pulsars, it is not a recent discovery. Instead, it was discovered by Mickaliger et al. (2012) in archived data from the Parkes Multibeam Pulsar Survey<sup>85</sup> (PMPS, Manchester et al., 2001). However, timing precision with Parkes/Murriyang was not precise enough to yield a good Shapiro delay measurement (Bates et al., 2015). Additionally the low eccentricity of the orbit,  $e = 1.15 \times 10^{-3}$ , implies a low GR-predicted rate of periastron of advance,  $\dot{\omega}$ , so it was not possible to measure it in early Parkes/Murriyang timing (Bates et al., 2015). Finally, the pulsar sky position was too far south for observation with the more powerful telescopes in the Northern Hemisphere. Therefore, the component masses in this system were left poorly constrained.

But then MeerKAT came along. With four times the sensitivity of Parkes/Murriyang, it promised a much improved measurement of the Shapiro delay parameters. Also, the accumulated time since discovery had allowed  $\omega$  to precess enough for a first  $\dot{\omega}$  measurement. That provided a golden opportunity for the first proper mass measurements in **PSR J1227–6208**, and adding an extra sample to the mass measurements in otherwise very scarce massive **PRP–ONeMg WD** systems. Therefore, in 2019, **PSR J1227–6208** was added into the **RelBin project** (Kramer et al., 2021b), a timing survey of several Galactic binaries belonging to the larger **MeerTIME project** (Bailes et al., 2020), a science program dedicated to the timing of Southern Sky pulsars with MeerKAT. Subsequently, this pulsar and project were handed to a very nice PhD student at the MPIfR (me) for mass measurements.

This chapter, following the same philosophy as the previous one, is a summary of Colom i Bernadich et al. (2024), A&A, 690, A253, the publication that came out of the timing project on **PSR J1227–6208**,

<sup>85</sup> This is what you can do if you do not delete your data every week: discovering new pulsars a decade after the observations.

and which is included in original form in Appendix C. In addition of the summary, I also provide some extra details for the methods used in this project. In particular, I describe the treatment of red timing noise, and how constraints on the longitude of the ascending node,  $\Omega_a$ , are performed. These methods are more standard in pulsar timing, and they are a direct extensions of the techniques already described in Section 2, so introducing them will not require too many leaps.

Regarding my contributions to the publication, as the first author of the paper, I wrote most of the text and performed all of the data analysis presented in it. That includes the measurement of ToAs, the timing analysis, the mass constraints, and the orbital geometry constraints. In addition, I also operated the Parkes/Murriyang telescope to record data for this pulsar and other RelBin pulsars under the observing program P1032, led by Vivek Venkatraman Krishnan. Vivek Venkatraman Krishnan is listed as a second author due to his very significant contributions regarding data preparation, with direct participation in the RFI excision and calibration of Parkes UWL observations, and for his technical and scientific guidance during this project. Accordingly, he also contributed to the writing of some of the text in Section 2.1 of the publication. The remaining co-authors are: David J. Champion, Paulo C. C. Freire, Michael Kramer, Thomas M. Tauris, Mathew Bailes, Matthew Bailes, Alessandr Ridolfi, Marcus E. Lower and Maciej Serylak. Each of them have made their own significant contributions: being builders of the MeerTIME/RelBin projects that enabled data collection in the first place, collecting and calibrating Parkes/Murriyang or MeerKAT data, and providing technical and scientific guidance during the project. Finally, the input from Thomas M. Tauris was more theoretical, with his understanding of binary evolution being essential to interpret the astrophysical implications of our timing measurements. His arguments convinced us against a neutron star nature of the companion despite its large mass due to the low orbital eccentricity of the system<sup>86</sup>. The main points of that discussion are now reflected in Section 5 of the publication. Now, with all of the proper introductions in place, let us delve into the amazing system in **PSR J1227–6208**.

## 5.2 Summary of the publication

### 5.2.1 Data used in the experiment

As described in Section 2 of the publication, this experiment was possible thanks to a long-term observing campaign performed across the years. Parkes data from the former 21 cm multibeam receiver goes as far back as 2011, while Parkes data from the newer UWL receiver was taken between 2019 and 2023. In addition, MeerKAT L-band data recording began in 2018 and continued until 2022. In May 2023, we also had a set of MeerKAT S-band observations in the S1 configuration, which observes in the  $f = 1968 - 2843$  MHz range<sup>87</sup> and makes this publication the first one with measurements from the MeerKAT MPIfR S-band receivers. The specific time and frequency coverage of each receiver set is also specified in Table 1 of the publication, but some specifications were also introduced in Section 2.2.3.

This collection of data is diverse both in nature and purpose. The long time-span of the Parkes multibeam data provides the much-needed timing baseline for measuring long-term PK effects such as  $\dot{\omega}$ . The MeerKAT L-band data, aside from continuing to increase the timing baseline, also increases

---

<sup>86</sup> To quote Thomas from one of our intense meeting between me, him, Paulo and Vivek, he said that “if the companion is a neutron star, I will eat my own COVID-19 mask”.

<sup>87</sup> See the [MeerKAT General documentation](#).



significantly the timing sensitivity and is the key to measuring the Shapiro delay parameters due to its lower ToA uncertainty. Indeed, with MeerKAT we achieved a typical ToA uncertainty of  $\sim 5 \mu\text{s}$ , while ToAs from Parkes observations had a typical ToA uncertainty of  $\sim 20 \mu\text{s}$ . However, that sensitivity is hampered by random time-variations of the DM, which are hard to model and bias our experiment. The Parkes UWL receiver allows us to track these variations empirically due to its large frequency range on which the DM value can be measured very precisely at every observation. Finally, the MeerKAT S-band data was taken to observe the pulsar at higher frequencies than in the L-band, where the amplitude of the ToA noise due to DM variations is reduced.

### 5.2.2 Profile measurements

Before going into timing, in Section 3 we present the measurement of scattering and spectral properties of the pulsed emission. We fit the pulse profile to equation 1.10 at different frequencies to measure the scattering time-scale  $\tau_s$  as a function of the observing frequency  $f$  in the time-integrated, frequency-resolved Parkes UWL and MeerKAT L-band profiles. We fit these measurements to equation (1.11) and derive  $\tau_{\text{m,ref}} \approx 1.2 \text{ ms}$  at the reference frequency of  $f = 1 \text{ GHz}$  and a scattering index of  $\alpha = 3.2\text{--}3.7$ . Similarly, we measure the mean flux density,  $F_m$ , at different frequencies  $f$  and fit them to equation (1.7), deriving  $F_{\text{m,ref}} = 530\text{--}620 \mu\text{Jy}$  at the reference frequency of  $f = 1 \text{ GHz}$  and a spectral index of  $\chi = -2.20 \pm 0.15$ . The scattering magnitude of **PSR J1227–6208** is typical to that of pulsars at similar DM values and its index is within the observed distribution of  $\alpha = 4.0 \pm 0.6$  values (Oswald et al., 2021), while its spectral index is steeper than most pulsars.

We detect a significant fraction of polarised emission from **PSR J1227–6208**, both with MeerKAT and Parkes. Approximately 15% of its emission is linearly and circularly polarised, as is shown in the integrated pulse profiles shown in Fig. 3 of the publication. It turns out that the linear and circularly polarised fractions do not coincide in spin phase: circularly polarised emission occurs  $\sim 0.025$  turns earlier in spin phase before linearly polarised emission. Nonetheless, we do not measure the angle of the linearly polarised emission due to significant scattering at the lower radio bands, and the low  $S/N$  at the higher Parkes radio bands. Because of that, we can not fit the rotating vecotor model from equation (1.6). Nonetheless, we measure  $\text{RM} \approx 47 \text{ rad cm}^{-1}$ , which indicates an average magnitude of the magnetic field parallel to the LOS of  $\sim 0.16 \mu\text{G}$ .

### 5.2.3 Timing models and timing noises

When it came to timing, this experiment became a bit trickier. In this thesis, I have yet to introduce to concept of timing noise. Dealing with timing noise is in theory universal in the field of pulsar timing, but in practice many timing experiments get away without worrying about it (for example, in the timing of **PSR J1208–5936**, Chapter 4). In the case of **PSR J1227–6208**, it became clear that the noise contributions were significant enough to bias our results.

There are two types of timing noise: white noise and red noise. White noise is as it sounds, Gaussian random noise in the timing residuals beyond what is expected from the ToA uncertainty. Its origin is both a genuine underestimation of the timing uncertainty and real random variations in the spin of the pulsar. It can shift the value of the  $\chi^2$  from equation (2.35) upwards, biasing our measurement uncertainties, but it is usually dealt with by re-scaling the uncertainties to a larger value until  $\chi_{\text{red}}^2 \approx 1$  is achieved in equation (2.50) is achieved (e.g. Lentati et al., 2014). However, while white noise has to be dealt with, the kind of noise that worries us the most is red noise.

### Red timing noise

Red timing noise is the real source of problems in the timing of **PSR J1227–6208**. The notion of red noise was already presented in Section 2.3.4 for uniformly sampled time series: long-term correlated baseline variations that manifest as increasing power in the lower Fourier frequencies. This definition is also applicable to the behaviour of ToA residuals: long-term correlated structures in the timing residuals that are not accounted for by our timing models (for some good examples, see Parthasarathy et al., 2019). It originates from real, physical long-term variations of the spin period of the pulsar, or in DM variations due to a changing ISM column along the LOS between the pulsar and the observer. If not accounted for, these variations are absorbed into the high-order derivatives in the spin frequency and the time evolution of DM via equations (2.27) and (2.29) or, even worse, into the orbital and PK parameters, leading to poor uncertainty estimates and even biased measurements.

Accounting for red noise in our timing models is therefore essential. We can model spin noise and DM noise independently from each other in the Fourier domain. We describe their Fourier power with

$$S(\nu) = A^2 \left( \frac{\nu}{\text{days}^{-1}} \right)^{-\gamma}, \quad (5.1)$$

where  $\nu$  here is the *noise frequency* in the Fourier domain<sup>88</sup> of the time variations measured in  $\text{days}^{-1}$ ,  $A$  is a dimensionless amplitude, and  $\gamma$  is a spectral index (Lentati et al., 2014). For spin noise originating from irregular pulsar rotation,  $A$  is constant across radio frequency, or *achromatic*. For DM noise originating in ISM variations,  $A$  is radio frequency-dependant or *chromatic*. Following equation (1.9). Therefore, we describe the ToA amplitude of DM noise scales with

$$A_{\text{DM},f} = A_{\text{DM}} \left( \frac{f}{1400 \text{ MHz}} \right)^{-2}, \quad (5.2)$$

where  $A_{\text{DM}}$  is measured at the reference radio frequency of  $f = 1400 \text{ MHz}$ .

These amplitudes and noise spectral indices describe the *correlation* between pairs of ToAs  $\Delta t_{i,j} = t_j - t_i$ . If  $\gamma$  is positive, as it usually is, the ToAs become more correlated the further apart they are. However, we do not model the noise to infinitely large noise frequencies. Above a chosen noise frequency,  $\nu_{\text{cut}}$ , the description is no longer applied and the noise is assumed to be well-described by uncorrelated white Gaussian noise (Lentati et al., 2014). Finally, to fully model the noise, a time-domain *noise realisation* has to be performed, where possible starting phases of each Fourier frequency are attempted. The one that minimizes the ToA residuals given a set of  $A$ ,  $\gamma$  and  $\nu_{\text{cut}}$  parameters is the one that is chosen.

Spin noise is a problem common of canonical pulsars and magnetars (see Parthasarathy et al., 2019) for a showcase of red spin noise in many young pulsars). In the case of recycled pulsars like **PSR J1227–6208**, their high timing stability typically makes them exempt from such issues. However, given its location deep in the Galactic plane in its high DM  $\approx 363 \text{ pc cm}^{-3}$ , **PSR J1227–6208** is victim of severe random DM variations that affect its timing. Therefore, modelling these DM variations became a key component of the experiment.

<sup>88</sup> We have three different frequencies defined in the thesis so far: radio frequency, spin frequency and noise frequency. At the same time, these frequencies are probed in the Fourier domain, where we use the term Fourier frequencies. Be careful to not get confused! I certainly get confused often.

### Measuring red noise with **TempoNEST**

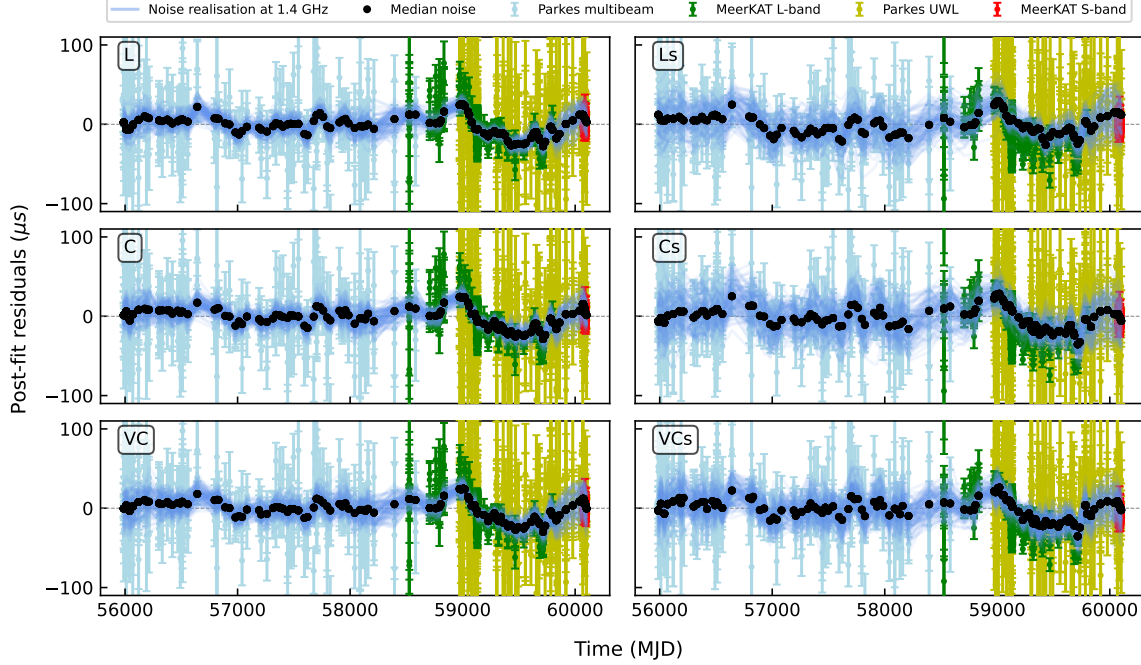


Figure 5.1: Timing residuals and time-domain realisations or red noise for the solutions listed in Table 2 of the publication. Since it is hard to depict the time-domain realisation of DM noise, I only plot the noise amplitudes at 1.4 GHz. Given that DM noise amplitude is  $\sim 100$  larger than that of spin noise, virtually all the contributions in the plots are from DM noise only.

If the timing noises can bias the measurement of timing parameters, how do we know what part of the ToAs residuals are caused by noise, and what part are caused by the PK effects we are interested in? The answer is that we do not know for certain, but that we can estimate the probability of one or the other and integrate this notion into the uncertainties of PK parameters.

We do this part with **TempoNEST**, a wrapper around TEMPO2 that fits for the timing model and the noise model (including white and red noises) at the same time with Bayesian methods (Lentati et al., 2014). As described in Section 4.1 of the publication, we implement up to six different noise models: three with white noise and red DM noise **VC**, **C**, and **L**, and three more with white noise, red DM noise and red spin noise **VCs**, **Cs**, and **LS**. In each of these models, we choose  $\nu_{\text{cut,DM}} = 30, 50$ , and  $100 \text{ days}^{-1}$  and  $\nu_{\text{cut,spin}} = 30, 100$ , and  $500 \text{ days}^{-1}$  (if spin noise applies). The quality of the resulting timing+noise models are displayed in Table 2 of the publication. We find that both models with and without spin noise adapt well to the data, and that in models with both spin and red noise, the best-fit amplitude of DM noise is larger than that of spin noise by almost two orders of magnitude. The implications is that DM noise is clearly dominant. We also find that models with larger  $\nu_{\text{cut}}$  values are more slightly favoured. The derived noise amplitudes, spectral index and quality fits are all listed in Table 2 as well. For visual aid, Fig. 5.1 in this chapter also depicts the timing noise in the ToA residuals in each derived timing model, along well as the time-domain noise realisations. In essence, all descriptions are a good match for the data, with all of them including significant DM variations across windows of  $\sim 100$  days.

### 5.2.4 PK parameters and mass constraints

We spend much of Section 4.2 discussing the effects of red noise on the measurement of PK parameters. Our conclusion is that individual noise models introduce specific variations in our measured PK parameters, but without a significant impact on the physics. We note that the orthometric Shapiro delay amplitude is sensitive to the  $\nu_{\text{cut}}$  values, with models going into larger frequencies yielding smaller  $h_3$  measurements. This is indicative that the Shapiro delay may be degenerate with the red noise model of choice. However, all PK measurements are consistent within  $1\sigma$  or  $2\sigma$  uncertainties across all noise models. To give uncertainty ranges that we are confident with, we report a measurement which encompasses all individual uncertainty estimates presented in Table 3 of the publication. As is shown in Table 3 of the publication, we report the values of  $\dot{\omega} = 0.0171(11) \text{ deg yr}^{-1}$ ,  $h_3 = 3.6 \pm 0.5 \mu\text{s}$  and  $\varsigma = 0.85 \pm 0.05$ , consider the wider uncertainty ranges that account for all noise models.

These are the first high-significance measurements of PK in this system, and they provide the following mass and inclination angle constraints:  $M_p = 1.54 \pm 0.11 M_\odot$ ,  $M_c = 1.40 \pm 0.01 M_\odot$  and  $i = 78.7 \pm 1.2 \text{ deg}$ , where the uncertainties are the 68.3% percentiles, as is reported in Section 4.3 of the publication. For reference, the mass and inclination constraints from both the individual PK parameters measured with the DDH model via equations (2.59) and (2.60) and the  $\chi^2$  mapping of DDGR solutions via equation (2.61) are plotted in Fig. 5 of the publication, while the timing residuals with noise contributions removed and the shape of the measured Shapiro delay are plotted in Fig. 4 of the publication.

Section 4.4 of the publication also includes an extensive analysis on the effect of including or removing the ToAs derived from the older Parkes multibeam data sets, and on how they contribute to the modelling of timing noise. Our conclusion is that they do not bring significant constraints to the DM noise due to their narrow frequency coverage, and that they may in fact be causing spurious spin noise measurements. Nonetheless, their inclusion is important for constraining the long-term PK effects.

### 5.2.5 Orbital geometry constraints

In addition to mass constraints, in Section 4.5 we perform the first preliminary constraints on the orbital geometry of this system besides the inclination angle  $i$ : we attempt to constrain the longitude of the ascending node,  $\Omega_a$ , from the limits on  $\dot{x}$  and the significant measurement of the proper motion vector,  $\vec{\mu} = (-6.1 \pm 0.3, 9.41 \pm 0.36) \text{ milliarcseconds yr}^{-1}$ .

Indeed, the timing models yield constraints on  $\dot{x}$  and  $\dot{P}_b$  (presented in Table 3 of the publication) with uncertainties similar to those predicted by the proper motion according to equation (2.57), and the Shklovskii and Galactic acceleration effects from equation (2.56). The GR contributions are expected to be orders of magnitude below these contributions, so if any measurement is to come soon it will be from proper motion. In addition,  $\dot{x}$  and  $\dot{P}_b$  show a consistent sign across all timing+noise models and their magnitudes are similar to their uncertainties, suggesting that we could be on the verge of a detection.

To test whether the  $\dot{x}$  constraints are actually affected by the proper motion and the orbital orientation of the system, we decided to perform another  $\chi^2$  mapping experiment, but instead of probing the masses, we probed the  $i - \Omega_a$  space while implementing consistency between  $\vec{\mu}$ ,  $i$ ,  $\Omega_a$ , and  $\dot{x}$  according to equation (2.57). In line to previous experiments of this kind (e.g. Geyer et al., 2023), that was done in a uniform grid of  $\cos i$  and  $\Omega_a$  values, where on each point we recorded the  $\chi^2$  of the solution, and

subsequently the probability distribution

$$\rho(\cos i, \Omega_a) \propto \exp \left[ -\frac{\chi^2(\cos i, \Omega_a) - \chi_{\min}^2}{2} \right], \quad (5.3)$$

is computed. The resulting map, shown in Fig. 6 of the publication, indicates that some regions of the  $i - \Omega_a$  space are more preferred than others, with the most preferred yielding  $\dot{x}$  predictions of the same sign as the constraints in Table 3 of the publication. Our interpretation is that, indeed, a detection of  $\dot{x}$  is almost peeking through even if we can not constrain  $\Omega_a$  just yet. Likely, this is also the case for  $\dot{P}_b$ .

### 5.2.6 Astrophysical implications

We discuss the astrophysical implications of our measurements in Section 5. They confirm the massive nature of the system, but also provide a large range of possible  $M_c$  values: in the extremes, it can be as low as  $M_c \approx 1.2 M_\odot$  and as high as  $M_c \approx 1.5 M_\odot$ . The truth is probably somewhere in the middle and right below the Chandrasekhar limit of  $M_{\text{WD}} = 1.38 M_\odot$ , but we can not know that for certain until we improve our uncertainties. When that happens, we may be able to say for certain whether this white dwarf is amongst the most massive known ones or not, accounting for pulsar binaries and other systems observed with optical telescopes. If it turns out to be extremely massive, we may even have a chance to probe the physics of the Chandrasekhar limit.

In addition, this experiments adds to the very few mass measurements in known **PRP–ONeMg WD** systems. A big take away from the publication that it has fallen above the  $M_c > 1.1 M_\odot$  line, and that no system seems to be in the region of  $M_c = 1.0 - 1.2 M_\odot$  yet, as is shown in Fig. 7 of the publication. As more measurements of this kind are published, the more or less certain we will be of this empirical gap that is not predicted by binary evolution theory. Regarding that, Fig. 7 also shows up to eight potential **PRP–CO/ONeMg WD** systems with no proper mass measurements whose minimum/median  $M_c$  ranges cross the  $M_c = 1.0 - 1.2 M_\odot$  region. This goes to show that many measurements of this kind are yet to be performed.

Finally, the massive nature of the pulsar confirms the existence of massive neutron stars without the need for accreting to much matter from a companion, a result that goes against the postulated notion of a bimodal mass distribution for pulsars presented in Section 1.4.1. Therefore, massive neutrons stars do not exist only in systems with recycled pulsars and light He WD companions where large amounts of mass transfers have occurred, but they can also be born massive. Results in the same direction has been found for other recycled pulsars with massive companions (e.g. Tauris et al., 2011; Cognard et al., 2017).

### 5.2.7 Future prospects

Measurements with MeerKAT have confirmed the interesting nature of this system, and it will continue to be exciting to measure it in the future. We estimate in Section 6 of the paper that continued timing of this system with MeerKAT will provide improved mass constraints thanks to the quick improvement of the  $\dot{\omega}$  measurement over the MeerKAT timing baseline. As stated in the publication, a propagation of the  $\dot{\omega}$  uncertainty to the future predicts a 65% reduction of mass uncertainties by 2026, helping us answer the questions posed above. However, improvements on the measurement of  $\dot{x}$  or  $\dot{P}_b$  will take many more years of continued timing.

Finally, despite the companion being a white dwarf, we do not expect an optical detection due to the likely large age of the system and the large distance from Earth. In this line, no counterpart has been found in optical surveys, and a potential future detection will only be possible with the most powerful optical/infrared telescopes available such as the *James Webb* and the *Hubble* in the best-case scenario that the companion is a young white dwarf. However, given the evolutionary history of the system, that is unlikely to be true.



---

## Conclusions and future prospects

---

After the many science methods and results given in the previous chapters, now comes the finalising chapter with an overview of all the projects I have worked on, and contextualize them in the broader field of pulsar astronomy. After that, I will give some future prospects on the future of radio surveys with MeerKAT, and on the pulsars whose discovery, follow-up and timing I have contributed to. Hopefully, you still have energy to go through this final pages. If you feel tired and exhausted of reading, imagine how it felt to write it! Jokes aside, let us jump into it.

### 6.1 Conclusions

In the projects of this thesis, we have expanded our understanding of Galactic pulsars with the MeerKAT radio-telescope, following the success of the Southern Sky radio surveys performed with Parkes/Murriyang. For that, I was lucky to join the collaborative effort of the MMGPS surveys, which had been in planning for years before the inauguration of MeerKAT. Under the TRAPUM umbrella of projects (Stappers and Kramer, 2016), MeerKAT has already discovered 250 pulsars (e.g. Ridolfi et al., 2021; Ridolfi et al., 2022; Caleb et al., 2022; Turner et al., 2024; Carli et al., 2024; Prayag et al., 2024), 82 of which belong to the MMGPS surveys. That is, a 7% increase of the known pulsar population, 2% from the MMGPS alone. And that only during the first five years of the telescope, without any possibility for re-processing in the MMGPS, which would likely double or triple the amount of pulsar discoveries as it did with the PMPS or HTRU Parkes/Murriyang surveys. It goes to show that MeerKAT really is the new spear of pulsar astronomy in the Southern Hemisphere, with the discovery of both faint and distant pulsars which were undetectable with previous surveys.

Within the MMGPS-L, the new discoveries amount to 78 (Padmanabh et al., 2023), 17 of the MMGPS-L discoveries are recycled pulsars in binary systems. These discoveries contribute to a variety of science cases regarding pulsar emission, binary evolution, and potentially tests of gravity. With the the discovery of nulling pulsars such as **PSR J1353–6341**, we increase the sample of pulsars where we can study the mechanism directly. The several binaries cover a wide range of binary evolution channels and masses. Some of them are quite valuable additions to rare populations with few known members before the start of the MMGPS surveys. The two fully recycled **MSP–CO WD** discoveries belong to a class with only three confirmed systems before the start of the survey, and of which only one had mass measurements from timing. The two massive **PRP–ONeMg WD** discoveries also belong to a

class with only four confirmed systems before the start of the MMGPS survey, with only one of them having published timing measurements in 2020. Of these, **PSR J1015–5358** in particular provides an excellent opportunity to study the upper end of the white dwarf mass distribution, it being potentially the most compact **PRP–ONeMg WD**, with an orbital period of eight hours. This systems is also very promising for testing theories of gravity with timing due to the the large expected magnitude of the PK parameters. Finally, several **MSP–He WD**, **PRP–CO WD** and DNS systems will also contribute to the understanding of the neutron star and white dwarf mass distributions and their binary evolution. Regarding DNSs, the MMGPS has already provided a relevant Galactic neutron star merger rate based on the sky area and depth covered by the MMGPS-L and the discovery of the merging DNS **PSR J1208–5936**.

In this survey, we have taken advantage of the interferometric nature of MeerKAT, a first in the field of pulsar astronomy. The multibeam nature of observations has introduced computational channels during the implementation of a pulsar search, but it also has enabled multiple beam detections of new discoveries and instant sky localisation, which significantly speeds up and optimises the follow-up and timing of new discoveries. We have implemented these multibeam tools for the first time in a blind pulsar survey, paving the road for future multibeam surveys. This will become even more relevant with future expansions of MeerKAT as it evolves into the SKA, which will increase the number of beams per observation. Thus, our multibeam implementations are only a first step towards an even more multibeam future.

MeerKAT has also opened the way for surveys at high frequencies in the Southern Sky with the installation of its S-band receivers. The targeted pulsar search at **SNR G340.6+0.3** was the first pulsar search at  $f > 2.0$  GHz with the MeerKAT telescope, testing the capability and functionality of the newly installed high-frequency receivers. Following that, the MMGPS-S and MMGPS-Sag A\* are the first blind or large-scale pulsar surveys at  $f > 2.0$  GHz from the Southern Hemisphere and with high sensitivity, seeking to untap a population of pulsars that is too dispersed or scattered to be detected at lower frequencies, and to faint feint to be detected at high frequencies with other telescopes.

Finally, the high sensitivity of MeerKAT allows the performance of experiments that were not possible with previous telescopes. In the case of MMGPS discovery **PSR J1208–5936**, its timing wielded mass measurements with less than one year of timing thanks to its high eccentricity that yields a high rate of precession of the periastron, but also to the the timing sensitivity of MeerKAT, which enables very useful constraints on the Shapiro delay and Einstein delay parameters. Similarly, it is only the MeerKAT-derived ToAs that result in the first meaningful Shapiro delay measurement in **PSR J1227–6208**, a feat that is not possible with any current telescope despite it having been discovered in Parkes/Murriyang. This has an important scientific impact, as for example **PSR J1227–6208** is one of the very few known massive **PRP–ONeMg WD** systems, which we would have not been able to measure otherwise. Similar results are expected to be seen from many of the MMGPS discoveries, which can only be timed and measured with MeerKAT. This makes it clear that MeerKAT is not only a discovery tool, but also the future of pulsar timing in the Southern Sky.

## 6.2 Future prospects

### 6.2.1 For the MMGPS surveys

In this thesis, we have only read about discoveries from the already completed MMGPS-L survey. But I have also shared some discoveries from the MMGPS-S, and we even already have one first discovery

from the MMGPS-UHF<sup>89</sup>. With only 22% of the MMGPS-S survey being complete, and with a discovery at the MMGPS-UHF from the first observing session, many more interesting new pulsars are bound to come during the following two years. In this regard, our MMGPS collaborator Vishnu Balakrishnan ran **PsrPopPy2** simulations and predicted the discovery of up to 57 canonical pulsars and 6 recycled pulsars by the end of the MMGPS-S, and at least 100 pulsars in the the MMGPS-UHF. This means that we are bound to have plenty of new interesting pulsars, both binary and isolated, waiting to be discovered, followed-up, and timed!

Regarding the already-discovered binaries, quite a few mass measurements are still pending from the MMGPS discoveries. So far, **PSR J1208–5936** and **PSR J1155–6529** are the only discoveries with proper mass constraints from pulsar timing. As these constraints improve over time, they will inevitably lead to four extra precise samplings of the neutron star mass distributions. While **PSR J1208–5936** has already been published in Colom i Bernadich et al. (2023) and discussed in extensive detail in this thesis, **PSR J1155–6529** is currently under analysis by my fellow MMGPS member Marina Berezina and will be soon reported in a publication. However, the vast majority of pulsars in Table 3.2 are still lacking this treatment.

Currently, work is ongoing to complete their analysis. In particular, all the massive **PRP–CO/ONeMg WD** binaries are quite promising due to their compact nature and the high mass of the companions. Some of these could have significantly measurable Shapiro delays and long-term PK effects such as the orbital decay and the rate of periastron advance (if there is some residual eccentricity), which would lead the measurement of both the pulsar and white dwarf masses. On my side, **PSR J1708–5843** already has a phase-connected timing solution and is waiting to be put on a continued timing campaign and perhaps even an orbital campaign.

Similarly, timing efforts are also ongoing for all the other massive binary discoveries, including the rare **MSP–CO WD** systems and the lighter **MSP–He WD** systems. Eventual mass measurements in these systems will continue to probe the the white dwarf mass distributions and the physics of binary evolution and mass accretion at the low end of the white dwarf mass distribution. My efforts in particular are focused on finding timing solutions for **PSR J1108–6329** and **PSR J1543–5439**, pulsars for which a plethora of MeerKAT data since their discovery is already available. For that, I am currently experimenting with the  $\chi^2$  penalisation’s described in Section 3.3.6 to account for the pulsar position during the implementation of the *dracula* algorithm. This implementation will be beneficial not only for the MMGPS discoveries, but also many future MeerKAT discoveries with sky localisation.

### 6.2.2 For PSR J1208–5936

**PSR J1208–5936** is under continued observations with MeerKAT, so future timing observations will lead to further constraints to the component masses in the following years. Given its eccentric orbit, the continued precession of the periastron will break the degeneracy between the Rømer delay and the Einstein delay, continuously improving the measurement of the amplitude of the Einstein delay. In turn, this will lead to a very significant improvement in the mass constraints.

To predict how quickly will these improvements arrive, I have simulated MeerKAT ToAs into the future years assuming the performance of two extra orbital campaigns of  $0.5 \text{ orbits yr}^{-1}$  in 2026 and

<sup>89</sup> I was not involved in that discovery because it happened while I was moving from Germany to Sardinia, so I missed that candidate inspection round.

2027. Then, I have measured the PK parameters on the simulated ToAs and checked the evolution of their uncertainties. The result of this simulation is presented on the left side of Fig. 6.1. While we barely have a  $3\sigma$  detection of  $\gamma_E$  from data up to May 2023 in Colom i Bernadich et al. (2023), we expect a  $10\sigma$  detection by the end of 2024, and a  $30\sigma$  detection by the end of 2025. Accounting for the two orbital campaigns, we may be getting close to a  $50\sigma$  detection by the end of 2027. This would result in mass uncertainties of  $\pm 0.02 M_\odot$ , and it would allow for example a very good placement of the companion of **PSR J1208–5936** in the  $e - M_c$  diagram, finally telling us if the large eccentricity of the systems correlates with a massive companion. Additionally, it will definitely improve samples to the neutron star mass distribution.




Figure 6.1: Future time-evolution of PK parameter measurement significance (measurement/uncertainty) from simulated ToAs for **PSR J1208–5936** (left) and **PSR J1227–6208** (right), assuming the same ToA quality as in the real timing experiments. The vertical axis shows the fractional uncertainty of a specific PK parameter measurement. The vertical dashed lines correspond to orbital campaigns. The horizontal axis is presented in units of MJD (below) or UTC years (above).

### 6.2.3 For PSR J1227–6208

**PSR J1227–6208** will continue to be observed with MeerKAT and Parkes as a continuation of the long-term timing campaign. The mass constraints we derived in Colom i Bernadich et al. (2024) show that it is indeed a very massive system, but also that there is still a lot of room for a precision improvement. As in the case of **PSR J1208–5936**, I have simulated ToAs into the future and predicted the time evolution of the most constraining PK measurements: the rate of advance of periastron and the Shapiro delay parameters, also accounting for two further orbital campaigns in 2025 and 2026. The results of this simulation are shown in the right side of Fig. 6.1, and they are quite clear: the parameter with the fastest improvement will be  $\dot{\omega}$ . While in Colom i Bernadich et al. (2024) we achieve a  $\sim 20\sigma$  measurement of  $\dot{\omega}$  in (depending on the noise model) with data from up to May 2023, after two possible orbital campaigns we expect an improvement to  $\sim 40\sigma$  by the end of 2026. At the same time, the precision in the measurement of  $h_3$  will increase slowly, reaching  $\sim 25\sigma$  by the end of 2028. This will translate into mass uncertainties of  $\pm 0.01 M_\odot$ , which will precisely tell us how close is the white dwarf companion of to the Chandrasekhar limit.

To make the most of this prediction, it will be essential to continue observing this pulsar with

the MeerKAT S-band receivers. While the L-band receivers can capture more flux due to the steep spectral index of the pulsed emission, the  $\sim 2.4$  GHz frequency is less affected by the noise introduced by the DM variations, allowing for timing measurements that are less sensitive to the chosen noise models. At the same time, continued observations with the Parkes UWL receiver will continue to constrain the DM evolution itself thanks to its wide frequency range.

#### 6.2.4 For pulsars in the era of MeerKAT+ and the SKA

There is no talking about the outlook of pulsar astronomy in the Southern Sky without considering what the future of MeerKAT will bring to the table. While MeerKAT has already discovered hundreds of new pulsars thanks to its increased sensitivity, the current configuration of the MeerKAT array is but a temporary stage in a series of consecutive improvements. In fact, MeerKAT is the embryonic stage of the future Square Kilometer Array at middle frequencies (**SKA-mid**), which aims to observe the sky with **almost 197 antennas (including the current 64) in the frequency range of 0.3 – 15 GHz** (Swart et al., 2022). That will vastly improve both its sensitivity and interferometric capabilities thanks to the extended antenna baseline of 160 km. Even now, the intermediate expansion stage from MeerKAT to **MeerKAT+** is **currently underway**, which will see the installation of 14 new antennas, providing a first sensitivity boost in the following few years.

Of course, it is not all butterflies and rainbows. To begin with, the expanded antenna array will pose a huge data challenge for pulsar surveys. With more antennas expanding the size of the current MeerKAT array, the increase of sensitivity will come with a decrease of the coherent beam size on the sky, meaning that if we already thought that 480 filterbanks per observation were a lot, we have to get ready for searching manifold more. Dealing with that would require another entire PhD thesis. Nonetheless, the collecting area of the telescope will increase almost by four, providing a very significant sensitivity boost to pulsars in the Southern Sky. It is predicted that, in conjunction of the SKA at low frequencies in Australia (**SKA-low**), thousands of previously unseen pulsars will be discovered in the Milky Way (Keane et al., 2015), of which  $\sim 10\%$  will be in binaries. That means that potentially tens DNS and massive **PRP–ONeMg** are waiting to be discovered and have their masses measured! How many more neutron star mass and white dwarf masses will we measure? And even better, how many weird and surprising systems will we find? In addition to new discoveries, the increased sensitivity will provide a significant boost to the timing precision, allowing for mass measurements in already-discovered systems. And all of that is without even starting to talk about targeted searches for pulsars in the the Galactic center, globular clusters, and nearby galaxies.

It is clear that, with MeerKAT and its successors, the future of pulsar astronomy in the Southern Sky looks brighter than ever.





---

## Epilogue

---

Wow, that was a long journey. Thank you for staying all the way until the end. It took longer to write than expected, and I guess than it also took longer to read than expected. It definitely took a lot of my supervisors' time to give feedback on these texts. I can not thank them enough for it. I hope that at least you had fun with it and learned something about pulsar astronomy! That definitely was the case for me, not only during writing, but also during all of those years in which I have been working with all of that precious radio data from MeerKAT and Murriyang. If I was given the option, I would choose to do it all over again without any trace of doubt. But alas, you only do a PhD in pulsar astronomy once. At least, I carry all of that experience and some of the uncompleted experiments on to the next stage of my career. Maybe you will read from me again, so let us look forward to that. Until next time and, as we say in Catalan, *i vet aquí un gat, i vet aquí un gos, aquest conte ja s'ha fos*.

*Gràcies per llegir la meva tesis doctoral i arreveure!*



---

## The many angles of pulsar astronomy

---

Astronomy is all about angles. In a bachelor's degree, it is the first topic that students deal with when it comes to astronomy. That is because the sky is described in *angular coordinates*, so they are needed even to describe the most basic constellations. On top of this, the description of stellar orbits relies on angles as well. Therefore, this thesis uses a large (excessive?) amount of angular coordinates, variables and parameters. While many of these are introduced in the main text, I thought it would be a good idea to introduce them more rigorously in this Appendix for non-expert readers, so that it can be used as a consistent reference. Three kinds of angular quantities are described here: sky coordinates in Section A.1, running angular orbital phase parameters in Section A.2.1, and static angular orbital orientation parameters in Section A.2.2.

### A.1 Sky coordinate systems

Sky coordinates systems are the grids we use to describe pulsars positions (or any other astronomical object) in the sky. The Universe is three-dimensional, and thus we need three parameters to describe an object's position in it. However, the distance information becomes irrelevant when it comes to pointing a telescope at them,. Thus, an astronomical chart is a two-dimensional projection of the entire universe on to an imaginary sphere around the observer, the *celestial sphere*. On this sphere, we place *sky coordinates*, which have an axis of *latitude* and an axis of *longitude*. In this thesis, three different coordinate systems are mentioned: *equatorial coordinates*, *ecliptic coordinates* and *Galactic coordinates*.

#### Equatorial coordinates

In equatorial coordinates, the equator of the celestial sphere is coincident with the equator of the Earth (United Kingdom Hydrographic Office, 2011). The longitude coordinate is the *right ascension* RA, expressed in hours. It increases eastward starting from the March equinox, the point in the equator that is crossed by the Sun at the start of the spring<sup>90</sup>, and it has a maximum value of 25 hours, after which it cycles back to 0. The latitude coordinate is the *declination* DEC, expressed in degrees. It

---

<sup>90</sup> The Earth's rotation axis is tilted by  $\sim 23.44$  deg, so the Sun crosses the celestial equator twice a year: on March 20th and on September 22nd

increases from 0 deg at the equator until 90 deg at the north celestial pole, and it decreases until  $-90$  deg at the south celestial pole.

Equatorial coordinates are the most natural sky coordinates for observers, as its polar axis coincides with the axis of rotation of the Earth. Therefore, it is the easiest frame of reference to work with for telescope operators on Earth. As a consequence, the sky coordinates of almost all astronomical objects are given in RA and DEC. That is also the case for pulsars, and in fact their names carry their sky location. For instance, **PSR J1208–5936** means "the pulsar at RA: 12 h, 08 min; DEC: -59 deg, 36 min".

The currently used realisation is the International Celestial Reference System (ICRS, Charlot et al., 2020), which is derived from the position of thousands of extragalactic radio sources, using the position of March equinox at the Julian date of 2000.0 ("J2000.0", 2000 Jan 1st 12:00 UTC). Hence, the majority of pulsar names include a "J" before their coordinates. However, you may have noticed that some pulsar names begin with "B" instead, such as **PSR B1919+21**. These are pulsars were discovered before JD 2000.0, and their names are based on an earlier realisation of the equatorial coordinates, based on the March equinox position in the Besselian date of 1950.0 ("B1950.0", 1949 Dec 31st 22:09 UTC), which became obsolete during the 1990's. However, all RA and DEC positions quoted in this thesis use the JD 2000.0 reference.

### Ecliptic coordinates

In ecliptic coordinates, the celestial equator is coincident with the orbital plane of Earth, which is offset by 23.44 deg from the equator of the Earth. Thus, ecliptic coordinates are essentially derived from the equatorial coordinates. Its coordinates are plain longitude  $l_e$  and latitude  $b_e$ , both expressed in degrees, with  $0 \leq l_e/\text{deg} < 360$ . The longitude is counted from the march equinox once again, and latitude works in the same way as DEC does the equatorial coordinate system (United Kingdom Hydrographic Office, 2011). This frame of reference is relatively uncommon, but it finds its niche in the tracking of Solar System bodies such as the Sun, planets, asteroids, and comets. It is also relevant in pulsar astronomy since the barycentric correction of spin periods and the Rømer delay from the orbital motion of the Earth are best understood using the ecliptic coordinates of pulsars. However, since the orbital motion of Earth is irregular, fitting programs compute these effects in the equatorial coordinate system nonetheless.

### Galactic coordinates

Galactic coordinates use the plane of the Milky Way as the frame of reference. They also use longitude  $l_G$  and latitude  $b_G$ , but the origin of coordinates (0,0) now lies in the direction of the Galactic centre, with  $l_G$  growing Eastward. Its current version was defined in 1959 by the International Astronomical Union (IAU, the same ones who demoted Pluto to the category of Dwarf Planet in 2006) in 1959 (Blaauw et al., 1960), and it is the natural system of coordinates for studying the Milky Way. It is particularly useful for astronomers working Galactic stellar populations, and even to those that do extra-Galactic astronomy. Therefore, it is also the coordinate systems in which is the most natural to study Galactic pulsars. For instance, the electron density models NE2001 (Cordes, 2004) and YMW16 ((Yao et al., 2017)), which predict the observed DM values of pulsars, are defined in Galactic coordinates. Thus, we pulsar astronomers also use these coordinates when simulating pulsar populations in the Milky Way and when planning pulsar surveys on the Galactic plane.

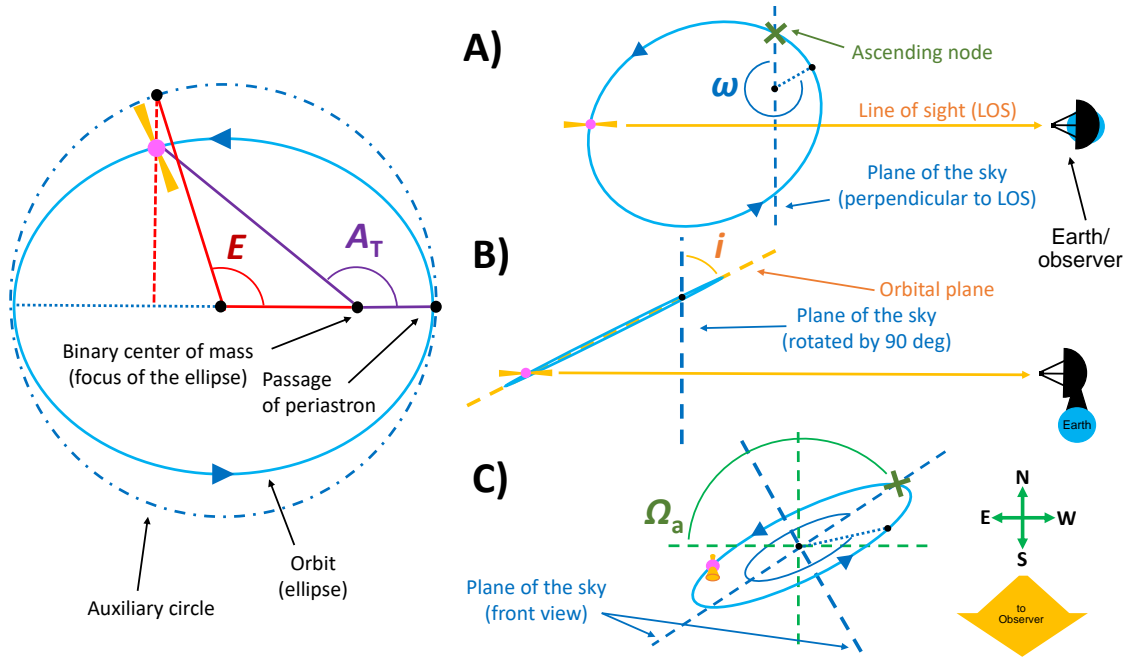


Figure A.1: Angles of pulsar binaries. **Left:** diagram of the Keplerian ellipse of binary orbits, tracing the eccentric anomaly,  $E$ , and the true anomaly,  $A_T$ , from the periastron in the position of the pulsar. **Right, A:** orientation of the ellipse within the orbital plane, defined by the angle of periastron  $\omega$ , measured from the sky plane to the Periastron. **Right, B:** same view as before, but rotated 90 degrees to show the inclination of the orbital plane with respect to the sky plane, measured by the inclination angle  $i$ . **Right, C:** same view as before, but from the observer's perspective, showing the orientation of the system, measured by the longitude of the ascending node,  $\Omega_a$ , from the East direction to the ascending node.

## A.2 Pulsar binary angles

There are many angles within a pulsar binary. I like to distinguish between two sets of them: the “running” angles relating to the orbital phase, which are independent of the observer, and the “static” angles describing the orbital geometry, which define the orientation of the orbit in relation to the observer. As explained in the main text, the latter are not truly static, as there are astrophysical and astrometric processes that can change their magnitude, but they would remain static if we ignored PK effects and the proper motion. We shall begin with the running angles, which are independent of the observer.

### A.2.1 Orbital phases: Running orbital angles

In pulsar astronomy, the orbital phase is expressed in three closely related angular parameters: the *mean anomaly*, the *eccentric anomaly* and the *true anomaly* (Roy, 2005; Lorimer and Kramer, 2005). All three of these angles are depicted in the left side of Fig. A.1, and they begin at 0 rad when the pulsar is at periastron, the tip of the orbital ellipse closest to the centre of mass of the system, and are cyclical until  $2\pi$ . However, they evolve differently as the pulsar moves along its orbit.

### Mean anomaly

The mean anomaly is proportional to the passage of time and is defined as

$$M = \frac{2\pi}{P_b} \left[ (t - T_0) - \frac{1}{2} \frac{\dot{P}_b}{P_b} (t - T_0)^2 \right], \quad (\text{A.1})$$

where  $P_b$  is the orbital period and  $T_0$  the time of passage of periastron. The only complication in the equation above is the presence of the orbital period derivative  $\dot{P}_b$ , which is only measurable in a few select systems. Other than that, it does not have much more of a secret than this: it is just time converted into an equivalent angle. Its usefulness lies in the ability to relate the passage of time to the other orbital angles. For circular systems where  $T_0$  is not properly defined, an alternative to the mean anomaly exists, that being the orbital phase measured from the time of passage of the ascending node,  $T_a$ , the point where the pulsar crosses the plain of the sky while moving away from the observer:

$$\Phi = \frac{2\pi}{P_b} \left[ (t - T_a) - \frac{1}{2} \frac{\dot{P}_b}{P_b} (t - T_a)^2 \right]. \quad (\text{A.2})$$

### Eccentric anomaly

The eccentric anomaly  $E$  is the first true angle. As shown in the left side of Fig. A.1, it is defined as the angle between the major axis of the orbital ellipse and the position of the pulsar in an imaginary auxiliary circle of radius equal to the semi-major axis of the ellipse, measured at the centre of the ellipse. It is related to  $M$  by the Kepler equations

$$E - e \sin E = M \quad \text{and} \quad \dot{E} = \frac{\dot{M}}{1 - e \cos E}, \quad (\text{A.3})$$

where  $e$  is the Keplerian orbital eccentricity of the pulsar. The eccentric anomaly  $E$  accounts for the eccentricity of the orbit and is used in the description of many orbital effects, such as the Rømer delay in the pulsar system, the Einstein delay and the Shapiro delay parameters. The equation written above also gives us a very nice physical insight:  $E$  evolves the fastest at periastron ( $E = M = 0$ ) and the slowest at apoastron ( $E = M = \pi$ ), which is coincident with the behaviour of the orbital velocity of the pulsar. It should be noted that if  $e = 0$ , this equation converses into  $E = M$ , meaning that the orbital velocity is the same at all orbital phases in circular orbits.

### True anomaly

The true anomaly  $A_T$  is the only physical angle that is measured directly by the pulsar. As shown in Fig. A.1, it is the angle between the major axis of the orbital ellipse and the pulsar position at the orbital ellipse as measured at the centre of mass of the system. It can be derived directly from  $E$  via

$$A_T = 2 \arctan \left[ \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2} \right], \quad (\text{A.4})$$

and its main use is to derive the real angle of the pulsar with respect to its companion. It can be interpreted as the angle of the imaginary line linking both components of the binary system, as the

system centre of mass is common to both of them, making it very useful to predict when conjunctions or eclipses will occur. If  $e = 0$ , it reduces to  $A_T = E = M$ .

### A.2.2 Orbital geometry: static orbital angles

The orbit of a pulsar is an ellipse with a three-dimensional orientation, but from Earth we can only measure its motion along the line-of-sight. This means that depending on how the orbit is oriented relating to us, its motion may manifest in different ways. There are three angles describing the orientation of a pulsar orbit: the angle of periastron,  $\omega$ , the inclination angle,  $i$ , and the longitude of the ascending node,  $\Omega_a$  (Roy, 2005; Lorimer and Kramer, 2005).

#### Longitude of periastron

The longitude of periastron  $\omega$  describes the orientation of the orbit within the orbital plane. As shown in the **A** plot in Fig. A.1, it is the angle between the ascending node and periastron. Running from 0 to 360 degrees, this angle already manifests on the observed spin evolution of the pulsar modulated by the Doppler effect and in the Rømer delay, and thus it is essential for the Keplerian description of the system.

If  $e = 0$ ,  $\omega$  becomes undefined, and we use the time at the ascending node  $T_a$  instead of  $T_0$ . The conversion between  $T_a$  and  $T_0$  is computed via

$$T_a = T_0 - \omega \frac{P_b}{2\pi}. \quad (\text{A.5})$$

The angle  $\omega$  is not truly fixed, as it rotates as predicted by general relativity, an effect that can become noticeable in the span of months or years via pulsar timing as changes in the shape of the Rømer delay. Thus, not only is  $\omega$  essential for the Keplerian description of a pulsar binary, but also in its post-Keplerian description.

#### Inclination angle

The inclination angle  $i$  measures the angle between the orbital plane and the sky plane, as shown in the **B** plot in Fig. A.1. Its magnitude decides which fraction of the velocity of the pulsar is projected onto the line-of-sight, with a direct effect on the measurement of the projected semi-major axis  $x$ . While  $i$  is not measurable from the Keplerian description of the system, it does change the apparent magnitude of the projected semi-major axis with  $x = a \times \sin i$ . This means that, at least in theory, every pulsar classified as isolated could be in a binary system with  $i = 0$  deg! However,  $i$  is directly measurable via the Shapiro delay, or indirectly through the mass function from mass measurements achieved with post-Keplerian parameters. As explained in the text, it is randomly distributed across observed binaries, with its median value being  $i = 60$  deg. In theory,  $i$  can run from 0 to 180 degrees, but for many systems  $i$  is degenerate with  $180 - i$ , as its effects often only show with  $\sin i$  or  $\tan i$  in the post-Keplerian parameters derived both from general relativity or proper motion. However, in nearby systems, the orbit of the Earth is large enough that the orbital orientation of the system seems to change periodically across the year, breaking this symmetry (Kopeikin, 1995).



### Longitude of ascending node

Honestly, this is the most subtle of the angles, and it took me a while to understand its significance. It is the angle between the east-ward direction in the plane of the sky and intersection of the orbital plane and the sky plane. Or, in other words, it is the angle of the ascending node from the east-ward direction, where the ascending node is the point where the pulsar crosses into the sky plane from the point of view of the observer (plot C in Fig. A.1). It has no effect whatsoever on the Keplerian description of the system nor on the post-Keplerian parameters derived from general relativity. However, it can affect  $\dot{x}$  and  $\dot{\omega}$  if the proper motion of the system is large, making it measurable after many years of timing in several systems (Arzoumanian et al., 1996; Kopeikin, 1996). It can take any orientation between 0 and 360 degrees, and it is measurable via the proper motion contributions of  $\dot{x}$  and  $\dot{\omega}$ . In the case of  $\dot{x}$ , if we transform  $i$  into  $180 - i$  degrees, and  $\Omega_a$  into  $\Omega_a + 180$  degrees, the post-Keplerian description of the system remains unaltered, so in the majority of systems there are two possible, equally valid  $(i, \Omega_a)$  solutions. However, this degeneracy does not exist for  $\dot{\omega}$ , and in systems where that effect is significant, or where there is significant parallax, this degeneracy can be broken (e.g. Guo et al., 2021).

### The distribution of $i$ and the mass function

The orientation of the orbital planes of all binary systems in the sky are uniformly distributed across an imaginary sphere with coordinates  $i$  and  $\Omega$ , where  $i$  acts as a latitude coordinate and  $\Omega$  as a longitude coordinate. This means that they follow a probability density distribution

$$\rho(i, \Omega_a) \delta i \delta \Omega_a = \frac{1}{4\pi} \sin i \delta i \delta \Omega_a, \quad (\text{A.6})$$

where  $0 \leq \Omega_a < 360$  deg and  $0 \leq i \leq 180$  deg. However, for the purposes of the mass function equation

$$\frac{4\pi^2}{T_\odot} \frac{x^3}{P_b^2} = \frac{(M_c \sin i)^3}{(M_p + M_c)^2}, \quad (\text{A.7})$$

only the  $i$  parameter is relevant, so we can integrate in  $\Omega_a$  and reduce the range of  $i$  to  $0 \leq i \leq 90$  deg. This results in a reduced probability distribution for  $i$

$$\rho(i) \delta i = \frac{1}{2\pi} \left[ \int_0^{2\pi} d\Omega_a \right] \sin i \delta i = \sin i \delta i. \quad (\text{A.8})$$

This distribution is zero at  $i = 0$  deg and increases steadily until  $i = 90$  deg. This means that it is actually very unlikely for a system to be (almost) completely face-on, and on the other hand it is more likely that it is closer to edge on. This is for the same reason as why we can fit every person on Earth in a line along the equator, but only one on each pole: many orbital orientation configurations result in  $i = 90$  deg, but only one results in  $i = 0$  deg. The median  $i$  value of this distribution is 60 deg, which is often used to quote the *median companion mass*.

---

### Paper on PSR J1208–5936

---

The paper Colom i Bernadich et al. (2023), A&A, **678**, A187, was published in the peer-reviewed journal of Astronomy & Astrophysics on 23/10/2023 under the Creative Commons CC-BY 4.0 license, and is reproduced in its original form below. This paper is part of a series of MMGPS papers, the first one being Padmanabh et al. (2023), MNRAS, **554**, 1, 1291–1315. More follow-up papers on MMGPS discoveries are to come in the future.

# The MPIfR-MeerKAT Galactic Plane Survey

## II. The eccentric double neutron star system PSR J1208–5936 and a neutron star merger rate update

M. Colomé Bernadich<sup>1</sup>, V. Balakrishnan<sup>1</sup>, E. Barr<sup>1</sup>, M. Berezina<sup>1,2</sup>, M. Burgay<sup>3</sup>, S. Buchner<sup>4</sup>,  
 D. J. Champion<sup>1</sup>, W. Chen<sup>1</sup>, G. Desvignes<sup>1</sup>, P. C. C. Freire<sup>1</sup>, K. Grunthal<sup>1</sup>, M. Kramer<sup>1</sup>, Y. Men<sup>1</sup>,  
 P. V. Padmanabh<sup>5,6,1</sup>, A. Parthasarathy<sup>1</sup>, D. Pillay<sup>1</sup>, I. Rammala<sup>1</sup>, S. Sengupta<sup>1</sup>, and V. Venkatraman Krishnan<sup>1</sup>

<sup>1</sup> Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany  
 e-mail: mcbnadich@mpi-fir-bonn.mpg.de

<sup>2</sup> Landessternwarte, Universität Heidelberg, Königstuhl 12, 69117 Heidelberg, Germany

<sup>3</sup> INAF – Osservatorio Astronomico di Cagliari, Via della Scienza 5, 09047 Selargius, CA, Italy

<sup>4</sup> South African Radio Astronomy Observatory, 2 Fir Street, Black River Park, Observatory 7925, South Africa

<sup>5</sup> Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstraße 38, 30167 Hannover, Germany

<sup>6</sup> Leibniz Universität Hannover, Welfengarten 1, 30167 Hannover, Germany

Received 19 May 2023 / Accepted 29 August 2023

### ABSTRACT

The MPIfR-MeerKAT Galactic Plane survey at *L*-band (MMGPS-L) is the most sensitive pulsar survey in the Southern Hemisphere, providing 78 discoveries in an area of 900 sq. deg. Here, we present a follow-up study of one of these new discoveries, PSR J1208–5936, a 28.71-ms recycled pulsar in a double neutron star system with an orbital period of  $P_b = 0.632$  days and an eccentricity of  $e = 0.348$ , merging within the Hubble time. Through timing of almost one year of observations, we detected the relativistic advance of periastron ( $\dot{\omega} = 0.918(1) \text{ deg yr}^{-1}$ ), resulting in a total system mass of  $M_t = 2.586(5) M_\odot$ . We also achieved low-significance constraints on the amplitude of the Einstein delay and Shapiro delay, in turn yielding constraints on the pulsar mass ( $M_p = 1.26^{+0.13}_{-0.25} M_\odot$ ), the companion mass ( $M_c = 1.32^{+0.25}_{-0.13} M_\odot$ ), and the inclination angle ( $i = 57 \pm 12^\circ$ ). This system is highly eccentric compared to other Galactic field double neutron stars with similar periods, possibly hinting at a larger-than-usual supernova kick during the formation of the second-born neutron star. The binary will merge within 7.2(2) Gyr due to the emission of gravitational waves, making it a progenitor of the neutron star merger events seen by ground-based gravitational wave observatories. With the improved sensitivity of the MMGPS-L, we updated the Milky Way neutron star merger rate to be  $\mathcal{R}_{\text{MW}}^{\text{new}} = 25^{+19}_{-9} \text{ Myr}^{-1}$  within 90% credible intervals, which is lower than previous studies based on known Galactic binaries owing to the lack of further detections despite the highly sensitive nature of the survey. This implies a local cosmic neutron star merger rate of  $\mathcal{R}_{\text{local}}^{\text{new}} = 293^{+222}_{-103} \text{ Gpc}^{-3} \text{ yr}^{-1}$ , which is consistent with LIGO and Virgo O3 observations. With this, we also predict the observation of  $10^{+8}_{-4}$  neutron star merger events during the LIGO-Virgo-KAGRA O4 run. We predict the uncertainties on the component masses and the inclination angle will be reduced to  $5 \times 10^{-3} M_\odot$  and  $0.4^\circ$  after two decades of timing, and that in at least a decade from now the detection of  $\dot{P}_b$  and the sky proper motion will serve to make an independent constraint of the distance to the system.

**Key words.** binaries: close – ephemerides – gravitational waves – stars: neutron – celestial mechanics – stars: fundamental parameters

### 1. Introduction

Double neutron star (DNS) binaries are the evolutionary end-point of massive binary systems ( $M_{\text{stars}} > 8 M_\odot$ ) that survive two supernovae while remaining bound (e.g. Tauris et al. 2017; Vigna-Gómez et al. 2018; Chattopadhyay et al. 2020). In this picture, the primary star undergoes a Type Ib/c core collapse supernova, becoming the first-born neutron star (NS). Subsequently, the system undergoes a common-envelope phase as the secondary evolves out of the main sequence (e.g. van den Heuvel 2019), after which the outer shells of the evolved secondary are expelled, turning the system into a circular, compact NS–He star binary (e.g. Chattopadhyay et al. 2020). As the secondary continues to evolve, a case BB Roche-lobe overflow (RLO) ensues in which the primary NS sustains partial recycling (Tauris et al. 2015), burying its magnetic field and spinning up to a rotation period of a few tens of milliseconds (Bhattacharya & van den Heuvel 1991). Finally, the sys-

tem must avoid the disruption during the second, ultra-stripped supernova in which the secondary NS is born with the expulsion of 0.1 to  $1 M_\odot$  from the system, and with a maximum kick velocity of  $\approx 100 \text{ m s}^{-1}$  (Tauris et al. 2015). The nature of this supernova is still a matter of discussion as it can be triggered by either electron capture in the degenerate ONeMg core or iron core collapse, depending on how much stripping has been suffered by the He star, with more massive remnants undergoing the latter channel and suffering greater mass losses and supernova kicks (Tauris et al. 2015). The distinction between these two formation mechanisms is blurry due to the randomness of the resulting orbital parameters arising from the unpredictability of the supernova kick direction and magnitude, but the remaining mass of the second-born NS may be a tell-tale sign, as it has been seen that higher eccentricity systems correlate with more massive second-born NSs and less rapidly spinning recycled pulsars (e.g. Sengar et al. 2022; Andrews & Mandel 2019; Faulkner et al. 2005; Tauris et al. 2017). Nonetheless, if

all of these steps are completed without disruption, then a new, eccentric DNS system is born.

Observable radio pulsars discovered in DNS systems are amongst the most useful astrophysical tools. With dedicated follow-up and observing campaigns using radio telescopes, their orbital parameters and component masses can be measured to high precision through the technique of pulsar timing (e.g. [Lorimer & Kramer 2005](#)), allowing for tests of formation channels and even of fundamental physics. This is demonstrated in the literature on formation mechanisms (e.g. [Tauris et al. 2017](#)), breakthrough tests of gravity (e.g. [Taylor et al. 1979](#); [Taylor & Weisberg 1982](#); [Kramer et al. 2021](#)), and the validation or exclusion of dense matter models (e.g. [Özel & Freire 2016](#); [Hu et al. 2020](#)). Beyond radio observations, the discovery of pulsars in DNS systems has also had an indirect impact on other fields of astrophysics. The observation of the orbital decay of B1913+16 ([Taylor et al. 1979](#)) experimentally demonstrated the existence of gravitational waves as predicted by general relativity (GR) for the first time and supported the prediction of its merger in 301 Myr, providing scientific justification for the construction of ground-based gravitational-wave observatories. Indeed, with knowledge of the pulsar DNS population in the Milky Way, their rates of orbital decay, and of the sensitivity of blind surveys on the sky, predictions of the observed cosmic rate of NS mergers can be made (e.g. [Kim et al. 2003, 2010, 2015](#); [Pol et al. 2019, 2020](#); [Grunthal et al. 2021](#)). The most recent estimate ([Grunthal et al. 2021](#)) provides an upper limit of  $\mathcal{R}_{\text{local}} \leq 597 \text{ yr}^{-1}$ , which is consistent with the observed rates at ground-based gravitational wave detectors ([Abbott et al. 2021](#)), highlighting the synergy between radio and gravitational wave observations. Finally, the first detection of a gravitational wave signal emitted by the coalescence of two NSs, GW170817 ([Abbott et al. 2017a](#)), opened a new era for multi-wavelength and multi-messenger astronomy, highlighting DNS systems among other astrophysical systems even further ([Abbott et al. 2017b](#)).

However, observable pulsars in DNS systems are a rarity. At the time of writing only 17 Galactic DNS systems are confirmed, of which only nine are competitive for tests of gravity and NS merger rate predictions (Table 1). Indeed, simulations predict a Galactic DNS formation rate of just 5–31  $\text{Myr}^{-1}$ , with only  $\sim 0.13\%$  of massive binaries surviving the DNS formation channels ([Vigna-Gómez et al. 2018](#)). The number of discoveries is also limited by observational reasons. With the exceptions of J0737–3039B, J1906+0746 and perhaps J1755–2550 (Table 1), typically only the first-born, recycled NS is observable due to their longer-lasting emission and larger beaming fractions, while the second-born NS is set to spin down and cross its pulsar death line a few Myr after birth ([Lorimer & Kramer 2005](#)). Limitations are also computational, as the extreme orbital motion of pulsars in compact, relativistic DNS systems hampers the effectiveness of traditional periodicity searches from pulsar surveys due to the Doppler smearing of the signal across an observation ([Johnston & Kulkarni 1991](#)). It is, therefore, necessary to implement computationally expensive and sophisticated pulsar searches to increase the sample of known DNS systems (and hence maximise their science output). These generally include acceleration searches (e.g. [Keith et al. 2010](#)), jerk searches (e.g. [Andersen & Ransom 2018](#); [Eatough et al. 2021](#); [Suresh et al. 2022](#)) or even template bank searches (e.g. [Allen et al. 2013](#); [Balakrishnan et al. 2022](#)).

The MPIfR–MeerKAT Galactic Plane survey at  $L$ -band (MMGPS-L; [Kramer et al. 2016](#); [Stappers & Kramer 2016](#); [Padmanabh et al. 2023](#)) has been highly successful in this objec-

tive. Designed to discover massive, compact binaries in the southern Galactic plane with a time resampling-based acceleration search, it has yielded the discovery of two previously unknown DNS systems, PSR J1155–6529 ([Berezina et al., in prep.](#)) and PSR J1208–5936 (J1208–5936 from now on). Additionally, the MMGPS-L is the most sensitive pulsar survey in the southern sky, making its large coverage area highly relevant to estimate the rate of NS mergers in the Milky Way. In this work, we present the follow-up study of J1208–5936, a 28.71-ms recycled pulsar in a highly eccentric orbit with another NS, with whom it is bound to merge in  $7.2 \pm 0.2 \text{ Gyr}$  due to the emission of gravitational waves. In addition, we also use this discovery and the improvement in depth of sky coverage provided by the MMGPS-L to recompute the predicted NS merger rate.

This paper is structured as follows: in Sect. 2, we present the discovery and the derivation of an early orbital solution. In Sect. 3, we show a phase-coherent timing solution and perform preliminary mass measurements with almost one year of timing data. In Sect. 4, we study its general properties, including its pulse profile, a comparison with other pulsars in DNS systems and a discussion of its formation channel. In Sect. 5, we search for pulsations from its companion. In Sect. 6, we investigate the implications of the discovery and the performance of the MMGPS-L for estimations of the local NS merger rate, and the detection of these events by ground-based gravitational-wave observatories. And finally, in Sect. 7 we discuss possible biases and prospects for future improvements.

## 2. Discovery and orbital solution

Led by the Max Planck Institute for Radioastronomy<sup>1</sup> (MPIfR) in collaboration with the South African Radio Astronomy Observatory<sup>2</sup> (SARAO), the MMGPS-L was a fast Fourier transform (FFT)-based pulsar search survey with the MeerKAT radio interferometer<sup>3</sup> ([Jonas & MeerKAT Team 2016](#)). Covering 900 sq. deg. on the southern Galactic plane with the primary aim to discover faint relativistic binaries, the MMGPS-L implemented a time-domain acceleration search ([Padmanabh et al. 2023](#)) with the PEASOUP<sup>4</sup> pipeline ([Barr 2020](#)).

J1208–5936 was discovered on 30 May 2021 in an observation from May 6th as a topocentric 28.706-ms signal in the FFT with signal-to-noise ratio  $S/N = 15.1$ , a line-of-sight acceleration of  $-8.86 \text{ m s}^{-2}$ ,  $\text{DM} = 344.2 \text{ pc cm}^{-3}$  and a PulsarX<sup>5</sup> ([Men et al. 2023](#)) fold of  $S/N \approx 19$ . Upon discovery, we performed a first 20 min-long follow-up observation at the same position on 4 June 2021. This yielded a re-detection with  $S/N \approx 15$ , and a PRESTO/prepfold<sup>6</sup> fold of  $S/N \approx 12$  with a changed period of  $P = 28.697 \text{ ms}$ , thus confirming both its existence and its binary nature.

The interferometric nature of MeerKAT, and therefore the multibeam pattern of the follow-up observation allowed for precise localisation of J1208–5936 with the SeeKAT<sup>7</sup> software ([Bezuidenhout et al., in prep.](#)), which compares the  $S/N$  of detections in different beams and does a maximum-likelihood estimate of the best position of the source, taking into account beam positions and their point-spread functions as derived by

<sup>1</sup> <https://www.mpifr-bonn.mpg.de/2169/en>

<sup>2</sup> <https://www.sarao.ac.za/>

<sup>3</sup> <https://www.sarao.ac.za/science/meerkat/about-meerkat/>

<sup>4</sup> <https://github.com/ewanbarr/peasoup>

<sup>5</sup> <https://github.com/ypmen/PulsarX>

<sup>6</sup> <https://github.com/scottransom/presto>

<sup>7</sup> <https://github.com/BezuidenhoutMC/SeeKAT>

**Table 1.** All known pulsars in Galactic DNS systems and candidates with their spin parameters, orbital parameters, and component masses ordered by merger times  $\tau_m$ , including J1208–5936 (highlighted).

PSR	$P_0$ (ms)	$\dot{P}$ ( $\text{s s}^{-1}$ )	$P_b$ (days)	$x$ (ls)	$e$	$M_t$ ( $M_\odot$ )	$M_c$ ( $M_\odot$ )	$M_p$ ( $M_\odot$ )	$\tau_m$ (Gyr)
J1946+2052 <sup>(1)</sup>	16.960	$9.20 \times 10^{-19}$	0.078	1.154	0.064	2.50(4)	>1.18	<1.31	$\sim 0.0455$ <sup>(a)</sup>
J1757–1854 <sup>(2)</sup>	21.497	$2.63 \times 10^{-18}$	0.184	2.238	0.606	2.732876(8)	1.3917(4)	1.3412(4)	0.0761
J0737–3039A <sup>(3)</sup>	22.699	$1.76 \times 10^{-18}$	0.102	1.415	0.088	2.587052(9)	1.248868(13)	1.338185(14)	0.0860
J0737–3039B <sup>(3)</sup>	2773.5	$8.92 \times 10^{-16}$	0.102	1.516	0.088	2.587052(9)	1.338185(14)	1.248868(13)	0.0860
B1913+16 <sup>(4)</sup>	59.030	$8.62 \times 10^{-18}$	0.323	2.342	0.617	2.828(1)	1.390(1)	1.438(1)	0.301
J1906+0746 <sup>(5)</sup>	144.07	$2.03 \times 10^{-14}$	0.166	1.42	0.085	2.6134(3)	1.322(11)	1.291(11)	0.308
J1913+1102 <sup>(6)</sup>	27.285	$1.61 \times 10^{-19}$	0.206	1.755	0.090	2.8887(6)	1.27(3)	1.62(3)	0.470
J0509+3801 <sup>(7)</sup>	76.541	$7.93 \times 10^{-18}$	0.38	2.051	0.586	2.805(3)	1.46(8)	1.34(8)	0.576
J1756–2251 <sup>(8)</sup>	28.462	$1.02 \times 10^{-18}$	0.32	2.756	0.181	2.56999(6)	1.230(7)	1.341(7)	1.66
B1534+12 <sup>(9)</sup>	37.904	$2.42 \times 10^{-18}$	0.421	3.729	0.274	2.678463(4)	1.3455(2)	1.3330(2)	2.73
<b>J1208–5936<sup>(10)</sup></b>	<b>28.714</b>	<b><math>&lt;4 \times 10^{-20}</math></b>	<b>0.632</b>	<b>4.257</b>	<b>0.348</b>	<b>2.586(6)<sup>(c)</sup></b>	<b>1.32<sup>+0.25</sup><sub>–0.13</sub><sup>(c)</sup></b>	<b>1.26<sup>+0.13</sup><sub>–0.25</sub><sup>(c)</sup></b>	<b>7.2(2)</b>
J1829+2456 <sup>(11)</sup>	41.010	$5.25 \times 10^{-20}$	1.176	7.238	0.139	2.60551(19)	1.299(4)	1.306(7)	55
J1325–6253 <sup>(12)</sup>	28.969	$4.80 \times 10^{-20}$	1.816	7.574	0.064	2.57(6)	>0.98	<1.59	$\sim 189$ <sup>(a)</sup>
J1411+2551 <sup>(13)</sup>	62.453	$9.56 \times 10^{-20}$	2.616	9.205	0.17	2.538(22)	>0.92	<1.62	$\sim 466$ <sup>(a)</sup>
J1759+5036 <sup>(14)</sup>	176.02	$2.43 \times 10^{-19}$	2.043	6.825	0.308	2.62(3)	>0.7006	<1.9194	$\sim 177$ <sup>(a)</sup>
J0453+1559 <sup>(15)</sup>	45.782	$1.86 \times 10^{-19}$	4.072	14.467	0.113	2.734(4)	1.174(4)	1.559(5)	1453
J1811–1736 <sup>(16)</sup>	104.18	$9.01 \times 10^{-19}$	18.779	34.783	0.828	2.57(10)	>0.93	<1.64	$\sim 1800$ <sup>(a)</sup>
J1518+4904 <sup>(17)</sup>	40.935	$2.72 \times 10^{-20}$	8.634	20.044	0.249	2.7183(7)	1.31(8)	1.41(8)	8844
J1018–1523 <sup>(18)</sup>	83.152	$1.09(6) \times 10^{-19}$	8.984	26.157	0.228	2.3(3)	>1.16	<1.1(3)	$\sim 1.4(3) \times 10^4$ <sup>(a)</sup>
J1930–1852 <sup>(19)</sup>	185.52	$1.80 \times 10^{-17}$	45.06	86.89	0.399	2.59(4)	>1.30	<1.32	$\sim 5.32 \times 10^5$ <sup>(a)</sup>
J1755–2550 <sup>(20)</sup>	315.20	$2.43 \times 10^{-15}$	9.696	12.284	0.089	...	>0.39	...	... <sup>(b)</sup>
J1753–2240 <sup>(21)</sup>	95.138	$9.70 \times 10^{-19}$	13.638	18.115	0.304	...	>0.4875	...	... <sup>(b)</sup>

**Notes.** The first solid line splits systems merging within the Hubble time from those that do not. The last two entries have no total mass measurements, and therefore their DNS nature is uncertain. Nonetheless, they are promising candidates due to their eccentricity. <sup>(a)</sup>Due to uncertain masses, merger times values are only estimates. <sup>(b)</sup>Merger time unavailable due to the lack of a total system mass measurement. <sup>(c)</sup>Total mass quoted from the direct DDGR fit, component masses quoted from the DDGR  $\chi^2$  mapping marginal one-dimensional likelihood distributions, shown in Fig. 2 and explained in Sect. 3.

**References.** (1) Stovall et al. (2018), (2) Cameron et al. (2023), (3) Kramer et al. (2021), (4) Weisberg & Huang (2016), (5) van Leeuwen et al. (2015), (6) Ferdman et al. (2020), (7) Lynch et al. (2018), (8) Ferdman et al. (2014), (9) Fonseca et al. (2014), (10) this work, (11) Haniewicz et al. (2021), (12) Sengar et al. (2022), (13) Martinez et al. (2017), (14) Agazie et al. (2021), (15) Martinez et al. (2015), (16) Corongiu et al. (2007), (17) last published work (Janssen et al. 2008) and Sect. 8.3 of Tauris et al. (2017), (18) Swiggum et al. (2023), (19) Swiggum et al. (2015), (20) Ng et al. (2018), and (21) Keith et al. (2009).

the Mosaic<sup>8</sup> software (Chen et al. 2021a). This improved the sky position to an uncertainty of just  $\sim 5$  arcsec, much better than the  $\sim 20$  arcsec precision provided by the beam size alone, and which was only 3.3 arcsec away from the timing-derived position (Table 2).

From 4 June 2021 to 31 August 2021, we typically scheduled dedicated sessions twice a week, consisting of two 20-min observations stored as filterbank data in the accelerated pulsar search user supplied equipment cluster in South Africa (APSUSE-mode; Padmanabh et al. 2023). Each session had the two observations conducted just a few hours apart. We folded each observation into archives with the dspsr<sup>9</sup>, where cycles of the pulsar are stacked every 8 or 10 seconds. We then cleaned these archives with the clfd<sup>10</sup> radio-frequency interference (RFI) excision software (Morello et al. 2019). These resulting data allow for accurate tracking of the evolution or drift of the pulse over time, leading to measurements of the barycentric period ( $P_{\text{bary}}$ ) at each observing epoch ( $t_{\text{obs}}$ ) with the PSRCHIVE/pdmp<sup>11</sup> software (Hotan et al. 2004), but also with the timing software TEMPO2<sup>12</sup> (Hobbs et al. 2006; Edwards et al. 2006) for observations with large line-of-sight acceleration or jerk.

We implemented a modified version of the roughness estimate algorithm (REA) from Bhattacharyya & Nityananda (2008) to solve for the orbital period, available in estimateOrbit.py<sup>13</sup>. Given adequate orbital coverage, the REA is more sensitive to orbital periods in systems with non-zero eccentricity than the Lomb–Scargle periodogram. It performs a search in the orbital period  $P_b$  space by folding the  $P_{\text{bary}}(t_{\text{obs}})$  series and evaluating the smoothness of the resulting curve with

$$R = \sum_{i=1}^n \left[ \frac{P_{\text{bary}}(t_{i+1}) - P_{\text{bary}}(t_i)}{\phi(t_{i+1}) - \phi(t_i)} \right]^2, \quad (1)$$

where  $t_i = t_{\text{obs},i}/P_{b,\text{trial}} - \text{mod}(t_{\text{obs},i}, P_{b,\text{trial}})$  are the folded observation epochs and  $\phi$  corresponds to the orbital phase (mean anomaly). We then select  $P_{b,\text{trial}}$  values that minimise  $R$ , as they produce a smooth fold and are likely to correspond to the true  $P_b$  value. The normalisation by the difference in orbital phase  $\Delta\phi$  is a modification of the original REA presented in Bhattacharyya & Nityananda (2008) that gives higher significance to data points that are closer in orbital phase, dealing with possible gaps in the orbital coverage.

<sup>8</sup> <https://github.com/wchenastro/Mosaic>

<sup>9</sup> <https://dspsr.sourceforge.net/>

<sup>10</sup> <https://github.com/v-morello/clfd>

<sup>11</sup> <http://psrchive.sourceforge.net/>

<sup>12</sup> <https://bitbucket.org/psrsoft/tempo2>

<sup>13</sup> <https://github.com/mcbernadich/CandyCracker/blob/main/estimateOrbit.py>



**Table 2.** Timing parameters resulting from the DDH and DDGR fits.


Data reduction parameters		
Binary model	DDH	DDGR
TEMPO2 wrapper	TempoNest	chi2Map.py
Solar System ephemeris	DE430	DE430
Timescale	TCB	TCB
Reference epoch for period and DM	59390	59390
Number of ToAs	384	384
Root mean squared of ToA residuals ( $\mu$ s)	36.99	36.93 <sup>(a)</sup>
Reduced $\chi^2$	0.949	0.944 <sup>(a)</sup>
Spin and astrometric parameters		
Right ascension, $\alpha$ (J2000, hh:mm:ss)	12:08:27.024(1)	12:08:27.023(1)
Declination, $\delta$ (J2000, dd:mm:ss)	−59:36:20.485(5)	−59:36:20.486(6)
Spin frequency, $\nu$ (Hz)	34.8263871091(1)	34.8263871091(2)
Spin-down rate, $\dot{\nu}$ ( $10^{-17}$ Hz s $^{-1}$ )	−3.6(1.1)	−3.1(1.3)
Dispersion measure, DM (pc cm $^{-3}$ )	344.427(6)	344.428(6)
First derivative of DM, DM1 ( $10^{-2}$ pc cm $^{-3}$ yr $^{-1}$ )	2.36(86)	2.40(98)
Keplerian orbital parameters		
Orbital period, $P_b$ (days)	0.631566177(3)	0.631566176(4)
Orbital eccentricity, $e$	0.347988(1)	0.347988(2)
Longitude of periastron, $\omega$ (deg)	247.98(1)	247.98(2)
Projected semi-major axis of the pulsar orbit, $x$ (ls)	4.2570(4)	4.2571(5)
Epoch of periastron, $T_0$ (MJD)	59340.210571(2)	59340.210572(2)
Post-Keplerian (PK) orbital parameters		
Rate of advance of periastron $\dot{\omega}$ (deg yr $^{-1}$ )	0.918(1)	[0.918(1)] <sup>(b)</sup>
Amplitude of Einstein delay, $\gamma_E$ (ms)	2.93(98)	{3.01 $^{+0.80}_{-0.40}$ }
Orthometric amplitude of Shapiro delay, $h_3$ ( $\mu$ s)	1.10(97)	{1.06 $^{+0.92}_{-0.52}$ }
Orthometric ratio of Shapiro delay, $\varsigma$	0.41(18)	{0.55 $^{+0.15}_{-0.13}$ }
Derivative of orbital period, $\dot{P}_b^{\text{GR}}$ ( $10^{-13}$ s s $^{-1}$ )	...	{−1.225 $^{+0.026}_{-0.009}$ }
Derivative of projected semi-major axis, $\dot{x}^{\text{GR}}$ ( $10^{-18}$ ls s $^{-1}$ )	...	{−6.37 $^{+0.14}_{-0.05}$ }
Mass measurements and derived orbital parameters		
Mass function, $f_M$ ( $M_\odot$ )	0.20758(6)	0.20760(7)
Total mass, $M_t$ ( $M_\odot$ )	[2.586(5)] <sup>(b)</sup>	2.586(6)
Companion mass, $M_c$ ( $M_\odot$ )	{1.36 $^{+0.18}_{-0.12}$ }	1.39(37)/{1.32 $^{+0.25}_{-0.13}$ } <sup>(c)</sup>
Pulsar mass, $M_p$ ( $M_\odot$ )	{1.24 $^{+0.12}_{-0.18}$ }	{1.26 $^{+0.13}_{-0.25}$ }
Inclination angle, $i$ (deg)	{55(9)}	{57(12)}
Merger time, $\tau_m$ (Gyr)		7.2(2)
Derived spin and astrometric parameters		
Galactic longitude, $l$ (deg)		297.512
Galactic latitude, $b$ (deg)		2.813
Spin period, $P_0$ (ms)	28.7138598921(1)	28.7138598921(2)
Spin period derivative, $\dot{P}$ ( $10^{-20}$ s s $^{-1}$ )	2.98(0.90)	2.57(1.04)
Characteristic age, $\tau_c$ (Gyr)		>10
Surface magnetic field strength, $B_{\text{surf}}$ ( $10^9$ G)		<30
NE2001 DM-derived distance, $d$ (kpc)		8.2(1.6)
YMW16 DM-derived distance, $d$ (kpc)		8.5(1.7)

**Notes.** Values without brackets come from the TempoNest fit of DDH or the direct TEMPO2 fit of DDGR. Values with square brackets [...] are derived from the TempoNest DDH parameters or the direct TEMPO2 DDGR parameters ( $M_t$  for DDH and  $\dot{\omega}$  for DDGR). Values with curly brackets {...} are derived from the multiplication of PK density probabilities in DDH or the  $\chi^2$  mapping probability in DDGR. All Keplerian, spin, position and DM parameters presented for DDGR are taken from the direct TEMPO2 fit. These fits also include a time jump between the APSUSE and PTUSE data-sets, as well as a DM jump (see Appendix A for details). <sup>(a)</sup>Taken from the direct TEMPO2 DDGR fit. Variations from the  $\chi^2$  mapping (Sect. 3) are only of the order of  $1/369 \approx 0.003$  (Reduced  $\chi^2 = \chi^2 / \text{degrees of freedom}$ ). <sup>(b)</sup>Derived through Eq. (2). In the DDGR case, this simple derivation has been chosen instead of the marginalisation of the  $\chi^2$  mapping due to it being much more constraining. <sup>(c)</sup>The first value is the direct DDGR fit, with a Gaussian uncertainty. The second one results from the  $\chi^2$  mapping of solutions (Sect. 3).

The REA found a significant signal at  $P_{b,\text{trial}} = 0.632$  days. We then used pyfitorbit<sup>14</sup> software to fit for the remaining

parameters using the best REA  $P_{b,\text{trial}}$  as a first guess, which resulted in a confirmation of the orbital period and a fit for the remaining Keplerian parameters (Fig. 1), with a eccentricity  $e = 0.34$  and a mass function  $f_M = 0.208 M_\odot$ .

<sup>14</sup> <https://github.com/gdesvignes/pyfitorbit>



**Fig. 1.** Orbital solution of J1208–5936 based on the  $P_{\text{bary}}(t_{\text{obs}})$  time series between 5 May 2021 and 31 August 2021. Observations with precise acceleration and/or jerk measurements were used to extract several data points (5 May:  $-10.5 \text{ m s}^{-2}$ ; 11 June:  $+22.4 \text{ m s}^{-2}$ , and 16 August:  $+7.2 \text{ m s}^{-2}$ ,  $-3.6 \times 10^{-3} \text{ m s}^{-3}$ ; at  $\phi \approx 0.72, 0.95$  and  $0.15$ ). Assuming a canonical pulsar mass of  $M_p = 1.35 M_\odot$ , the minimum companion mass computed from the mass function is  $M_c \gtrsim 1.1 M_\odot$  (see Table 2).

### 3. Timing and mass measurements

After August 2021, follow-up observations were scheduled to focus on the timing analysis. The observing time was dropped to monthly 5-min-long observations, but in exchange for the reduced observational time, observations were moved to being recorded with the dedicated MeerKAT pulsar timing user supplied equipment (PTUSE; Bailes et al. 2020) and APSUSE simultaneously. The advantages of PTUSE over APSUSE are GPU-based coherent de-dispersion, recording of full-Stokes information, real-time data folding and a finer sampling resolution of  $9 \mu\text{s}$  at the same radio band, thus enabling high precision pulsar timing.

To achieve a phase-coherent timing solution, we used `dracula2.py`<sup>15</sup>, an implementation of the `dracula` algorithm (Freire & Ridolfi 2018) with `python` and `TEMPO2`. `dracula` searches for timing solutions assuming different combinations of phase wraps in between the observations, finding the unique solution that accounts for all the rotations of the pulsar from start to end. For this, we used PTUSE pulsar archives from the newer observations and APSUSE archives produced from older search data with `dspsr`, both of them with 128 phase bins across the profile. Those archives were de-dispersed and scrunches in frequency with the `PSRCHIVE/pam` command, and three Times of Arrival (ToAs) were produced per observation with `PSRCHIVE/paas`, using a single, narrow von Mises function fitted with `PSRCHIVE/pat` as an initial timing template. `dracula.py` quickly converged into a single solution without the need of fitting for the sky position in the process, indicative of the quality of the SeeKAT multibeam localisation.

Given the massive, eccentric and compact nature of the system, we performed a semi-coherent orbital campaign with 14 observations at selected orbital phases from 2 March 2022 to 7 March 2022, accumulating a total of 11 h. Using our phase-connected timing model as a predictor, the orbital phases of the observations were chosen to cover features of the unabsorbed

signal of a potential Shapiro delay signal (Freire & Wex 2010) and the passage of periastron.

Together with an additional observation from 8 April 2022, the added PTUSE folded pulse profile with 1024 bins presents  $S/N \approx 200$ , while the added 512-bin APSUSE profile presents  $S/N \approx 170$ , the main difference in  $S/N$  coming from the improvement provided by coherent de-dispersion of PTUSE data. These two high  $S/N$  profiles were scrunches into four frequency sub-bands for frequency-resolved timing of the entire data set in order to account for DM evolution. Frequency-resolved analytical timing templates were produced with `paas` both for the PTUSE and APSUSE data sets. Subsequently and with `dspsr`, all PTUSE data were folded into 1024-bin, four-channel archives, and APSUSE data (when PTUSE data were not available) were folded into 512-bin, four-channel archives, with sub-integration times of at most 10 min in length in both cases. These files had been excised of RFI a priori at full time and frequency resolution using first the `clfd`<sup>16</sup> software (Morello et al. 2019) and then manually with `PSRCHIVE/paz` upon inspection of the remaining RFI. Finally, `pat` was used with the Fourier Phase Gradient algorithm (Taylor 1992) to generate 384 frequency-resolved ToAs.

Assuming GR to be the correct theory of gravity, we analysed the produced ToAs to constrain the pulsar mass ( $M_p$ ), the companion mass ( $M_c$ ) and the inclination angle of the orbit ( $i$ ) through two different methods. The first one is by measuring the theory-independent post-Keplerian (PK) parameters that arise from relativistic corrections of the orbital motion, and then deriving probability densities for  $M_p$ ,  $M_c$  and  $i$  on which we quote median values and  $1\sigma$  uncertainties from the two adjacent 34.1% percentiles (68.2% credible intervals). In this work, we measure the advance of periastron passage ( $\dot{\omega}$ ), the amplitude of the Einstein delay ( $\gamma_E$ ), and the orthometric parameters of Shapiro delay: the third-order orthometric amplitude ( $h_3$ ) and the orthometric ratio ( $\varsigma$ ) as defined in Freire & Wex (2010). In the first post-Newtonian approximation of GR,  $\dot{\omega}$  arises from the rotation of the Keplerian ellipse of the orbit in the direction of the orbital motion, depending on the total mass  $M_t = M_p + M_c$  as

$$\dot{\omega} = 3 \left( \frac{G}{c^3} \right)^{2/3} \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{M_t^{2/3}}{1 - e^2}, \quad (2)$$

while  $h_3$  and  $\varsigma$  parameterise the unabsorbed part of the Shapiro delay (Freire & Wex 2010). These describe the delay of the pulses propagating through the companion's gravitational field, depending on  $M_c$  and  $i$  as

$$h_3 = \frac{GM_c}{c^3} \left( \frac{1 - \cos i}{1 + \cos i} \right)^{3/2} \quad (3)$$

and


$$\varsigma = \left( \frac{1 - \cos i}{1 + \cos i} \right)^{1/2}, \quad (4)$$

for  $\cos i > 0$ , which is applicable to our study as we do not have any constrain on the angle of the ascending node ( $\Omega_A$ , see Sect. 7.2 for a discussion on this). It is worth noting that we use the  $h_3$  and  $\varsigma$  parametrisation introduced in Freire & Wex (2010) instead of the classic range ( $r$ ) and shape ( $s$ ) parameters used in other works (e.g. Kramer et al. 2021) because they better describe systems with low  $i$  and are less correlated with each

<sup>15</sup> <https://github.com/mcbernadich/CandyCracker/blob/main/dracula2.py>

<sup>16</sup> <https://github.com/v-morello/clfd>





**Fig. 2.** Corner plot showing the constraints on the inclination angle and component masses from PK parameters and the  $\chi^2$  mapping of DDGR solutions. *Central plots:* mass–mass diagrams portraying the DDH PK parameters constraints, with each color corresponding to a different parameter as indicated in the legend (solid lines: nominal value, dashed lines:  $1\sigma$  limits), and the  $1\sigma$  limits from DDGR  $\chi^2$  mapping in black. The shaded grey area on the right plot represents the region excluded by the mass function ( $i > 90^\circ$ ), while the shaded areas on the left plot represent the areas outside the prior for  $M_p$  (outside of  $0.6 < M_p < 1.5 M_\odot$ ). The explored regions of the  $M_p$  and  $M_c$  were decided based on the limits given by the mass function and  $M_t$  (from  $\dot{\omega}$ ). *Outer plots:* marginalised one-dimensional probability densities for  $M_p$ ,  $M_c$  and  $\cos i$  from DDGR  $\chi^2$  mapping, showcasing the median value (black solid line) and the 31.4%, 47.4% and 49.9% percentiles on both sides (shaded areas). The resulting mass constraints are consistent with a pair of NSs.

other. And finally, the Einstein delay is caused by the periodic modulation of the relativistic time dilation due to the pulsar’s changing orbital velocity and its motion across the companion’s gravitational field, its amplitude  $\gamma_E$  being modelled by

$$\gamma_E = \left(\frac{G}{c^3}\right)^{2/3} \left(\frac{P_b}{2\pi}\right)^{1/3} \frac{M_c(M_c + M_t)}{M_t^{4/3}} e. \quad (5)$$

These parameters are implemented in DDH, which is an extended version of the Damour-Deruelle pulsar timing model (DD; Damour & Deruelle 1986) that implements the orthometric parameters  $h_3$  and  $\zeta$  instead of  $r$  and  $s$ . The fit was done with TEMPO2/TempoNEST<sup>17</sup>, a multi-nested Bayesian sampling plugin to TEMPO2 (Lentati et al. 2014), in order to find stable values for both  $h_3$  and  $\zeta$  with realistic uncertainties.

The other method is to assume GR from the start with the DDGR model (Taylor & Weisberg 1989), which measures  $M_t$  and  $M_c$  directly to model the PK effects. For this measurement we implemented a common likelihood approach (first introduced by Splaver et al. 2002, see therein for more details) with chi2Map.py<sup>18</sup>, computing the likely constraints from the quality of the TEMPO2 fits of the DDGR model in a uniformly spaced grid on the  $M_t$ – $\cos i$  plane, which produces agnostic prior that follows the random distribution of  $i$  values of binary systems in the sky.

The resulting probability distributions, derived from multiplying the measured PK parameters from DDH and the  $\chi^2$  values from DDGR, were then marginalised into one-dimensional

probability densities for  $M_p$ ,  $M_c$  and  $\cos i$ , on which we quote median values and  $1\sigma$  uncertainties from the two adjacent 34.1% percentiles (68.2% credible intervals). The resulting TempoNEST fit of the DDH model and the chi2Map.py exploration of the DDGR  $\chi^2$  space are consistent with each other. The most significant PK parameter is  $\dot{\omega}$ , detected in DDH with  $\approx 900\sigma$  significance, yielding a highly precise measurement of the total system mass at

$$M_t^{\text{DDH}} = 2.586 \pm 0.005 M_\odot. \quad (6)$$

This is in excellent agreement with the

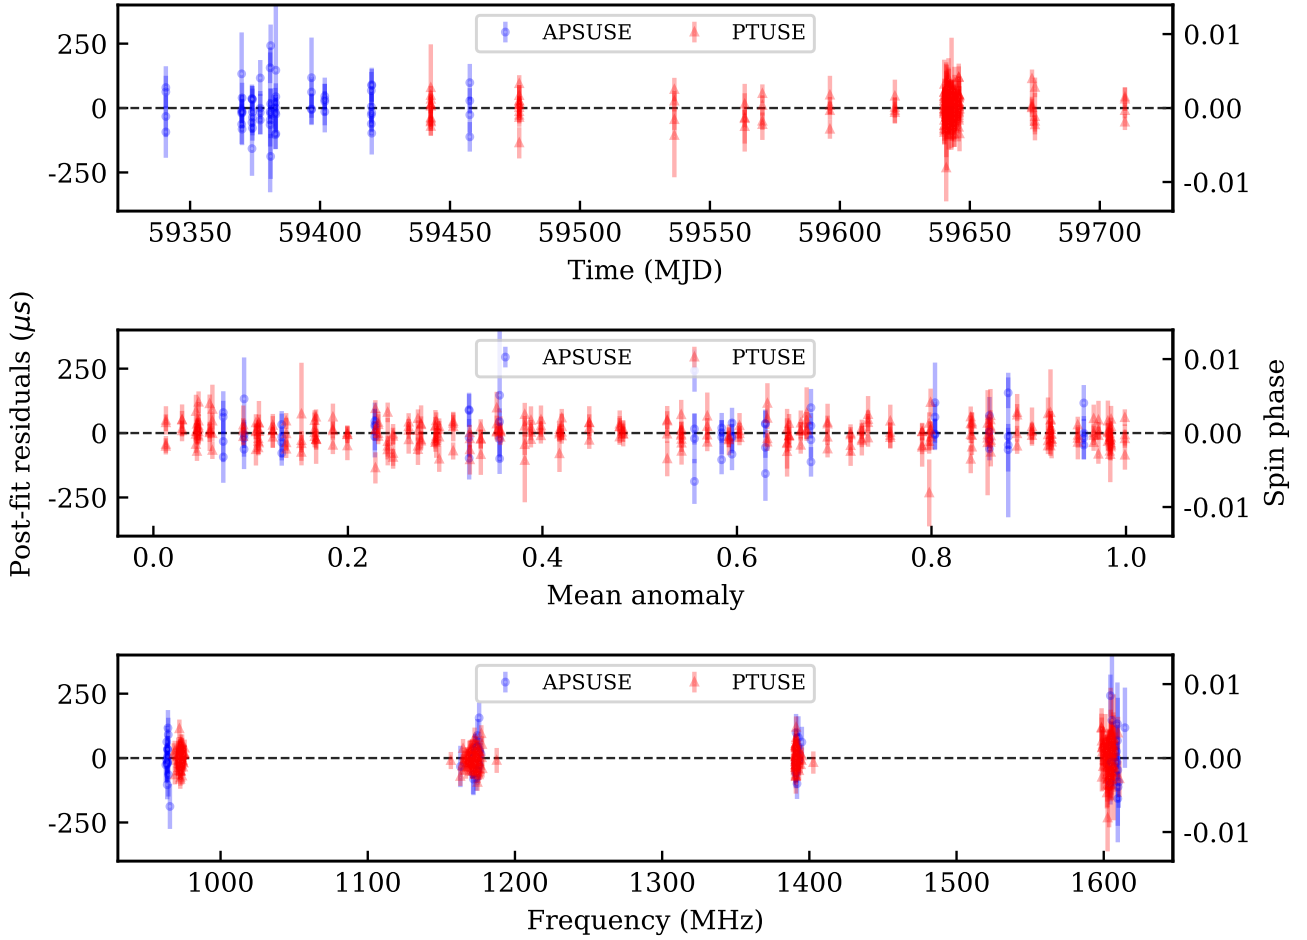
$$M_t^{\text{DDGR}} = 2.586 \pm 0.006 M_\odot \quad (7)$$

given by the direct DDGR TEMPO2 fit. We note that  $\zeta$  and  $\gamma_E$  are detected with low significance, and there is an important upper limit of Shapiro delay amplitude  $h_3$ . The derived  $1\sigma$  constraints from each DDH PK parameter and the  $\chi^2$  DDGR fits are overlaid in Fig. 2, with the DDGR contour tracing very strictly the limits imposed by  $\dot{\omega}$ , and being consistent with the loose limits imposed by  $\gamma_E$ ,  $h_3$  and  $\zeta$ . The most likely  $M_p$ ,  $M_c$  and  $i$  values derived from the DDH and DDGR models are in very good consistency with each other, with

$$\begin{aligned} M_p^{\text{DDH}} &= 1.24^{+0.12}_{-0.18} M_\odot, \\ M_c^{\text{DDH}} &= 1.36^{+0.18}_{-0.12} M_\odot, \text{ and} \\ i^{\text{DDH}} &= 59 \pm 9 \text{ deg} \end{aligned} \quad (8)$$

<sup>17</sup> <https://github.com/LindleyLentati/TempoNest>

<sup>18</sup> <https://github.com/mcbernadich/mass-diagrams>



**Fig. 3.** Timing residuals against observing time, mean anomaly, and radio frequency of the ToAs extracted from the TempoNEST DDH fit, presented in Table 2. No significant trends are seen in the data, which is indicative of the quality of the fit. The DDGR fit presents virtually the same residuals.

from multiplying probability distributions given by the PK parameter limits, and

$$\begin{aligned} M_p^{\text{DDGR}} &= 1.26^{+0.13}_{-0.25} M_\odot, \\ M_c^{\text{DDGR}} &= 1.32^{+0.25}_{-0.13} M_\odot, \text{ and} \\ i^{\text{DDGR}} &= 57 \pm 12 \text{ deg} \end{aligned} \quad (9)$$

from the  $\chi^2$  mapping of DDGR. Additionally, as a final consistency check, we also re-derived the likely values of  $\gamma_E$ ,  $h_3$  and  $\zeta$  by calculating their marginal one-dimensional probability densities from the DDGR  $\chi^2$  mapping experiment, on which we quote median values and  $1\sigma$  uncertainties from the two adjacent 34.1% percentiles, and attest that they are in good consistency with their DDH limits (Table 2).

We also attempted to fit any orbital period derivative  $\dot{P}_b$  or apparent change of the projected semi-major axis  $\dot{x}$ , but none of them yield significant limits with DDH. From the DDGR  $\chi^2$  mapping limits on  $M_p$ ,  $M_c$  and  $i$  and assuming GR, we predict them to have values of  $\dot{P}_b^{\text{GR}} = -1.225^{+0.026}_{-0.009} \times 10^{-13} \text{ s s}^{-1}$  and  $\dot{x}^{\text{GR}} = -6.37^{+0.14}_{-0.05} \times 10^{-18} \text{ ls s}^{-1}$ , but the current timing sensitivity is not high enough to yield a proper measurement. Nonetheless,  $\dot{P}_b$  is not likely to be detected without the addition of another decade of timing and it may be contaminated by effects introduced by the galactic acceleration field and proper motion in the sky, while  $\dot{x}$  is likely to be dominated by proper motion effects (Sect. 7.1, Table 5).

## 4. Properties of J1208–5936

### 4.1. General characteristics


Table 2 presents the timing parameters resulting from fits to the ToAs with the DDH and the DDGR models, and Fig. 3 shows the timing residuals. Both models are a good description of the data, with reduced  $\chi^2 \approx 0.95$ , and with spin, astrometric, Keplerian and PK parameters in very good consistency.

The results of the timing analysis are consistent with J1208–5936 being the first-born NS in a DNS system. It is a mildly recycled pulsar ( $P_0 = 28.71 \text{ ms}$ ,  $\dot{P} \lesssim 4 \times 10^{-20} \text{ s s}^{-1}$ ; Fig. 4) in an eccentric, compact orbit ( $e = 0.3480$ ,  $P_b = 0.6316 \text{ days}$ ) in a binary with both masses being consistent with those of NSs ( $M_t = 2.586(6) M_\odot$ ,  $M_p = 1.26^{+0.13}_{-0.25} M_\odot$ ,  $M_c = 1.32^{+0.25}_{-0.13} M_\odot$ ). Similar to most massive Galactic pulsar binaries, the system lies close to the Galactic plane, at a galactic latitude of  $b = 2.813^\circ$ . With  $\text{DM} \approx 344.4 \text{ cm}^{-3} \text{ pc}$ , and assuming the NE2001<sup>19</sup> (Cordes 2004) or the YMW16<sup>20</sup> (Yao et al. 2017) models of Galactic electron density, the corresponding DM-inferred distances from Earth are  $d \approx 8.2(1.6) \text{ kpc}$  or  $d \approx 8.5(1.7) \text{ kpc}$  respectively and with 20% uncertainties, placing it as possibly the furthest known Galactic DNS system.

Due to the low significance of the spin frequency derivative  $\dot{\nu}$ , both the characteristic age  $\tau_c$  and surface magnetic field

<sup>19</sup> <https://pypi.org/project/pyne2001/>

<sup>20</sup> <https://www.atnf.csiro.au/research/pulsar/ymw16/>



**Fig. 4.**  $P_0$ – $\dot{P}$  diagram showcasing all isolated and binary pulsars, and highlighting all known DNS systems and candidates (Table 1). The dashed lines indicate the characteristic age  $\tau_c$ , the surface magnetic field  $B_{\text{surf}}$  and the spin-down luminosity. J1208–5936 falls at the bottom of the mildly recycled population. All data points except J1208–5936 have been retrieved from the Australia Telescope National Facility (ATNF) website.

$B_{\text{surf}}$  of J1208–5936 are poorly constrained. Furthermore, the observed magnitude of  $\dot{\nu}$  is likely dominated by contributions from the Shklovskii effect and the Galactic acceleration field (Sect. 7.1, Table 5). However, the current constraint points towards J1208–5936 possibly having the lowest spin-down rate among all DNS systems, indicating a weak magnetic field.

With the constrained masses and orbital parameters, we numerically integrate Eqs. (5.6) and (5.7) in Peters (1964) to compute a merger time of  $\tau_m = 7.2 \pm 0.2$  Gyr, where the uncertainty arises from the ranges of the individual  $M_p$  and  $M_c$  values. Therefore, J1208–5936 joins the family of DNS systems merging within the Hubble time due to the orbital decay because of gravitational wave radiation (Table 1).

#### 4.2. Profile properties

The emission of J1208–5936 exhibits no detectable linear polarisation. Considering possible Faraday smearing, we searched the Rotation Measure (RM) space in the range of  $-20\,000 < \text{RM} < 20\,000$  rad m $^{-2}$  using a 1024-channel frequency-resolved profile from the PTUSE data of the semi-coherent orbital campaign. We used both `psrchive/rmfit` and `RMcalc`<sup>21</sup> but found no detection, which may be caused by the scattering of the pulse at low frequencies.

As seen in Fig. 5, a secondary component trails the brighter, primary component of the pulse at high frequencies with a relative separation between peaks of approximately 12.7 deg. This secondary component disappears in the lower half of the MeerKAT band, being absorbed into the scattered tail of the

main component, but becomes more prominent at high frequencies ( $>1700$  MHz). In addition to the MeerKAT observations, we also have eight fold-mode Parkes/Murriyang<sup>22</sup> observations totaling more than 15 h. They were recorded with the Ultra-Wideband (UWL) receiver, which covers the band between 704 MHz and 4032 MHz (Hobbs et al. 2020). We calibrated the bandpass with the `psrchive/pac` command and the standard candle observations provided by CSIRO<sup>23</sup>, cleaned them from RFI with `clfd`, and then took the pulse profile between 1715 and 2305 MHz by adding all the observations. Figure 6 shows the resulting pulse profile, where the secondary component is seen to gain prominence.

We find good evidence for scattering in the low-frequency band of MeerKAT observations ( $<1000$  MHz). We divide the MeerKAT band into 32 sub-bands and fit the profile in the 10 lowest frequency sub-bands with a Gaussian function convolved with an exponential


$$S_b = \int A \times \exp\left(\frac{(b' - b_0)^2}{2 \times \Delta b^2}\right) \times \exp\left(-\frac{b - b'}{\tau_s}\right) db', \quad (10)$$

where  $b$  stands for bins,  $b_0$  for the Gaussian centre,  $\Delta b$  for the standard deviation thereof, and  $\tau_s$  for the scattering time-scale. Table 3 shows the  $\tau_s$  value for each frequency, and Fig. 7 shows the fit at  $f = 877$  MHz as an example, where it is clear that the pulse shape is well described by a scattered Gaussian function.


<sup>22</sup> <https://www.csiro.au/en/about/facilities-collections/atnf/parkes-radio-telescope-murriyang>

<sup>23</sup> [https://www.parkes.atnf.csiro.au/observing/Calibration\\_and\\_Data\\_Processing\\_Files.html](https://www.parkes.atnf.csiro.au/observing/Calibration_and_Data_Processing_Files.html)

<sup>21</sup> <https://gitlab.mpifr-bonn.mpg.de/nporayko/RMcalc>



**Fig. 5.** De-dispersed ( $DM = 344.258 \text{ pc cm}^{-3}$ ) PTUSE profiles at the upper and lower half of the bands of the emission of J1208–5936, with full Stokes resolution at  $RM = 0 \text{ rad m}^{-2}$ , derived with `psrchive/psrplot`. No linearly polarised emission has been detected.



**Fig. 6.** De-dispersed ( $DM = 344.308 \text{ pc cm}^{-3}$ ) UWL profile from scrunching the emission of J1208–5936 between 1715 and 2305 MHz. Only total intensity data is available.

This results in a significant trend of  $\tau_s$  decreasing in frequency  $f$  in the shape of the power law


$$\tau_s = 693 \pm 12 \mu\text{s} (f/\text{GHz})^{-2.8 \pm 0.2}. \quad (11)$$

The scattering index  $2.8 \pm 0.2$  is lower than the 4.0 typically measured in single-component scattered pulsars (e.g. Oswald et al. 2021), which is likely indicative of a bias introduced by the double-component nature of J1208–5936. Nonetheless, our analysis provides robust evidence in favour of the presence of scattering in the pulse. Finally, from the

**Table 3.** Gaussian width  $\Delta b$  and scattering timescale  $\tau_s$ , with both units of bins (out of 1024) and ms.

$f$ (MHz)	$\Delta b$ (bins)	$\tau_s$ (bins)	$\tau_s$ (ms)
1110	$11.0 \pm 0.7$	$19.4 \pm 1.5$	$0.54 \pm 0.04$
1083	$10.0 \pm 0.6$	$20.8 \pm 1.3$	$0.58 \pm 0.04$
1057	$11.2 \pm 0.6$	$22.0 \pm 1.4$	$0.62 \pm 0.04$
1030	$12.5 \pm 0.6$	$21.4 \pm 1.4$	$0.60 \pm 0.04$
1003	$11.3 \pm 0.6$	$25.4 \pm 1.3$	$0.71 \pm 0.04$
976	$14.4 \pm 0.6$	$21.6 \pm 1.4$	$0.60 \pm 0.04$
955	$9.7 \pm 0.8$	$28.1 \pm 2.0$	$0.79 \pm 0.05$
922	$14.3 \pm 0.7$	$31.0 \pm 1.5$	$0.87 \pm 0.04$
896	$14.9 \pm 0.7$	$33.7 \pm 1.6$	$0.95 \pm 0.05$
877	$16.0 \pm 1.2$	$39.7 \pm 2.9$	$1.11 \pm 0.08$

**Notes.** Aside from a couple of outliers, there is a clear increase of  $\tau_s$  with decreasing  $f$ .




**Fig. 7.** Fit of Eq. (10) (red line) at the bottom-most sub-band of 877 MHz (cyan line). The fit is effective at getting the general shape of the pulse profile, with little structure left in the residuals (dark blue line). Table 3 lists the best scattering timescales and Gaussian widths for all fitted frequencies.

PTUSE 1024-channel profile, the best DM measurement from the 2 March 2022 to 8 March 2022 observational campaign is  $344.298 \pm 0.054 \text{ pc cm}^{-3}$  from the `pdmp` fit at full frequency resolution.


#### 4.3. Comparison with the known DNS population

J1208–5936 falls on the lower side of the expected  $P_0$  and  $P_b$  relationship, having a spin period below the average (Fig. 8), but it is still consistent with the rest of the DNS population. This relationship is explained through the less efficient recycling of the primary NS resulting from a delayed RLO accretion onset in longer-orbit progenitor NS–He star systems (Tauris et al. 2017).


A relationship between  $e$  and  $P_b$  has been postulated on a similar basis (Tauris et al. 2015, 2017). In this case, wide NS–He star progenitors undergo a reduced mass transfer due to the delayed RLO onset, which results in a less stripped He star and therefore in an increased mass-loss during the second supernova. Therefore, large orbital periods are associated with large eccentricities. Figure 9 shows that J1208–5936 is a high-eccentricity outlier amongst DNS systems with  $e < 0.5$



**Fig. 8.**  $P_0$ – $P_b$  diagram of the known recycled pulsars in DNS systems population. The lines represent linear regression fits to the data points in the log space, added to aid in the visualisation of the trend.



**Fig. 10.**  $e$ – $P_0$  diagram of the known recycled pulsars in DNS systems population. The spread of values in observed systems is larger than predicted in some simulations (Dewi et al. 2005), with J1208–5936 staying slightly above the expectation.



**Fig. 9.**  $e$ – $P_b$  diagram of the known pulsars in DNS systems population. The lines represent linear fits to the data points at  $e < 0.5$  and  $e > 0.5$ , with the extended fit including unconfirmed DNS systems.

with short orbital periods. Such outliers are to be expected due to the large spread in outcome eccentricities introduced by supernova kicks (Tauris et al. 2017), but it could also be a hint towards larger supernova kick (see Sect. 4.4).


Finally, we also look at the postulated  $e$ – $P_0$  relationship in Galactic DNS systems. This one goes in parallel with the  $e$ – $P_b$  relationship, and is explained from the  $P_0$ – $P_b$  relationship: pulsars in longer orbital periods undergo less efficient recycling, the He star undergoes greater mass loss during the supernova, and therefore high eccentricities should imply longer spin periods. This was first observed by Faulkner et al. (2005), and then theorised by Dewi et al. (2005) under the assumption of supernova kicks smaller than  $50 \text{ km s}^{-1}$ . Figure 10 compares the updated DNS population with the simulated results from Dewi et al. (2005), and it shows that J1208–5936 is once again at the high-eccentricity end but still lies within the distribution. As recently reported by Sengar et al. (2022), the simulations from Dewi et al. (2005) do not coincide with the observed high-eccentricity end of the Galactic DNS population, which would favour the different formation channel hypothesis brought forward in Andrews & Mandel (2019).

#### 4.4. Formation channel

We note in Fig. 9 that J1208–5936 has a particularly high orbital eccentricity compared to other DNS systems with  $e < 0.5$ . This could be indicative that J1208–5936 has been formed through a different channel than other systems in this group. We also pay attention to the division between  $e > 0.5$  and  $e < 0.5$  systems in the Galactic DNS populations in the  $e$ – $P_b$  space (Fig. 9). According to Andrews & Mandel (2019), the high-eccentricity population can be explained by larger supernova kicks from heavier He stars progenitors of the second-born NS. Tauris et al. (2015) also explores the possible evolutionary origin of this division. In most known DNS systems, the He star loses enough mass during binary interaction to stop the possibility of nuclear fusion in the ONeMg core. In these cases, the He star with a core mass  $\leq 1.43 M_\odot$  would implode through the electron-capture process instead of following shell burning until reaching the iron core collapse (Tauris et al. 2015). On the other hand, He stars that keep a core with  $\geq 1.43 M_\odot$  would be able to reach iron core collapse, possibly leaving a heavier NS behind with a larger supernova kick and with greater mass loss (Tauris et al. 2015). In this picture, most known DNS systems that do not get disrupted form through the electron-capture channel, constituting the  $e < 0.5$  population, while the DNS systems forming through the iron core collapse channel could create a heavier  $e > 0.5$  population. This picture is in principle corroborated by Fig. 11, where the high eccentricity systems are also shown to contain more massive second-born NSs, consistent with the idea of larger supernova kicks being associated with larger masses.

It is tempting to see J1208–5936 as a system bridging the gap between these two postulated populations. While its current eccentricity is consistent with the tail of the eccentricity distribution with a second supernova kick of  $50 \text{ km s}^{-1}$  and a progenitor He stars mass of  $3 M_\odot$  like other low-eccentricity DNS systems (see Fig. 10 in Tauris et al. 2017), it is also plausible that its companion has formed through the iron core collapse channel, or that it has at least suffered a stronger supernova kick than in other  $e < 0.5$  DNS systems formed through the electron capture process owing to a more massive He star progenitor. However, in Fig. 11 it is shown that the current uncertainties on  $M_c$  are too large to discriminate whether J1208–5936 lies within any of these two populations. Further timing in the following years will





**Fig. 11.**  $M_c$ - $e$  diagram of second-born NSs with measured masses (quoting the mass derived from the  $\chi^2$  mapping of DDGR solutions for J1208–5936). The line represents linear fits to the data points excluding the companion of J1208–5936, which does not yet have a sufficiently constrained mass. For PSR J1906–0746 we quote the pulsar mass instead of the companion mass because it is believed to be the second-born, un-recycled NS.

constrain the mass of the companion of J1208–5936, providing clarification on its formation channel.

## 5. Search for pulsations from the companion

We also searched for pulsations from the companion of J1208–5936. Detecting them would not only increase the sample of known young pulsars in DNS systems, but it would also provide much more precise mass measurements and even gravity tests in the future through the inclusion of the mass ratio as an extra constraint on top of the PK parameters, such as in the case of PSR J0737–3039A/B (Lyne et al. 2004).

We performed deep searches on de-modulated APSUSE data from 30-min and 60-min long observations from the Shapiro delay semi-coherent orbital campaign. We began by cleaning the data with `presto/rfifind`, de-dispersing it and integrating it into a barycentric time series at  $DM = 344.43 \text{ pc cm}^{-3}$  with `presto/prepdata`. This DM value is slightly offset from our current best estimate presented in Sect. 4.2, owing to the fact that we are dealing with APSUSE data instead of PTUSE. Nonetheless, small discrepancies in DM do not affect the search for pulsations at  $P > 10 \text{ ms}$ . We de-modulated the time series with python software `pysolator`<sup>24</sup>, which undoes the orbital motion of the companion assuming a chosen mass ratio, allowing us to search for pulsations in its rest frame of the companion. For this, we assumed several system mass ratios chosen from 27 different inclination angles between  $i = 45^\circ$  and  $i = 85^\circ$  uniformly spaced in  $\sin i$ , assuming the total mass of  $M_t = 2.586 M_\odot$  from the measurement of  $\dot{\omega}$  (Sect. 3, Table 2). The spacing in  $\sin i$  was chosen on the basis of introducing a maximum line-of-sight accelerations discrepancy of at most  $1 \text{ m s}^{-2}$  between the different de-modulations based on different  $\sin i$  values.

Subsequently, two different searches were performed on this data set. Firstly, we computed the FFT of each de-modulated time series with `presto/realfreq` and then de-reddened them with `presto/rednoise`. The de-reddened FFT spectra were

searched with `presto/accelsearch` with  $z_{\text{max}} = 2$ . For a 30-min-long and a 60-min-long observations, this corresponds to acceleration ranges of  $\pm 1.85 \text{ m s}^{-2}$  and  $\pm 0.46 \text{ m s}^{-2}$  for 10-ms pulsar, with increasing ranges for increasing spin period. All candidates were then refolded with `presto/prepfold`.

The second search implemented a fast Folding algorithm (FFA) instead, a phase-coherent periodicity search technique on evenly sampled time series (Staelin 1969). The FFA has the advantage of being more sensitive to long-period, narrow signals than the FFT (e.g. Cameron et al. 2017). Additionally, since it does not compute a Fourier Transform, de-reddening is typically performed with a running-median filter on the time series, which is less harmful towards real low-frequency signals than removing power from Fourier bins (Singh et al. 2022). Therefore, the FFA is well-suited for searching for a non-recycled second-born pulsar in the properly de-modulated time series. We restricted our FFA search to periods above 100 ms, on the basis that shorter periods would have easily been detected by the PRESTO FFT search. We used the RIPTIDE<sup>25</sup> software (Morello et al. 2019) to perform a non-accelerated FFA search on the de-modulated time series and fold the resulting candidates. For that we implemented the following steps with our `presto` wrapper pipelines `demodulate-search`<sup>26</sup>.

No significant pulsations from the companion were found. This suggests that the companion has either crossed its pulsar death line, that its radio emission is beamed away from Earth, or that the pulses are too faint to be detected with the available data and telescope sensitivity.

## 6. Implications for NS merger rate

The discovery of J1208–5936 and, more significantly, the increase in explored depth of the southern Galactic plane provided by the sensitivity of the MMGPS-L survey are a great opportunity to update the prediction of the observed NS coalescence events with gravitational wave observatories such as LIGO and Virgo. This is possible assuming that the observed population of pulsars in DNS systems is representative of the broad Galactic DNS population, with each new survey and discovery providing a more accurate sampling of it.

Following the methodology established in Kim et al. (2003) and also implemented in Kim et al. (2010, 2015), Pol et al. (2019, 2020), Grunthal et al. (2021), we use `PsrPopPy2`<sup>27</sup> to populate the Galactic field with DNS systems and simulate blind surveys on the sky. In this case, we simulate the known DNS systems merging within the Hubble time. For each known DNS system  $j$ , we seek to find the proportionality constant  $\alpha_j$  in

$$\lambda_j = \alpha_j N_{\text{tot},j}, \quad (12)$$

where  $\lambda_j$  is the number of discovered systems and  $N_{\text{tot},j}$  is the number of simulated systems with similar characteristics each. This  $\alpha_j$  value is then used to compute a probability distribution for the merger rate  $\mathcal{R}_j$

$$P(\mathcal{R}_j) = \left( \frac{\alpha_j \tau_{\text{life},j}}{f_{b,j}} \right)^2 \mathcal{R}_j \times \exp \left( - \frac{\alpha_j \tau_{\text{life},j}}{f_{b,j}} \mathcal{R}_j \right), \quad (13)$$

where  $\tau_{\text{life},j} = \tau_{\text{age}} + \tau_{\text{obs}}$  is the lifetime of the system, composed by its age and future observable time, and  $f_b$  is beaming factor of the pulsar (inverse of the beam sweep fraction). Typically,

<sup>25</sup> <https://github.com/v-morello/riptide>

<sup>26</sup> <https://github.com/mcbernadich/demodulate-search>

<sup>27</sup> <https://github.com/devanshkv/PsrPopPy2>

<sup>24</sup> <https://github.com/alex88ridolfi/pysolator>




$\tau_{\text{age}}$  is defined either by the characteristic age ( $\tau_c$ ) of the pulsar or the time after it has exited a recycling fiducial line, while  $\tau_{\text{obs}}$  consists either on the remaining time until the pulsar crosses the death-line or the merger time ( $\tau_m$ ). Finally, the rate of NS merger events in the Milky Way ( $\mathcal{R}_{\text{MW}}$ ) is then the sum of all the  $\mathcal{R}_j$  computed for all DNS systems merging within the Hubble time, which for the probability distribution translates to the convolution of all probability distributions.

The latest update on the NS merger rate prediction (Grunthal et al. 2021) includes the following blind pulsar surveys: the Pulsar Arecibo L-band Feed Array survey (PALFA; Cordes et al. 2006), the Low-latitude High Time-Resolution Universe pulsar survey (HTRU-Low; Keith et al. 2010), the Parkes High-latitude pulsar survey (Burgay et al. 2006), the Parkes Multibeam Survey (Manchester et al. 2001), the PSR B1534+12 discovery survey with Arecibo (Wolszczan 1991) and the Green Bank North Celestial Cap survey (GBNCC; Stovall et al. 2014). These surveys used to provide a realistic coverage of the explored pulsar sky, and they all have in common that they were performed with single-dish telescopes and therefore single-beam pointings on the sky, which are straightforward to model in PsrPopPy2 as circular beams with a Gaussian sensitivity function. Now we update this picture with the MMGPS-L, the most sensitive survey in the southern Galactic plane. Its interferometric nature requires an extra layer of complexity in the simulation of pointings in the sky, which we implement in PsrPopPy2 with the addition of an extra degradation factor for coherent beams (see Appendix B for details).

The biggest uncertainty in accounting for J1208–5936 in the simulations is the computation of  $\tau_{\text{life}, \text{J1208}}$ . Thanks to the preliminary mass estimates presented in Sect. 3, we compute a merger time of  $\tau_m = 7.2 \pm 0.2$  Gyr for J1208–5936. However, unlike with the rest of pulsars merging within the Hubble time, there is little information on the characteristic age  $\tau_c$  of J1208–5936 due to the poorly constrained  $\dot{P}$ . Despite this, a reasonable assumption for  $\tau_{\text{age}}$  can be made without sacrificing accuracy. Due to its large merger time, J1208–5936 is already the pulsar with the smallest contribution to the estimated galactic merger rate, and therefore we do not expect the choice to have a significant impact when all contributions are added. Given its small  $\dot{P}$ , it is likely a pulsar comparable to PSR J1913+1102 (Table 1). We therefore pick a realistic recycling age of 2.5 Gyr. For the beaming factor, we choose  $f_{b, \text{J1208}} = 4.59$ , consistent with the average of pulsars in DNS systems with actually measured  $f_{b, i}$  values (Pol et al. 2019), and model the pulse duty fraction of  $\delta_{\text{J1208}} = 8\%$  from the part of the profile that was visible above the noise during the discovery. Finally, we compute its Doppler degeneration factor  $0 \leq \gamma_{2m} \leq 1$  in all surveys, assuming acceleration search for each of them (see Appendix B for more details).

For consistency with Pol et al. (2019, 2020), Grunthal et al. (2021), we simulate pulsars with a mean luminosity of  $\langle \log_{10}(L/[\text{mJy kpc}^2]) \rangle = -1.1$  with a standard deviation of  $\sigma_{\log_{10} L} = 0.9$ , and a mean spectral index of  $\langle \Gamma \rangle = 1.4$  with a deviation of  $\sigma_{\Gamma} = 1$ . The height scale of the simulated population is set to  $z_0 = 0.33$  kpc, with pulsars following a density distribution above and below the Galactic plane of  $f(z) = \exp(-z/z_0)$ . For PSR J1906+0746, we use the beam shape modelling implemented in Sect. 4 of Grunthal et al. (2021). From each DNS system  $j$ , we start with simulating only  $N_{\text{tot}, j} = 100$  pulsars on the sky and simulate the surveys to see the number of discoveries  $\lambda_j$ . Then, on each step, the number of pulsars is increased by  $\Delta N_{\text{tot}, j} = 100$  and the surveys are repeated until we reach  $N_{\text{tot}, j} = 4000$ . Within each step, the population is simulated 100 different times, and in each simulation, the surveys are per-



**Fig. 12.** Updated probabilities distribution of the merger rates after the inclusion of J1208–5936 and the MMGPS-L survey. *Top:* individual contributions from each DNS population, derived from the values in Table 4 and Eq. (13). The 90% credible intervals are quoted in Table 4. *Bottom:* galactic merger rate probability distribution from the convolution of individual DNS distributions, with a highlighted 90% credible interval.

formed 100 different times, producing a loop with 10 000 iterations. All of this is done to get an averaged measurement of  $\alpha_j$  from Eq. (12).

The updated contributions to the Galactic merger rate, along with the parameters used in the simulations and computations, are presented in Fig. 12 and Table 4. PSR J1906+0746 remains the most impactful contribution due to its short lifetime, while the contribution of J1208–5936 is the smallest due to its large lifetime. The added Galactic merger rate results in

$$\mathcal{R}_{\text{MW}}^{\text{new}} = 25_{-9}^{+19} \text{ Myr}^{-1} \quad (14)$$


with a 90% confidence interval. This result is shifted downwards with respect to the  $\mathcal{R}_{\text{MW}}^{2021} = 32_{-9}^{+19} \text{ Myr}^{-1}$  presented by Grunthal et al. (2021), owing to the inclusion of the MMGPS-L survey, because despite it being the most sensitive blind pulsar survey on the southern Galactic plane, it has only added one DNS system merging within Hubble time, therefore implying a reduction of the expected number of unseen merging DNS systems in the Milky Way from the lack of new detections.

We pay attention to the increasingly constrained NS merger rates in the Milky Way based on the observation of Galactic binaries, as well as the decreasing trend of the estimated NS merger rates in the Milky Way based on the observation of Galactic binaries. As shown in Fig. 13, the discovery of PSR J0737–3039 with its relatively short merger time brought

**Table 4.** Pulsars used in the NS merger rate computation simulation, along with their used lifetime, beaming fraction, pulse duty ( $\delta$ , used only during the PsrPopPy2 runs), and the derived  $\alpha_i$  and individual merger rate contributions.

PSR <sub><i>j</i></sub>	$\tau_{\text{life},j}$ (Myr)	$f_{b,j}$	$\delta_j$ (%)	$\alpha_j$	$\mathcal{R}_j$ (Myr <sup>-1</sup> )
J1946+2052	293	4.59 <sup>(c)</sup>	6	0.00609	2.6 <sup>+11.7</sup> <sub>-1.88</sub>
J1757-1854	235	4.59 <sup>(c)</sup>	6	0.00916	2.15 <sup>+9.67</sup> <sub>-1.49</sub>
J0737-3039A	244	2.0 <sup>(a)</sup>	27	0.00429	1.92 <sup>+8.7</sup> <sub>-1.41</sub>
B1913+16	377	5.7 <sup>(a)</sup>	16.9	0.00914	1.66 <sup>+7.55</sup> <sub>-1.15</sub>
J1906+0746	60	1.0 <sup>(b)</sup>	1	0.00368	4.49 <sup>+20.71</sup> <sub>-3.26</sub>
J1913+1102	3125	4.59 <sup>(c)</sup>	6	0.01010	0.14 <sup>+0.67</sup> <sub>-0.09</sub>
J0509+3801	729	4.59 <sup>(c)</sup>	18	0.01008	0.62 <sup>+2.86</sup> <sub>-0.83</sub>
J1756-2251	2086	4.59 <sup>(c)</sup>	3	0.01290	0.17 <sup>+0.77</sup> <sub>-0.11</sub>
B1534+12	2908	6.0 <sup>(a)</sup>	4	0.01415	0.14 <sup>+0.67</sup> <sub>-0.09</sub>
J1208-5936	9700	4.59 <sup>(c)</sup>	8	0.01157	0.04 <sup>+0.19</sup> <sub>-0.03</sub>

**Notes.** For PSR J1906+0746, the lifetime is shorter than the merger time due to it crossing the pulsar deadline in ~60 Myr. <sup>(a)</sup>Beaming fractions obtained from Table 2 in Kim et al. (2015). <sup>(b)</sup>Beaming fraction set to 1.0 as we implement explicit modelling of its variation in time. See Grunthal et al. (2021) for more details. <sup>(c)</sup>Beaming factor taken from the average of J0737-3039A, B1913+16 and B1534+12.



**Fig. 13.** Evolution of estimated NS merger rates in the Milky Way based on observable Galactic binaries during the last two decades, quoting the 90% or 95% confidence intervals (Kim et al. 2003, 2010, 2015; Kalogera et al. 2004; Pol et al. 2019, 2020; Grunthal et al. 2021).

forward a drastic increase of the predicted Galactic merger rate (Burgay et al. 2003; Kalogera et al. 2004; Kim et al. 2010) with respect to previous estimates based only on PSR B1913+16 and PSR B1534+12 (e.g., Kim et al. 2003). However, further discoveries have not increased the estimated rate in recent years, as better modelling of beam shapes and sky coverage have allowed for an increase in the accuracy of the estimates. Pol et al. (2019) presented  $\mathcal{R}_{\text{MW}}^{2019} = 42^{+30}_{-14}$  Myr<sup>-1</sup> at the 90% confidence with the inclusion of at-the-time newly discovered systems such as of PSR J1757-1854 and the highly relativistic PSR J1946+2052, the most constrained estimate up to that date owing to the increased coverage of the sky. Then, in Pol et al. (2020) the sensitivity of the newly included the GBNCC survey (Stovall et al. 2014) outweighed the inclusion of the newly discovered relativistic J0509+3801 DNS system, decreasing the estimated number of unseen binaries in the sky and reducing the N merger rate

estimate to  $\mathcal{R}_{\text{MW}}^{2020} = 37^{+24}_{-11}$  Myr<sup>-1</sup>, with more constrained uncertainties. Grunthal et al. (2021) implemented a modified beam shape in their simulations, leading to another reduction of the total value and uncertainties. Our estimate presents a continuation of this trend, being our estimate once again reduced in comparison to the most recent estimates owing to the sensitivity of the MMPGS-L, the most important newly contributing factor. Therefore, our work confirms that NS merger rate estimates based on known electromagnetic binaries are converging into more constrained, lower values as pulsar surveys scout the sky at larger depths.

We transform this rate into a local cosmic merger rate density and a prediction for the LIGO event detection rates by assuming that the number of DNS systems merging within the Hubble time in a galaxy is proportional to its total *B*-band luminosity, as it is a tracer of the star-formation rate (Kopparapu et al. 2008). This results in a local merger rate density of

$$\mathcal{R}_{\text{local}}^{\text{new}} = 293^{+222}_{-103} \text{ Gpc}^{-3} \text{ yr}^{-1}, \quad (15)$$

which results on an upper limit of  $\mathcal{R}_{\text{local}} \leq 515 \text{ Gpc}^{-3} \text{ yr}^{-1}$ . Assuming that LIGO has sensitivity to a range distance of  $D_r = 130$  Mpc during the O3 run<sup>28</sup> (see Chen et al. 2021b for a definition of  $D_r$ ), we predict a LIGO NS merger detection rate of

$$\mathcal{R}_{\text{LIGO,O3}}^{\text{new}} = 2.76^{+2.10}_{-0.97} \text{ yr}^{-1}. \quad (16)$$

Based on the observed events between April 1st, 2019 and October 1st, 2019 catalogued in the Gravitational-Wave Transient Catalogue-2.1, a NS merger rate density of  $\mathcal{R}_{\text{local}}^{\text{GRTC-2.1}} = 256^{+510}_{-237} \text{ Gpc}^{-3} \text{ yr}^{-1}$  is computed, which for detections translates into  $\mathcal{R}_{\text{LIGO,O3}}^{\text{GRTC-2.1}} = 2.36^{+4.69}_{-2.18} \text{ yr}^{-1}$  (Abbott et al. 2021), in good consistency with our estimate. This shows that the rates computed from electromagnetic binaries are in good consistency with gravitational wave observatories (see also Pol et al. 2020; Grunthal et al. 2021). With this in mind, we assume  $D_r = 175$  Mpc to make a prediction of the rate NS merger events seen by the LIGO-Virgo-KAGRA O4 run<sup>28</sup>, resulting in

$$\mathcal{R}_{\text{LIGO,O4}}^{\text{new}} = 6.73^{+5.12}_{-2.37} \text{ yr}^{-1}, \quad (17)$$

which implies a prediction of the detection of  $10^{+8}_{-4}$  events during the 18 months of the O4 run<sup>28</sup>, or at most 18 events within 90% credible intervals.

## 7. Biases and future prospects

### 7.1. Proper motion and galactic field

There are three main contributions that could be biasing our mass constraints. Firstly, we consider the Shklovskii and the Galactic acceleration field effects (Shklovskii 1970; Damour & Taylor 1991). These two combine into an apparent evolution of any periodicity or length  $L = \{P_0, P_b, x\}$ , expressed with the equation

$$\frac{\dot{L}}{L} = -\left(\frac{\dot{D}}{D}\right)_{\text{Shkl}} - \left(\frac{\dot{D}}{D}\right)_{\text{Gal}} = \frac{1}{c} \left( \frac{V_t^2}{d} + K_0(a_{\text{PSR}} - a_{\text{SSB}}) \right), \quad (18)$$

where  $D$  is the Doppler factor,  $K_0$  is the unit vector from the Solar System barycentre (SSB) to the pulsar system,  $a_{\text{PSR}}$  and  $a_{\text{SSB}}$  are the Galactic acceleration vectors at the pulsar system and at the SSB,  $V_t$  is the magnitude of the transverse velocity of the system with respect to the Sun, and  $d$  is the distance to

<sup>28</sup> <https://observing.docs.ligo.org/plan/>

**Table 5.** Possible biases onto spin a PK parameters measurements introduced by the Shklovskii effect and the acceleration in the Galactic field (Eq. (18)) and from proper motion (Eqs. (19) and (20)), assuming  $d = 8.2$  kpc and  $|\mu_t| = 6$  mas s<sup>-1</sup> (see text in Sect. 5).

Parameter	Shklovskii	Galactic	Proper motion
$\dot{P}$ (s s <sup>-1</sup> )	$2 \times 10^{-20}$	$-2 \times 10^{-20}$	...
$\dot{P}_b$ (s s <sup>-1</sup> )	$4 \times 10^{-14}$	$-4 \times 10^{-14}$	...
$\dot{x}$ (ls s <sup>-1</sup> )	$3 \times 10^{-18}$	$-3 \times 10^{-18}$	$3 \times 10^{-15}$
$\dot{\omega}$ (deg yr <sup>-1</sup> )	...	...	$2 \times 10^{-6}$

the source. Secondly, we also consider the effects on  $\dot{\omega}$  and  $\dot{x}$  introduced by the proper motion (PM) vector  $\mu_t = (\mu_{RA}, \mu_{Dec})$ , along with the angle of the ascending node  $\Omega_a$  and inclination angle  $i$ , expressed in Kopeikin (1996) as

$$\dot{\omega}^{\text{PM}} = 2.78 \times 10^{-7} \csc i (\mu_{RA} \cos \Omega_a + \mu_{Dec} \sin \Omega_a) \quad (19)$$

and


$$\left(\frac{\dot{x}}{x}\right)^{\text{PM}} = 1.54 \times 10^{-16} \cot i (-\mu_{RA} \sin \Omega_a + \mu_{Dec} \cos \Omega_a). \quad (20)$$

In our case, since we do not have a measurement of the proper motion or  $\Omega_a$ , these effects are not yet quantifiable, but we can nonetheless estimate the likely magnitude of the contributions. From the Galactic coordinates and DM distance of J1208–5936, we take the Galactic rotation velocity curve from Sofue (2020) and predict  $V_t \approx 240$  km s<sup>-1</sup> and  $|\mu_t| \approx 6$  mas s<sup>-1</sup>. This assumption is taken based on the small magnitude of the supernova kicks introduced during the formation of the system, as otherwise the binary would have been disrupted (Tauris et al. 2017), making the Galactic flow the more dominant component of the acceleration. For the Galactic acceleration field, we use the Galactic mass distribution model presented in McMillan (2017) and extract  $a_{\text{PSR}}$  and  $a_{\text{SSB}}$ .

Our estimation of the magnitude of these contributions are listed in Table 5. We expect the Shklovskii and Galactic contributions to have similar orders of magnitude but of opposite sign, likely cancelling each other out. However, the exact values will only be computable when precise measurements of PM and  $d$  are available, and it may well be possible that either of them is dominant over the other. Comparing it with the values listed in Table 2 ( $\dot{P} = 2.6 \pm 1.0$  s s<sup>-1</sup>), we see that the measured  $\dot{P}$  may also be heavily contaminated by either contribution. Nonetheless, it is also clear that our mass constraints should not be biased by the proper motion contribution to  $\dot{\omega}$ , as the maximum expected contribution is still three orders of magnitude smaller than our current measurement uncertainty, while the effect on  $\dot{P}_b$  will eventually be able to constrain the distance through comparison with the prediction given by GR (Table 2,  $\dot{P}_b^{\text{GR}} = 1.225^{+0.026}_{-0.009}$  s s<sup>-1</sup>).

## 7.2. Future prospects for timing

We estimate the prospects for improved mass measurements of J1208–5936 by simulating new ToAs with the TEMPO2/fake plug-in. We emulate the observing cadence outside of the Shapiro delay campaign (4 ToAs/month), our current timing precision (37  $\mu$ s) and add an assumed orbital decay of  $\dot{P}_b = -1.247533 \times 10^{-13}$  s s<sup>-1</sup> from the crossing of the current values of  $\dot{\omega}$  and  $\gamma_E$ . We also keep  $h_3$  and  $\zeta$  fixed to their current values. As seen in Fig. 2,  $\dot{\omega}$  is already providing a very tight constraint,



**Fig. 14.** Fractional uncertainty evolution with timing baseline of the measured  $\dot{\omega}$ ,  $\gamma_E$  and  $\dot{P}_b$  PK parameters from the TEMPO2/fake simulation (solid lines: uncertainty/measurement, cyan dashed line: uncertainty/simulation input value, grey dashed lines: significance thresholds). For  $\dot{P}_b$ , the dashed line is taken as a reference due to the low significance of the measurement.  $\dot{\omega}$  and  $\gamma_E$  scale with  $T^{-3/2}$ , while  $\dot{P}_b$  scales with  $T^{-5/2}$ .

so the evolution of the uncertainties on  $M_p$ ,  $M_c$  and  $i$  is expected to be dependant on the evolution of the  $\gamma_E$  measurement. From Fig. 14, we expect the precision of  $\gamma_E$  to surpass a significance of  $10\sigma$  at  $T = 3$  yr, decreasing the individual uncertainties of  $M_p$  and  $M_c$  to  $\pm 0.1 M_\odot$ , and that of  $i$  to  $\pm 5$  deg. Beyond this point, the uncertainties on PK parameters are likely to become dominated by correlated spin and DM noise, but assuming only white noise or a good modelling with Bayesian sampling of correlated noise as offered by TempoNest (Lentati et al. 2014), the first reliable measurement of  $\dot{P}_b$  at  $3\sigma$  is reached at  $T = 11$  yr, while  $5\sigma$  is reached at  $T = 14$  yr, and  $10\sigma$  at  $T = 19$  yr. At this time, the effects of Shklovskii and the Galactic acceleration field will already have an impact on the measurement, which would only be accountable from a good measurement of the PM. Depending on the magnitude of the PM, an independent estimate of the distance may be performed at that time. The uncertainties on the masses at those times are  $\pm 0.01$ ,  $\pm 0.008$  and  $\pm 0.005 M_\odot$ , and for  $i$  they are  $\pm 1$  deg,  $\pm 0.5$  deg and  $\pm 0.4$  deg. The parameters  $\dot{\omega}$  and  $\gamma_E$  always remain the most constraining ones, with  $\gamma_E$  dominating the uncertainty and  $\dot{P}_b$  providing an independent estimate of true distance by forcing it to be consistent with GR and the measured PM, which in turn will help constrain the true value of  $\dot{P}$ .

It is unlikely that the geometry of the system will be well determined in the future. Even at  $T = 20$  yr, the proper motion contribution to  $\dot{\omega}$  (Eq. (19), Table 5) is expected to have a maximum possible value which is one order of magnitude below the uncertainty with our assumed PM magnitude. The value of  $\dot{x}$  is also unlikely to be constrained in the following decades, as it is highly degenerate with  $\gamma_E$ , resulting in a spurious measurement until the precession angle of  $\omega$  is large enough to break the degeneracy (for a discussion of this phenomenon, the reader is encouraged to consult Ridolfi et al. 2019). Furthermore, even a detection of the PM contribution to  $\dot{x}$  (Eq. (20), Table 5), which is expected to be three orders of magnitude larger than the GR-predicted one (Table 2), will not break the sign degeneracy of  $i$  owing to the large distance to the system.



### 7.3. Future profile changes

The two components present in the profile (Sect. 4.2, Fig. 5) are likely to change both in relative amplitude and in phase separation due to the geodetic precession of the pulsar spin axis around the orbital angular momentum vector during the following years (Damour & Ruffini 1974). In GR and assuming  $M_p = M_c$ , the rate of the spin-orbit coupling-induced precession is proportional to  $\dot{\omega}$  (Barker & O’Connell 1975) as in

$$\Omega_g = \frac{7}{24} \dot{\omega}, \quad (21)$$

which in the case of J1208–5936 results in an expected value of  $\Omega_g \approx 0.268 \text{ deg yr}^{-1}$ . Assuming a cone-shaped pulse from which we are seeing the maximal cross-section, a spin axis perpendicular to the angular momentum and the pulse being emitted from the equator, this would give us a minimum timescale of  $\approx 24 \text{ yr}$  before the pulses precess out of view. However, a different spin-orbit orientation would only extend this range upwards and the primary component will most likely be visible for a longer period of time. A detailed study of profile changes during the following years will enable a more accurate prediction from the detection of any change in the profile or lack thereof.

## 8. Conclusions

In this work we report the discovery and follow-up study of the MMGPS-L discovery J1208–5936. Spinning at 28.71 ms and in close orbit with another NS, this is in the tenth known Galactic DNS merging within the Hubble time. We have constrained the masses and inclination angle to  $M_p = 1.26^{+0.13}_{-0.25} M_\odot$ ,  $M_c = 1.32^{+0.25}_{-0.13} M_\odot$  and  $i = 57 \pm 12 \text{ deg}$  from the mapping  $\chi^2$  mapping of DDGR solutions, with the tightest constraint coming from the  $900\sigma$  measurement of the periastron advance  $\dot{\omega}$ . The measurement of  $\dot{P} < 4 \times 10^{-4} \text{ s s}^{-1}$  is consistent with a mildly recycled pulsar and makes J1208–5936 the pulsar in a DNS with the smallest period derivative. However, the value is likely to be biased by the Shklovskii and Galactic acceleration field contamination. Its high eccentricity is still consistent with the tail of eccentricity distribution arising from a  $50 \text{ km s}^{-1}$  supernova kick during the formation of the companion NS (Tauris et al. 2017), but at the same time it could be indicative of a larger supernova kick caused by a massive He star. A more precise measurement of  $M_c$  in the future may clarify which is the case and help confirm the idea of two main formation channels for Galactic DNS systems depending on the supernova type. We have been unable to detect polarised emission from J1208–5936, but we have observed a faint, secondary leading component to the main pulse that becomes more prominent at high frequencies, and significantly scattered at low frequencies. We have also found robust evidence for scattering, with a scattering index of  $2.8 \pm 0.2$ , even though a better modelling may be able to provide a better measurement of the scattering index.

The merger time of  $\tau_m = 7.2 \pm 0.2 \text{ Gyr}$  adds J1208–5936 to the family of DNS systems merging within the Hubble time, therefore making it a progenitor of NS merger events seen by gravitational-wave observatories such as the landmark GW170817 event (Abbott et al. 2017a), making it relevant for predictions of the cosmic NS merger rate based on Galactic binaries. The performance of the MMGPS-L, the most sensitive survey in the southern sky, encouraged us to revisit these predictions. The end result provides an updated merger rate of  $\mathcal{R}_{\text{MW}}^{\text{new}} = 25^{+19}_{-9} \text{ Myr}^{-1}$  and local cosmic merger rate of  $\mathcal{R}_{\text{local}}^{\text{new}} = 293^{+222}_{-103} \text{ Gpc}^{-3} \text{ yr}^{-1}$  within a 90% confidence interval, smaller

than the limits provided by previous studies on Galactic DNS systems (Grunthall et al. 2021; Pol et al. 2019, 2020) owing to the fact that despite the high sensitivity of the MMGPS-L only one new system merging within the Hubble time has been discovered, reducing the expected number of unseen DNS systems. This continues the trend of more constrained, decreasing estimates over time as the depth of pulsar surveys increase and the modelling of pulsar beam shapes improves. The resulting prediction for the LIGO-Virgo-KAGRA O4 run is the observation of  $10^{+8}_{-4}$  NS merger events within 90% credible intervals.

We expect the mass constraints in this system to improve significantly in the following years. Through simulations, we predict the masses and inclination angle uncertainty to be reduced to  $\pm 0.1 M_\odot$  and  $\pm 5 \text{ deg}$  with only two extra years of timing. After two decades, mass and inclination angle uncertainties can be reduced down to  $\pm 0.005 M_\odot$  and  $\pm 0.4 \text{ deg}$ , with the uncertainty always being dominated by the precision in the measurement of the Einstein delay amplitude. An eventual detection of  $\dot{P}_b$  will help constrain the true distance to the system by forcing consistency with GR and the PM.

Deep surveys with new sensitive facilities such as MeerKAT, FAST or the SKA in the future will continue to provide new systems similar to J1208–5936, and increase the discovery rate of DNS systems with improved sensitivity and search algorithms. The MPIfR-MeerKAT Galactic Plane survey at S-band (Kramer et al. 2016; Padmanabh et al. 2023) will probe deep into southern Galactic plane at high radio frequencies, allowing the discovery of even more distant and faint compact pulsar binaries without being hampered by propagation effects introduced by the interstellar medium. Further in time, space-based gravitational-wave observatories like LISA will probe tens of electromagnetically dark DNS systems with orbital periods of one hour or less (Lau et al. 2020), constraining estimates of the merger rate in the Milky Way even further.

**Acknowledgements.** The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. The Parkes radio telescope is part of the ATNF, which is funded by the Australian Government for operation as a National Facility managed by the Commonwealth Scientific and Industrial Research Organisation. We acknowledge the Wiradjuri people as the Traditional Owners of the Observatory site. SARAO acknowledges the ongoing advice and calibration of GPS systems by the National Metrology Institute of South Africa (NMISA) and the time space reference systems department of the Paris Observatory. Observations used the FBFUSE and APSUSE computing clusters for data acquisition, storage and analysis. These clusters were funded and installed by the MPIfR and the Max-Planck-Gesellschaft (MPG). All authors affiliated with the MPG acknowledge its constant support. Marina Berezina acknowledges support from the Bundesministerium für Bildung und Forschung D-MeerKAT award 05A17VH3 (Verbundprojekt D-MeerKAT). Marta Burgay acknowledges support through the research grant ‘iPeska’ (PI: A. Possenti) funded under the INAF national call Prin-SKA/CTA approved with the Presidential Decree 70/2016. Vivek Venkatraman Krishnan acknowledges financial support from the European Research Council (ERC) starting grant ‘COMPACT’ (grant agreement number: 101078094). We also thank Alessandro Ridolfi for providing a working version of `pysolator.py` and for his input in Sect. 5, Norbert Wex for his comments on the interpretations of the NS merger rate results in Sect. 6, and Livia Silva Rocha and Robert Main for their general feedback on this manuscript. The data underlying this work will be shared on reasonable request to the MMGPS Collaboration.

## References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *Phys. Rev. Lett.*, **119**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *ApJ*, **848**, L12
- Abbott, R., Abbott, T. D., Acernese, F., et al. 2021, ArXiv e-prints [arXiv:2108.01045]
- Agazie, G. Y., Mingyar, M. G., McLaughlin, M. A., et al. 2021, *ApJ*, **922**, 35

- Allen, B., Knispel, B., Cordes, J. M., et al. 2013, *ApJ*, **773**, 91
- Andersen, B. C., & Ransom, S. M. 2018, *ApJ*, **863**, L13
- Andrews, J. J., & Mandel, I. 2019, *ApJ*, **880**, L8
- Bagchi, M., Lorimer, D. R., & Wolfe, S. 2013, *MNRAS*, **432**, 1303
- Bailes, M., Jameson, A., Abbate, F., et al. 2020, *PASA*, **37**, e028
- Balakrishnan, V., Champion, D., Barr, E., et al. 2022, *MNRAS*, **511**, 1265
- Barker, B. M., & O’Connell, R. F. 1975, *Phys. Rev. D*, **12**, 329
- Barr, E. 2020, Astrophysics Source Code Library [record ascl:2001.014]
- Bhattacharyya, B., & Nityananda, R. 2008, *MNRAS*, **387**, 273
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *Phys. Rep.*, **203**, 1
- Burgay, M., D’Amico, N., Possenti, A., et al. 2003, *Nature*, **426**, 531
- Burgay, M., Joshi, B. C., D’Amico, N., et al. 2006, *MNRAS*, **368**, 283
- Cameron, A. D., Barr, E. D., Champion, D. J., Kramer, M., & Zhu, W. W. 2017, *MNRAS*, **468**, 1994
- Cameron, A. D., Bailes, M., Champion, D. J., et al. 2023, *MNRAS*, **523**, 5064
- Chatteropadhyay, D., Stevenson, S., Hurley, J. R., Rossi, L. J., & Flynn, C. 2020, *MNRAS*, **494**, 1587
- Chen, W., Barr, E., Karuppusamy, R., Kramer, M., & Stappers, B. 2021a, *J. Astron. Instrum.*, **10**, 2150013
- Chen, H.-Y., Holz, D. E., Miller, J., et al. 2021b, *Class. Quant. Grav.*, **38**, 055010
- Cordes, J. M. 2004, *ASP Conf. Ser.*, **317**, 211
- Cordes, J. M., Freire, P. C. C., Lorimer, D. R., et al. 2006, *ApJ*, **637**, 446
- Corongiu, A., Kramer, M., Stappers, B. W., et al. 2007, *A&A*, **462**, 703
- Damour, T., & Deruelle, N. 1986, *Ann. Inst. Henri Poincaré Phys. Théor.*, **44**, 263
- Damour, T., & Ruffini, R. 1974, *C. R. Acad. Sc. Paris Serie A*, **279**, 971
- Damour, T., & Taylor, J. H. 1991, *ApJ*, **366**, 501
- Dewi, J. D. M., Podsiadlowski, P., & Pols, O. R. 2005, *MNRAS*, **363**, L71
- Eatough, R. P., Torne, P., Desvignes, G., et al. 2021, *MNRAS*, **507**, 5053
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, **372**, 1549
- Faulkner, A. J., Kramer, M., Lyne, A. G., et al. 2005, *ApJ*, **618**, L119
- Ferdman, R. D., Stairs, I. H., Kramer, M., et al. 2014, *MNRAS*, **443**, 2183
- Ferdman, R. D., Freire, P. C. C., Perera, B. B. P., et al. 2020, *Nature*, **583**, 211
- Fonseca, E., Stairs, I. H., & Thorsett, S. E. 2014, *ApJ*, **787**, 82
- Freire, P. C. C., & Ridolfi, A. 2018, *MNRAS*, **476**, 4794
- Freire, P. C. C., & Wex, N. 2010, *MNRAS*, **409**, 199
- Grunthal, K., Kramer, M., & Desvignes, G. 2021, *MNRAS*, **507**, 5658
- Haniewicz, H. T., Ferdman, R. D., Freire, P. C. C., et al. 2021, *MNRAS*, **500**, 4620
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
- Hobbs, G., Manchester, R. N., Dunning, A., et al. 2020, *PASA*, **37**, e012
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, **21**, 302
- Hu, H., Kramer, M., Wex, N., Champion, D. J., & Kehl, M. S. 2020, *MNRAS*, **497**, 3118
- Janssen, G. H., Stappers, B. W., Kramer, M., et al. 2008, *A&A*, **490**, 753
- Johnston, H. M., & Kulkarni, S. R. 1991, *ApJ*, **368**, 504
- Jonas, J., & MeerKAT Team 2016, *MeerKAT Science: On the Pathway to the SKA*, 1
- Kalogera, V., Kim, C., Lorimer, D. R., et al. 2004, *ApJ*, **601**, L179
- Keith, M. J., Kramer, M., Lyne, A. G., et al. 2009, *MNRAS*, **393**, 623
- Keith, M. J., Jameson, A., van Straten, W., et al. 2010, *MNRAS*, **409**, 619
- Kim, C., Kalogera, V., & Lorimer, D. R. 2003, *ApJ*, **584**, 985
- Kim, C., Kalogera, V., & Lorimer, D. 2010, *New Astron. Rev.*, **54**, 148
- Kim, C., Perera, B. B. P., & McLaughlin, M. A. 2015, *MNRAS*, **448**, 928
- Kopeikin, S. M. 1996, *ApJ*, **467**, L93
- Kopparapu, R. K., Hanna, C., Kalogera, V., et al. 2008, *ApJ*, **675**, 1459
- Kramer, M., Menten, K., Barr, E. D., et al. 2016, *MeerKAT Science: On the Pathway to the SKA*, 3
- Kramer, M., Stairs, I. H., Manchester, R. N., et al. 2021, *Phys. Rev. X*, **11**, 041050
- Lau, M. Y. M., Mandel, I., Vigna-Gómez, A., et al. 2020, *MNRAS*, **492**, 3061
- Lentati, L., Alexander, P., Hobson, M. P., et al. 2014, *MNRAS*, **437**, 3004
- Lorimer, D. R., & Kramer, M. 2005, *Handbook of Pulsar Astronomy* (Cambridge: Cambridge University Press)
- Lynch, R. S., Swiggum, J. K., Kondratiev, V. I., et al. 2018, *ApJ*, **859**, 93
- Lyne, A. G., Burgay, M., Kramer, M., et al. 2004, *Science*, **303**, 1153
- Manchester, R. N., Lyne, A. G., Camilo, F., et al. 2001, *MNRAS*, **328**, 17
- Martinez, J. G., Stovall, K., Freire, P. C. C., et al. 2015, *ApJ*, **812**, 143
- Martinez, J. G., Stovall, K., Freire, P. C. C., et al. 2017, *ApJ*, **851**, L29
- McMillan, P. J. 2017, *MNRAS*, **465**, 76
- Men, Y. P., Barr, E., Clark, C. J., Carli, E., & Desvignes, G., 2023, *A&A*, in press <https://doi.org/10.1051/0004-6361/202347356>,
- Morello, V., Barr, E. D., Cooper, S., et al. 2019, *MNRAS*, **483**, 3673
- Ng, C., Kruckow, M. U., Tauris, T. M., et al. 2018, *MNRAS*, **476**, 4315
- Oswald, L. S., Karastergiou, A., Posselt, B., et al. 2021, *MNRAS*, **504**, 1115
- Özel, F., & Freire, P. 2016, *ARA&A*, **54**, 401
- Padmanabh, P. V., Barr, E. D., Sridhar, S. S., et al. 2023, *MNRAS*, **524**, 1291
- Peters, P. C. 1964, *Phys. Rev.*, **136**, 1224
- Pol, N., McLaughlin, M., & Lorimer, D. R. 2019, *ApJ*, **870**, 71
- Pol, N., McLaughlin, M., & Lorimer, D. R. 2020, *Res. Notes Am. Astron. Soc.*, **4**, 22
- Ridolfi, A., Freire, P. C. C., Gupta, Y., & Ransom, S. M. 2019, *MNRAS*, **490**, 3860
- Sengar, R., Balakrishnan, V., Stevenson, S., et al. 2022, *MNRAS*, **512**, 5782
- Shklovskii, I. S. 1970, *Sov. Astron.*, **13**, 562
- Singh, S., Roy, J., Panda, U., et al. 2022, *ApJ*, **934**, 138
- Sofue, Y. 2020, *Galaxies*, **8**, 37
- Splaver, E. M., Nice, D. J., Arzoumanian, Z., et al. 2002, *ApJ*, **581**, 509
- Staelin, D. H. 1969, *IEEE Proc.*, **57**, 724
- Stappers, B., & Kramer, M. 2016, *MeerKAT Science: On the Pathway to the SKA*, 9
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, **791**, 67
- Stovall, K., Freire, P. C. C., Chatterjee, S., et al. 2018, *ApJ*, **854**, L22
- Suresh, A., Cordes, J. M., Chatterjee, S., et al. 2022, *ApJ*, **933**, 121
- Swiggum, J. K., Rosen, R., McLaughlin, M. A., et al. 2015, *ApJ*, **805**, 156
- Swiggum, J. K., Pleunis, Z., Parent, E., et al. 2023, *ApJ*, **944**, 154
- Tauris, T. M., Langer, N., & Podsiadlowski, P. 2015, *MNRAS*, **451**, 2123
- Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *ApJ*, **846**, 170
- Taylor, J. H. 1992, *Philos. Trans. R. Soc. A*, **341**, 117
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, **253**, 908
- Taylor, J. H., & Weisberg, J. M. 1989, *ApJ*, **345**, 434
- Taylor, J. H., Fowler, L. A., & McCulloch, P. M. 1979, *Nature*, **277**, 437
- van den Heuvel, E. P. J. 2019, *IAU Symp.*, **346**, 1
- van Leeuwen, J., Kasian, L., Stairs, I. H., et al. 2015, *ApJ*, **798**, 118
- Vigna-Gómez, A., Neijssel, C. J., Stevenson, S., et al. 2018, *MNRAS*, **481**, 4009
- Weisberg, J. M., & Huang, Y. 2016, *ApJ*, **829**, 55
- Wolszczan, A. 1991, *Nature*, **350**, 688
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, **835**, 29

## Appendix A: Implementation of the DM jump fit

One of the most relevant differences between APSUSE and PTUSE is the implementation of coherent de-dispersion during recording for PTUSE. This leads to a discrepancy between the best DM in the two data sets, as APSUSE data suffers from intra-channel smearing. While this hampers the  $S/N$  of APSUSE-derived ToAs, it also introduces a best-DM discrepancy between the APSUSE-derived ToAs and the PTUSE-derived ToAs, which leads to spurious DM trends in the time series when combined, and therefore to biased fit parameters and reduced quality fit. We implement a DM jump between the two data sets to enforce the continuity of the DM variations in the data, and to ensure a good quality fit of the DM and DM1 parameters in the DDH and DDGR models, taking advantage that in some epochs the data sets overlap.

Such DM jump is not implemented in TEMPO2, and therefore we implement a manual fitting instead. The process is as follows: firstly, a global fit is implemented on the entire data set, including DM and DM1. Such fit has an unreliable DM1 value because APSUSE data dominates the beginning of the time series, while PTUSE dominates the middle and later half. Then, we take the resulting model and fit only DM and DM1 for the APSUSE and PTUSE data sets individually. This creates two individual DM models, one for each data set, represented as

$$\begin{aligned} DM^{\text{APSUSE}}(t) &= DM_0^{\text{APSUSE}} + t \times DM_p^{\text{APSUSE}} \quad \text{and} \\ DM^{\text{PTUSE}}(t) &= DM_0^{\text{PTUSE}} + t \times DM_p^{\text{PTUSE}}, \end{aligned} \quad (\text{A.1})$$


where  $DM_0^i$  and  $DM_p^i$  are the DM value at a reference time and its constant derivative, and  $t$  the instantaneous time, which covers a different range for each data set. Then, the average of  $\Delta DM = DM^{\text{PTUSE}} - DM^{\text{APSUSE}}$  is computed in the overlapping range, and the DM value of the APSUSE template is shifted by  $-\Delta DM$ . Subsequently, the APSUSE ToAs are re-processed with this fix. This cycle is repeated several times until the  $\Delta DM$  value is seen to approach zero. In our implementation, the profiles were de-dispersed originally with  $DM = 344.312 \text{ pc cm}^{-3}$ , and the value that sets  $\Delta DM \sim 0$  was found to be  $DM^{\text{APSUSE}} = 344.336 \text{ pc cm}^{-3}$ .

The setup for the DDH and DDGR fits and their measurements are listed in Table 2. The timing residuals are plotted in Fig. 3.

## Appendix B: Implementation of the MerKAT coherent beam pattern in PsrPopPy2

Unlike the pulsar surveys included in the previous NS merger rate estimates, the MMGPS-L makes use of the MeerKAT multi-dish telescope. It is, therefore, necessary to implement the interferometric nature of its observations, which manifest as a tiling of coherently phased beams within the main survey beam defined by the resolution of individual antennas (Chen et al. 2021a; Padmanabh et al. 2023).

Each MMGPS-L pointing consists of a survey beam with a Full Width at Half Maximum (FWHM) of 29.85 arcmin, within which a tiling of 480 coherently phased beams is embedded (Chen et al. 2021a; Padmanabh et al. 2023). The size and orientations of these coherent beams depends on the local sky position and the antenna array configuration, and therefore they can be considered independent of the RA and DEC coordinates (Chen et al. 2021a). Thus, we consider them random in nature, with a probability distribution in semi-major and semi-minor axis that mimics the true distribution of MMGPS-L size values.



**Fig. B.1.** Randomly computed coherent beam sizes for the PsrPopPy2 simulation of the MMGPS-L survey for the NS-NS merger rate computation (orange) against the true distribution of coherent beam sizes of the bulk of the MMGPS-L pointings (blue).

The probability distributions are derived by fitting skewed Gaussian functions to the full set of MMGPS-L pointing size values, and upon simulation of random values, we see that this method is successful in mimicking typical MMGPS-L coherent beam sizes (Fig. B.1).

A second aspect of randomness in the simulation of MMGPS-L pointings is the size of the coherent beam tiling itself. Before October 5th, 2021, the coherent tiling was circular and its radius  $r_c$  depended on the local sky position (Padmanabh et al. 2023). In most of the pointings, it was held that  $r_c < \text{FWHM}/2$ , but occasionally the opposite was true. We find that a sum of two normal Gaussian functions with standard deviations of  $\sigma_1 = 0.898$  and  $\sigma_2 = 2.264$  arcmin, central values of  $\langle x_1 \rangle = 12.546$  and  $\langle x_2 \rangle = 12.761$  arcmin and a height ratio of  $A_1/A_2 = 2.942$  describes the distribution of the distance of the furthest coherent beam from the centre of the tiling, and take it as a distribution for  $r_c$ . After October 5th, 2021, the coherent beam tiling was instead forced to adopt an hexagonal shape with an inner radius equal to the survey beam to provide a complete coverage of the sky, which allows us to model the tiling consistently across pointings (Padmanabh et al. 2023).

This parametrisation is implemented to PsrPopPy2 by modifying the `doSurvey.py` script with the computation of an extra degradation factor  $0 \leq D_{\text{coh}} \leq 1$  for the coherent beams on top of the Gaussian degradation factor of the survey beam  $0 \leq D_{\text{sur}} \leq 1$  used for single-dish surveys<sup>29</sup>. The process goes as following: for each computed  $D_{\text{sur}}$ , we randomly select whether the pointing should be considered under the older or newer rules for tiling filling (before or after October 5th, 2021) with a 0.39 or 0.61 chance, corresponding to the fraction of pointings under the old and new tiling rules, respectively. In the case that the newer tiling rules are assigned, then we assume an efficient hexagonal tiling which takes an area of

$$A = (29.85/2 \text{ arcmin})^2 \times 6 \times \frac{\sin 30^\circ}{\cos 30^\circ}, \quad (\text{B.1})$$

where the 480 coherent beams are uniformly distributed. Otherwise, the area taken by the tiling is

$$A = \pi \times (0.95 \times r_c \text{ arcmin})^2, \quad (\text{B.2})$$

<sup>29</sup> <https://github.com/mcbernadich/DNS-merger-rate-2022>



where the 480 coherent beams are also uniformly distributed, and where the 0.95 factor comes from the fact that the outermost coherent beam is typically an outlier slightly outside of the main radius. This area is then divided between the 480 beams and a random position for a potential pulsar discovery is drawn within the square of the corresponding size. Then,  $D_{\text{coh}}$  is a second degradation factor computed from a 2D Gaussian function with the random dimensions drawn from the coherent beam tiling size distributions (Fig. B.1) centered at the square and aligned with its sides. However, for cases from the older tiling rules in which  $r_c < \text{FWHM}/2$ , if the offset from the centre of the survey beam is found to be larger than  $1.05 \times r_c$ , then we set  $D_{\text{coh}} = 0$  instead.

Finally, like in the previous works (Grunthal et al. 2021; Pol et al. 2020), each pulsar in a DNS merging within the Hubble

time has a computed Doppler degeneration factor that depends on the orbital parameters, the spin period, the duration of the observations and the applied search method based on the definitions given by Bagchi et al. (2013). Like the other surveys, the MMGPS-L implements an acceleration search, and therefore we computed the parameter  $0 \leq \gamma_{2m} \leq 1$  from equation (11) in Bagchi et al. (2013) using the code<sup>30</sup> prepared in Pol et al. (2019), under the assumption of 500-second long observations. Therefore, the  $S/N$  of a simulated pulsar in PsrPopPy2 is computed as

$$S/N_{\text{pulsar}} = S/N_{\text{input}} \times D_{\text{sur}} \times D_{\text{coh}} \times \gamma_{2m}^2. \quad (\text{B.3})$$

For the MMGPS-L, we only consider as discoveries signals with  $S/N_{\text{pulsar}} \geq 9$ .

<sup>30</sup> [https://github.com/NihanPol/SNR\\_degradation\\_factor\\_for\\_BNS\\_systems](https://github.com/NihanPol/SNR_degradation_factor_for_BNS_systems)



---

### Paper on PSR J1227–6208

---

The paper Colomi Bernadich et al. (2024), A&A, **690**, A253, was published in the peer-reviewed journal of Astronomy & Astrophysics on 15/10/2024 under the Creative Commons CC-BY 4.0 license, and is reproduced in its original form below. This paper is part of many other RelBin papers on relativistic binaries measured with MeerKAT and Parkes/Murriyang.

# PSR J1227–6208 and its massive white dwarf companion: Pulsar emission analysis, timing update, and mass measurements

Miquel Colom i Bernadich<sup>1,\*</sup>, Vivek Venkatraman Krishnan<sup>1,\*</sup>, David J. Champion<sup>1</sup>, Paulo C. C. Freire<sup>1</sup>,  
Michael Kramer<sup>1</sup>, Thomas M. Tauris<sup>2,1</sup>, Matthew Bailes<sup>3</sup>, Alessandro Ridolfi<sup>4,1</sup>,  
Marcus E. Lower<sup>5</sup>, and Maciej Serylak<sup>6,7</sup>

<sup>1</sup> Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

<sup>2</sup> Dept. of Materials and Production, Aalborg University, DK-9220 Aalborg Øst, Denmark

<sup>3</sup> INAF – Osservatorio Astronomico di Cagliari, via della Scienza 5, 09047 Selargius (CA), Italy

<sup>4</sup> Centre for Astrophysics and Supercomputing, Swinburne University of Technology, P.O. Box 218, Hawthorn, Vic 3122, Australia

<sup>5</sup> SKA Observatory, Jodrell Bank, Lower Withington, Macclesfield SK11 9FT, United Kingdom

<sup>6</sup> Australia Telescope National Facility, CSIRO, Space and Astronomy, PO Box 76, Epping, NSW 1710, Australia

<sup>7</sup> Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town 7535, South Africa

Received 14 May 2024 / Accepted 16 July 2024

## ABSTRACT

PSR J1227–6208 is a 34.53-ms recycled pulsar with a massive companion. This system has long been suspected to belong to the emerging class of massive recycled pulsar–ONeMg white dwarf systems such as PSR J2222–0137, PSR J1528–3146, and J1439–5501. Here, we present an updated emission and timing analysis with more than 11 years of combined Parkes and MeerKAT data, including 19 hours of high-frequency data from the newly installed MeerKAT S-band receivers. We measure a scattering timescale of 1.22 ms at 1 GHz with a flat scattering index of  $3.33 < \beta < 3.62$ , and a mean flux density of 0.53–0.62 mJy at 1 GHz with a steep spectral index of  $2.06 < \alpha < 2.35$ . Around 15% of the emission is linearly and circularly polarised, but the polarisation angle does not follow the rotating vector model. Thanks to the sensitivity of MeerKAT, we successfully measure a rate of periastron advance of  $\dot{\omega} = 0.0171(11) \text{ deg yr}^{-1}$ , and a Shapiro delay with an orthometric amplitude of  $h_3 = 3.6 \pm 0.5 \mu\text{s}$  and an orthometric ratio of  $\varsigma = 0.85 \pm 0.05$ . The main source of uncertainty in our timing analysis is chromatic correlated dispersion measure noise, which we model as a power law in the Fourier space thanks to the large frequency coverage provided by the Parkes UWL receiver. Assuming general relativity and accounting for the measurements across all the implemented timing noise models, the total mass, companion mass, pulsar mass, and inclination angle are constrained at  $2.3 < M_t/M_\odot < 3.2$ ,  $1.21 < M_c/M_\odot < 1.47$ ,  $1.16 < M_p/M_\odot < 1.69$ , and  $77.5 < i/\text{deg} < 80.3$ . We also constrain the longitude of ascending node to either  $\Omega_a = 266 \pm 78 \text{ deg}$  or  $\Omega_a = 86 \pm 78 \text{ deg}$ . We argue against a neutron star nature of the companion based on the very low orbital eccentric of the system ( $e = 1.15 \times 10^{-3}$ ), and instead classify the companion of PSR J1227–6208 as a rare, massive ONeMg white dwarf close to the Chandrasekhar limit.

**Key words.** binaries: general – stars: evolution – stars: fundamental parameters – stars: neutron – stars: individual: PSR J1227-6208 – white dwarfs

## 1. Introduction

The Chandrasekhar limit of white dwarf (WD) masses ( $M_{\text{WD}}$ ) is a topic of ongoing research in astrophysics and theoretical physics. Studies show that maximally rotating rigid WDs could sustain up to  $M_{\text{WD}} \approx 1.48 M_\odot$  (Yoon & Langer 2005), but several effects have been speculated to allow them to exist beyond this limit, such as differential rotation (Yoon & Langer 2005) or high magnetisation (Kundu & Mukhopadhyay 2012), and super-Chandrasekhar masses have been inferred indirectly from some Type Ia supernovae (Tomaschitz 2018). Additionally, extensions or modifications of general relativity (GR) must incorporate a prediction for the upper mass limit of WDs (e.g. Gregoris & Ong 2023; Mathew & Nandy 2021). To test these postulations, the empirical measurement of  $M_{\text{WD}}$  values close to the Chandrasekhar limit is required. However, WDs with  $M_{\text{WD}} \gtrsim 1.3 M_\odot$  are a rarity in the Galactic field. One of their windows of study is optical photometry, which has enabled the

characterisation of their masses based on emission models (e.g. Caiazzo et al. 2021; Hollands et al. 2020; Külebi et al. 2010; Miller et al. 2023; Pshirkov et al. 2020). The other window of study is mass measurements of binary radio pulsars with massive WD companions obtained via pulsar timing. This technique allows for the measurement of relativistic effects in the orbital motion and in the light propagation time; these are quantified by the post-Keplerian (PK) parameters. Under the assumption of GR, such measurements can result in the precise measurement of the WD and pulsar masses (Lorimer & Kramer 2005).

Timing measurements of pulsars with massive WD companions are relevant not only for probing the Chandrasekhar limit, but also for testing binary evolution and fundamental physics. In recycled pulsar binaries, we can test how binary interaction affects the resulting WD and neutron star (NS) mass distribution (for general insights into binary evolution and pulsar recycling, see Tauris & van den Heuvel 2023). For instance, there is a bi-modality in the  $M_{\text{WD}}$  distribution, with  $M_{\text{WD}} \lesssim 0.5 M_\odot$  WDs being found with fully recycled pulsars and  $M_{\text{WD}} \gtrsim 0.7 M_\odot$  WDs with mildly recycled pulsars (McKee et al. 2020; Shamohammadi et al. 2023). This divide is well understood

\* Corresponding authors; mcbnadich@mpi-fr-bonn.mpg.de;  
vkrishnan@mpi-fr-bonn.mpg.de

thanks to models of binary evolution (Lazarus et al. 2014; Tauris et al. 2012; Tauris & van den Heuvel 2023), but the upper end of the  $M_{\text{WD}}$  distribution ( $M_{\text{WD}} \gtrsim 1.1 M_{\odot}$ ) remains relatively unsampled. The birth NS mass distribution can also be directly sampled in these kinds of systems given the small amount of mass accretion occurring during recycling (e.g. Lazarus et al. 2014; Cognard et al. 2017). For instance, the measurement of a pulsar mass of  $M_p = 1.831(10) M_{\odot}$  in PSR J2222–0137 (Guo et al. 2021) is evidence that NSs can be born massive (heavier than  $1.4 M_{\odot}$ ) instead of acquiring large amounts of extra mass via accretion (Cognard et al. 2017). Furthermore, the NS mass distribution is a probe of the physics of matter under conditions of extreme density (e.g. Özel & Freire 2016; Fonseca et al. 2021; Hu et al. 2020). Finally, in the most compact systems, the timing of pulsars with WD companions also provides extremely precise tests of gravity theories (e.g. Voisin et al. 2020), as is shown by the strict constraints on dipolar gravitational wave emission from PSR J2222–0137 (Guo et al. 2021), which have ruled out the phenomenon of spontaneous scalarisation predicted by some alternative gravity theories (Zhao et al. 2022).

As of today, only four pulsars with WD companions at the upper end of the mass distribution have been characterised. The most studied one is PSR J2222–0137, a 32.8-ms recycled pulsar in a circular 2.45-day orbit with a  $M_{\text{WD}} = 1.319(4) M_{\odot}$  companion (Boyles et al. 2013; Cognard et al. 2017; Guo et al. 2021). The system distance of 268 pc, measured via very-long-baseline interferometry observations (Deller et al. 2013; Guo et al. 2021), and the lack of an optical detection imply a cold WD ( $T < 3000$  K) with a cooling age of at least several Gyr (Kaplan et al. 2014). A similar case is PSR J1528–3146, a 60.8-ms recycled pulsar in a circular 3.18-day orbit with a  $M_{\text{WD}} = 1.33^{+0.08}_{-0.07} M_{\odot}$  companion (Jacoby et al. 2006, 2007; Berthreau et al. 2023). In this case, the optical detection implies a cooling age between 1.5 and 3.2 Gyr (Jacoby et al. 2006), consistent with the pulsar characteristic age of 3.9 Gyr estimated by Berthreau et al. (2023). A third system is PSR J1439–5501, with a 28.6-ms pulsar in a 2.12-day orbit with a  $M_{\text{WD}} = 1.27^{+0.14}_{-0.12} M_{\odot}$  optically detected companion with a cooling age of 0.1–0.5 Gyr (Jang et al., in prep. Faulkner et al. 2004; Lorimer et al. 2006; Pallanca et al. 2013). The final system is PSR B2303+46, with a young 0.937-s pulsar in a highly eccentric 12.34-day orbit with a  $M_{\text{WD}} = 1.34^{+1.08}_{-0.15} M_{\odot}$  companion (Thorsett et al. 1993; Thorsett & Chakrabarty 1999). The WD companion is hot and young, implying a system age of 30 Myrs (van Kerkwijk & Kulkarni 1999).

In this work, we present a detailed study of PSR J1227–6208 (J1227–6208 from now on). Discovered independently by three different studies in data from the Murriyang Parkes telescope<sup>1</sup> (Bates et al. 2015; Knispel et al. 2013; Mickaliger et al. 2012), it is a 34.5-ms mildly recycled pulsar in orbit with a massive companion. With an orbital period of  $P_b = 6.72$  days and a projected semimajor axis of  $x = 23.2$  light seconds (ls), it has a high mass function of  $f_M = 0.297 M_{\odot}$ . Assuming a pulsar mass  $M_p = 1.35 M_{\odot}$ , this leads to a minimum companion mass of  $M_c > 1.27 M_{\odot}$ . Its low orbital eccentricity ( $e = 1.15 \times 10^{-3}$ ) makes the possibility of a NS companion unlikely, giving more weight to a massive WD hypothesis instead. However, owing to the low orbital eccentricity, and large timing uncertainties, precise measurements of PK parameters in this system have been impossible until now.

It is for these reasons that the system was included in the Relativistic Binary program (RelBin, Kramer et al. 2021) of Meer-

TIME (Bailes et al. 2020), a large science project that takes advantage of the high sensitivity of the MeerKAT telescope<sup>2</sup> (Jonas & MeerKAT Team 2016) to improve the timing precision of southern pulsars. The RelBin project is designed to measure masses and test theories of gravity in 25 selected binary pulsar systems (Kramer et al. 2021), including J1227–6208. The updated timing analysis of J1227–6208 presented here includes a decade-long timing baseline of Parkes/Murriyang observations and two years of dedicated MeerKAT observations. This is the first timing experiment to include data from the newly commissioned S-band MeerKAT receivers (Barr 2018).

This paper is structured as follows. Section 2 details the observations used in this study and the data reduction for timing. Section 3 reports the emission study, including the modelling of the profile, scattering, and spectral properties, and the detection of polarised light. Section 4 reports the timing analysis, including the modelling of several timing noises, constraints on PK parameters, mass measurements, and an exploration of the orbital orientation. Section 5 discusses the astrophysical implications of our measurements. In Section 6 we outline future prospects and lines of action for the study of this system. Finally, Section 7 concludes the paper and summarises the key aspects of our measurements and discussions.

## 2. Data acquisition and reduction

### 2.1. Parkes/Murriyang

Table 1 shows the four data sets used in this work. Observations were performed first with the central beam of the 21-cm multibeam receiver (Staveley-Smith et al. 1996) and later with the ultra-wide bandwidth low-frequency receiver (UWL, Hobbs et al. 2020). The multibeam data set provides the longest baseline for the measurement of secular variations in the Keplerian parameters, while the UWL observations are particularly useful for constraining the dispersion measure (DM) evolution due to their overlap with the MeerKAT observations and their very large bandwidth. UWL observations also include three dense orbital campaigns on the dates of 4–10 October 2020, 20–26 July 2022, and 3–9 May 2023, with seven observations each, accumulating 5.74, 10.41, and 17.63 hours each. These campaigns have two aims: aiding in the measurement of the time delay of the pulses as they propagate through the gravitational field of the companion (Shapiro delay, Shapiro 1964) and constraining DM evolution within a single orbit. In line with this objective, the first of the orbital campaign was coordinated with the MeerKAT orbital campaign with the L-band receivers (Section 2.2) so that observations alternate each other.

The multibeam receiver data were recorded by the Center for Astronomy Signal Processing and Electronics Research (CASPER) Parkes Swinburne Recorder (CASPSR, Sarkissian et al. 2011) backend. The data were folded with 512 frequency channels, four polarisation channels, 1024 phase bins, and with coherent de-dispersion at  $\text{DM} \approx 363 \text{ cm}^{-3} \text{ pc}$ . Each observation was accompanied by a noise-diode observation for polarisation calibration. Calibration was performed on each file with the `pac` command from the PSRCHIVE<sup>3</sup> software package (Hotan et al. 2004). The archives were manually excised of radio-frequency-interference (RFI) with the PSRCHIVE/`pazi` interface, and the 80 bottom channels and 32 top channels were zero-weighted with the PSRCHIVE/`paz` command to remove

<sup>2</sup> <https://www.sarao.ac.za/science/meerkat/about-meerkat/>

<sup>3</sup> <https://psrchive.sourceforge.net/>

<sup>1</sup> <https://www.parkes.atnf.csiro.au/>

**Table 1.** Summary of the data sets used in this analysis.

Receiver(s)	Freq. (MHz)	Bandw. (MHz)	Temp. K	Gain K Jy <sup>-1</sup>	# obs	Time (hours)	First obs.	Last obs.	# bands	# ToAs	Median err. ( $\mu$ s)
Parkes multibeam	1382	400	~21	~0.9	93	33	1 March 2012	8 April 2019	4	365	26.4
MeerKAT <i>L</i> -band	1284	856	4–7	~2.6	37	26	12 February 2019	25 May 2023	8	508	5.39
Parkes UWL	2368	3328	~22	~0.9	56	72	4 May 2020	6 June 2023	8	444	19.5
MeerKAT <i>S</i> -band	2406	875	4–7	~2.3	10	19	12 May 2023	28 May 2023	4	228	5.43

**Notes.** The data were recorded in folding mode with four polarisation channels and 1024 phase bins of 28.28 of  $\mu$ s in length. The system temperatures are drawn from the [MeerKAT General documentation](#), the [Multibeam Receiver Description webpage](#) and [Hobbs et al. \(2020\)](#). The gain is estimated from the collecting area, accounting that only 56 antennas are available in the MeerKAT *S*-band configuration. The last three columns indicate timing data reduction parameters: number of sub-bands used in multi-frequency timing, number of ToAs per data set, and the median ToA uncertainty.

Gaussian noise caused by the loss of sensitivity at the edge of the bandpass. The bands were then scrunched to four frequency sub-bands, a single intensity channel, and a single subintegration per observation with the PSRCHIVE/pam command for the production of frequency-resolved times of arrival (ToAs). To derive timing templates for each band, we obtained a frequency-resolved standard profile resulting from the time integration of all multibeam CAPSR observations (with the exclusion of three heavily RFI-affected observations) into a single rotational phase cycle. The analytic timing templates were then produced for each sub-band by fitting a combination of von Mises functions with PSRCHIVE/paas program. The ToAs were produced with the PSRCHIVE/pat command using the FDM algorithm with the two-dimensional, frequency-resolved timing template, where each subintegration resulted a ToA from each of the four sub-bands.

The UWL receiver data were recorded by the Medusa cluster ([Hobbs et al. 2020](#)) with 3328 frequency channels, four polarisation channels, 1024 phase bins, and with coherent dedispersion at  $DM \approx 363 \text{ cm}^{-3} \text{ pc}$ . In addition, the data were also recorded with real-time folding based on an early pulsar ephemeris. The data were processed using the psrpye processing pipeline<sup>4</sup>. This pipeline carries out flux and polarisation calibration, and also automatically removes RFI using c1fd<sup>5</sup>. It results in data that is cleansed of RFI, calibrated, and broken down into various time, frequency, and polarisation resolutions. The RFI excision step also excises a standard set of frequencies, on top of whichever part of the data that the algorithm considers to be contaminated. These frequency sets were decided based on the knowledge of the known transmitter frequencies that routinely affect the data. To boost the quality of ToAs, the data was scrunched into total intensity, eight frequency channels, and subintegrations of two hours in length with the PSRCHIVE/pam command. The ToAs were produced with the PSRCHIVE/pat command using the FDM algorithm. Analogously to the multibeam CASPSR data set, the timing template was modelled with PSRCHIVE/paas from the integration of all of the files into a two-dimensional frequency-resolved standard profile, with the exclusion of three severely RFI-affected observations.

## 2.2. MeerKAT

The MeerKAT data sets provide the most precise ToAs owing to the large sensitivity of the telescope. Most of the data were recorded with the *L*-band receivers (26 hours, 856–1712 MHz), but 19 hours were recorded with the new *S*-band receivers in

the *S1* configuration (1968–2843 MHz, [Barr 2018](#)) on 12–28 May 2023 to ensure a measurement of the Shapiro delay with a significant reduction of DM noise (see Table 1). The data were recorded with the pulsar timing user-supplied equipment (PTUSE, [Bailes et al. 2020](#)) machines as part of the MeerTIME science program ([Bailes et al. 2020](#)), and included an orbital campaign on the dates of 4–11 October 2020 that accumulated 13 hours over nine *L*-band observations, with a dedicated five-hour-long observation at superior conjunction on 11 October 2020, aimed at constraining the Shapiro delay.

Both data sets were recorded with 1024 frequency channels, four polarisation channels, and coherent dedispersion at  $DM \approx 363 \text{ cm}^{-3} \text{ pc}$ . The *L*-band data were processed, cleaned of RFI, and calibrated with the meerpipeline<sup>6</sup>, which resulted in the trimming of the outer edges, leaving only 926 channels of useful data with a bandwidth of 775.5 MHz. Any remaining RFI was manually removed with the PSRCHIVE/paz command. Observations were then fully summed in polarisation, and scrunched into eight frequency channels and 20-min-long subintegrations to extract frequency-resolved ToAs with the PSRCHIVE/pat command. Analogous to the Parkes data sets (Section 2.1), the frequency-resolved timing template was derived using the PSRCHIVE/paas on frequency-resolved standard profiles resulting from the time integration of the observations in the October 4–11 October 2020 orbital campaign. Unfortunately, calibration files were not available for the *S*-band data set, as it was taken jointly with commissioning data, and therefore it could not be calibrated nor band-pass corrected. This, however, is unlikely to affect our timing precision. *S*-band data were cleaned of RFI with the c1fd software, and it was sub-banded to four frequency channels for frequency-resolved timing, with the profile being created from the integration of all of the observations, following the same steps as the *L*-band data.

## 3. Emission analysis

### 3.1. Profile evolution


J1227–6208 suffers from significant scattering and has a steep flux density spectrum. Fig. 1 shows time-integrated, frequency-resolved standard profiles from the four different receivers listed in Table 1. The most obvious feature is scattering at the MeerKAT *L*-band and the bottom of the Parkes UWL band, resulting in an exponential, frequency-dependent tail of the pulse profile. In addition, the pulse brightness fades quickly at the top of the Parkes UWL band. It should be noted that the dark bands in the Parkes UWL profile are a result of the RFI excision processes, and are not intrinsic to the pulsar emission. To measure

<sup>4</sup> <https://github.com/vivekvenkris/psrpye>

<sup>5</sup> <https://github.com/v-morello/c1fd>

<sup>6</sup> <https://zenodo.org/records/7961071>





**Fig. 1.** Time-integrated pulse profiles of J1227–6208 as recorded with the different receivers (from left to right: Parkes multibeam, MeerKAT *L*-band, Parkes UWL, and MeerKAT *S*-band). The top plots show the intensity of the integrated emission, while the bottom plots are frequency-resolved. Only the MeerKAT *L*-band and Parkes UWL data sets are flux calibrated. The plots were made with PSRCHIVE/psrplot.

the scattering timescale,  $\tau_s$ , we modelled the pulse at each frequency,  $f$ , as a single Gaussian function convolved with a scattering exponential tail,

$$S_b = \int A \times \exp\left(\frac{(b' - b_0)^2}{2 \times \Delta b^2}\right) \times \exp\left(-\frac{b - b'}{\tau_s}\right) db', \quad (1)$$

where  $b$  is the spin phase in bins, and  $b_0$  and  $\Delta b$  stand for the Gaussian centre and the standard deviation, respectively. We measured the scattering index,  $\alpha$ , by fitting the evolution of  $\tau_s$  as a power law function of  $f$ ,

$$\tau_s(f) = \tau_{\text{GHz}} \left(\frac{f}{\text{GHz}}\right)^{-\alpha}, \quad (2)$$

where  $\tau_{\text{GHz}}$  is the reference value at  $f = 1$  GHz. For the flux density spectral index,  $\beta$ , we measured the mean flux density,  $F_m$ , across the pulse phase at each sub-band, deriving the uncertainty from the off-pulse baseline noise, and fitted another power law,

$$F_m(f) = F_{\text{GHz}} \left(\frac{f}{\text{GHz}}\right)^{-\beta}, \quad (3)$$

where  $F_{\text{GHz}}$  is the reference value at  $f = 1$  GHz.

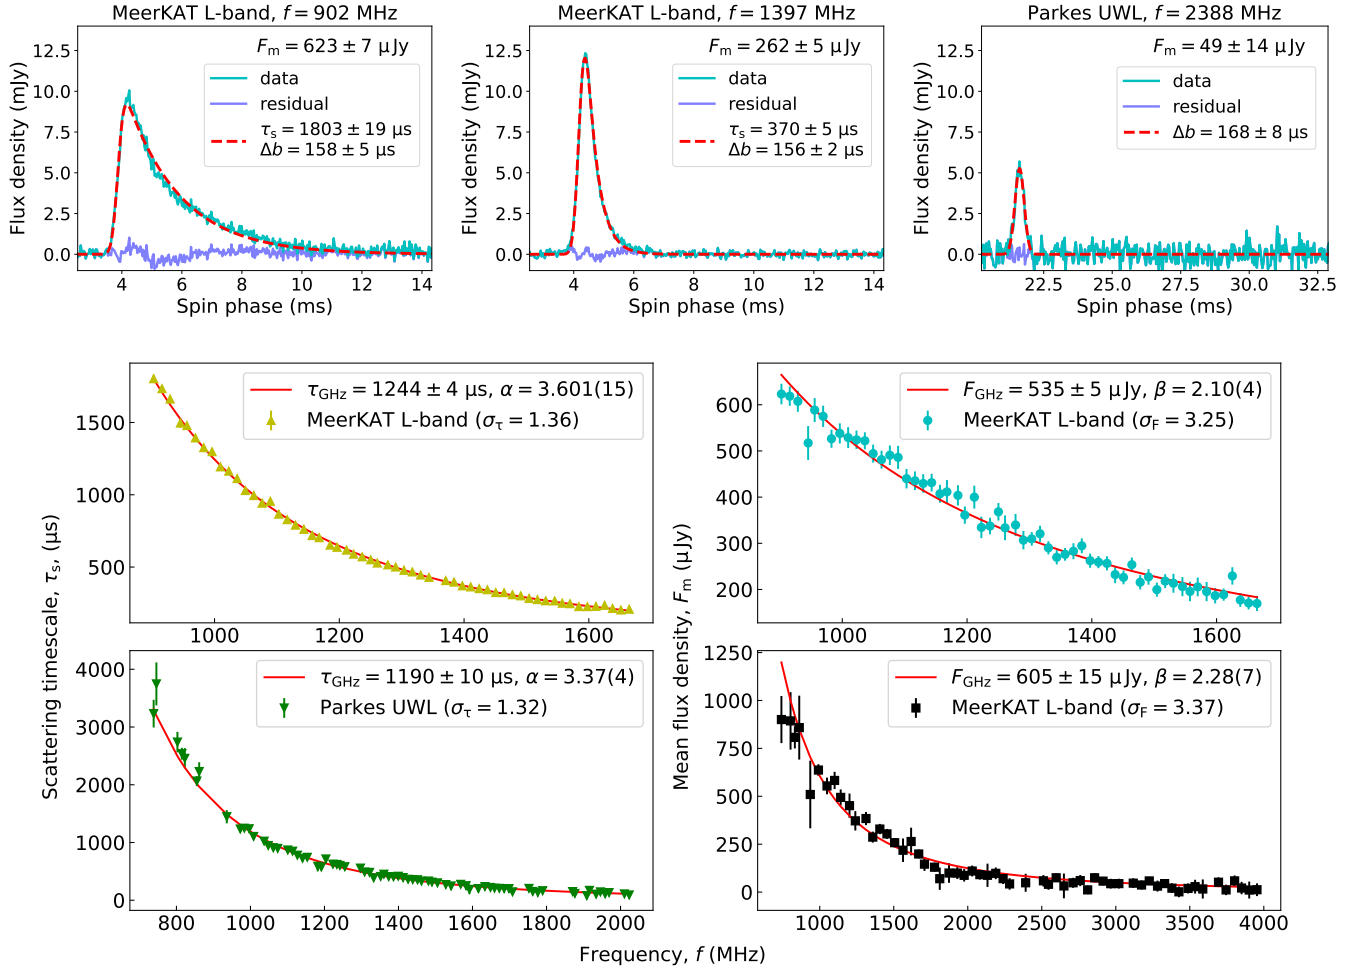
We performed least- $\chi^2$  fits of equations (1), (2), and (3) on the frequency-resolved, time-integrated MeerKAT *L*-band and Parkes UWL profiles (second and third plots in Fig. 1) with the python module *scipy* (Virtanen et al. 2020). As is shown in the upper part of Fig. 2, the scattered Gaussian modelling of the profile adjusts well to the data, with only some minor residual structure after subtraction of the model. The intrinsic pulse width stays consistent at  $150 < \Delta b < 170 \mu\text{s}$  across the whole band, and the parameter  $\tau_s$  becomes redundant in the modelling of the pulse at  $f > 2$  GHz in the Parkes UWL dataset, where a simple Gaussian function provides an adequate description. Thus, we restricted the measurement of  $\tau_{\text{GHz}}$  and  $\alpha$  at  $f < 2$  GHz. To ensure a proper uncertainty estimation on the scattering and spectral parameters, we multiplied the individual measurement uncertainties of  $\tau_s$  and  $F_m$  by the error factors  $\sigma_\tau$  and  $\sigma_F$  ( $\delta' = \sigma \times \delta$ , where  $\delta$  is the set of measurement uncertainties). Subsequently, we tuned the values of  $\sigma_\tau$  and  $\sigma_F$  to achieve a reduced  $\chi^2 = 1$  in the fit of equations (2) and (3).

Our fits confirmed the steep spectrum of J1227–6208 and measure a scattering index of  $\alpha < 4$ . The lower part of Fig. 2 shows the fit parameters for  $\tau_{\text{GHz}}$ ,  $\alpha$ ,  $F_m$ , and  $\beta$  from equations (2) and (3) on the measurements of  $\tau_s$  and  $F_m$  in the

MeerKAT *L*-band and Parkes UWL datasets. The power law describes  $\tau_s$  accurately for the MeerKAT *L*-band data set, with  $\tau_{\text{GHz}} = 1244 \pm 4 \mu\text{s}$  and  $\alpha = 3.601(15)$ . However, this is different from the measurement performed in the Parkes UWL profile, which resulted in  $\tau_{\text{GHz}} = 1190 \pm 10 \mu\text{s}$  and  $\alpha = 3.37(15)$ . The difference likely comes from the different frequency coverage and the low quality of measurements at  $f < 900$  MHz. Nonetheless, both fits agree on  $\alpha < 4$ , far from the  $\alpha = 4.4$  predicted from a Kolmogorov medium and the  $\alpha = 4$  from Gaussian inhomogeneities (Rickett 1977; Romani et al. 1986), as already seen in some pulsars behind complex ISM environments (e.g. Löhmer et al. 2004; Krishnakumar et al. 2019; Oswald et al. 2021). On the other hand, our  $F_m$  measurements are not as precise, as indicated by the  $\sigma_F > 3$  values. The most likely explanation is the lower *S/N* of the Parkes standard profiles due to its lower gain and the large degree of zapping of frequency channels during RFI excision. Nonetheless, both the MeerKAT *L*-band and Parkes UWL agree on a very steep spectrum, with  $\beta = 2.10(4)$  and  $\beta = 2.28(7)$ , respectively, on the steeper side the typically observed values of  $\beta = 1.4 \pm 1.0$  (Bates et al. 2013). The derived reference flux densities are  $F_{\text{GHz}} = 535 \pm 5 \mu\text{Jy}$  and  $F_{\text{GHz}} = 605 \pm 15 \mu\text{Jy}$ . Regarding the peak flux density, the combined effect of scattering and spectral behaviour results in the brightest observed peak flux density of 13–14 mJy -and highest *S/N* at  $f \approx 1200$ –1300 MHz.

### 3.2. Polarised emission

We performed a rotation measure (RM) fit on the MeerKAT *L*-band and Parkes UWL profiles with PSRCHIVE/rmfit, finding the maximum of integrated linearly polarised flux density at  $\text{RM} = 47.7$  and  $46.2 \text{ rad cm}^{-1}$ , with full widths at half maximums of 50 and  $30 \text{ rad cm}^{-1}$ , respectively. These two values are consistent with each other, and they imply an average magnetic field parallel to the line of sight (LOS) of  $0.16 \mu\text{G}$ . Fig. 3 shows the polarised emission of J1227–6208 at  $f = 1574$ –1670 MHz from the MeerKAT *L*-band profile, and at  $f = 2368$ –4029 MHz from the Parkes UWL profile. Approximately 15% of the total intensity is linearly polarised, and a similar fraction of circular polarisation is detected. Linear polarisation is more prominent in the first half of the pulse, peaking before the total intensity, while circular polarisation peaks along with the total intensity. The position angle (PA) of the linear polarisation shows an increase before the linearly polarised emission peaks, and a decrease afterwards, instead of following the rotating vector model from Radhakrishnan & Cooke (1969). In the MeerKAT *L*-band,



**Fig. 2.** Frequency evolution of the scattering and mean flux density, and their modelling. Top: pulse profiles of J1227–6208 at  $f = 902$  MHz and  $f = 1397$  MHz as is seen by the MeerKAT *L*-band receiver (left and middle plots), and at  $f = 2388$  MHz as is seen by the Parkes UWL receiver (right plot). The data (cyan lines) is shown against the best model (red dashes, fitted with equation 1) and the fit residuals (blue lines). The fit parameters  $\tau_s$  and  $\Delta b$  from equation (1), and the mean flux density  $F_m$  are quoted on each profile. Bottom: measurements of  $\Delta\tau$  and  $F_m$  at each frequency on the MeerKAT *L*-band and Parkes UWL profiles (dots with uncertainties), and the best fits of the power laws (2) and (3) (red lines). The data points are displayed accounting for the error factors  $\sigma_\tau$  and  $\sigma_F$ , and the relevant  $\tau_{\text{GHz}}$ ,  $\alpha$ ,  $F_{\text{GHz}}$ , and  $\beta$  values are quoted in the legend.

scattering drags some of the polarised emission into the scattering tail, resulting in some spurious PA angle measurements towards the end of the pulse. In the Parkes data set, an extra feature of the profile was revealed at the beginning of the pulse, making it slightly asymmetric. This feature is coincident with the off-centre linearly polarised emission.

## 4. Timing analysis

### 4.1. Noise modelling

We fitted the entire ToA time series with the DDH model, a modified version of the Damour-Deurelle timing model (DD, Damour & Deruelle 1986) that includes the orthometric parametrisation of the Shapiro delay (Freire & Wex 2010). The fit itself was performed with the pulsar timing software TEMPO2<sup>7</sup> (Edwards et al. 2006; Hobbs et al. 2006) and its plug-in TEMPO2/TempoNEST<sup>8</sup>, which implements multi-nested Bayesian sampling of the highly multidimensional space of the timing model, including the pulsar spin and astrometric parameters, the

DM evolution, the five Keplerian parameters, and five independent PK parameters. In addition, TEMPO2/TempoNEST also fits for Gaussian white noise and correlated red noise.


The ToA time series is affected by Gaussian instrumental noise, red spin noise, and DM noise, which can significantly contaminate the measurement of other parameters in the timing model. In particular, there is a significant DM variability as measured in different observations, an effect that becomes very prominent in the MeerKAT *L*-band and Parkes UWL data sets. Following the method first described in van Haasteren et al. (2009), we modelled the presence of red noise as a power spectral density in the Fourier domain described with a power law,

$$S(\nu) = A^2 \left( \frac{\nu}{\text{yr}^{-1}} \right)^{-\gamma}, \quad (4)$$

where  $A = \{A_{\text{spin}}, A_{\text{DM}}\}$  is the dimensionless amplitude of the correlation matrix,  $\nu$  is the spectral frequency, and  $\gamma = \{\gamma_{\text{spin}}, \gamma_{\text{DM}}\}$  is the power law index. We accounted for two possible sources of red noise: spin noise caused by the rotational variations originating within the pulsar, and DM noise originated by variations within the interstellar medium (ISM) along the LOS.

<sup>7</sup> <https://bitbucket.org/psrsoft/tempo2>

<sup>8</sup> <https://github.com/LindleyLentati/TempoNest>



**Fig. 3.** Parkes UWL (top) and MeerKAT L-band (bottom) polarisation profiles of J1227–6208, RM-corrected at  $47.3 \text{ rad cm}^{-1}$ . Total flux density (black lines), linearly polarised intensity (red lines), circularly polarised intensity (blue lines), and the PA angle (black error bars) are plotted as a function of the pulse phase.

The spin noise is frequency-independent (achromatic), while the DM noise is defined by its frequency dependence (chromatic), with

$$A_{\text{DM},f} = A_{\text{DM}} \times \left( \frac{f}{1,400 \text{ MHz}} \right)^{-2}, \quad (5)$$

where  $A_{\text{DM}}$  corresponds to the amplitude of red DM noise at  $f = 1400 \text{ MHz}$ . To account for uncorrelated Gaussian (white) instrumental and pulsar noise for each telescope backend  $k$  (Section 2), we add the  $\text{EQUAD}_k$  and  $\text{EFAC}_k$  parameters that re-scale the ToA uncertainties,  $\sigma_k$ , into

$$\hat{\sigma}_k^2 = (\text{EFAC}_k \sigma_k)^2 + \text{EQUAD}_k^2, \quad (6)$$

where  $k$  is specific to each telescope backend. As a starting point, TEMPO2/TempoNEST only requires the extra assumption of a cut-off frequency for the correlated noises as a user-given input. We performed six different TEMPO2/TempoNEST runs: half of them include white noise and red DM noise, while the other half also include red spin noise. Each of these runs has different noise cut-offs, chosen from  $\nu_{\text{cut}}^{-1} = 30, 50, 100$ , and  $500$  days. The resulting noise models, listed in Table 2, were named according to the cut-offs and inclusion or exclusion of red spin noise: loose (L), loose

with spin (Ls), constrained (C), constrained with spin (Cs), very constrained (VC), and very constrained with spin (VCs).

All of the noise models provide an adequate description of the data, including those that do not include red spin noise, and red DM noise is always dominant over red spin noise when both are included. As is shown in Table 2, models that include both noises are favoured by the nested importance sampling global log-evidence ( $\log E$ ), and  $A_{\text{DM}}$  is two orders of magnitude larger than  $A_{\text{spin}}$ . That is not surprising, as J1227–6208 is a recycled pulsar with an expected high rotational stability, and it shows that the main source of timing noise is indeed the ISM. Furthermore, Table 2 also shows that models with shorter  $\nu_{\text{cut}}^{-1}$  values are favoured by the  $\log E$ . Nonetheless, the resulting  $\chi^2 < 0.95$  values for the VC and VCs models hint towards the possibility of over-fitting when accounting for large frequencies in the red noises.

#### 4.2. Post-Keplerian parameter constraints

We fitted five PK parameters to constrain the component masses and the effects of proper motion. These are the rate of periastron advance,  $\dot{\omega}$ , the orthometric amplitude and ratio parameters of the Shapiro delay,  $h_3$  and  $\varsigma$ , as defined in Freire & Wex (2010), the orbital period derivative,  $\dot{P}_b$ , and the projected semi-major axis derivative,  $\dot{x}$ . From the preliminary timing solution, we expect  $\dot{\omega}$  to be dominated by the relativistic precession of the Keplerian orbit as predicted by GR,

$$\dot{\omega} = 3 \left( \frac{P_b}{2\pi} \right)^{-5/3} \frac{(T_\odot M_t)^{2/3}}{1 - e^2}, \quad (7)$$

where the total system mass,  $M_t$ , is expressed in units of  $M_\odot$ , and  $T_\odot = GM_\odot/c^3 = 4.92549094764 \mu\text{s}$  as defined in Prša et al. (2016). The Shapiro delay describes the periodic time delay of the pulses as they propagate through the gravitational field of the companion (Shapiro 1964). However, as its periodicity is equal to  $P_b$ , part of its signal is degenerate with the Rømer delay for systems with low inclination angles. Therefore, we used  $h_3$  and  $\varsigma$  to model the residual unabsorbed component of the Shapiro delay instead (Freire & Wex 2010). The orthometric amplitude,

$$h_3 = T_\odot M_c \left( \frac{1 - \cos i}{1 + \cos i} \right)^{3/2}, \quad (8)$$

describes the amplitude of the unabsorbed component, which depends on the companion mass,  $M_c$ , and the inclination angle,  $i$ , while the orthometric ratio,

$$\varsigma = \left( \frac{1 - \cos i}{1 + \cos i} \right)^{1/2}, \quad (9)$$

describes the shape of the delay in the orbital phase, depending only on  $i$ .

The remaining PK parameters,  $\dot{P}_b$  and  $\dot{x}$ , have the estimated GR contributions of  $\dot{P}_b \sim 10^{-15} \text{ s s}^{-1}$  and  $\dot{x} \sim 10^{-20} \text{ ls s}^{-1}$ , but much more dominant contributions are those of the Galactic acceleration field and the proper motion. For  $\dot{P}_b$ , we expected the Shklovskii effect and the Galactic acceleration field to be the dominant contributors (Shklovskii 1970; Damour & Taylor 1991), which introduce a derivative based on the Doppler factor ( $D$ ) derivatives:

$$\frac{\dot{P}_b}{P_b} = - \left( \frac{\dot{D}}{D} \right)_{\text{Shkl}} - \left( \frac{\dot{D}}{D} \right)_{\text{Gal}} = \frac{1}{c} \left[ |\boldsymbol{\mu}|^2 d + \mathbf{K}_0 (\mathbf{a}_{\text{PSR}} - \mathbf{a}_{\text{SSB}}) \right], \quad (10)$$

**Table 2.** Resulting noise models as derived by TEMPO2/TempoNEST.

Model	$\nu_{\text{cut,spin}}^{-1}$ (days)	$\nu_{\text{cut,DM}}^{-1}$ (days)	$\log E$	$\chi^2$	$\log A_{\text{spin}}$	$\gamma_{\text{spin}}$	$\log A_{\text{DM}}$	$\gamma_{\text{DM}}$
<b>L</b>	...	100	16310.61(4)	1.013	...	...	-10.02	1.148
<b>Ls</b>	500	100	16387.3(1)	0.973	-11.71	1.485	-9.99	1.320
<b>C</b>	...	50	16326.21(8)	0.983	...	...	-10.09	1.113
<b>Cs</b>	100	50	16416.80(2)	0.966	-11.75	1.499	-10.01	1.573
<b>VC</b>	...	30	16329.1(2)	0.944	...	...	-10.09	1.432
<b>VCs</b>	30	30	16430.9(5)	0.932	-11.81	1.097	-10.04	1.307

**Notes.** The best-fit reduced  $\chi^2$  is derived from a subsequent TEMPO2 fit.

**Table 3.** Constraints on the PK parameters and the component masses from the global fits.

Noise model	$\dot{\omega}$ (deg yr <sup>-1</sup> )	$h_3$ ( $\mu\text{s}$ )	$\zeta$	$\dot{P}_b$ $\times 10^{-13} \text{ s s}^{-1}$	$\dot{x}$ $\times 10^{-15} \text{ ls s}^{-1}$	$M_p$ ( $M_\odot$ )	$M_c$ ( $M_\odot$ )	$i$ (deg)
<b>L</b>	0.0171(9)	3.8(3)	0.84(4)	$0.02 \pm 2.4$	$-1.6 \pm 6.6$	1.58(11)	1.41(5)	79.2(6)
<b>Ls</b>	0.0173(8)	3.7(3)	0.84(3)	$2.9 \pm 2.3$	$-3.5 \pm 6.5$	1.59(10)	1.42(5)	78.6(7)
<b>C</b>	0.0170(9)	3.7(3)	0.85(3)	$1.0 \pm 2.6$	$-5.0 \pm 6.9$	1.53(11)	1.38(5)	79.3(7)
<b>Cs</b>	0.0171(8)	3.5(2)	0.86(3)	$3.7 \pm 2.5$	$-9.1 \pm 7.2$	1.53(10)	1.39(5)	78.4(7)
<b>VC</b>	0.0171(9)	3.6(3)	0.86(3)	$1.0 \pm 2.6$	$-3.7 \pm 7.0$	1.54(11)	1.39(5)	79.2(7)
<b>VCs</b>	0.0170(8)	3.4(2)	0.85(3)	$4.6 \pm 2.7$	$-6.5 \pm 8.0$	1.51(11)	1.38(5)	78.2(7)
<b>extended</b>	0.0171(11)	3.6(5)	0.85(5)	...	...	1.54(15)	1.40(7)	$78.7 \pm 1.2$

**Notes.** The PK parameters were measured in a global fit with TEMPO2 under the assumption of the noise models derived by TEMPO2/TempoNEST (Table 2). The component masses and the inclination angle are derived from the  $\chi^2$  mapping of DDGR solutions (Section 4.3). The **extended** row shows comprehensive uncertainties accounting for all measurements across the different noise models.

where  $\mu = (\mu_{\text{RA}}, \mu_{\text{DEC}})$  is the sky proper motion of the pulsar,  $\mathbf{K}_0$  is the Solar System barycentre (SSB) to pulsar system unit vector, and  $\mathbf{a}_{\text{PSR}}$  and  $\mathbf{a}_{\text{SSB}}$  are the Galactic acceleration field measured at the pulsar system and the SSB, respectively. On the other hand, the contribution to  $\dot{x}$  was expected to be dominated by the geometric effect introduced by the proper motion of the system on the sky (Kopeikin 1996):

$$\frac{\dot{x}}{x} = 1.54 \times 10^{-16} \text{ s}^{-1} \cot i \left( -\frac{\mu_{\text{RA}}}{\text{mas yr}^{-1}} \sin \Omega_a + \frac{\mu_{\text{DEC}}}{\text{mas yr}^{-1}} \cos \Omega_a \right), \quad (11)$$

where  $\Omega_a$  is the longitude of ascending node. Finally, we did not include the amplitude of the Einstein delay  $\gamma_E$ , as for circular systems like J1227–6208 it is expected to be highly degenerate with  $\dot{x}$ , and its inclusion would lead to the non-detection of  $\dot{x}$  instead (see Ridolfi et al. 2019 for a detailed discussion of this phenomenon).

To measure the PK parameters, we assumed the noise models found by TEMPO2/TempoNEST and performed a global re-fit of all model parameters with TEMPO2. We quote the fit values and  $1\sigma$  uncertainties as reported by the least- $\chi^2$  TEMPO2 fit, which are consistent with the ones reported by TEMPO2/TempoNEST but with slightly larger uncertainties. We chose this because the TEMPO2 fits are more consistent across noise models, because it provides a more conservative uncertainty estimate, and because it is consistent with the use of TEMPO2 in the mass measurements presented in the following sections.

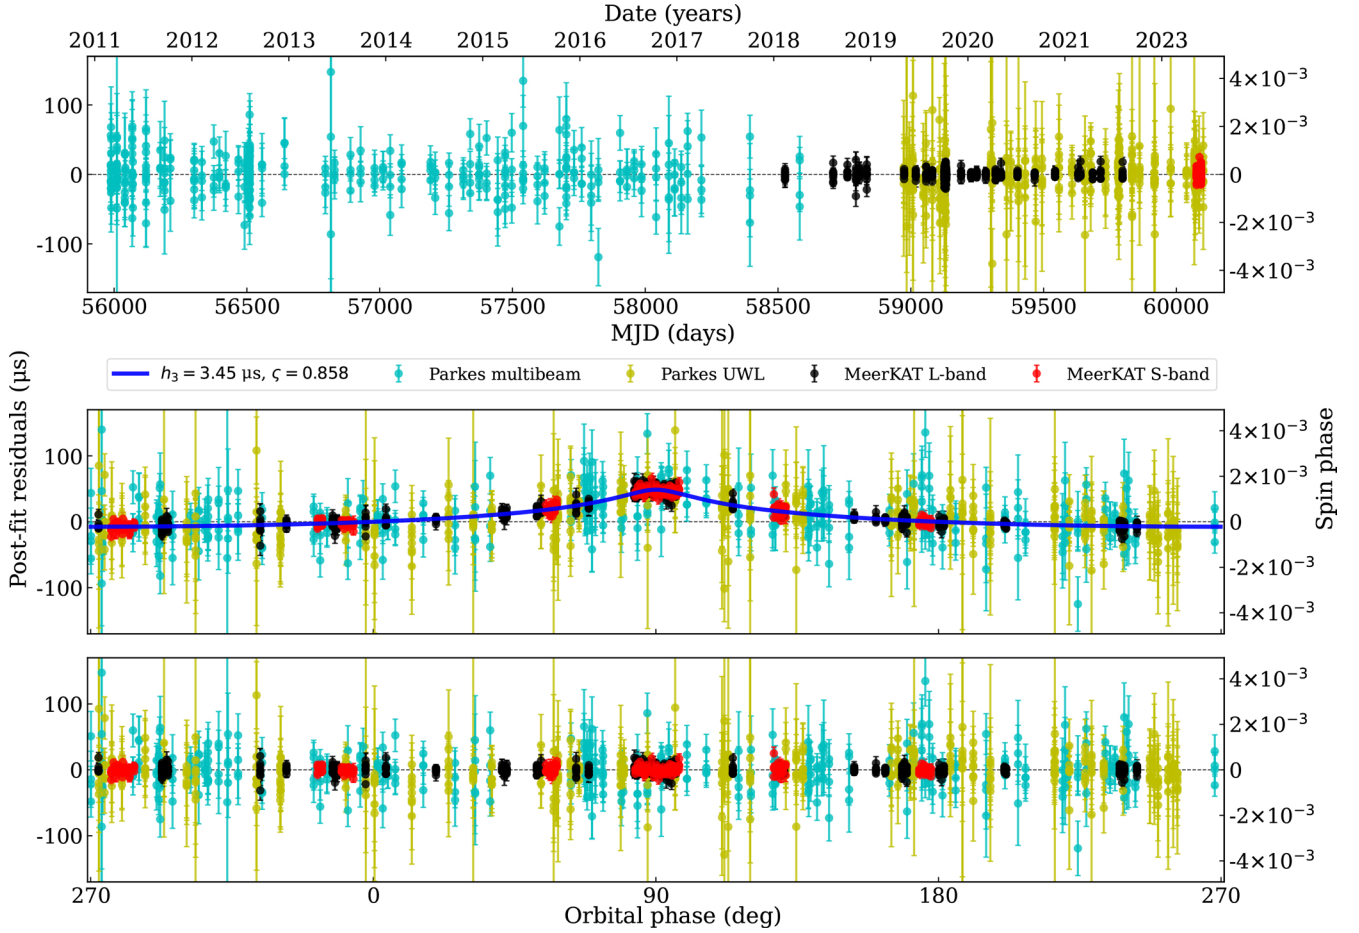
The measured PK parameters are consistent across all noise models. Table 3 presents the measured value of each PK parameter for each assumed noise model in the global fit. The rate of periastron advance presents a very consistent value of  $\dot{\omega} =$

0.0171(11) deg yr<sup>-1</sup> if we consider extended uncertainties from all of the measurements listed in Table 3, which results in a  $\sim 15\sigma$  detection and indicates a total system mass of  $M_t = 2.9 \pm 0.3 M_\odot$ . The Shapiro delay parameters  $h_3 = 3.6 \pm 0.5 \mu\text{s}$  and  $\zeta = 0.85 \pm 0.05$  were measured with high significance for the first time, presenting  $\sim 7\sigma$  and  $\sim 17\sigma$  detections if we consider extended uncertainties from across all the noise models.

The red DM noise is a major source of uncertainty in our measurements. The parameters most affected by the choice of the noise model is  $h_3$ , ranging from 3.8(3)  $\mu\text{s}$  in the **L** model to 3.4(2)  $\mu\text{s}$  in the **VCs**. In general, we measured lower values in models with high-frequency cut-offs, and with both spin and DM noise. A possible explanation is that unmodelled red noise can bias the Shapiro delay measurement to higher amplitude values, especially achromatic spin noise. A similar phenomenon can occur in the opposite direction, with noise models removing power from the Shapiro delay signal when the orbital period is close to  $\nu_{\text{cut}}$ . That may be a potential explanation for the  $\chi^2 < 0.95$  value in the **VC** and **VCs** models (Table 2). Nonetheless, the measured PK values of all parameters are consistent across all noise models within the  $1\sigma$  uncertainty ranges.

The MeerKAT data set dominates the measurement of the orthometric Shapiro delay parameters. Fig. 4 shows the timing residuals as a function of the orbital phase under the assumption of the **Cs** noise model. Despite some small gaps in orbital coverage, only the MeerKAT-derived ToAs have enough precision to detect the Shapiro delay with high significance, and therefore to constrain the component masses in this system. To corroborate this, we have attempted to fit the  $h_3$  and  $\zeta$  parameters on each ToA dataset. The MeerKAT *L*-band ToAs reproduced the measurement from the global fit, while the *S*-band ToAs yielded only a slight loss of precision owing to the sparser orbital





**Fig. 4.** Timing residuals under the assumption of the **Cs** noise model. Top: ToA residuals as a function of MJD, showing flat, Gaussian residuals. Middle: ToA residuals as a function of the orbital phase (orbital position from the ascending node, true anomaly + periastron) excluding the  $h_3$  and  $\zeta$  parameters from the timing model, showing the full amplitude of the Shapiro delay signal. The continuous blue line depicts the predicted Shapiro delay described by  $h_3 = 3.45 \mu\text{s}$  and  $\zeta = 0.858$ . Bottom: same as the middle plot, but the orthometric Shapiro delay parameters are now included in the model, resulting in flat, Gaussian residuals.

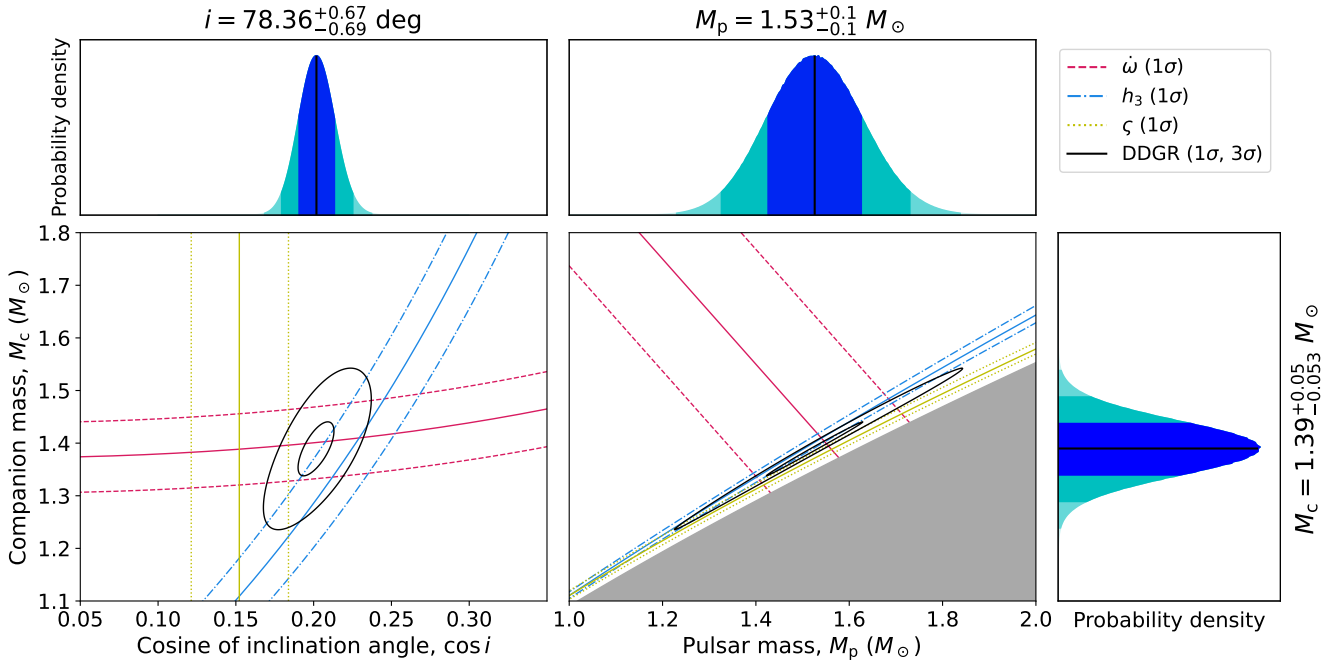
coverage. On the other hand, the Parkes multibeam data set provided only very loose constraints, as is discussed in Section 4.4, and the Parkes UWL dataset was unable to converge on significant values.

The two remaining PK parameters,  $\dot{P}_b$  and  $\dot{x}$ , did not yield a significant detection with the current timing baseline but still offer useful constraints. Table 5 presents the values for all the timing parameters of the global fit using the **Cs** noise model, as well as the expected contributions to  $\dot{P}_b$  and  $\dot{x}$ . The DM-derived distance from the NE2001 (Cordes 2004) and YMW16 (Yao et al. 2017) electron density models are  $d = 8.3$  kpc and  $d = 8.5$  kpc, respectively. With the detected proper motion vector of  $\mu = (-6.1 \pm 0.3, 0.41 \pm 0.36) \text{ mas yr}^{-1}$  and the Galactic gravitational potential from McMillan (2017) in eq. (10), we predict an order of magnitude of  $\dot{P}_b \sim 10^{-13} \text{ s s}^{-1}$ . Likewise, assuming  $i \approx 79^\circ$  (Section 4.3) in equation (11), we predict a maximum value of  $|\dot{x}_{\text{max}}| \approx 4.2 \times 10^{-15} \text{ ls s}^{-1}$  in case of a favourable  $\Omega_a$  value. These estimates are of the same order of magnitude as the constraints on  $\dot{P}_b$  and  $\dot{x}$  presented in Table 3, which also present a consistent sign across all noise models, suggesting that the constraints can provide some physical information. Taking this into account, in Section 4.5 we translate the constraint on  $\dot{x}$  into constraints on the orbital geometry of the system, and in Section 4.6

we discuss the possible implications of the  $\dot{P}_b$  constraints for the measurement of the spin period derivative  $\dot{P}_s$ .

#### 4.3. Mass constraints

In a common Bayesian approach (e.g. Splaver et al. 2002), we enforced the consistency of the PK parameters with GR and derived constraints on  $M_p$ ,  $M_c$ , and  $i$ . We use the DDGR model, a modified version of the DD model that implements the PK effects directly from  $M_t$  and  $M_c$  as predicted by GR (Taylor & Weisberg 1989). This allows us to explore a two-dimensional space of parameters to constrain the mass components and orbital inclination. For each noise model derived by TEMPO2/TempoNEST, we performed TEMPO2 fits with the DDGR model in a uniform grid in the  $M_t - \cos i$  space, where the uniform spacing on  $\cos i$  was chosen to achieve a uniform sampling of the possible orbital geometries of the system. At each point in the grid, we registered the resulting  $\chi^2$  and transformed it into a probability value, deriving a likelihood distribution. Subsequently, the distribution was integrated into marginal one-dimensional probability distributions for  $M_p$ ,  $M_c$ , and  $i$ . Given that  $\dot{P}_b$  and  $\dot{x}$  are influenced by the proper motion of J1227–6208 and the Galactic acceleration field, we included them as



**Fig. 5.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model. The constraints have been derived in the global fit and under the assumption of the Cs noise model (Table 2). Central plots: two-dimensional constraints on the  $M_c - \cos i$  and  $M_c - M_p$  spaces. The solid coloured lines represent the nominal values of the PK parameters, the dashed coloured lines their  $1\sigma$  limits, and the solid black lines the  $1\sigma$  and  $3\sigma$  limits from DDGR. The shaded grey area in the right plot is the region excluded by the mass function ( $i > 90$  deg). Corner plots: marginalised one-dimensional probability densities for  $M_p$ ,  $M_c$ , and  $\cos i$  from DDGR  $\chi^2$  mapping, showcasing the median value (solid black line) and the 31.4%, 47.4%, and 49.9% percentiles on both sides (shaded areas under the curve).

independent excess parameters that do not need to be consistent with GR.

The resulting values of  $M_p$ ,  $M_c$ , and  $i$  are listed in Table 3. They are consistent across all noise models, but there is some tension between the DDGR constraints and the  $\zeta$  parameter. Fig. 5 depicts the constraints on the  $M_c - \cos i$  and  $M_c - M_p$  spaces as derived both with the DDGR  $\chi^2$  mapping method and the constraints imposed by the independently measured  $\dot{\omega}$ ,  $h_3$ , and  $\zeta$  parameters, in both cases under the assumption of the Cs noise model. The resulting DDGR constraints are  $1\sigma$  consistent with the measurement of  $\dot{\omega}$  and  $h_3$ , and  $2\sigma$  consistent with  $\zeta$ . This is also observed with the noise models Ls, C, VC, and VCs. Only the assumption of the L noise model results in self-consistent constraints. The median values and the 68.2% percentiles of the marginal one-dimensional distributions for  $M_p$ ,  $M_c$ , and  $i$  are quoted as the measurements and their  $1\sigma$  uncertainties in Table 3, where it is seen that all noise models present  $1\sigma$  consistent constraints. We note that models with high-frequency cut-offs resulted in slightly lower mass ranges, going from  $M_p = 1.59(10) M_\odot$  and  $M_c = 1.42(5) M_\odot$  in Ls (Fig. A.1) to  $M_p = 1.51(11) M_\odot$  and  $M_c = 1.38(5) M_\odot$  in VCs (Fig. A.2). That is consistent with the variation in the  $h_3$  values across noise models as exposed in Section 4.2, and it shows how the DM noise is the main limiting factor in the precision of our mass measurements. On the other hand, the constraints on  $i$  are more sensitive to the inclusion or exclusion of red spin noise, with  $i = 79.2(7)$  deg for the L, C, and VC models and ranging from  $i = 78.6(7)$  deg to  $i = 78.2(7)$  deg from Ls to VCs models. Using extended uncertainty ranges from all the values listed in Table 3, we quote  $M_t = 2.9 \pm 0.3 M_\odot$ ,  $M_p = 1.54(15) M_\odot$ ,  $M_c = 1.40(7) M_\odot$ , and  $i = 78.7 \pm 1.2$  deg, which constitute  $10\sigma$ ,  $10\sigma$ ,  $20\sigma$ , and  $65\sigma$  constraints.

#### 4.4. Discrepancies between post-Keplerian parameters and General Relativity

As is shown in Fig. 5,  $\zeta$  is only  $2\sigma$  consistent with the DDGR constraints, and the  $h_3$  and  $\dot{\omega}$  parameters. To explore the cause of this slight discrepancy, we analysed the individual contributions to the PK parameters from segregated data sets. Keeping the noise models derived by TEMPO2/TempoNEST, we created two new isolated ToA time series: one with the ToAs derived from Parkes multibeam observations (multibeam), and another one with the ToAs derived from MeerKAT and the Parkes UWL observations (MeerKAT+UWL). These two sets differ significantly both in timing baseline and in ToA quality, with the Parkes multibeam-derived ToAs providing the longest baseline, but with the MeerKAT+UWL ToAs being much more precise and providing observing frequency information (Figs. 1 and 4). Subsequently, we repeated the TEMPO2 fits on each side.

The Shapiro delay measurement is not significantly affected by the split. In the multibeam fits, the orthometric Shapiro delay parameters were measured with larger uncertainties across all noise models ( $h_3 = 4.1\text{--}4.7 \pm 2.1 \mu\text{s}$ ,  $\zeta = 0.69\text{--}0.76 \pm 0.29$ ), but they are still consistent with the measurements from the global fit. On the other hand, the MeerKAT+UWL  $h_3$  and  $\zeta$  values listed in Table 6 are almost identical to the ones from the global fit quoted in Table 3. That is consistent with the Shapiro delay being primarily constrained by the MeerKAT ToAs owing to their higher precision (see Section 4.2).

The tension between the PK parameters was reduced in the MeerKAT+UWL fits. In the multibeam fits, the periastron advance  $\dot{\omega}$  increased to  $\dot{\omega} = 0.026(5) \text{ deg yr}^{-1}$ , which is more than  $1\sigma$  away from the global fit. Not only is this an unrealistic value, as it implies total system mass of  $M_t = 5.5 \pm 1.7 M_\odot$ , but it also suggests that the multibeam data set is somehow



**Table 4.** Data reduction, model fit, spin, DM, and astrometric parameters from the DDH TEMPO2/TempoNEST fits.

Data set	Global	Multibeam	MeerKAT+UWL
Data reduction parameters			
Number of ToAs	1545	365	1180
First ToA (MJD)	55 987.54	55 987.54	58 526.27
Last ToA (MJD)	60 101.30	58 581.61	60 101.30
Solar System ephemeris	DE430	DE430	DE430
Timescale	TCB	TCB	TCB
Shortest correlated spin noise timescale, $\nu_{\text{spin}}^{-1}$ (days)	100	100	100
Shortest correlated DM noise timescale, $\nu_{\text{DM}}^{-1}$ (days)	50	50	50
Model fit parameters			
Log correlated spin noise amplitude, $\log A_{\text{spin}}$ ( $\mu\text{s}$ )	-11.7524	-12.0384	-11.8007
Log correlated DM noise amplitude, $\log A_{\text{DM}}$ ( $\mu\text{s}$ )	-10.0075	-11.7218	-9.993
Spin noise spectral index, $\gamma_{\text{spin}}$ ( $\mu\text{s}$ )	1.49911	0.281109	1.0046
DM noise spectral index, $\gamma_{\text{DM}}$ ( $\mu\text{s}$ )	1.57333	-1.80875	1.45055
Weighted root mean square of the timing residuals ( $\mu\text{s}$ )	6.995	23.571	6.094
Reduced $\chi^2$	0.9650	1.0298	0.9643
Spin and astrometric parameters			
Reference epoch for spin, position, and DM (MJD)	55 991.22	55 991.22	58 526
Right ascension, RA (J2000, hh:mm:ss)	12:27:00.4414(4)	12:27:00.4412(4)	12:27:00.441(2)
Declination, DEC (J2000, dd:mm:ss)	-62:08:43.791(3)	-62:08:43.788(2)	-62:08:43.80(1)
Proper motion in RA ( $\mu_{\text{RA}}$ , mas yr $^{-1}$ )	-6.1(3)	-5.2(7)	-5.9 $\pm$ 1.1
Proper motion in DEC ( $\mu_{\text{DEC}}$ , mas yr $^{-1}$ )	0.41 $\pm$ 0.36	-0.4(7)	1.4 $\pm$ 1.2
Spin frequency, $F_0$ (Hz)	28.962140551274(9)	28.96214055129(2)	28.96214051691(1)
Spin frequency derivative, $F_1$ ( $10^{-16}$ Hz s $^{-1}$ )	-1.5677(4)	-1.570(2)	-1.5636(14)
Dispersion measure, $\text{DM}_0$ (pc cm $^{-3}$ )	362.816(12)	362.808(11)	362.895(7)
First derivative of DM, $\text{DM}_1$ ( $10^{-2}$ pc cm $^{-3}$ yr $^{-1}$ )	1.72(5)	2.8(9)	-1.86(9)
Second derivative of DM, $\text{DM}_2$ ( $10^{-3}$ pc cm $^{-3}$ yr $^{-2}$ )	-2.3(9)	-3.1 $\pm$ 2.1	3.1 $\pm$ 2.1
Derived parameters			
Galactic longitude, $l$ (deg)	300.082	...	...
Galactic latitude, $b$ (deg)	0.591	...	...
Spin period, $P_s$ (ms)	34.527834647774(10)	34.52783464775(2)	34.527834688742(19)
Spin period derivative, $\dot{P}_s$ ( $10^{-19}$ s s $^{-1}$ )	1.8690(5)	1.872(3)	1.8640(17)
NE2001 DM-derived distance, $d$ (kpc)	8.3	...	...
YMW16 DM-derived distance, $d$ (kpc)	8.5	...	...
Shklovskii contribution to $\dot{P}_s$ ( $10^{-19}$ s s $^{-1}$ )	$\sim 0.26$	...	...
Galactic potential contribution to $\dot{P}_s$ ( $10^{-19}$ s s $^{-1}$ )	$\sim (-0.26)$	...	...
Characteristic age, $\tau_c$ (Gyr)	2.9	...	...
Spin-down luminosity, $ \dot{E} $ (erg s $^{-1}$ )	$6.3 \times 10^{32}$	...	...
Surface magnetic field strength, $B_{\text{surf}}$ (G)	$2.6 \times 10^9$	...	...

**Notes.** The derived parameters are derived only for the global fit, the most constraining solution, and left blank for the multibeam and MeerKAT+UWL fit. The value of  $\dot{P}_s$  has not been corrected for the possible Shklovskii and Galactic potential contributions.

biased towards a larger  $\dot{\omega}$  value. However, the MeerKAT+UWL fits give values much more consistent with the global fits, with  $\dot{\omega} = 0.0163(14) \text{ deg yr}^{-1}$  from the extended uncertainties. The median value of  $\dot{\omega}$  was slightly reduced within the  $1\sigma$  uncertainties consistently across all noise models (Table 6) with only a small increase in uncertainty compared to the global fit ( $\dot{\omega} = 0.0171(11) \text{ deg yr}^{-1}$ , Table 3). With this subtle change,  $\dot{\omega}$ ,  $h_3$ , and  $\zeta$  became  $1\sigma$  consistent with each other across all noise models.

In an attempt to understand the origin of the apparent biases introduced by the multibeam data set, we performed two extra TEMPO2/TempoNEST runs with  $\nu_{\text{spin}}^{-1} = 100$  days and  $\nu_{\text{DM}}^{-1} = 50$  days on the multibeam and MeerKAT+UWL data sets, and compared the resulting noise and timing parameters with those

of the Cs model derived in the global fit (Table 2). Tables 4 and 5 show the global fit, the multibeam fit and the MeerKAT+UWL fit results side to side. In Table 4, it becomes readily evident that TEMPO2/TempoNEST converged on similar timing noise models for the global and MeerKAT+UWL data sets. Table 5 also shows that the MeerKAT+UWL PK parameter measurements are also virtually identical to those derived with the assumption of the Cs model, presented in Table 6. However, the noise model diverges in the isolated multibeam data set, and it was unable to converge on realistic values, showing that the Parkes multibeam data does not have enough timing precision and bandwidth to constrain the spin and DM timing noise. This could explain why the multibeam gives a discrepant  $\dot{\omega}$  measurement, and why it could be biasing the global fit measurements as well.

**Table 5.** Orbital Keplerian and PK parameters from the DDH TEMPO2/TempoNEST fits, and mass constraints from the DDGR  $\chi^2$  mapping fits.

Data set	Global	Multibeam	MeerKAT+UWL
Keplerian orbital parameters			
Orbital period, $P_b$ (days)	6.7210192(3)	6.72102(2)	6.7210189(4)
Orbital eccentricity, $e$	0.00114918(4)	0.0011492(4)	0.00114918(5)
Longitude of periastron, $\omega$ (deg)	27.090(9)	27.07(3)	27.211(12)
Projected semi-major axis of the pulsar orbit, $x$ (ls)	23.200666(3)	23.20066(3)	23.200665(3)
Epoch of periastron, $T_0$ (MJD)	55 991.6995(2)	55 991.6991(6)	58 525.52384(9)
Post-Keplerian (PK) orbital parameters			
Rate of periastron advance $\dot{\omega}$ (deg yr $^{-1}$ )	0.01710(83)	0.0252(48)	0.0163(11)
Orthometric amplitude of Shapiro delay, $h_3$ ( $\mu$ s)	3.45(23)	$3.5 \pm 1.8$	3.44(23)
Orthometric ratio of Shapiro delay, $\varsigma$	0.858(27)	0.76(27)	0.853(30)
Derivative of orbital period, $\dot{P}_b$ ( $10^{-13}$ s s $^{-1}$ )	$3.7 \pm 2.5$	$-10 \pm 30$	$8 \pm 40$
Estimated Shklovskii contribution to $\dot{P}_b$ ( $10^{-13}$ s s $^{-1}$ )	$\sim 4.5$	...	...
Estimated Galactic potential contribution to $\dot{P}_b$ ( $10^{-13}$ s s $^{-1}$ )	$\sim (-4.5)$	...	...
Derivative of projected semi-major axis, $\dot{x}$ ( $10^{-15}$ ls s $^{-1}$ )	$-9.1 \pm 7.2$	$-88 \pm 42$	$-17 \pm 11$
Estimated maximum proper motion contribution to $\dot{x}$ , $ \dot{x}_{\max} $ ( $10^{-15}$ ls s $^{-1}$ )	$\sim 4.2$	...	...
Mass and orbital geometry constraints			
Mass function, $f_M$ ( $M_\odot$ )	0.297	0.297	0.297
Total mass, $M_t$ ( $M_\odot$ )	$3.1^{+0.6}_{-0.8}$	$5.5 \pm 1.7$	$3.0^{+0.7}_{-1.0}$
Companion mass, $M_c$ ( $M_\odot$ )	1.39(5)	...	1.31(7)
Pulsar mass, $M_p$ ( $M_\odot$ )	1.53(10)	...	1.37(14)
Inclination angle, $i$ (deg)	78.4(7)	...	78.9(8)
Longitude of ascending node, $\Omega_a$ (deg)	$86/266^{+76}_{-78}$	...	...

**Notes.** The non-GR contributions to  $\dot{P}_b$  and  $\dot{x}$  are derived only from the global fit constraints, where the proper motion measurement is the most significant (Table 4). Mass and orbital geometry constraints are derived from the  $\chi^2$  mapping of DDGR solutions, as exposed in Sections 4.3, 4.4, and 4.5.

**Table 6.** Constraints on the PK parameters and the component masses from the MeerKAT+UWL fits.

Model	$\dot{\omega}$ (deg yr $^{-1}$ )	$h_3$ ( $\mu$ s)	$\varsigma$	$M_p$ ( $M_\odot$ )	$M_c$ ( $M_\odot$ )	$i$ (deg)
<b>L</b>	0.0163(12)	3.9(3)	0.82(4)	1.42(14)	1.33(7)	79.6(8)
<b>Ls</b>	0.0166(11)	3.7(3)	0.84(3)	1.43(14)	1.34(7)	79.2(8)
<b>C</b>	0.0163(12)	3.7(3)	0.84(4)	1.39(15)	1.32(7)	79.4(8)
<b>Cs</b>	0.0164(11)	3.4(3)	0.86(3)	1.37(14)	1.31(7)	78.9(8)
<b>VC</b>	0.0162(12)	3.7(3)	0.84(4)	1.38(15)	1.32(8)	79.4(9)
<b>VCs</b>	0.0160(11)	3.4(2)	0.85(3)	1.31(13)	1.28(7)	78.9(8)
<b>extended</b>	0.0163(14)	3.7(5)	0.84(6)	1.36(20)	1.31(10)	$79.2 \pm 1.1$

**Notes.** The PK parameters were measured in the MeerKAT+UWL fit with TEMPO2 under the assumption of the noise models as derived by TEMPO2/TempoNEST (Table 2). The component masses were derived from the  $\chi^2$  mapping of DDGR solutions (Section 4.4). The **extended** row shows comprehensive uncertainties accounting for all measurements across the different noise models.


In light of this, we repeated the  $\chi^2$  mapping of DDGR solutions with the MeerKAT+UWL data set only and with the assumption of the noise models derived in the global TEMPO2/TempoNEST fits (Table 2), deriving the  $M_p$ ,  $M_c$ , and  $i$  constraints quoted in Table 6. Similar to the global fits,  $\dot{\omega}$  and  $h_3$  dominate the constraints, but this time, the DDGR limits were  $1\sigma$  consistent with all of the PK parameters. Like in the global fits, the more constraining noise models resulted in lower mass ranges, going from  $M_p = 1.31(31) M_\odot$  and  $M_c = 1.28(7) M_\odot$  with the **VCs** noise model (Fig. A.3) to  $M_p = 1.43(14) M_\odot$  and  $M_c = 1.34(7) M_\odot$  with the **Ls** noise model (Fig. A.4). For reference, Fig. A.5 in the Appendix (A) shows the counterpart for Fig. 5, with the assumption of the **Cs** noise model on the MeerKAT+UWL data set. Nonetheless, the uncertainty ranges are still  $1\sigma$  consistent with the global fits presented in Table 3. Accounting for extended uncertainties from all measurements in

Table 6, we quote  $M_t = 2.7 \pm 0.4 M_\odot$ ,  $M_p = 1.36(21) M_\odot$ ,  $M_c = 1.30(9) M_\odot$ , and  $i = 79.3 \pm 1.1$  deg from the MeerKAT+UWL global fits, which are  $7\sigma$ ,  $6\sigma$ ,  $14\sigma$ , and  $77\sigma$  measurements, and are in  $1\sigma$  consistency with the ones quoted in Section 4.3.

Finally, we attested that  $\dot{P}_b$  and  $\dot{x}$  can only be constrained in the global fit. Both the multibeam and the MeerKAT+UWL fits were unable to constrain to  $\dot{P}_b$  and  $\dot{x}$ . As is shown in Table 5, the uncertainties on both parameters increase by an order of magnitude. Therefore, it is evident that only the accumulated baseline of the global fit can constrain them.

#### 4.5. Orbital geometry constraints

In Section 4.2, it is noted that the uncertainty in the constraints of  $\dot{x}$  across have a similar size as the  $|\dot{x}_{\max}| \approx 4.2 \times 10^{-15}$  expected from the proper motion, as quoted in Table 5. There-



**Fig. 6.** Constraints on the orbital orientation from the  $\chi^2$  mapping of DDGR solutions consistent with Eq. (11). The constraints have been derived in the global fit and assuming the Cs noise model (Table 2). Central plots: two-dimensional constraints in the  $\Omega_a - \cos i$  space. The solid contours represent the explored space ( $\cos i > 0$ ), while the dashed contours are derived from symmetries in eq. (11) ( $\cos i < 0$ ). Corner plots, top: marginalised one-dimensional probability densities for  $\cos i$ , showcasing the median value (solid black lines) and the 31.4%, 47.4%, and 49.9% percentiles (shaded areas under the curve, blue for  $\cos i > 0$  and purple for  $\cos i < 0$ ). Corner plot, right: marginalised one-dimensional probability densities for  $\Omega_a$  from  $\cos i > 0$  and  $\cos i < 0$  (solid and dashed lines, respectively) and the addition of both (shaded cyan area).

fore, we investigated how the DDGR solutions are affected upon enforcement of consistency between  $\dot{x}$ ,  $\mu$ , and the orbital geometry of the system. Assuming the noise models derived with TEMPO2/TempoNEST, we implemented the same likelihood approach from Section 4.3 with two major modifications. First, the mapping was done in a uniform three-dimensional grid on the  $M_t - \cos i - \Omega_a$  space. Second, equation (11) was implemented by forcing the excess  $\dot{x}$  value to be consistent with the  $i$ ,  $\Omega_a$ , and  $\mu = (\mu_{RA}, \mu_{Dec})$  values. To compensate for the extra dimensionality, the grid resolution was reduced to avoid a manifold increase of the computational running time of this experiment, with only 64 points along the  $M_t$  axis across  $10\sigma$ , 60 points along the  $\cos i$  across  $12\sigma$ , and 180 points across the  $0 \leq \Omega_a < 360$  range. In addition, equation (11) is symmetric with respect to the transformation  $i \rightarrow 180 \text{ deg} - i$  and  $\Omega_a \rightarrow \Omega_a + 180 \text{ deg}$ . Therefore, the  $\cos i < 0$  side of the explored space was derived by implementing this transformation. This last point is true only because the distance to J1227–6208 is large enough so that the orbital motion of Earth does not introduce year-long periodic contributions to  $\dot{x}$ . Finally, we integrated the likelihood distribution along the  $M_t$  axis to derive two-dimensional likelihood distribution on the  $\Omega_a - \cos i$  space, from which the marginal probability distributions for  $\Omega_a$  and  $\cos i$  were derived.

The constraints on  $\cos i$  and  $\Omega_a$  are drawn on Fig. 6. As expected, the constraints on  $i$  are consistent with those derived in Section 4.3. The value of  $\Omega_a$  is not constrained owing to the large overlap of the  $\cos i > 0$  and  $\cos i < 0$  probability distributions, but the probability density contours show that two regions in the  $\cos i - \Omega_a$  are preferred. For  $\cos i > 0$ , the longitude of ascending node is constrained at  $\Omega_a = 266 \pm 78 \text{ deg}$ , while for  $\cos i < 0$  it is constrained at  $\Omega_a = 86 \pm 78 \text{ deg}$ . From equation 11, using the values of  $x$ ,  $i$ ,  $\mu_{RA}$ , and  $\mu_{Dec}$  presented in Table 4 and  $\Omega_a = 266 \text{ deg}$ , we predict a value of  $\dot{x} \approx -4.2 \times 10^{-15} \text{ ls s}^{-1}$ , in

full consistency with the values presented in Table 3. Therefore, it is clear that the measured excess  $\dot{x}$  originates from the proper motion of the system in the sky.

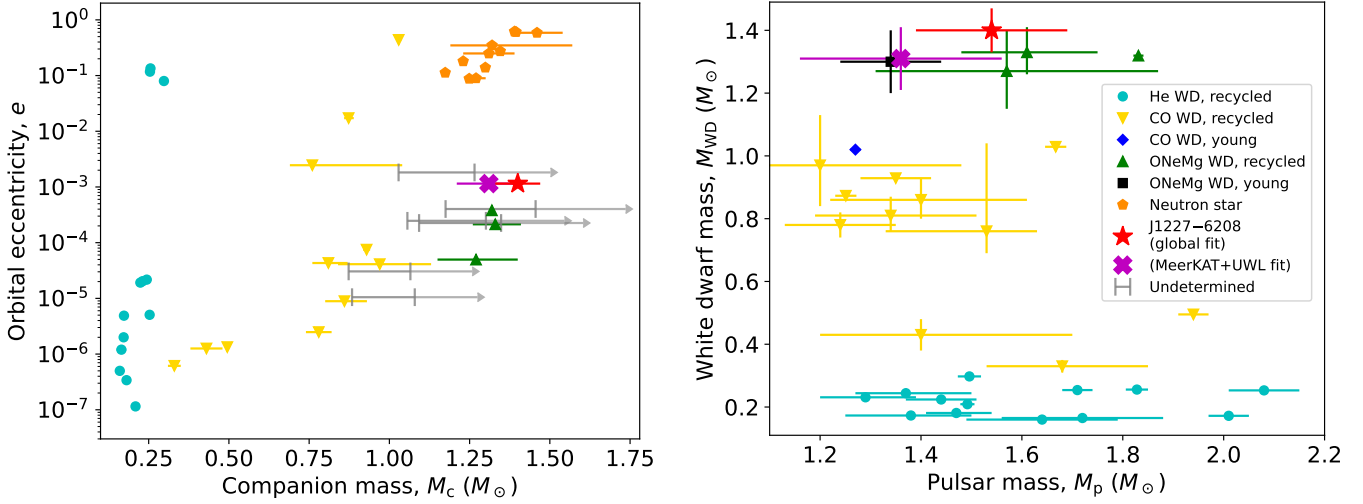
#### 4.6. Spin and orbital period evolution constraints

As is shown in Table 4, in the global fit we measured the spin evolution parameters  $P_s$  and  $\dot{P}_s$  with high significance, showing that J1227–6208 is indeed a mildly recycled pulsar. However, the Shklovskii effect and Galactic potential field affect  $\dot{P}_s$  in the same way as they affect  $\dot{P}_b$ , following equation (10). The estimated values of these contributions are listed at the end of Table 4, making it evident that albeit they are smaller than the measurement ( $\sim 10^{-20} \text{ s s}^{-1}$  against  $\dot{P}_s = 1.8690(5) \times 10^{-19} \text{ s s}^{-1}$ ), they can nonetheless be significant. Thus, we should be cautious of pulsar parameters derived from  $P_s$  and  $\dot{P}_s$ . Since a parallax distance measurement is out of the question, the best prospect for constraining the true combined contributions of the Shklovskii effect and the Galactic potential field is to improve the  $\dot{P}_b$  measurement, on which these effects are orders of magnitude larger than the GR prediction. Therefore, a measurement of  $\dot{P}_b$  will be direct measurement of the Doppler factor derivative. Nonetheless, with the current observations, the consistently positive sign of  $\dot{P}_b$  across all noise models (Table 3) already suggests a dominance of the Shklovskii effect over the Galactic potential contribution.

## 5. Astrophysical Implications

### 5.1. J1227–6208 as a DNS system

The mass ranges derived in Sections 4.3 and 4.4 and presented in Tables 3 and 6 are consistent with the mass distribution of the



**Fig. 7.** Comparison of orbital and mass constraints between Galactic pulsar–WD systems, Galactic DNS systems, and J1227–6208. For J1227–6208, the extended uncertainties from the global (red star, Table 3) and the MeerKAT+UWL fits (magenta cross, Table 6) are quoted. Left: orbital eccentricity against companion mass from recycled binary pulsars in the Galactic field, distinguished by the nature of their companions. For systems without timing mass measurements (“Undetermined” in the legend), the lower constraints on the minimum and median  $M_c$  values derived from the mass function (assuming  $i = 90, 60$  deg) are plotted. Right: WD mass against pulsar mass of PSR–WD systems in the Galactic field. References: constraints on systems besides J1227–6208 are taken from [https://www3.mpifr-bonn.mpg.de/staff/pfreire/NS\\_masses.html](https://www3.mpifr-bonn.mpg.de/staff/pfreire/NS_masses.html) and references therein, and from Camilo et al. (2001) (PSR J1435–6100), Cruces et al. (2021) (PSR J2338+4818), Edwards & Bailes (2001) (PSR J1157–5112), Gautam et al. (2022) (PSR J1952+2630), Jang et al., in prep. (PSR J1439–5501), Martinez et al. (2019) (PSR J0709+0458), Parent et al. (2019) (PSR J1932+1756), and Tan et al. (2020) (PSR J1658+3630).

double-NS (DNS) population, with  $M_c$  being significantly above the current lower mass limit of a NS ( $1.17 M_\odot$ , Martinez et al. 2015), making a NS nature a hypothesis in need of consideration.

Recycled pulsars with massive WD and NS companions follow a similar evolutionary path. With original companions more massive than  $5 M_\odot$  (Lazarus et al. 2014), these systems go through a intermediate-mass or high-mass X-ray binary (IMXB or HMXB) stage, resulting in a dynamically unstable mass transfer and the formation of a common envelope (CE, e.g. van den Heuvel 2019). After the hydrogen-rich CE is expelled, the system becomes a circular PSR–He naked star. Stable mass accretion on to the NS occurs when the companion leaves the He main sequence and the system enters a short-lived (less than 100 kyr) Case BB RLO stage, leading to the partial recycling of the pulsar (Lazarus et al. 2014). Afterwards, and if the stripped He star has retained enough mass ( $\geq 1.45 M_\odot$ ), the companion will undergo an electron-capture supernova, forming a second NS (Tauris et al. 2015, 2017). This process entails the loss  $0.2–0.4 M_\odot$  and a supernova kick of  $< 50 \text{ km s}^{-1}$  that either disrupts the system or introduces a significant orbital eccentricity ( $e \geq 0.1$ , Tauris et al. 2017). On the other hand, if the stripped companion is unable to trigger a supernova, it becomes a massive WD instead (Lazarus et al. 2014; Tauris et al. 2012).

We argue that J1227–6208 is very unlikely to be a DNS system based on its orbital eccentricity. The left plot of Fig. 7 depicts the orbital eccentricity  $e$  of recycled pulsars against  $M_c$  and it shows that, while the  $M_p$  in J1227–6208 is consistent with the DNS population, its  $e$  value is two orders of magnitude below those of known DNS systems. Instead, it lays among other massive recycled PSR–CO/ONeMg WD systems. While Tauris et al. (2017) show that very low eccentricities ( $e < 0.01$ ) in DNS are technically possible, they also argue that this scenario is extremely unlikely, requiring an extraordinary fine-tuning of the supernova kick magnitude and direction. That is strong evidence against the companion of J1227–5936 having undergone a supernova, making it a massive WD instead.

## 5.2. J1227–6208 as a massive PSR–WD system

J1227–6208 thus belongs to an emerging class of massive recycled PSR–ONeMg WD systems. The right side of Fig. 7 maps measured PSR–WD component masses in a  $M_c - M_p$  diagram, showing that J1227–6208 lies among the well-studied massive systems PSR J2222–0137, PSR J1528–3146, and J1439–5501, represented by the green triangles. J1227–6208 also shares with these three systems similar spin properties ( $P_0 = 34.52 \text{ ms}$  against 32.82, 60.82, and 28.64 ms), binary period ( $P_b = 6.72 \text{ days}$  against 2.45, 3.18, and 2.12 days), and orbital eccentricity ( $e = 1.15 \times 10^{-3}$  against  $4.65 \times 10^{-4}$ ,  $2.13 \times 10^{-4}$ , and  $4.99 \times 10^{-5}$ ), suggesting a similar nature. Therefore, it is very plausible that J1227–6208 has followed the evolutionary path of recycled PSR–ONeMg WD systems ( $M_{WD} \gtrsim 1.1 M_\odot$ ). It should be noted that the young system PSR B2303+46, shown as a black square in Fig. 7, is the result of an exotic evolution where the WD formed before the pulsar and thus no recycling has occurred in it (van Kerkwijk & Kulkarni 1999; Tauris & Sennels 2000).

As we argue in favour of a WD nature for the companion, the possibility of  $M_c$  laying beyond the Chandrasekhar limit (Section 4.3) is brought into question. Theoretical models predict that fast-rotating WDs may exist in the  $1.38 < M_{WD}/M_\odot < 1.48$  mass range without collapsing in a supernova (e.g. Yoon & Langer 2005). However, these conditions can only be the result of accretion and the spin-up in the WD, or be the product of a WD merger. Both of these scenarios are difficult to reconcile with the evolutionary model proposed above and the low orbital eccentricity, and therefore the mass of the companion of J1227–6208 is likely lower than  $1.38 M_\odot$ . Continued timing in the future will be essential to increase the precision of the mass constraints so that a more definitive statement can be made on this aspect.

J1227–6208 and its companion are the fourth partially recycled pulsar with a  $M_{WD} > 1.2 M_\odot$  ONeMg WD com-



panion with mass measurements from timing. As is seen on the right plot of Fig. 7, an apparent gap in  $M_{\text{WD}}$  values is arising at  $M_{\text{WD}} \approx 1.1 M_{\odot}$ , in addition to the already-known separation between massive CO WD systems ( $0.7 \lesssim M_{\text{WD}} \lesssim 1.0 M_{\odot}$ ) and the fully recycled millisecond pulsars with light He and CO WD companions ( $M_{\text{WD}} \lesssim 0.5 M_{\odot}$ ) McKee et al. (2020), Shamohammadi et al. (2023), implying a trimodality in the  $M_{\text{WD}}$  distribution. However, such gap is not predicted by theory and the left plot of Fig. 7 shows that a handful of known massive systems without mass measurements can potentially fill this gap: PSR J1435–6100 (Camilo et al. 2001), PSR J2338+4818 (Cruces et al. 2021), PSR J1157–5112 (Edwards & Bailes 2001), PSR J1952+2630 (Jang et al., in prep.), PSR J0709+0458 (Martinez et al. 2019), PSR J1932+1756 (Parent et al. 2019), and PSR J1658+3630 (Tan et al. 2020). Thus, this gap could be spurious, arising from the reduced size of the sample of mass measurements in massive systems. New mass measurements and the discovery and follow-up of further massive partially recycled PSR–ONeMg WD systems will be key in confirming whether the  $M_{\text{WD}}$  distribution is trimodal or if this notion is just a statistical fluke.


### 5.3. The pulsar mass of J1227–6208

Theoretical models for recycled pulsars with the most massive WD companions suggest that these systems went through a short phase of accretion during post-CE Case BB RLO from their naked He-star companion, resulting in observed spin-up periods of  $P_s \gtrsim 20$  ms. According to Tauris et al. (2012), the amount of mass accreted by a recycled pulsar given a  $P_s$  can be estimated with the following equation:

$$\Delta M_p \approx 0.22 M_{\odot} \frac{(M_p/M_{\odot})^{1/3}}{(P_s/\text{ms})^{4/3}}. \quad (12)$$

Assuming its current spin period of  $P_s \approx 34.5$  ms results in the accretion of  $0.0022 M_{\odot}$  from its companion during recycling. On the other hand, assuming a possible initial spin-up period of  $P_s \approx 20$  ms results in the accretion of at most  $0.0045 M_{\odot}$  from its companion. In both cases, the implication is that the current  $M_p$  measurement is close its birth mass. Therefore a more precise measurement of  $M_p$  will contribute to the statistics of NS birth masses, like in the similar systems PSR J2222–0137, PSR J1534–3146, and PSR J1439–5501. In fact, the measurement of  $M_p = 1.76(6) M_{\odot}$  in PSR J2222–0137 by Cognard et al. (2017) ( $M_p = 1.831(10) M_{\odot}$  in Guo et al. 2021) has already expanded such distribution on its upper end.

If the amount of accreted matter estimated from theory is correct, the mass-transfer rate was likely Eddington or super-Eddington. For a NS of  $1.5 M_{\odot}$ , the Eddington luminosity limit is  $L \approx 1.5 \times 10^{38} \text{ erg s}^{-1}$ , and the accretion-to-luminosity efficiency is  $L \approx 0.1 \times \dot{M} c^2$  (Shakura & Sunyaev 1973; Poutanen et al. 2007), resulting in a maximum accretion rate of  $\dot{M} \approx 0.03 M_{\odot} \text{ Myr}^{-1}$ . Simulations find that the duration of the accretion stage during the Case BB RLO is  $\lesssim 100\,000$  yr for massive PSR–WD systems (Lazarus et al. 2014; Tauris et al. 2015, e.g.), and Cognard et al. (2017) even found that the accretion time was  $\lesssim 20\,000$  yr for the massive PSR–ONeMg WD system in PSR J2222–0137. The Eddington limit would allow for the accretion of at most  $0.003 M_{\odot}$  in the span of  $100\,000$  yr, and  $5 \times 10^{-4} M_{\odot}$  in  $20\,000$  yr, implying that if the recycling process occurred on a similar time span for J1227–6208, then it has accreted either close to the Eddington limit or significantly above it like in PSR J2222–0137. This is a plausible scenario, as



**Fig. 8.** ToA uncertainty against observing frequency for each telescope receiver. The highest and lowest frequencies of the Parkes multibeam and MeerKAT S-band data sets are biased upwards because the signal drops off at the edge of the bands (see Fig. 1).

super-Eddington accretion onto NSs has been directly observed in extra-galactic ultra-luminous X-ray sources such as NGC 5907 X-1 (Israel et al. 2017), NGC 300 ULX1 (Carpano et al. 2018), and M82 X-2 (Bachetti et al. 2014). However, an alternative explanation is that the spin-up efficiency is larger than what is predicted by Tauris et al. (2012), requiring less mass accretion to reach a similar  $P_s$ .

## 6. Future prospects

### 6.1. Prospects for timing

Future timing observations of J1227–6208 should seek to strike a balance between its steep spectrum and the ISM-induced effects. Fig. 8 portrays how ToA uncertainty evolves with  $f$  across different receivers, and it shows that scattering affects the ToA quality negatively at  $f < 1300$  MHz, leading to an increase of timing uncertainty. The best ToA quality is found at the  $f = 1.3$ – $2.2$  GHz range, while at  $f > 2200$  MHz the steep spectrum starts hampering the  $S/N$ , leading to another increase in timing uncertainty. Another factor to take into account is that, as reported in Section 4.1, red DM noise is very prominent, making it preferable to observe at high frequencies where its amplitude is lower. However, doing so comes at the price of achieving a sub-optimal  $S/N$  and ToA quality due to the steep spectral index reported in Section 3. Therefore, MeerKAT observations with the S-band receivers at the S0 configuration (1750–2625 MHz) will offer the best compromise between a high ToA precision and a reduced ISM influence, with  $A_{\text{DM},1700} = 0.68 \times A_{\text{DM}}$  and  $A_{\text{DM},2600} = 0.28 \times A_{\text{DM}}$ , thus enabling a precision increase on the measurement of PK parameters. In complement to this, UWL observations will be essential to track the DM evolution to properly model the still-present DN noise in the MeerKAT S-band ToAs.

The significance of the PK parameters will improve with an increasing baseline. The most significant improvement in the mass constraints will come from an improved measurement of  $\dot{\omega}$ . Assuming that the uncertainty on  $\dot{\omega}$  scales with the timing baseline according to the power law (Damour & Taylor 1992)

$$\Delta \dot{\omega}' \approx \Delta \dot{\omega} \times \left( \frac{T}{4.31 \text{ yr}} \right)^{-3/2}, \quad (13)$$

where 4.31 yr is the timing baseline of the MeerKAT+UWL data set, we expect  $\Delta\dot{\omega}' \approx \Delta\dot{\omega}/2$  with just two extra years of timing. This improvement will reduce the total mass measurement uncertainty by 65% if we assume the current MeerKAT ToA uncertainty, aiding a more robust settlement or confirmation of any discrepancies between the PK parameters in the global fit.

Tighter constraints on  $\dot{P}_b$  and  $\dot{x}$  will also lead to an independent estimate of the distance of J1227–6208 and a better constraint on its orbital geometry. However, since these measurements are dependent on the global 11.26 year baseline, reducing uncertainty will require more time. Assuming once again a power law evolution of the uncertainties  $\dot{P}_b$  (Damour & Taylor 1992)

$$\Delta\dot{P}_b' \approx \Delta\dot{P}_b \times \left( \frac{T}{11.26 \text{ yr}} \right)^{-5/2}, \quad (14)$$

then another decade of timing will be required to achieve a  $3\sigma$  detection of  $\dot{P}_b$ . Constraints on the orbital geometry will also require an improvement of the measurement of the proper motion. Following the power law evolution of the uncertainty on  $\dot{x}$  (Damour & Taylor 1992)

$$\Delta\dot{x}' \approx \Delta\dot{x} \times \left( \frac{T}{11.26 \text{ yr}} \right)^{-3/2}, \quad (15)$$

two more decades of timing will be necessary to achieve a  $3\sigma$  detection on  $\dot{x}$ . However, as is seen in Section 4.5, constraints on the  $\Omega_a - \cos i$  can already be achieved even with a low-significance detection. Furthermore, the proper motion vector measurement will continue to improve alongside it, leading to more significant constraints in the  $\Omega_a - \cos i$  space before that time.

## 6.2. Prospects for an optical detection

There would be obvious benefits to having an optical detection of the companion of J1227–6208. Besides the confirmation of its WD nature, a measurement of its colour index and magnitude could be used to constrain its age with the help of WD evolutionary models. This would provide an independent measurement of the mass and age of the companion of J1227–6208 like other massive PSR–WD systems, such as PSR J1528–3146 (Jacoby et al. 2006) and PSR J1439–5501 (Pallanca et al. 2013). That would allow, for example, for a true estimate the initial spin-up period.

As of now, we are unable to report an optical detection. *Gaia* DR3 is sensitive to sources with magnitude  $G \approx 20.7$  (Gaia Collaboration 2016, 2023). Assuming the YMW16-derived DM distance of  $d = 8.5$  kpc, detection would require an absolute magnitude of at most  $M_G \approx 6$  to be detected, too bright for any WD. The closest reported *Gaia* DR3 source<sup>9</sup> is 0.29 arcsec away from the timing position, with a colour index of  $B - R = 1.59$  and a magnitude of  $G = 17.4$ . Its parallax distance is  $d = 5.3 \pm 2.0$ , giving it an absolute magnitude of  $M_G = 3-5$ , and putting it either on the main sequence or the giant branch. Therefore, this source is most likely not associated with J1227–6208.

A potential detection will be possible only with the most powerful optical and infrared telescopes available. Assuming  $d = 8.5$  kp, we use the cooling tracks provided by Bergeron et al. (2011)<sup>10</sup> for a  $M_{WD} = 1.3 M_\odot$  with a He atmosphere to predict

visible and infrared magnitudes given an age. If the age of the system is consistent with the characteristic age of  $\tau_c = 2.93$  Gyr, then the companion will have a optical, red, and infrared magnitudes of  $V, R, I > 30$  and it will be undetectable. But if, like in PSR J1439–5501, the true age is significantly younger than  $\tau_c$  for instance only 100 Myr old (Pallanca et al. 2013), the magnitudes could reach  $V, R, I < 27$ , making is potentially detectable in long exposure by sensitive optical and infrared telescopes such as the Very Large Telescope, the *Hubble* Space Telescope or the *James Webb* Space Telescope.

## 7. Conclusions

This study adds J1227–6208 to the rare class of confirmed mildly recycled pulsars with massive ONeMg WD companions ( $M_{WD} > 1.1 M_\odot$ ), along with PSR J2222–0137, PSR J1528–3146, and PSR J1439–5501. In addition, we have provided a detailed study of the pulsar's emission and the origin of the PK parameters measured in pulsar timing. The study includes 33 hours of Parkes multibeam data since 2012, 72 hours of Parkes ULW data since 2020, 25 hours of MeerKAT data since 2019 and 19 hours of MeerKAT S-band data from 2023, producing a timing baseline of more than 11 years.

Taking advantage of the large frequency coverage offered by these receivers to study the profile evolution of J1227–6208, we measure a steep flux density spectral index of  $2.06 < \alpha < 2.35$  and a flat scattering index  $3.33 < \beta < 3.62$ , with a scattering timescale of  $\tau_s \approx 1.2$  ms and a mean flux density of  $F_m = 530-620 \mu\text{Jy}$  at the reference frequency of 1 GHz. Due to the combined effect of scattering and spectral properties, the largest peak flux density is observed at the frequency of 1.2–1.3 GHz, with  $F_{\max} = 13-14$  mJy. Around 15% of the emission of J1227–6208 is linearly polarised, with a behaviour that does not resemble the rotating vector model, and a further  $\sim 15\%$  of the emission is circularly polarised.

The timing of J1227–6208 suffers from red spin noise and, more prominently, red DM noise. We have implemented a Bayesian nested sampling algorithm to measure the parameters of the DDH timing model, including the spectral modelling of red DM noise, red spin noise, and white noise, and the timing model parameters. We derived six different noise models: three with correlated spin noise and three without it, using different frequency cut-offs in all cases. We find that the red DM noise amplitude at 1400 MHz is almost two orders of magnitude larger than the red spin noise amplitude, and that models with both correlated spin and correlated DM noise with high-frequency cut-offs are favoured by the sampling evidence. Nonetheless, all of the resulting PK parameters are consistent across models within  $1\sigma$  uncertainty. The Shapiro delay orthometric parameters,  $h_3$  and  $\zeta$ , and the periastron advance  $\dot{\omega}$  are dominated by the MeerKAT-derived ToAs, while the constraints of the orbital period derivative  $\dot{P}_b$  and the projected semi-major axis derivative  $\dot{x}$  are possible only with the accumulated timing baseline.

We argue that the constraints on  $\dot{P}_b$  and  $\dot{x}$  are physical despite their low significance.  $\dot{P}_b$  presents a consistent sign across all noise models and is about the same order of magnitude as the combined prediction of the Shklovskii effect given by the DM distance and the proper motion vector and the Galactic acceleration fields ( $\dot{P}_b \sim 10^{-13} \text{ s s}^{-1}$ ), therefore implying that a more significant detection will occur in the following years. This detection will also clarify the magnitude of the Shklovskii effect and the Galactic acceleration on the spin period derivative  $\dot{P}_s$ . The situation is similar for  $\dot{x}$ , as we predict  $|\dot{x}_{\max}| \approx 4.2 \times 10^{-15} \text{ ls s}^{-1}$  for from the proper motion vector,  $i = 79$  deg and a favourable

<sup>9</sup> Designation: *Gaia* DR3 6054663745667894016.

<sup>10</sup> Tables available here: <https://www.astro.umontreal.ca/~bergeron/CoolingModels/>



longitude of ascending node angle  $\Omega_a$ , which is of the same order of magnitude as the timing constraints. The mapping of solutions that implement consistency with the excess  $\dot{x}$ , the proper motion,  $i$  and  $\Omega_a$  indicate that the resulting  $\chi^2$  is indeed sensitive to the location in the  $\Omega_a - \cos i$  grid, with a preference for either  $\Omega_a = 266 \pm 78$  deg and  $i = 79$  deg, or  $\Omega_a = 86 \pm 78$  deg and  $i = 101$  deg. This parameter is also likely to get a more significant detection in the following years.

With the noise models from the DDH fits, we measure  $M_p$ ,  $M_c$ , and  $i$  by performing a mapping of solutions under the assumption of GR across all noise models, and attest that they are consistent with the PK parameters. In the global fit, the  $h_3 = 3.6 \pm 0.5 \mu\text{s}$  and  $\dot{\omega} = 0.0171(11) \text{ deg yr}^{-1}$  measurements are  $1\sigma$  consistent with the assumption of GR, while  $\zeta = 0.85 \pm 0.05$  is  $2\sigma$  consistent with GR, resulting in  $M_p = 1.54(15) M_\odot$ ,  $M_c = 1.40(7) M_\odot$ , and  $i = 78.7 \pm 1.2$  deg in a conservative uncertainty range that includes all noise models. However, if we exclude the Parkes multibeam observations from the analysis, the median of the periastron advance is reduced to  $\dot{\omega} = 0.0163(14) \text{ deg yr}^{-1}$ , resulting in  $1\sigma$  consistency between the three PK parameters and GR, and the mass constraints of  $M_p = 1.36(20) M_\odot$ ,  $M_c = 1.31(10) M_\odot$ , and  $i = 79.2 \pm 1.1$  deg. All of these measurements are within  $1\sigma$  uncertainty overlap, and they result in the total ranges of possible values of  $2.3 < M_t/M_\odot < 3.2$ ,  $1.16 < M_p/M_\odot < 1.69$ ,  $1.21 < M_c/M_\odot < 1.47 M_\odot$ , and  $75.5 < i/\text{deg} < 80.3$ . This wide range of possible values demonstrates that DM variability and its modelling as red DM noise is the main limiting factor of the precision of our timing analysis.

Despite the companion mass uncertainty range extending into the regime of known NS masses, the probability of it being a NS is extremely low due to the very small orbital eccentricity ( $e = 1.15 \times 10^{-3}$ ), which is much more consistent with the massive recycled PSR–CO/ONeMg WD population. The  $M_{\text{WD}}$  measured in this work is significant as it overlaps with the companion masses measured for three other similar systems. This measurement explores the potential multi-modality of the  $M_{\text{WD}}$  distribution noted in previous works, which suggests that the  $M_{\text{WD}}$  distribution is not continuous, but characterised by discrete clumps. Whether WDs with  $M_{\text{WD}} \gtrsim 1.2 M_\odot$  are segregated from their  $0.7 \lesssim M/M_\odot \lesssim 1.0$  counterparts or not will be decided by further measurements of systems in these mass ranges. Regarding the pulsar mass, given the relatively large spin period of 34 ms, J1227–6208 has likely accreted no more than  $0.0045 M_\odot$  during the recycling process, that a precise measurement of  $M_p$  will contribute to the NS birth mass distribution. However, if accretion occurred in a time period significantly shorter than 100,000 years like in the case of PSR J2222–0137 (Cognard et al. 2017), this would imply either super-Eddington accretion or a larger spin-up efficiency than predicted in Tauris et al. (2012).

Future timing of this system will be key for constraining the PK parameters and the mass constraints even further. The most precise ToAs are produced in the 1.3–2.2 GHz frequency range due to the balance between the steep spectral index and scattering. However, the presence of red DM noise requires that future observations are taken at higher frequencies, making the 1.7–2.6 GHz MeerKAT S-band receiver on the S0 configuration the best observing band in the future. Optical detection of the WD companion may be possible in long exposures with the most powerful optical/infrared telescopes if the system is younger than a few hundred million years, leading to an independent confirmation of the nature of the system, or otherwise to a lower limit of its true age.

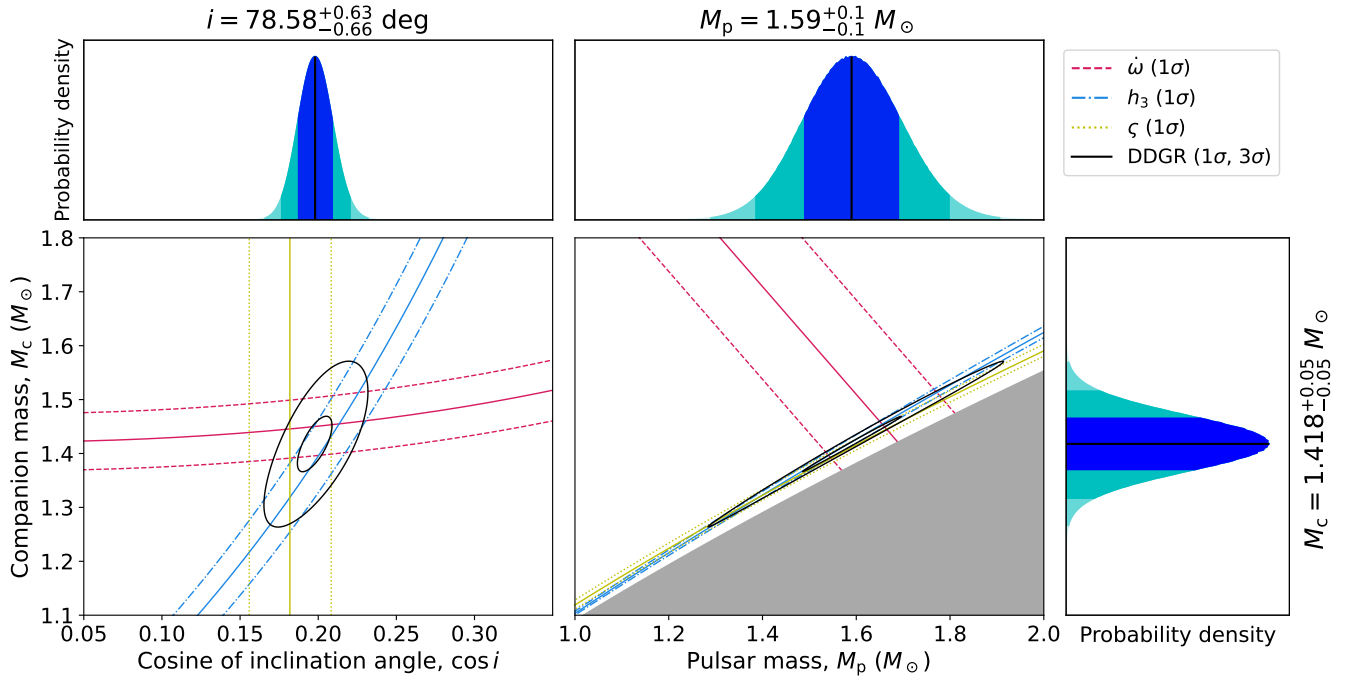
**Acknowledgements.** The MeerKAT telescope is operated by the South African Radio Astronomy Observatory (SARAO), which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. SARAO acknowledges the ongoing advice and calibration of GPS systems by the National Metrology Institute of South Africa (NMISA) and the time space reference systems department of the Paris Observatory. The Parkes radio telescope is part of the Australia Telescope National Facility (ATNF), which is funded by the Australian Government for operation as a National Facility managed by the Commonwealth Scientific and Industrial Research Organisation (CSIRO). We acknowledge the Wiradjuri people as the Traditional Owners of the Observatory site. This work used the OzSTAR national facility at Swinburne University of Technology. OzSTAR is funded by Swinburne University of Technology and the National Collaborative Research Infrastructure Strategy (NCRIS). All authors affiliated with the Max Planck Institute for radioastronomy (MPIfR) acknowledge the continuing valuable support from the Max-Planck Society (MPG). Vivek Venkatraman Krishnan acknowledges financial support from the European Research Council (ERC) starting grant ‘COMPACT’ (grant agreement number: 101078094). Alessandro Ridolfi is supported by the Italian National Institute for Astrophysics (INAF) through an ‘IAF - Astrophysics Fellowship in Italy’ fellowship (Codice Unico di Progetto: C59J21034720001; Project ‘MINERS’). We acknowledge the team of P789 and P1032 Parkes observers and the SARAO operators for the data acquisition. We give our sincere thanks to the anonymous referee and Dr. Huanchen Hu for their comments on this article.

## References

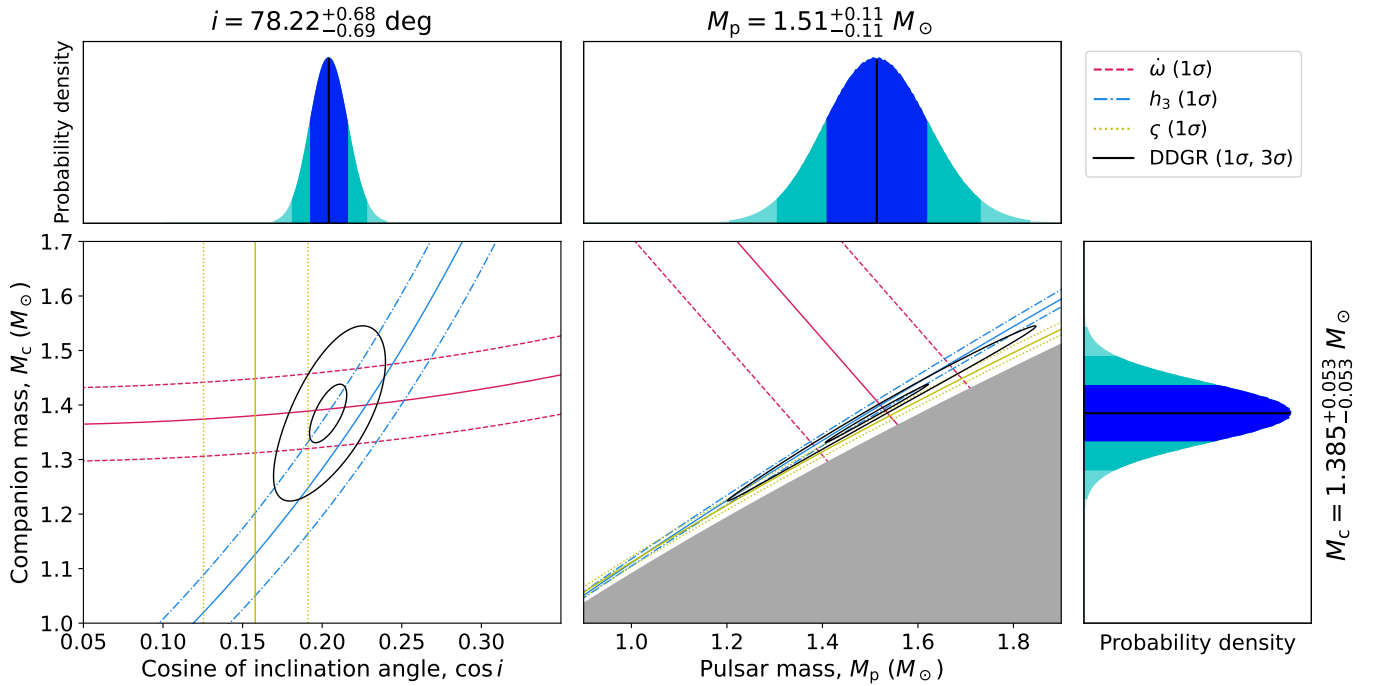
- Bachetti, M., Harrison, F. A., Walton, D. J., et al. 2014, *Nature*, **514**, 202  
 Bailes, M., Jameson, A., Abbate, F., et al. 2020, *PASA*, **37**, e028  
 Barr, E. D. 2018, *IAU Symp.*, **337**, 175  
 Bates, S. D., Lorimer, D. R., & Verbiest, J. P. W. 2013, *MNRAS*, **431**, 1352  
 Bates, S. D., Thornton, D., Bailes, M., et al. 2015, *MNRAS*, **446**, 4019  
 Bergeron, P., Wesemael, F., Dufour, P., et al. 2011, *ApJ*, **737**, 28  
 Berthreau, A., Guillemot, L., Freire, P. C. C., et al. 2023, *A&A*, **674**, A71  
 Boyles, J., Lynch, R. S., Ransom, S. M., et al. 2013, *ApJ*, **763**, 80  
 Caiazzo, I., Burdge, K. B., Fuller, J., et al. 2021, *Nature*, **595**, 39  
 Camilo, F., Lyne, A. G., Manchester, R. N., et al. 2001, *ApJ*, **548**, L187  
 Carpano, S., Haberl, F., Maitra, C., & Vasilopoulos, G. 2018, *MNRAS*, **476**, L45  
 Cognard, I., Freire, P. C. C., Guillemot, L., et al. 2017, *ApJ*, **844**, 128  
 Cordes, J. M. 2004, *ASP Conf. Ser.*, **317**, 211  
 Cruces, M., Champion, D. J., Li, D., et al. 2021, *MNRAS*, **508**, 300  
 Damour, T., & Deruelle, N. 1986, *Ann. Inst. Henri Poincaré Phys. Théor.*, **44**, 263  
 Damour, T., & Taylor, J. H. 1991, *ApJ*, **366**, 501  
 Damour, T., & Taylor, J. H. 1992, *Phys. Rev. D*, **45**, 1840  
 Deller, A. T., Boyles, J., Lorimer, D. R., et al. 2013, *ApJ*, **770**, 145  
 Edwards, R. T., & Bailes, M. 2001, *ApJ*, **553**, 801  
 Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, **372**, 1549  
 Faulkner, A. J., Stairs, I. H., Kramer, M., et al. 2004, *MNRAS*, **355**, 147  
 Fonseca, E., Cromartie, H. T., Pennucci, T. T., et al. 2021, *ApJ*, **915**, L12  
 Freire, P. C. C., & Wex, N. 2010, *MNRAS*, **409**, 199  
 Gaia Collaboration (Prusti, T., et al.) 2016, *A&A*, **595**, A1  
 Gaia Collaboration (Vallenari, A., et al.) 2023, *A&A*, **674**, A1  
 Gautam, T., Freire, P. C. C., Batrakov, A., et al. 2022, *A&A*, **668**, A187  
 Gregoris, D., & Ong, Y. C. 2023, *Ann. Phys.*, **452**, 169287  
 Guo, Y. J., Freire, P. C. C., Guillemot, L., et al. 2021, *A&A*, **654**, A16  
 Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655  
 Hobbs, G., Manchester, R. N., Dunning, A., et al. 2020, *PASA*, **37**, e012  
 Hollands, M. A., Tremblay, P. E., Gänsicke, B. T., et al. 2020, *Nat. Astron.*, **4**, 663  
 Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, **21**, 302  
 Hu, H., Kramer, M., Wex, N., Champion, D. J., & Kehl, M. S. 2020, *MNRAS*, **497**, 3118  
 Israel, G. L., Belfiore, A., Stella, L., et al. 2017, *Science*, **355**, 817  
 Jacoby, B. A., Chakrabarty, D., van Kerkwijk, M. H., Kulkarni, S. R., & Kaplan, D. L. 2006, *ApJ*, **640**, L183  
 Jacoby, B. A., Bailes, M., Ord, S. M., Knight, H. S., & Hotan, A. W. 2007, *ApJ*, **656**, 408  
 Jonas, J., & MeerKAT Team 2016, *MeerKAT Science: On the Pathway to the SKA*, **1**  
 Kaplan, D. L., Boyles, J., Dunlap, B. H., et al. 2014, *ApJ*, **789**, 119  
 Knispel, B., Eatough, R. P., Kim, H., et al. 2013, *ApJ*, **774**, 93  
 Kopeikin, S. M. 1996, *ApJ*, **467**, L93  
 Kramer, M., Stairs, I. H., Venkatraman Krishnan, V., et al. 2021, *MNRAS*, **504**, 2094  
 Krishnakumar, M. A., Maan, Y., Joshi, B. C., & Manoharan, P. K. 2019, *ApJ*, **878**, 130  
 Külebi, B., Jordan, S., Nelan, E., Bastian, U., & Altmann, M. 2010, *A&A*, **524**, A36

- Kundu, A., & Mukhopadhyay, B. 2012, *Mod. Phys. Lett. A*, **27**, 1250084
- Lazarus, P., Tauris, T. M., Knispel, B., et al. 2014, *MNRAS*, **437**, 1485
- Löhmer, O., Mitra, D., Gupta, Y., Kramer, M., & Ahuja, A. 2004, *A&A*, **425**, 569
- Lorimer, D. R., & Kramer, M. 2005, *Handbook of Pulsar Astronomy* (Cambridge University Press)
- Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, *MNRAS*, **372**, 777
- Martinez, J. G., Stovall, K., Freire, P. C. C., et al. 2015, *ApJ*, **812**, 143
- Martinez, J. G., Gentile, P., Freire, P. C. C., et al. 2019, *ApJ*, **881**, 166
- Mathew, A., & Nandy, M. K. 2021, *R. Soc. Open Sci.*, **8**, 210301
- McKee, J. W., Freire, P. C. C., Berezina, M., et al. 2020, *MNRAS*, **499**, 4082
- McMillan, P. J. 2017, *MNRAS*, **465**, 76
- Mickaliger, M. B., Lorimer, D. R., Boyles, J., et al. 2012, *ApJ*, **759**, 127
- Miller, D. R., Caiazzo, I., Heyl, J., et al. 2023, *ApJ*, **956**, L41
- Oswald, L. S., Karastergiou, A., Posselt, B., et al. 2021, *MNRAS*, **504**, 1115
- Özel, F., & Freire, P. 2016, *ARA&A*, **54**, 401
- Pallanca, C., Lanzoni, B., Dalessandro, E., et al. 2013, *ApJ*, **773**, 127
- Parent, E., Kaspi, V. M., Ransom, S. M., et al. 2019, *ApJ*, **886**, 148
- Poutanen, J., Lipunova, G., Fabrika, S., Butkevich, A. G., & Abolmasov, P. 2007, *MNRAS*, **377**, 1187
- Prša, A., Harmanec, P., Torres, G., et al. 2016, *AJ*, **152**, 41
- Pshirkov, M. S., Dodin, A. V., Belinski, A. A., et al. 2020, *MNRAS*, **499**, L21
- Radhakrishnan, V., & Cooke, D. J. 1969, *Astrophys. Lett.*, **3**, 225
- Rickett, B. J. 1977, *ARA&A*, **15**, 479
- Ridolfi, A., Freire, P. C. C., Gupta, Y., & Ransom, S. M. 2019, *MNRAS*, **490**, 3860
- Romani, R. W., Narayan, R., & Blandford, R. 1986, *MNRAS*, **220**, 19
- Sarkissian, J. M., Carretti, E., & van Straten, W. 2011, in *Am. Inst. Phys. Conf. Ser.*, **1357**, 351
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, 337
- Shamohammadi, M., Bailes, M., Freire, P. C. C., et al. 2023, *MNRAS*, **520**, 1789
- Shapiro, I. I. 1964, *Phys. Rev. Lett.*, **13**, 789
- Shklovskii, I. S. 1970, *Sov. Astron.*, **13**, 562
- Splaver, E. M., Nice, D. J., Arzoumanian, Z., et al. 2002, *ApJ*, **581**, 509
- Staveley-Smith, L., Wilson, W. E., Bird, T. S., et al. 1996, *PASA*, **13**, 243
- Tan, C. M., Bassa, C. G., Cooper, S., et al. 2020, *MNRAS*, **492**, 5878
- Tauris, T. M., & Sennels, T. 2000, *A&A*, **355**, 236
- Tauris, T. M., & van den Heuvel, E. P. J. 2023, *Physics of Binary Star Evolution. From Stars to X-ray Binaries and Gravitational Wave Sources* (Princeton University Press)
- Tauris, T. M., Langer, N., & Kramer, M. 2012, *MNRAS*, **425**, 1601
- Tauris, T. M., Langer, N., & Podsiadlowski, P. 2015, *MNRAS*, **451**, 2123
- Tauris, T. M., Kramer, M., Freire, P. C. C., et al. 2017, *ApJ*, **846**, 170
- Taylor, J. H., & Weisberg, J. M. 1989, *ApJ*, **345**, 434
- Thorsett, S. E., & Chakrabarty, D. 1999, *ApJ*, **512**, 288
- Thorsett, S. E., Arzoumanian, Z., McKinnon, M. M., & Taylor, J. H. 1993, *ApJ*, **405**, L29
- Tomaschitz, R. 2018, *Phys. Stat. Mech. Appl.*, **489**, 128
- van den Heuvel, E. P. J. 2019, *IAU Symp.*, **346**, 1
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, *MNRAS*, **395**, 1005
- van Kerkwijk, M. H., & Kulkarni, S. R. 1999, *ApJ*, **516**, L25
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Voisin, G., Cognard, I., Freire, P. C. C., et al. 2020, *A&A*, **638**, A24
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, **835**, 29
- Yoon, S. C., & Langer, N. 2005, *A&A*, **435**, 967
- Zhao, J., Freire, P. C. C., Kramer, M., Shao, L., & Wex, N. 2022, *CQG*, **39**, 11LT01

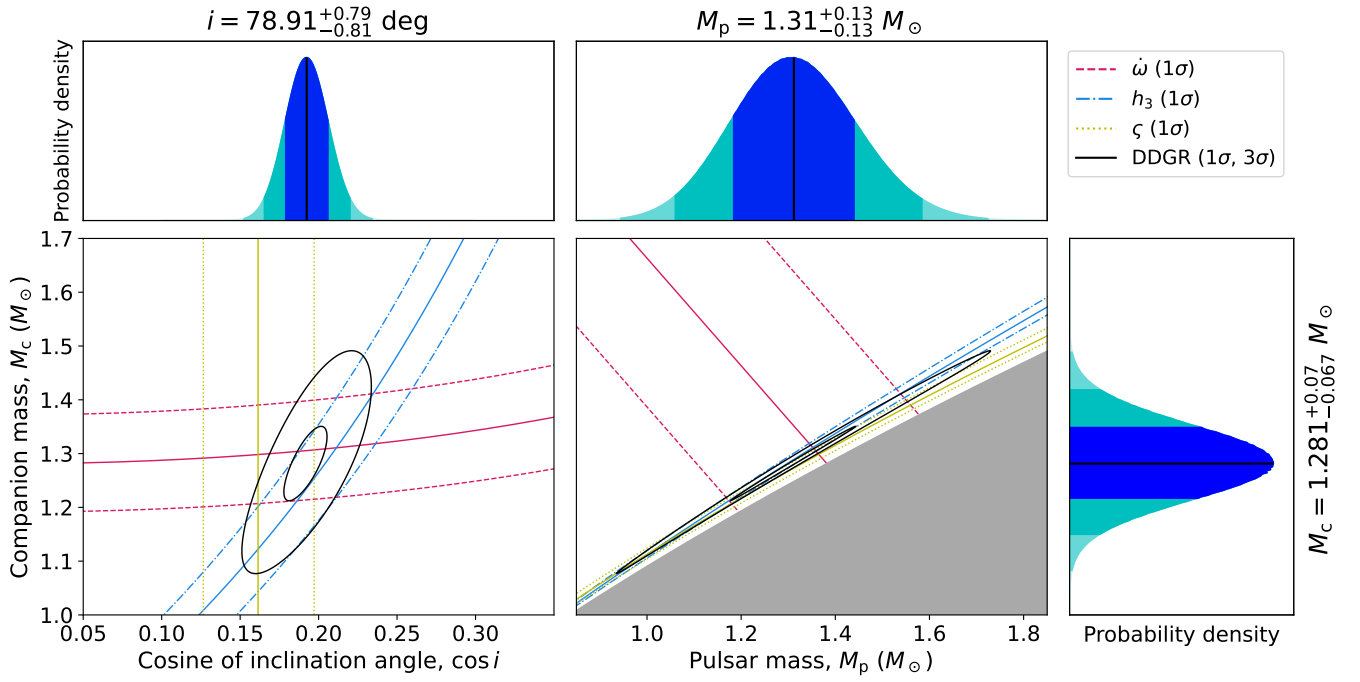
## Appendix A: Additional mass constraints plots




**Fig. A.1.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model, both from the global fit and under the assumption of the **Ls** noise model (Table 2). Contours have been drawn following the same principles as in Fig. 5.



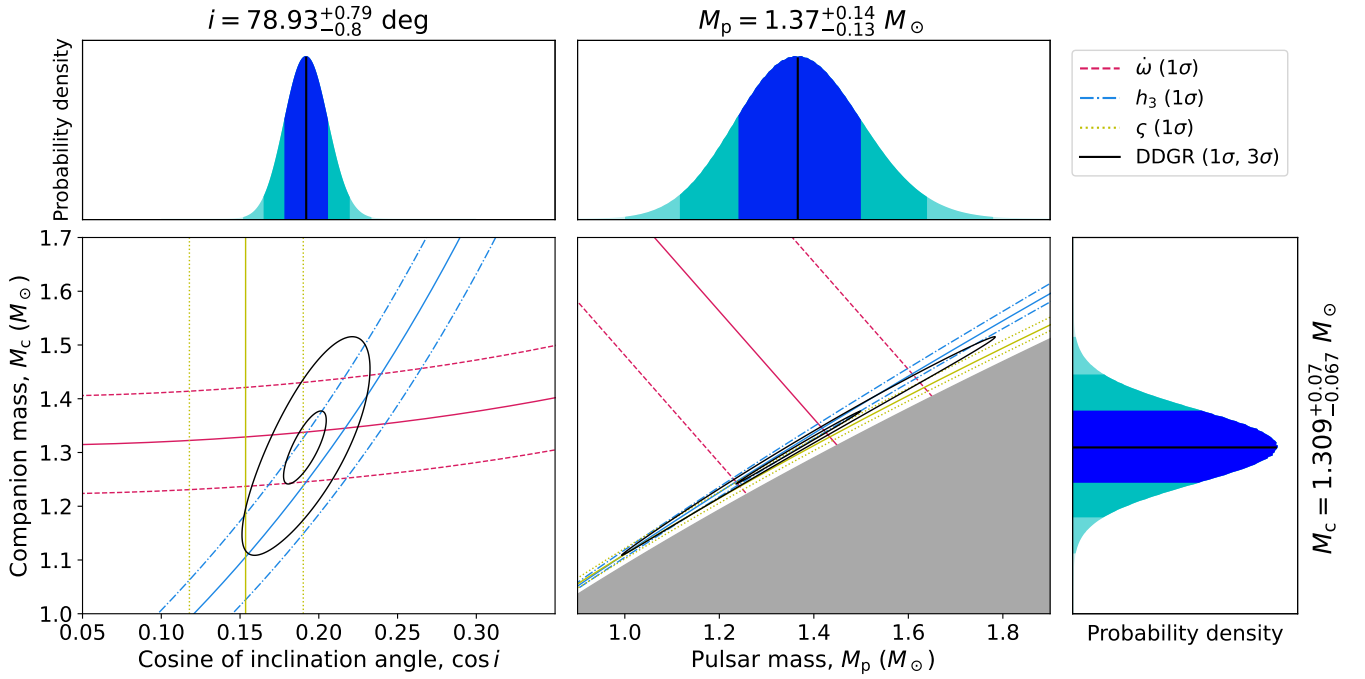
**Fig. A.2.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model, both from the global fit and under the assumption of the **VCs** noise model (Table 2). Contours have been drawn following the same principles as in Fig. 5.



**Fig. A.3.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model, both from the MeerKAT+UWL fit and under the assumption of the VC noise model (Table 2). Contours have been drawn following the same principles as in Fig. 5.



**Fig. A.4.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model, both from the MeerKAT+UWL fit and under the assumption of the Ls noise model (Table 2). Contours have been drawn following the same principles as in Fig. 5.



**Fig. A.5.** Mass and inclination angle constraints from the DDH PK measurements and the  $\chi^2$  mapping with the DDGR model, both from the MeerKAT+UWL fit and under the assumption of the Cs noise model (Table 2). Contours have been drawn following the same principles as in Fig. 5.





---

## Music to spin to

---

One may think that pulsars have no personality, that there is no way to relate to them as a human being. I thank that, like with everything else, we assign our own meaning to everything that we discover in life, even if we are not aware of it. But in the case of astronomical objects, we are never short of meaning. If we assign meaning to constellations, I can also assign my own meaning to pulsars.

This Appendix has no scientific content at all. It is just me sharing my thoughts and feelings about the pulsars I have discovered in the MMGPS surveys. In particular, I want to share the music I was listening to when a pulsar appeared on my CandyJar screen. The inspection of a batch of candidates always takes a few minutes at the very least, so it was my habit to put on my headphones when the time to scroll through CandyJar came upon. This means that, from my point of view, each pulsar has a set of feelings and personality assigned by the music playing through my headphones all the time. An let me tell you, I have a lot to say about that. In this appendix, I list the song that was playing during the discovery of *every pulsar*, and share some meaning about it. It is deeply personal, so ignore it if you do not feel like reading it.

### D.1 Music to spin to, the list

**PSR J1306–6043:** *Stickerbush Symphony*, by David Wise, from *Donkey Kong Country 2 OST*. Stickerbush Symphony has to be one of the most iconic pieces of videogame soundtrack ever made. It plays for the first time in the *Bramble Blast* level of *Donkey Kong Country 2*, released in 1995 for the SNES. In this game, you infiltrate the Kremlings' island to rescue Donkey Kong. This island is a dark and perilous swampland, so when I reach the heights of the world map and am hit by this smooth melody in a bright, sky-themed level, I feel an unexpected peace. The peace and tranquility it transmits is so powerful, that it became a tradition for people to share their life experiences in the comment sections of videos where this song is played. In particular, experiences where one got out of rough personal situations and is looking towards the future with hope. I would say it is extremely fitting for the first discovery brought in by the MMGPS survey: after thousands of noise plots, RFI signals and unconfirmed T2 candidates, **PSR J1306–6043** appeared like a bright, high- $S/N$  beacon of hope.

**PSR J1353–6341:** *Nausicaä of the Valley of Wind, opening sequence*, by Joe Hisaishi. *Nausicaä*, directed by Hayao Miyazaki, is *the* Studio Ghibli movie before Studio Ghibli was founded. The movie was so successfully that when it released in 1984 it gave the team the resources to create the official

Studio Ghibli that we all know and love today. It tells a story about humanity repeating the mistakes of the past, and it does so with an **opening sequence** that depicts a ravaging war that destroyed the world in ancient times, where this OST plays. The song is slow and beautiful, as if the apocalypse was just an ancient tale, and is on good pace with the simple lives that the descendants of the survivors enjoy, where the remains of the disaster are no more than a few forgotten ruins. But in doing so, it gives a hint towards the much grander and powerful past that precedes the time when the movie is set. I found that to be extremely fitting for the gentle **PSR J1353–6341**, a beautifully aged pulsar with a 2-s spin period, which lives in tranquility in start constraint to the supernova explosion that gave it birth hundredths of millions of years before our time.

**PSR J1208–5936: *Cara Mia Addio*, by Mike Morasky (composer) and Ellen McLain (singer), from Portal 2.** At first glance, this song is just a joke **played before the end credits** of the videogame Portal 2, released in 2011. And it really is just that: an opera sung by robot turrets, rather dumb, comic-relief enemies that appear in various part of the game. But they also sing goodbye to the player as they leave the facility they have been trapped in for years. It is emotive and evoking. I like to think of pulsars this way: they are lumps of dead matter, but they still sing to us. **PSR J1208–5936** in particular was singing a beautiful binary melody that was calls for us, waiting to be discovered and solved.

**PSR J1449–6339: *Hyrule Castle, exterior*, by Manaka Kataoka, Hajime Wakai, and Yasuaki Iwata, from the Legend of Zelda Breath of the Wild OST.** Released in 2017, the legend of Zelda Breath of the Wild is the most calm Zelda experience out there. It is set in a mostly unpopulated Hyrule where the player is given more freedom to explore at their leisure than ever. The OST of the game reflects this by providing **subdued ambient music** that takes a step back and blends in with the sounds of wind, birds and water streams, unlike in other entries where it provided a grandiose sense of epic adventure. That is, with one big exception. You see, the whole game revolves around exploring the tranquility of the world to power up your character at your pace, but with the ever-present objective to storm Hyrule castle and fight the evil within. After months of enjoying this game, when you finally decide you have explored enough and believe yourself to be strong enough, you decide to finally brave-up **and tackle the final dungeon**. Then, as you cross the main gates, the ambience turns dark and hostile, and this song starts playing as you push through hordes of enemies. Everything has lead to this point, this is no longer exploring around, it is the moment of truth, the final end-game, and the makers of the game made this song as a clear statement. It gives a clear sense of urgency and purpose. And in a stroke of geniality, after the minute 1:20, this theme even remixes the **original Legend of Zelfda theme** from 1986 by the legendary Koji Kondo, and it is the only occasion in which the classic Zelda theme appears through the whole game. **PSR J1449–6339** is an energetic young pulsar which happened to be discovered at the same time I was beating Breath of the Wild, so it will forever be associated with this theme.

**PSR J1708–4843: *Alpine Incline*, by David Wise, from Donkey Kong Country: Tropical Freeze OST.** The man himself, David Wise, came back as the main composer for Donkey Kong Country: Tropical Freeze in 2014 after its last contribution in Donkey Cong Country 2. This game's OST is so good I used to play it in playlists as I looked for pulsars. It would be hard to pic one single theme that is the best one, but the one playing during the **level of Alpine Incline** was the one to win the pulsar lotery. This theme is special because it changes as your progress through the level. You start at the base of a mountain under a rather moody weather, and climbing up until you are bouncing on aerostatic balloons over the cloud tops at the end of the level. The music also begins with moody tune, but it becomes more colorful and lively as you get closer to the heights, ending with a floaty joyful

tune that rewards you for your effort. In the final parts of the level, the aerostatic balloons follow predictable “orbits” around other platforms, which is not unfitting given the binary nature of **PSR J1708–4843**. In a way, I like to think of this song as it rewarding our efforts for looking for new pulsars, with the bright music rewarding for the discoveries after all that effort.

**PSR J1134–6207: *Lords of Iron*, by Antti Martikainen.** Everyone has had a metal musical period. I discovered this song during that time. While I no longer listen to much music of this stile, this one song I keep coming back to because it motivates me to push forward. It transmits the strength and determination necessary to complete a job. It is also almost 10-minutes long, making it ideal for the inspecting large batches of candidates. Because of this, I associate **PSR J1134–6207** with a will of iron. Many pulsars are made from the collapse of the iron cores of their progenitor stars, so it is not unfitting at all.

**PSR J1108–6329: *Cloudy Court Galaxy*, by Mahito Yokota, from Super Mario Galaxy 2 OST.** The Super Mario Galaxy Games are amongst the best-reviewed in history. That is for good reason: they took the famous plumper out of the Mushroom Kingdom and brought him to space. In the games, Mario goes through levels made of tiny planets and platforms in space, jumping from one to the next via slingshots that launch him into brief-lasting orbits. There is nothing like those slingshot transitions, they give a sense of freedom and wonder that is, in my opinion, yet to be matched by any other videogame. But if there is one thing that gave these games the last push into excellence, it was the OST. In a bold choice, Nintendo went fully orchestral, and it was the most fitting decision, as now an orchestra follows Mario every time he flies between planets. Super Mario Galaxy 2 is slightly more lighthearted than the first entry. The first game took the space theming to a heart, having for example *levels set in battle station in orbit around a realistic planet or in a cluster of molten asteroids between two colliding planets*. On the other hand, Super Mario Galaxy 2 has more abstract levels like *Fluffy Bluff Galaxy* or *Cloudy Court Galaxy*, where Mario goes through impossible, clearly man-made gardens floating in an infinite sky-blue expanse filled with elegant fluffy clouds. If taken seriously, when seeing these levels one should ask: “Why is there a sky without a planet? Why are there gardens floating in it? Who paved these grounds and made these bush sculptures? Why are there giant musical instruments floating around?” But the developers likely never had that in mind. Instead, they focused on making a grandiose and whimsical journey about riding the winds and jumping over clouds. And in that, they were successful, especially thanks to the great music that accompanies this. In my opinion, it fits **PSR J1108–6329** perfectly, a binary pulsar that traces an elegant orbit around its companion, just like Mario’s jumps between the green gardens amongst the clouds.

**PSR J1231–5929: *Voyager*, by Key Puncher.** This song belongs to *synthwave*, a micro-genre of retro electronic music that is inspired by the OSTs of action and science-fiction movies from the 1980’s. It is, therefore, a nostalgia genre, although a weird kind of nostalgia. Rather than a glorious past, it reminds us of the imaginative universe of the 1980s childhood of their composers, when clunky computers, flying cars, neon lights and sunglasses were depicted as wonders of the future in analog TVs. These songs really do indeed remind us of a better, wondrous time from before the chores of adulthood. And even though I was born in 1996, I somehow resonate with it<sup>91</sup>. I discovered this genre in the first year of my PhD, thanks to the recommendations of my dear college and friend Jonah Wagenveld, so I will forever associate this genre with my time in Bonn and, as a consequence, with my early exploration of the pulsar world. I think that is very fitting. Indeed, the discovery of pulsars came

<sup>91</sup> This particular song played by chance, but my favourite song of this genre would be *Decades* by Skjalg A. Skagen. Nonetheless, many synthwave songs are fitting for **PSR J1231–5929**

with radio, with a boom in the 1970s and 1980s. Pulsars were a wondrous discovery extremely fitting with the boom of analog media, another sign of the greatness that the brand new technologies were bringing us. And, let's not lie, being a periodic radio sourced discovered by the first radio telescopes, they are the most *retro* and analogue objects of astronomy out there, right on par to time reminisced by synthwave. I would say that **PSR J1231–5929** is a most-classical synthwave pulsar, with a spin period of 410 ms and a narrow profile, it would have fit right into the boom of pulsar astronomy in the 1970s and 1980s, making this song extremely fitting for it.

**PSR J1352–6141: Lake Lamode, by Shiho Fujii.** I will be honest: I think it is a bit of a pity this pulsar got assigned this piece of OST. Super Mario Odyssey, released in 2017, is a videogame with an excellent OST evoking a sense of adventure, epicness, wonder, and exploration. I would have much rather been listening to, for example, *Fossil Falls* when this pulsar appeared on my screen, featuring a vibrant orchestra for a *level made of vibrant grasslands, rocky cliffs, giant waterfalls, and majestic dinosaur fossils*. But alas, it was Lake Lamode instead, a soft, relaxing melody for a *level set on a whimsical, sparkling lake deep in a gorge under an aurora-filled night sky*. It is nice and well-crafted, but not what I would associate with a 4.74-ms MSP spitting radiation all around. Or most importantly, it is not a song that evokes any significant emotions on me. However, if one thing can be saved it would be its orbit. Its orbit of 87 days could be seen like a slow dance the stars are participating in. Perhaps, it would not be unfitting to see these stars from afar in the skies of Lake Lamonde. Perhaps the auroras could even come from the pulsar winds.

**PSR J1543–5439: Parsec, by Scann-Tec.** Now, this is a real piece of music. It is an electronic piece belonging to the genre of *PsyDub*. This type of music is designed to put you in a psychedelic trance. When I listen to it, I feel transported to another world thanks to its layers and layers of shifting beats and motifs. This genre is even inspired by the more mystical aspects of astronomy and cosmology: many of its playlists on youtube are decorated with astral imagery, evoking to the connection between the individual and the Cosmos. It also means it is damned good for long session of studying, coding, and writing about astronomy. Parsec is an excellent representation of it, even the song of the name is on the nose. I found this song as part of a playlist of equally high quality songs of the same genre called *Gravitational Balance*. Just listening to Parsec and all of the other songs meshed in this masterpiece of a playlist makes me *feel* the thousands of spinning pulsars and their orbits that exist in the Milky Way, eternal an unending. And that could have not been any more fitting for an wide-orbit MSP like **PSR J1543–5439**.

**PSR J1413–5936: Audiotool Day 2016, by Xtract.** We go back to synthwave from this one. However, this is an evolved, matured kind of synthwave. In particular, it could classify as *chillwave* or *dreamwave*, which focus on an ethereal sense of relaxation or “chill”, with nostalgia for an abstract childhood time. The first composers of synthwave drew nostalgia for the 1980's media, but over time the younger composers have grown somewhat detached from these roots, and started drawing from their own 1990's or even 2000's experiences. And I think that AT Day 2016, composed in 2016 (duh) is a great example of that. Somehow, this song is somewhat more positive and has another *feel* to it. It is also very fitting for **PSR J1413–5936**. Synthwave always goes well with pulsars, but this one in particular is extremely favoured by it. Perhaps, it is for the feeling of *déjà vu* I have when thinking about it.

**PSR J1652–5154: Memory Upload, Tonebox.** This is another synthwave classic, belonging perhaps to chillwave. This song is omnipresent in many playlists. Somehow, it also ends up always appearing in my automatic playlists. It is standard modern synth-wave, designed to make you fly with your senses. Its second half is very subdued compared to its first half, making me feel adrift in an

ocean of nostalgia, and it ends with a very nice fade-out. That is fitting for a polite, a most-classical 600-ms drifting canonical pulsar like **PSR J1652–5154**.

**PSR J1015–5358: Monorail**, by Gereon. More syntwave pleasure for my ears. However, this one is a bit different. It lies closer to the genre of *darkwave*, which evokes darker, more sinister emotions. If classic synthwave was born from nostalgia from the imagined future in the science fiction of 1980's, this sub-genre draws from its darkest, most dystopian fantasies. Think of *Blade Runner* or the *Matrix*. This song fills me with purpose. It makes me imagine myself as a cyborg secret agent in a dark, futuristic, neon-lit city ruled by ruthless internet corporations. And I am not talking about the convenient internet we have nowadays, but about the evil “net” from 1990's fictions, where evil, power-thirsty sentient viruses sought to escape and conquer to the real world in every corner. **PSR J1015–5358** does not quite deserve this reputation, but it is a **bad boy** nonetheless. It is a recycled pulsar in a 8-hour long orbit around what could be one of the most massive companions we know of. It oozes dystopian synthwaves vibes from both of its magnetic poles.

**PSR J0917–4413: ERENthe標**, by SawanoHiroyuki, from the *Attack on Titan Season 2 OST*. Where to even begin with the soundtrack of *Attack on Titan*. Where to even begin with *Attack on Titan* itself! This is an action-packed Japanese animated series that made the jump into the mainstream audiences of the West. It depicts a dystopian world where giant humanoid man-eating monsters have driven the last remnants of humanity into a survival corner. Set in a rustic world taken straight out of XIXth century Europe, it transmits a sense of surreal, eery beauty. Despite its bleak premise and plot, it is filled with impressive open landscapes, beautiful skies, melancholic sunsets, and emotional moments. The OST does not fall any short in this regard, and in fact it is even a great contributor to it. The piece I share here has two halves. The first transmits epicness and urgency. The second one transmits a sense of calm, hopeful melancholy that transitions into more epic, future-looking tunes as it moves on. It plays during *the climatic fight of the second season, when giant warriors try to kidnap the main character and flee his chasing friends*. But the part that hits home for me the most is not when they fight, but when the former group of friends asks one of the reluctant villains for their motives. In it, the characters show their disbelief for the actions of the villains, who are part of a much grander scheme they can not yet comprehend. During this interrogation *the guilt-ridden villain breaks in tears and asks not for forgiveness, but understanding*, despite well knowing his former friends can never give it to him. In a way, it is a goodbye to the apparent simplicity the series had before that scene. **PSR J0917–4413** is probably a most-classical young canonical pulsar no different than many others we have found, but it is the last one I saw in the MMGPS, and a fitting goodbye and look-for-the future in this thesis.





---

## Bibliography

---

- Andersen, B. C. and S. M. Ransom (2018), *A Fourier Domain “Jerk” Search for Binary Pulsars*, *ApJ* **863**, L13 L13, arXiv: 1807.07900 [astro-ph.HE] (cit. on p. 128).
- Andrews, J. J. and I. Mandel (2019), *Double Neutron Star Populations and Formation Channels*, *ApJ* **880**, L8 L8, arXiv: 1904.12745 [astro-ph.HE] (cit. on pp. 6, 35, 36, 134).
- Antoniadis, J., C. G. Bassa, N. Wex, M. Kramer and R. Napiwotzki (2011), *A white dwarf companion to the relativistic pulsar PSR J1141-6545*, *MNRAS* **412** 580, arXiv: 1011.0926 [astro-ph.SR] (cit. on pp. 2, 23, 31, 38).
- Antoniadis, J. et al. (2016), *The millisecond pulsar mass distribution: Evidence for bimodality and constraints on the maximum neutron star mass*, arXiv e-prints, arXiv:1605.01665 arXiv:1605.01665, arXiv: 1605.01665 [astro-ph.HE] (cit. on pp. 34, 35).
- Arzoumanian, Z., K. Joshi, F. A. Rasio and S. E. Thorsett (1996), “Orbital Parameters of the PSR B1620-26 Triple System”, *IAU Colloq. 160: Pulsars: Problems and Progress*, ed. by S. Johnston, M. A. Walker and M. Bailes, vol. 105, Astronomical Society of the Pacific Conference Series 525, arXiv: astro-ph/9605141 [astro-ph] (cit. on pp. 86, 162).
- Ascenzi, S., V. Graber and N. Rea (2024), *Neutron-star Measurements in the Multi-messenger Era*, arXiv e-prints, arXiv:2401.14930 arXiv:2401.14930, arXiv: 2401.14930 [astro-ph.HE] (cit. on pp. 7, 35).
- Avakyan, A. et al. (2023), *XRBcats: Galactic low-mass X-ray binary catalogue*, *A&A* **675**, A199 A199, arXiv: 2303.16168 [astro-ph.HE] (cit. on p. 27).
- Baade, W. and F. Zwicky (1934), *Cosmic Rays from Super-novae*, *Proceedings of the National Academy of Science* **20** 259 (cit. on p. 3).
- Bachetti, M. et al. (2014), *An ultraluminous X-ray source powered by an accreting neutron star*, *Nature* **514** 202, arXiv: 1410.3590 [astro-ph.HE] (cit. on p. 28).
- Backer, D. C., S. R. Kulkarni, C. Heiles, M. M. Davis and W. M. Goss (1982), *A millisecond pulsar*, *Nature* **300** 615 (cit. on p. 1).
- Bahramian, A. and N. Degenaar (2023), “Low-Mass X-ray Binaries”, *Handbook of X-ray and Gamma-ray Astrophysics. Edited by Cosimo Bambi and Andrea Santangelo* 120 120 (cit. on p. 28).
- Bailes, M. et al. (1994), *Discovery of Three Binary Millisecond Pulsars*, *ApJ* **425** L41 (cit. on p. 24).
- Bailes, M. et al. (2020), *The MeerKAT telescope as a pulsar facility: System verification and early science results from MeerTime*, *PASA* **37**, e028 e028, arXiv: 2005.14366 [astro-ph.IM] (cit. on pp. 2, 47, 105, 141).
- Balakrishnan, V. et al. (2023), *Missing for 20 yr: MeerKAT Redetects the Elusive Binary Pulsar M30B*, *ApJ* **942**, L35 L35, arXiv: 2301.04983 [astro-ph.HE] (cit. on p. 16).

- Barr, E. D. (2018), “An S-band Receiver and Backend System for MeerKAT”, *Pulsar Astrophysics the Next Fifty Years*, ed. by P. Weltevrede, B. B. P. Perera, L. L. Preston and S. Sanidas, vol. 337 175 (cit. on p. 47).
- Barr, E. (2020), *Peasoup: C++/CUDA GPU pulsar searching library*, Astrophysics Source Code Library, record ascl:2001.014 (cit. on p. 93).
- Barsdell, B. R., M. Bailes, D. G. Barnes and C. J. Fluke (2012), *Accelerating incoherent dedispersion*, *MNRAS* **422** 379, arXiv: 1201.5380 [astro-ph.IM] (cit. on p. 93).
- Bates, S. D., D. R. Lorimer, A. Rane and J. Swiggum (2014), *PSRPOPPY: an open-source package for pulsar population simulations*, *MNRAS* **439** 2893, arXiv: 1311.3427 [astro-ph.IM] (cit. on p. 138).
- Bates, S. D., D. R. Lorimer and J. P. W. Verbiest (2013), *The pulsar spectral index distribution*, *MNRAS* **431** 1352, arXiv: 1302.2053 [astro-ph.SR] (cit. on p. 13).
- Bates, S. D. et al. (2011), *The High Time Resolution Universe Pulsar Survey - II. Discovery of five millisecond pulsars*, *MNRAS* **416** 2455, arXiv: 1101.4778 [astro-ph.SR] (cit. on p. 3).
- Bates, S. D. et al. (2012), *The High Time Resolution Universe Pulsar Survey — VI. An artificial neural network and timing of 75 pulsars*, *MNRAS* **427** 1052, arXiv: 1209.0793 [astro-ph.SR] (cit. on p. 3).
- Bates, S. D. et al. (2015), *The High Time Resolution Universe survey - XI. Discovery of five recycled pulsars and the optical detectability of survey white dwarf companions*, *MNRAS* **446** 4019, arXiv: 1411.1288 [astro-ph.SR] (cit. on pp. 3, 141).
- Bejger, M. and P. Haensel (2002), *Moments of inertia for neutron and strange stars: Limits derived for the Crab pulsar*, *A&A* **396** 917, arXiv: astro-ph/0209151 [astro-ph] (cit. on p. 18).
- Bell, J. F., M. S. Bessell, B. W. Stappers, M. Bailes and V. M. Kaspi (1995), *PSR J0045-7319: A Dual-Line Binary Radio Pulsar*, *ApJ* **447** L117, arXiv: astro-ph/9505044 [astro-ph] (cit. on p. 29).
- Benvenuto, O. G., M. A. De Vito and J. E. Horvath (2014), *Understanding the Evolution of Close Binary Systems with Radio Pulsars*, *ApJ* **786**, L7 L7, arXiv: 1402.7338 [astro-ph.SR] (cit. on pp. 30, 31).
- Berthereau, A. et al. (2023), *Radio timing constraints on the mass of the binary pulsar PSR J1528–3146*, *A&A* **674**, A71 A71, arXiv: 2304.06578 [astro-ph.HE] (cit. on pp. 33, 113).
- Bezuidenhout, M. C. et al. (2023), *Tied-array beam localization of radio transients and pulsars*, *RAS Techniques and Instruments* **2** 114, arXiv: 2302.09812 [astro-ph.HE] (cit. on p. 104).
- Bhat, N. D. R., M. Bailes and J. P. W. Verbiest (2008), *Gravitational-radiation losses from the pulsar white-dwarf binary PSR J1141 6545*, *Phys. Rev. D* **77**, 124017 124017, arXiv: 0804.0956 [astro-ph] (cit. on p. 31).
- Bhattacharya, D. and E. P. J. van den Heuvel (1991), *Formation and evolution of binary and millisecond radio pulsars*, *Phys. Rep.* **203** 1 (cit. on pp. 2, 22, 27).
- Bhattacharya, D. (2002), *Evolution of Neutron Star Magnetic Fields*, *Journal of Astrophysics and Astronomy* **23** 67 (cit. on p. 27).
- Bhattacharyya, B. and R. Nityananda (2008), *Determination of the orbital parameters of binary pulsars*, *MNRAS* **387** 273, arXiv: 0803.1907 [astro-ph] (cit. on p. 65).
- Blaauw, A., C. S. Gum, J. L. Pawsey and G. Westerhout (1960), *The new I. A. U. system of galactic coordinates (1958 revision)*, *MNRAS* **121** 123 (cit. on p. 158).
- Blandford, R. and S. A. Teukolsky (1976), *Arrival-time analysis for a pulsar in a binary system.*, *ApJ* **205** 580 (cit. on p. 85).

- 
- Bond, H. E. et al. (2017), *The Sirius System and Its Astrophysical Puzzles: Hubble Space Telescope and Ground-based Astrometry*, *ApJ* **840**, 70 70, arXiv: [1703.10625 \[astro-ph.SR\]](#) (cit. on p. 5).
- Boyles, J. et al. (2013), *The Green Bank Telescope 350 MHz Drift-scan Survey. I. Survey Observations and the Discovery of 13 Pulsars*, *ApJ* **763**, 80 80, arXiv: [1209.4293 \[astro-ph.HE\]](#) (cit. on p. 33).
- Bracewell, R. N. (2000), *The Fourier transform and its applications* (cit. on p. 52).
- Breen, A. and D. McCarthy (1995), *A Re-evaluation of the Eastern and Western Records of the Supernova of 1054*, *Vistas in Astronomy* **39** 363 (cit. on p. 21).
- Breton, R. P. et al. (2008), *Relativistic Spin Precession in the Double Pulsar*, *Science* **321** 104, arXiv: [0807.2644 \[astro-ph\]](#) (cit. on p. 87).
- Brosch, N. (2008), *Sirius Matters*, Springer (cit. on p. 5).
- Burgay, M. et al. (2003), *An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system*, *Nature* **426** 531, arXiv: [astro-ph/0312071 \[astro-ph\]](#) (cit. on p. 33).
- Burgay, M. et al. (2006), *The Parkes High-Latitude pulsar survey*, *MNRAS* **368** 283 (cit. on pp. 2, 138).
- Burgay, M. et al. (2013), *The High Time Resolution Universe Pulsar Survey - VII. Discovery of five millisecond pulsars and the different luminosity properties of binary and isolated recycled pulsars*, *MNRAS* **433** 259, arXiv: [1307.7629 \[astro-ph.HE\]](#) (cit. on p. 3).
- Burgay, M. et al. (2019), *The High Time Resolution Universe Pulsar Survey - XV. Completion of the intermediate-latitude survey with the discovery and timing of 25 further pulsars*, *MNRAS* **484** 5791, arXiv: [1902.05571 \[astro-ph.HE\]](#) (cit. on p. 3).
- Burke-Spolaor, S., M. Bailes, R. Ekers, J.-P. Macquart and I. Crawford Fronefield (2011), *Radio Bursts with Extragalactic Spectral Characteristics Show Terrestrial Origins*, *ApJ* **727**, 18 18, arXiv: [1009.5392 \[astro-ph.CO\]](#) (cit. on pp. 3, 58).
- Burrows, A. and D. Vartanyan (2021), *Core-collapse supernova explosion theory*, *Nature* **589** 29, arXiv: [2009.14157 \[astro-ph.SR\]](#) (cit. on p. 6).
- Caiazzo, I. et al. (2021), *A highly magnetized and rapidly rotating white dwarf as small as the Moon*, *Nature* **595** 39, arXiv: [2107.08458 \[astro-ph.SR\]](#) (cit. on p. 36).
- Caleb, M. et al. (2022), *Discovery of a radio-emitting neutron star with an ultra-long spin period of 76 s*, *Nature Astronomy* **6** 828, arXiv: [2206.01346 \[astro-ph.HE\]](#) (cit. on pp. 7, 149).
- Cameron, A. D., E. D. Barr, D. J. Champion, M. Kramer and W. W. Zhu (2017), *An investigation of pulsar searching techniques with the fast folding algorithm*, *MNRAS* **468** 1994, arXiv: [1703.05581 \[astro-ph.IM\]](#) (cit. on pp. 52, 56).
- Cameron, A. D. et al. (2018), *The High Time Resolution Universe Pulsar Survey - XIII. PSR J1757-1854, the most accelerated binary pulsar*, *MNRAS* **475** L57, arXiv: [1711.07697 \[astro-ph.HE\]](#) (cit. on p. 3).
- Cameron, A. D. et al. (2020), *The High Time Resolution Universe Pulsar Survey - XVI. Discovery and timing of 40 pulsars from the southern Galactic plane*, *MNRAS* **493** 1063, arXiv: [2001.01823 \[astro-ph.HE\]](#) (cit. on p. 3).
- Camilo, F., R. N. Manchester, B. M. Gaensler and D. R. Lorimer (2002), *Heartbeat of the Mouse: A Young Radio Pulsar Associated with the Axisymmetric Nebula G359.23-0.82*, *ApJ* **579** L25, arXiv: [astro-ph/0209480 \[astro-ph\]](#) (cit. on p. 6).

- Carli, E. et al. (2024), *The TRAPUM Small Magellanic Cloud pulsar survey with MeerKAT - I. Discovery of seven new pulsars and two Pulsar Wind Nebula associations*, *MNRAS* **531** 2835, arXiv: [2405.12029 \[astro-ph.HE\]](#) (cit. on pp. [47](#), [89](#), [99](#), [149](#)).
- Carpano, S., F. Haberl, C. Maitra and G. Vasilopoulos (2018), *Discovery of pulsations from NGC 300 ULX1 and its fast period evolution*, *MNRAS* **476** L45, arXiv: [1802.10341 \[astro-ph.HE\]](#) (cit. on p. [28](#)).
- Carroll, B. W. and D. A. Ostlie (1996), *An Introduction to Modern Astrophysics*, Cambridge University Press (cit. on pp. [4](#), [6](#)).
- Caswell, J. L., R. F. Haynes, D. K. Milne and K. J. Wellington (1983), *Radio maps revealing shell structures in five supernova remnants.*, *MNRAS* **203** 595 (cit. on p. [129](#)).
- Chadwick, J. (1932), *Possible existence of a neutron*, *Nature* **129** 312 (cit. on p. [3](#)).
- Champion, D. et al. (2020), *Spin-evolution of the new magnetar J1818.0-1607*, *The Astronomer's Telegram* **13559** 1 (cit. on p. [18](#)).
- Chandrasekhar, S. (1931), *The Maximum Mass of Ideal White Dwarfs*, *ApJ* **74** 81 (cit. on p. [5](#)).
- Chanlaridis, S. et al. (2022), *Thermonuclear and electron-capture supernovae from stripped-envelope stars*, *A&A* **668**, A106 A106, arXiv: [2201.00871 \[astro-ph.HE\]](#) (cit. on pp. [5](#), [36](#)).
- Charlot, P. et al. (2020), *The third realization of the International Celestial Reference Frame by very long baseline interferometry*, *A&A* **644**, A159 A159, arXiv: [2010.13625 \[astro-ph.GA\]](#) (cit. on p. [158](#)).
- Chawner, H. et al. (2020), *A complete catalogue of dusty supernova remnants in the Galactic plane*, *MNRAS* **493** 2706, arXiv: [2001.05504 \[astro-ph.GA\]](#) (cit. on p. [127](#)).
- Chen, H.-L., X. Chen, T. M. Tauris and Z. Han (2013), *Formation of Black Widows and Redbacks—Two Distinct Populations of Eclipsing Binary Millisecond Pulsars*, *ApJ* **775**, 27 27, arXiv: [1308.4107 \[astro-ph.SR\]](#) (cit. on pp. [30](#), [31](#)).
- Chen, H.-Y., S. Vitale and F. Foucart (2021a), *The Relative Contribution to Heavy Metals Production from Binary Neutron Star Mergers and Neutron Star-Black Hole Mergers*, *ApJ* **920**, L3 L3, arXiv: [2107.02714 \[astro-ph.HE\]](#) (cit. on pp. [38](#), [135](#)).
- Chen, W., E. Barr, R. Karuppusamy, M. Kramer and B. Stappers (2021b), *Wide Field Beamformed Observation with MeerKAT*, *Journal of Astronomical Instrumentation* **10**, 2150013-178 2150013, arXiv: [2110.01667 \[astro-ph.IM\]](#) (cit. on pp. [47](#), [93](#)).
- Clark, C. J. et al. (2023), *The TRAPUM L-band survey for pulsars in Fermi-LAT gamma-ray sources*, *MNRAS* **519** 5590, arXiv: [2212.08528 \[astro-ph.HE\]](#) (cit. on pp. [47](#), [89](#), [99](#)).
- Cognard, I. et al. (2017), *A Massive-born Neutron Star with a Massive White Dwarf Companion*, *ApJ* **844**, 128 128, arXiv: [1706.08060 \[astro-ph.HE\]](#) (cit. on pp. [28](#), [33](#), [35](#), [147](#)).
- Collins George W., I., W. P. Claspy and J. C. Martin (1999), *A Reinterpretation of Historical References to the Supernova of A.D. 1054*, *PASP* **111** 871, arXiv: [astro-ph/9904285 \[astro-ph\]](#) (cit. on p. [21](#)).
- Colom i Bernadich, M. et al. (2023), *The MPIfR-MeerKAT Galactic Plane Survey. II. The eccentric double neutron star system PSR J1208–5936 and a neutron star merger rate update*, *A&A* **678**, A187 18, arXiv: [2308.16802 \[astro-ph.HE\]](#) (cit. on pp. [131](#), [151](#), [152](#), [163](#)).
- Colom i Bernadich, M. et al. (2024), *PSR J1227 –6208 and its massive white dwarf companion: pulsar emission analysis, timing update, and mass measurements*, arXiv: [2407.13593 \[astro-ph.HE\]](#) (cit. on pp. [141](#), [152](#), [183](#)).
- Comella, J. M., H. D. Craft, R. V. E. Lovelace and J. M. Sutton (1969), *Crab Nebula Pulsar NP 0532*, *Nature* **221** 453 (cit. on pp. [1](#), [4](#)).



- 
- Condon, J. J. and S. M. Ransom (2016), *Essential Radio Astronomy* (cit. on pp. 42–46).
- Cooley, J. W. and J. W. Tukey (1965), *An algorithm for the machine calculation of complex Fourier series*, *Mathematics of computation* **19** 297 (cit. on p. 52).
- Cordes, J. M. (1978), *Observational limits on the location of pulsar emission regions.*, *ApJ* **222** 1006 (cit. on p. 13).
- Cordes, J. M. (2004), *NE2001: A New Model for the Galactic Electron Density and its Fluctuations*, *ASP Conference Series* **317** 211 (cit. on pp. 49, 158).
- Cordes, J. M., A. Pidwerbetsky and R. V. E. Lovelace (1986), *Refractive and Diffractive Scattering in the Interstellar Medium*, *ApJ* **310** 737 (cit. on p. 15).
- Cordes, J. M. et al. (2006), *Arecibo Pulsar Survey Using ALFA. I. Survey Strategy and First Discoveries*, *ApJ* **637** 446, arXiv: [astro-ph/0509732](#) [[astro-ph](#)] (cit. on pp. 95, 138).
- CPTA Collaboration (2023), *Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I*, *Res. Astron. Astrophys.* **23**, 075024 075024, arXiv: [2306.16216](#) [[astro-ph.HE](#)] (cit. on p. 38).
- Crawford, F. et al. (2006), *A Survey of 56 Midlatitude EGRET Error Boxes for Radio Pulsars*, *ApJ* **652** 1499, arXiv: [astro-ph/0608225](#) [[astro-ph](#)] (cit. on pp. 32, 112).
- Damour, T. and N. Deruelle (1986), *General relativistic celestial mechanics of binary systems. II. The post-Newtonian timing formula.*, *Ann. Inst. Henri Poincaré Phys. Théor* **44** 263 (cit. on p. 74).
- Damour, T. and J. H. Taylor (1991), *On the Orbital Period Change of the Binary Pulsar PSR 1913+16*, *ApJ* **366** 501 (cit. on p. 86).
- Davis, P. J., L. Siess and R. Deschamps (2013), *Mass transfer in eccentric binary systems using the binary evolution code BINSTAR*, *A&A* **556**, A4 A4, arXiv: [1305.6092](#) [[astro-ph.SR](#)] (cit. on p. 28).
- Deneva, J. S. et al. (2013), *Goals, Strategies and First Discoveries of AO327, the Arecibo All-sky 327 MHz Drift Pulsar Survey*, *ApJ* **775**, 51 51, arXiv: [1307.8142](#) [[astro-ph.SR](#)] (cit. on p. 34).
- Desvignes, G. et al. (2016), *High-precision timing of 42 millisecond pulsars with the European Pulsar Timing Array*, *MNRAS* **458** 3341, arXiv: [1602.08511](#) [[astro-ph.HE](#)] (cit. on pp. 24, 60).
- Desvignes, G. et al. (2019), *Radio emission from a pulsar’s magnetic pole revealed by general relativity*, *Science* **365** 1013, arXiv: [1909.06212](#) [[astro-ph.HE](#)] (cit. on p. 11).
- Dewey, R. J., J. H. Taylor, J. M. Weisberg and G. H. Stokes (1985), *A search for low-luminosity pulsars.*, *ApJ* **294** L25 (cit. on p. 31).
- Dhurandhar, S. V. and A. Vecchio (2001), *Searching for continuous gravitational wave sources in binary systems*, *Phys. Rev. D* **63**, 122001 122001, arXiv: [gr-qc/0011085](#) [[gr-qc](#)] (cit. on p. 62).
- Di Salvo, T. and A. Sanna (2022), “Accretion Powered X-ray Millisecond Pulsars”, *Astrophysics and Space Science Library*, ed. by S. Bhattacharyya, A. Papitto and D. Bhattacharya, vol. 465, *Astrophysics and Space Science Library* 87 (cit. on p. 27).
- Dubner, G. M., E. B. Giacani, W. M. Goss, D. A. Moffett and M. Holdaway (1996), *VLA Observations of Nine Galactic Supernova Remnants*, *AJ* **111** 1304 (cit. on pp. 127, 129).
- Eatough, R. P. et al. (2013), *A strong magnetic field around the supermassive black hole at the centre of the Galaxy*, *Nature* **501** 391, arXiv: [1308.3147](#) [[astro-ph.GA](#)] (cit. on pp. 15, 17, 21, 49).
- Edwards, R. T., M. Bailes, W. van Straten and M. C. Britton (2001), *The Swinburne intermediate-latitude pulsar survey*, *MNRAS* **326** 358, arXiv: [astro-ph/0105126](#) [[astro-ph](#)] (cit. on p. 2).

- Edwards, R. T., G. B. Hobbs and R. N. Manchester (2006), *TEMPO2, a new pulsar timing package - II. The timing model and precision estimates*, *MNRAS* **372** 1549, arXiv: [astro-ph/0607664](#) [[astro-ph](#)] (cit. on pp. [72](#), [118](#)).
- Einstein, A. (1915a), *Die Feldgleichungen der Gravitation*, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften 844 (cit. on p. [82](#)).
- Einstein, A. (1915b), *Erklärung der Perihelbewegung des Merkur aus der allgemeinen Relativitätstheorie*, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften 831 (cit. on p. [83](#)).
- EPTA and InPTA Collaborations (2023), *A&A* **678**, A50 A50, arXiv: [2306.16214](#) [[astro-ph.HE](#)] (cit. on p. [38](#)).
- Faulkner, A. J. et al. (2004), *The Parkes Multibeam Pulsar Survey - V. Finding binary and millisecond pulsars*, *MNRAS* **355** 147, arXiv: [astro-ph/0408228](#) [[astro-ph](#)] (cit. on pp. [2](#), [32](#), [112](#)).
- Fonseca, E. et al. (2021), *Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620*, *ApJ* **915**, L12 L12, arXiv: [2104.00880](#) [[astro-ph.HE](#)] (cit. on pp. [2](#), [34](#), [38](#)).
- Fortin, F., F. García, A. Simaz Bunzel and S. Chaty (2023), *A catalogue of high-mass X-ray binaries in the Galaxy: from the INTEGRAL to the Gaia era*, *A&A* **671**, A149 A149, arXiv: [2302.02656](#) [[astro-ph.HE](#)] (cit. on p. [27](#)).
- Freire, P. C. C., M. Kramer and A. G. Lyne (2001), *Determination of the orbital parameters of binary pulsars*, *MNRAS* **322** 885, arXiv: [astro-ph/0010463](#) [[astro-ph](#)] (cit. on p. [64](#)).
- Freire, P. C. C. and A. Ridolfi (2018), *An algorithm for determining the rotation count of pulsars*, *MNRAS* **476** 4794, arXiv: [1802.07211](#) [[astro-ph.IM](#)] (cit. on pp. [78](#), [80](#), [81](#)).
- Freire, P. C. C. and N. Wex (2010), *The orthometric parametrization of the Shapiro delay and an improved test of general relativity with binary pulsars*, *MNRAS* **409** 199, arXiv: [1007.0933](#) [[astro-ph.IM](#)] (cit. on pp. [77](#), [84](#)).
- Freire, P. C. C. and N. Wex (2024), *Gravity experiments with radio pulsars*, *Living Reviews in Relativity* **27**, 5 5, arXiv: [2407.16540](#) [[gr-qc](#)] (cit. on p. [87](#)).
- Gaensler, B. M. et al. (2004), *The Mouse that Soared: High-Resolution X-Ray Imaging of the Pulsar-powered Bow Shock G359.23-0.82*, *ApJ* **616** 383, arXiv: [astro-ph/0312362](#) [[astro-ph](#)] (cit. on p. [6](#)).
- Gangadhara, R. T. (2010), *Circular Polarization in Pulsars Due to Curvature Radiation*, *ApJ* **710** 29, arXiv: [1001.2671](#) [[astro-ph.HE](#)] (cit. on p. [13](#)).
- Gautschi, W. (2011), *Numerical analysis*, Springer Science & Business Media (cit. on pp. [66](#), [136](#)).
- Geyer, M. et al. (2023), *Mass measurements and 3D orbital geometry of PSR J1933-6211*, *A&A* **674**, A169 A169, arXiv: [2304.09060](#) [[astro-ph.HE](#)] (cit. on pp. [32](#), [112](#), [146](#)).
- Ghosh, P. and F. K. Lamb (1979), *Accretion by rotating magnetic neutron stars. III. Accretion torques and period changes in pulsating X-ray sources.*, *ApJ* **234** 296 (cit. on p. [27](#)).
- Ghosh, P. and F. K. Lamb (1992), “Diagnostics of disk-magnetosphere interaction in neutron star binaries.”, *X-Ray Binaries and Recycled Pulsars*, ed. by E. P. J. van den Heuvel and S. A. Rappaport, vol. 377, NATO Advanced Study Institute (ASI) Series C 487 (cit. on p. [27](#)).
- Goedhart, S. et al. (2024), *The SARA MeerkAT 1.3 GHz Galactic Plane Survey*, *MNRAS* **531** 649, arXiv: [2312.07275](#) [[astro-ph.GA](#)] (cit. on p. [129](#)).
- Goldreich, P. and W. H. Julian (1969), *Pulsar Electrodynamics*, *ApJ* **157** 869 (cit. on p. [9](#)).
- Grunthal, K. et al. (2024), *Triple trouble with PSR J1618-3921: Mass measurements and orbital dynamics of an eccentric millisecond pulsar*, arXiv: [2409.03615](#) [[astro-ph.HE](#)] (cit. on p. [86](#)).



- 
- Grunthal, K., M. Kramer and G. Desvignes (2021), *Revisiting the Galactic Double Neutron Star merger and LIGO detection rates*, *MNRAS* **507** 5658, arXiv: 2107.13307 [astro-ph.HE] (cit. on pp. 136, 138, 139).
- Guo, Y. J. et al. (2021), *PSR J2222–0137. I. Improved physical parameters for the system*, *A&A* **654**, A16 A16, arXiv: 2107.09474 [astro-ph.HE] (cit. on pp. 32, 33, 35, 113, 162).
- Gupta, T. et al. (2021), *New binary pulsar constraints on Einstein-æther theory after GW170817*, *Classical and Quantum Gravity* **38**, 195003 195003, arXiv: 2104.04596 [gr-qc] (cit. on p. 39).
- Han, J. L. et al. (2021), *The FAST Galactic Plane Pulsar Snapshot survey: I. Project design and pulsar discoveries*, *Research in Astronomy and Astrophysics* **21**, 107 107, arXiv: 2105.08460 [astro-ph.HE] (cit. on p. 90).
- Hankins, T. H. and B. J. Rickett (1975), *Pulsar signal processing.*, *Methods in Computational Physics* **14** 55 (cit. on p. 61).
- Harding, A. K. (2018), “Pulsar Emission Physics: The First Fifty Years”, *Pulsar Astrophysics the Next Fifty Years*, ed. by P. Weltevrede, B. B. P. Perera, L. L. Preston and S. Sanidas, vol. 337 52, arXiv: 1712.02409 [astro-ph.HE] (cit. on p. 9).
- Hessels, J. W. T. et al. (2006), *A Radio Pulsar Spinning at 716 Hz*, *Science* **311** 1901, arXiv: astro-ph/0601337 [astro-ph] (cit. on p. 7).
- Hewish, A., S. J. Bell, J. D. H. Pilkington, P. F. Scott and R. A. Collins (1968), *Observation of a Rapidly Pulsating Radio Source*, *Nature* **217** 709 (cit. on pp. 1, 3, 15, 47).
- Heywood, I. et al. (2022), *The 1.28 GHz MeerKAT Galactic Center Mosaic*, *ApJ* **925**, 165 165, arXiv: 2201.10541 [astro-ph.GA] (cit. on pp. 127, 129).
- Hiramatsu, D. et al. (2021), *The electron-capture origin of supernova 2018zd*, *Nature Astronomy* **5** 903, arXiv: 2011.02176 [astro-ph.HE] (cit. on p. 6).
- Hobbs, G. et al. (2004), *The Parkes multibeam pulsar survey - IV. Discovery of 180 pulsars and parameters for 281 previously known pulsars*, *MNRAS* **352** 1439, arXiv: astro-ph/0405364 [astro-ph] (cit. on p. 2).
- Hobbs, G. B., R. T. Edwards and R. N. Manchester (2006), *TEMPO2, a new pulsar-timing package - I. An overview*, *MNRAS* **369** 655, arXiv: astro-ph/0603381 [astro-ph] (cit. on p. 118).
- Hobbs, G. et al. (2020), *An ultra-wide bandwidth (704 to 4 032 MHz) receiver for the Parkes radio telescope*, *PASA* **37**, e012 e012, arXiv: 1911.00656 [astro-ph.IM] (cit. on p. 46).
- Hollands, M. A. et al. (2020), *An ultra-massive white dwarf with a mixed hydrogen-carbon atmosphere as a likely merger remnant*, *Nature Astronomy* **4** 663, arXiv: 2003.00028 [astro-ph.SR] (cit. on p. 36).
- Hotan, A. W., M. Bailes and S. M. Ord (2005), *PSR J0737-3039A: baseband timing and polarimetry*, *MNRAS* **362** 1267 (cit. on p. 68).
- Hotan, A. W., W. van Straten and R. N. Manchester (2004), *PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis*, *PASA* **21** 302, arXiv: astro-ph/0404549 [astro-ph] (cit. on p. 94).
- Hu, H. et al. (2022), *Gravitational signal propagation in the double pulsar studied with the MeerKAT telescope*, *A&A* **667**, A149 A149, arXiv: 2209.11798 [astro-ph.HE] (cit. on pp. 39, 47).
- Hulse, R. A. and J. H. Taylor (1975), *Discovery of a pulsar in a binary system.*, *ApJ* **195** L51 (cit. on pp. 1, 22, 24).
- Humphreys, R. M. and J. A. Larsen (1995), *The Sun’s Distance Above the Galactic Plane*, *AJ* **110** 2183 (cit. on p. 6).

- Israel, G. L. et al. (2017), *An accreting pulsar with extreme properties drives an ultraluminous x-ray source in NGC 5907*, *Science* **355** 817, arXiv: [1609.07375 \[astro-ph.HE\]](#) (cit. on p. 28).
- Jackson, J. D. (1998), *Classical Electrodynamics, 3rd Edition* (cit. on p. 19).
- Jacoby, B. A., M. Bailes, S. M. Ord, R. T. Edwards and S. R. Kulkarni (2009), *A Large-Area Survey for Radio Pulsars at High Galactic Latitudes*, *ApJ* **699** 2009 (cit. on p. 2).
- Jacoby, B. A., M. Bailes, S. M. Ord, H. S. Knight and A. W. Hotan (2007), *Discovery of Five Recycled Pulsars in a High Galactic Latitude Survey*, *ApJ* **656** 408, arXiv: [astro-ph/0609448 \[astro-ph\]](#) (cit. on pp. 2, 33, 113).
- Jacoby, B. A., D. Chakrabarty, M. H. van Kerkwijk, S. R. Kulkarni and D. L. Kaplan (2006), *Optical Detection of Two Intermediate-Mass Binary Pulsar Companions*, *ApJ* **640** L183, arXiv: [astro-ph/0602501 \[astro-ph\]](#) (cit. on p. 33).
- Janka, H.-T., A. Wongwathanarat and M. Kramer (2022), *Supernova Fallback as Origin of Neutron Star Spins and Spin-kick Alignment*, *ApJ* **926**, 9 9, arXiv: [2104.07493 \[astro-ph.HE\]](#) (cit. on pp. 6–8, 35).
- Jankowski, F. et al. (2018), *Spectral properties of 441 radio pulsars*, *MNRAS* **473** 4436, arXiv: [1709.08864 \[astro-ph.HE\]](#) (cit. on p. 13).
- Jansky, K. G. (1933), *Electrical phenomena that apparently are of interstellar origin*, *Popular Astronomy* **41** 548 (cit. on p. 41).
- Johnston, H. M. and S. R. Kulkarni (1991), *On the Detectability of Pulsars in Close Binary Systems*, *ApJ* **368** 504 (cit. on pp. 54, 55).
- Johnston, S., R. N. Manchester, A. G. Lyne, L. Nicastro and J. Spyromilio (1994), *Radio and Optical Observations of the PSR:B1259-63 / SS:2883 Be-Star Binary System*, *MNRAS* **268** 430 (cit. on p. 29).
- Johnston, S. et al. (2023), *The Thousand-Pulsar-Array programme on MeerKAT - XI. Application of the rotating vector model*, *MNRAS* **520** 4801, arXiv: [2212.03988 \[astro-ph.HE\]](#) (cit. on p. 12).
- Johnston, S. and M. Kerr (2018), *Polarimetry of 600 pulsars from observations at 1.4 GHz with the Parkes radio telescope*, *MNRAS* **474** 4629, arXiv: [1711.10092 \[astro-ph.HE\]](#) (cit. on p. 11).
- Johnston, S. et al. (1992), *A high-frequency survey of the southern galactic plane for pulsars.*, *MNRAS* **255** 401 (cit. on p. 29).
- Jonas, J. and MeerKAT Team (2016), “The MeerKAT Radio Telescope”, *MeerKAT Science: On the Pathway to the SKA 1* 1 (cit. on p. 2).
- Kaplan, D. L. et al. (2014), *A 1.05  $M_{\odot}$  Companion to PSR J2222-0137: The Coolest Known White Dwarf?*, *ApJ* **789**, 119 119, arXiv: [1406.0488 \[astro-ph.SR\]](#) (cit. on p. 33).
- Karastergiou, A. et al. (2024), *The Thousand-Pulsar-Array programme on MeerKAT - XV. A comparison of the radio emission properties of slow and millisecond pulsars*, *MNRAS* **532** 3558, arXiv: [2407.06836 \[astro-ph.HE\]](#) (cit. on p. 11).
- Kaspi, V. M. and D. J. Helfand (2002), “Constraining the Birth Events of Neutron Stars”, *Neutron Stars in Supernova Remnants*, ed. by P. O. Slane and B. M. Gaensler, vol. 271, Astronomical Society of the Pacific Conference Series 3, arXiv: [astro-ph/0201183 \[astro-ph\]](#) (cit. on p. 19).
- Kaspi, V. M. et al. (2000), *Discovery of a Young Radio Pulsar in a Relativistic Binary Orbit*, *ApJ* **543** 321, arXiv: [astro-ph/0005214 \[astro-ph\]](#) (cit. on pp. 2, 31).
- Kaspi, V. M. and A. M. Beloborodov (2017), *Magnetars*, *ARA&A* **55** 261, arXiv: [1703.00068 \[astro-ph.HE\]](#) (cit. on p. 21).

- 
- Keane, E. et al. (2015), “A Cosmic Census of Radio Pulsars with the SKA”, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)* 40 40, arXiv: [1501.00056 \[astro-ph.IM\]](#) (cit. on p. 153).
- Keith, M. J. et al. (2010), *The High Time Resolution Universe Pulsar Survey - I. System configuration and initial discoveries*, *MNRAS* **409** 619, arXiv: [1006.5744 \[astro-ph.HE\]](#) (cit. on pp. 2, 46, 91, 138).
- Keith, M. J. et al. (2012), *The High Time Resolution Universe Pulsar Survey - IV. Discovery and polarimetry of millisecond pulsars*, *MNRAS* **419** 1752, arXiv: [1109.4193 \[astro-ph.SR\]](#) (cit. on p. 3).
- Kepler, S. O., D. Koester and G. Ourique (2016), *A white dwarf with an oxygen atmosphere*, *Science* **352** 67 (cit. on p. 5).
- Kim, C., V. Kalogera and D. R. Lorimer (2003), *The Probability Distribution of Binary Pulsar Coalescence Rates. I. Double Neutron Star Systems in the Galactic Field*, *ApJ* **584** 985, arXiv: [astro-ph/0207408 \[astro-ph\]](#) (cit. on p. 136).
- Kim, V., I. Izmailova and Y. Aimuratov (2023), *Catalog of the Galactic Population of X-Ray Pulsars in High-mass X-Ray Binary Systems*, *ApJS* **268**, 21 21, arXiv: [2308.09295 \[astro-ph.HE\]](#) (cit. on p. 27).
- Kippenhahn, R. and A. Weigert (1967), *Entwicklung in engen Doppelsternsystemen I. Massenaustausch vor und nach Beendigung des zentralen Wasserstoff-Brennens*, *ZAp* **65** 251 (cit. on p. 28).
- Kluźniak, W., M. Ruderman, J. Shaham and M. Tavani (1988), *Nature and evolution of the eclipsing millisecond binary pulsar PSR1957 + 20*, *Nature* **334** 225 (cit. on pp. 31, 127).
- Knispel, B. et al. (2013), *Einstein@Home Discovery of 24 Pulsars in the Parkes Multi-beam Pulsar Survey*, *ApJ* **774**, 93 93, arXiv: [1302.0467 \[astro-ph.GA\]](#) (cit. on p. 2).
- Komesaroff, M. M. (1970), *Possible Mechanism for the Pulsar Radio Emission*, *Nature* **225** 612 (cit. on p. 9).
- Kopeikin, S. M. (1995), *On Possible Implications of Orbital Parallaxes of Wide Orbit Binary Pulsars and Their Measurability*, *ApJ* **439** L5 (cit. on p. 161).
- Kopeikin, S. M. (1996), *Proper Motion of Binary Pulsars as a Source of Secular Variations of Orbital Parameters*, *ApJ* **467** L93 (cit. on pp. 86, 162).
- Kopparapu, R. K. et al. (2008), *Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events*, *ApJ* **675** 1459, arXiv: [0706.1283 \[astro-ph\]](#) (cit. on p. 139).
- Korol, V. et al. (2017), *Prospects for detection of detached double white dwarf binaries with Gaia, LSST and LISA*, *MNRAS* **470** 1894, arXiv: [1703.02555 \[astro-ph.HE\]](#) (cit. on p. 38).
- Kothes, R. and S. M. Dougherty (2007), *The distance and neutral environment of the massive stellar cluster Westerlund 1*, *A&A* **468** 993, arXiv: [0704.3073 \[astro-ph\]](#) (cit. on pp. 127, 129).
- Kramer, M., A. G. Lyne, J. T. O’Brien, C. A. Jordan and D. R. Lorimer (2006), *A Periodically Active Pulsar Giving Insight into Magnetospheric Physics*, *Science* **312** 549, arXiv: [astro-ph/0604605 \[astro-ph\]](#) (cit. on pp. 108, 109).
- Kramer, M. et al. (2003), *The Parkes Multibeam Pulsar Survey - III. Young pulsars and the discovery and timing of 200 pulsars*, *MNRAS* **342** 1299, arXiv: [astro-ph/0303473 \[astro-ph\]](#) (cit. on p. 2).
- Kramer, M. et al. (2021a), *Strong-Field Gravity Tests with the Double Pulsar*, *Physical Review X* **11**, 041050 041050, arXiv: [2112.06795 \[astro-ph.HE\]](#) (cit. on pp. 2, 18, 33, 39, 83–85, 87).
- Kramer, M. et al. (2021b), *The relativistic binary programme on MeerKAT: science objectives and first results*, *MNRAS* **504** 2094, arXiv: [2102.05160 \[astro-ph.HE\]](#) (cit. on p. 141).

- Kramer, M. (1998), *Determination of the Geometry of the PSR B1913+16 System by Geodetic Precession*, *ApJ* **509** 856, arXiv: [astro-ph/9808127](#) [[astro-ph](#)] (cit. on p. 11).
- Kramer, M. et al. (1998), *The Characteristics of Millisecond Pulsar Emission. I. Spectra, Pulse Shapes, and the Beaming Fraction*, *ApJ* **501** 270, arXiv: [astro-ph/9801177](#) [[astro-ph](#)] (cit. on p. 11).
- Krishnakumar, M. A., Y. Maan, B. C. Joshi and P. K. Manoharan (2019), *Multi-frequency Scatter-broadening Evolution of Pulsars. II. Scatter-broadening of Nearby Pulsars*, *ApJ* **878**, L30, arXiv: [1905.01651](#) [[astro-ph.HE](#)] (cit. on p. 17).
- Lange, C. et al. (2001), *Precision timing measurements of PSR J1012+5307*, *MNRAS* **326** 274, arXiv: [astro-ph/0102309](#) [[astro-ph](#)] (cit. on p. 76).
- Large, M. I., A. E. Vaughan and B. Y. Mills (1968), *A Pulsar Supernova Association?*, *Nature* **220** 340 (cit. on p. 101).
- Lattimer, J. M. and M. Prakash (2001), *Neutron Star Structure and the Equation of State*, *ApJ* **550** 426, arXiv: [astro-ph/0002232](#) [[astro-ph](#)] (cit. on p. 35).
- Lau, M. Y. M. et al. (2020), *Detecting double neutron stars with LISA*, *MNRAS* **492** 3061, arXiv: [1910.12422](#) [[astro-ph.HE](#)] (cit. on p. 38).
- Lazarus, P. et al. (2014), *Timing of a young mildly recycled pulsar with a massive white dwarf companion*, *MNRAS* **437** 1485, arXiv: [1310.5857](#) [[astro-ph.SR](#)] (cit. on pp. 29, 30, 32).
- Lazarus, P. et al. (2020), *CoastGuard: Automated timing data reduction pipeline*, Astrophysics Source Code Library, record ascl:2003.008 (cit. on p. 60).
- Lentati, L. et al. (2014), *TEMPONEST: a Bayesian approach to pulsar timing analysis*, *MNRAS* **437** 3004, arXiv: [1310.2120](#) [[astro-ph.IM](#)] (cit. on pp. 143–145).
- Levin, L. et al. (2013), *The High Time Resolution Universe Pulsar Survey -VIII. The Galactic millisecond pulsar population*, *MNRAS* **434** 1387, arXiv: [1306.4190](#) [[astro-ph.SR](#)] (cit. on p. 3).
- LIGO and Virgo Collaborations (2016), *Observation of Gravitational Waves from a Binary Black Hole Merger*, *Phys. Rev. Lett.* **116**, 061102, arXiv: [1602.03837](#) [[gr-qc](#)] (cit. on p. 83).
- LIGO and Virgo Collaborations (2017), *Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A*, *ApJ* **848**, L13, arXiv: [1710.05834](#) [[astro-ph.HE](#)] (cit. on pp. 38, 135).
- LIGO Scientific Collaboration et al. (2017), *Multi-messenger Observations of a Binary Neutron Star Merger*, *ApJ* **848**, L12, arXiv: [1710.05833](#) [[astro-ph.HE](#)] (cit. on p. 135).
- LISA Consortium (2023), *Waveform Modelling for the Laser Interferometer Space Antenna*, arXiv e-prints, arXiv:[2311.01300](#) arXiv:[2311.01300](#), arXiv: [2311.01300](#) [[gr-qc](#)] (cit. on p. 38).
- Liu, K., R. P. Eatough, N. Wex and M. Kramer (2014), *Pulsar-black hole binaries: prospects for new gravity tests with future radio telescopes*, *MNRAS* **445** 3115, arXiv: [1409.3882](#) [[astro-ph.GA](#)] (cit. on p. 39).
- Löhmer, O., D. Mitra, Y. Gupta, M. Kramer and A. Ahuja (2004), *The frequency evolution of interstellar pulse broadening from radio pulsars*, *A&A* **425** 569, arXiv: [astro-ph/0406601](#) [[astro-ph](#)] (cit. on p. 17).
- Lomb, N. R. (1976), *Least-Squares Frequency Analysis of Unequally Spaced Data*, *Ap&SS* **39** 447 (cit. on p. 64).
- Lorimer, D. R. and M. Kramer (2005), *Handbook of pulsar astronomy*, Cambridge University press (cit. on pp. 41–43, 48–50, 52, 56, 58, 67, 69, 74, 159, 161).
- Lorimer, D. R. et al. (2006), *Arecibo Pulsar Survey Using ALFA. II. The Young, Highly Relativistic Binary Pulsar J1906+0746*, *ApJ* **640** 428, arXiv: [astro-ph/0511523](#) [[astro-ph](#)] (cit. on p. 33).



- 
- Lorimer, D. R. et al. (2007), *VizieR Online Data Catalog: Parkes Multibeam Pulsar Survey. VI. (Lorimer+, 2006)*, VizieR Online Data Catalog, J/MNRAS/372/777 J/MNRAS/372/777 (cit. on pp. 2, 3, 18).
- Lorimer, D. R. et al. (2015), *The Parkes multibeam pulsar survey - VII. Timing of four millisecond pulsars and the underlying spin-period distribution of the Galactic millisecond pulsar population*, *MNRAS* **450** 2185, arXiv: 1501.05516 [astro-ph.IM] (cit. on p. 3).
- Lower, M. E. et al. (2024), *A MeerKAT view of the double pulsar eclipses. Geodetic precession of pulsar B and system geometry*, *A&A* **682**, A26 A26, arXiv: 2311.06445 [astro-ph.HE] (cit. on p. 87).
- Lynch, R. S. et al. (2018), *The Green Bank North Celestial Cap Pulsar Survey. III. 45 New Pulsar Timing Solutions*, *ApJ* **859**, 93 93, arXiv: 1805.04951 [astro-ph.HE] (cit. on p. 34).
- Lyne, A. G. and R. N. Manchester (1988), *The shape of pulsar radio beams.*, *MNRAS* **234** 477 (cit. on p. 11).
- Lyne, A. G. et al. (2004), *A Double-Pulsar System: A Rare Laboratory for Relativistic Gravity and Plasma Physics*, *Science* **303** 1153, arXiv: astro-ph/0401086 [astro-ph] (cit. on pp. 2, 23, 33).
- Lyne, A. G. et al. (2017), *Two Long-Term Intermittent Pulsars Discovered in the PALFA Survey*, *ApJ* **834**, 72 72, arXiv: 1608.09008 [astro-ph.HE] (cit. on pp. 108, 109).
- Manchester, R. N., G. B. Hobbs, A. Teoh and M. Hobbs (2005), *The Australia Telescope National Facility Pulsar Catalogue*, *AJ* **129** 1993, arXiv: astro-ph/0412641 [astro-ph] (cit. on p. 20).
- Manchester, R. N. et al. (2001), *The Parkes multi-beam pulsar survey - I. Observing and data analysis systems, discovery and timing of 100 pulsars*, *MNRAS* **328** 17, arXiv: astro-ph/0106522 [astro-ph] (cit. on pp. 2, 46, 91, 138, 141).
- Martinez, J. G. et al. (2015), *Pulsar J0453+1559: A Double Neutron Star System with a Large Mass Asymmetry*, *ApJ* **812**, 143 143, arXiv: 1509.08805 [astro-ph.HE] (cit. on pp. 2, 34).
- Martinez, J. G. et al. (2019), *The Discovery of Six Recycled Pulsars from the Arecibo 327 MHz Drift-Scan Pulsar Survey*, *ApJ* **881**, 166 166, arXiv: 1906.05071 [astro-ph.HE] (cit. on p. 11).
- Mayor, M. and D. Queloz (1995), *A Jupiter-mass companion to a solar-type star*, *Nature* **378** 355 (cit. on p. 60).
- McConnell, D. et al. (1991), *Radio pulsars in the Magellanic clouds.*, *MNRAS* **249** 654 (cit. on p. 29).
- McKee, J. W. et al. (2020), *A precise mass measurement of PSR J2045 + 3633*, *MNRAS* **499** 4082, arXiv: 2009.12283 [astro-ph.HE] (cit. on pp. 32, 37).
- McLaughlin, M. A. et al. (2006), *Transient radio bursts from rotating neutron stars*, *Nature* **439** 817 (cit. on p. 109).
- McMillan, P. J. (2017), *The mass distribution and gravitational potential of the Milky Way*, *MNRAS* **465** 76, arXiv: 1608.00971 [astro-ph.GA] (cit. on p. 86).
- Mdzinarishvili, T. G. and G. I. Melikidze (2004), *On the z-distribution of pulsars*, *A&A* **425** 1009 (cit. on p. 7).
- Melrose, D. (2003), “What Causes the Circular Polarization in Pulsars?”, *Radio Pulsars*, ed. by M. Bailes, D. J. Nice and S. E. Thorsett, vol. 302, Astronomical Society of the Pacific Conference Series 179 (cit. on p. 13).
- Men, Y., E. Barr, C. J. Clark, E. Carli and G. Desvignes (2023), *PulsarX: A new pulsar searching package. I. A high performance folding program for pulsar surveys*, *A&A* **679**, A20 A20, arXiv: 2309.02544 [astro-ph.IM] (cit. on pp. 60, 94).

- Mereghetti, S. et al. (2005), *An XMM-Newton View of the Soft Gamma Repeater SGR 1806-20: Long-Term Variability in the Pre-Giant Flare Epoch*, *ApJ* **628** 938, arXiv: [astro-ph/0502417](#) [[astro-ph](#)] (cit. on p. 18).
- Mickaliger, M. B. et al. (2012), *Discovery of Five New Pulsars in Archival Data*, *ApJ* **759**, 127 127, arXiv: [1206.2895](#) [[astro-ph.SR](#)] (cit. on p. 141).
- Misner, C. W., K. S. Thorne and J. A. Wheeler (1973), *Gravitation* (cit. on p. 82).
- Miyaji, S., K. Nomoto, K. Yokoi and D. Sugimoto (1980), *Supernova triggered by electron captures.*, *PASJ* **32** 303 (cit. on p. 6).
- Morello, V., E. D. Barr, B. W. Stappers, E. F. Keane and A. G. Lyne (2020), *Optimal periodicity searching: revisiting the fast folding algorithm for large-scale pulsar surveys*, *MNRAS* **497** 4654, arXiv: [2004.03701](#) [[astro-ph.IM](#)] (cit. on p. 128).
- Morello, V., E. D. Barr, B. W. Stappers, E. F. Keane and A. G. Lyne (2023), *riptide: Pulsar searching with the Fast Folding Algorithm*, Astrophysics Source Code Library, record ascl:2310.010 (cit. on p. 128).
- Morello, V. et al. (2019), *The High Time Resolution Universe survey - XIV. Discovery of 23 pulsars through GPU-accelerated reprocessing*, *MNRAS* **483** 3673, arXiv: [1811.04929](#) [[astro-ph.IM](#)] (cit. on pp. 3, 60, 93).
- Morris, D. J. et al. (2002), *The Parkes Multibeam Pulsar Survey - II. Discovery and timing of 120 pulsars*, *MNRAS* **335** 275, arXiv: [astro-ph/0204238](#) [[astro-ph](#)] (cit. on p. 2).
- NANOGrav Collaboration (2023), *The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background*, *ApJ* **951**, L8 L8, arXiv: [2306.16213](#) [[astro-ph.HE](#)] (cit. on p. 38).
- Ng, C. et al. (2014), *The High Time Resolution Universe pulsar survey - X. Discovery of four millisecond pulsars and updated timing solutions of a further 12*, *MNRAS* **439** 1865, arXiv: [1401.3003](#) [[astro-ph.HE](#)] (cit. on p. 3).
- Ng, C. et al. (2015), *The High Time Resolution Universe Pulsar Survey - XII. Galactic plane acceleration search and the discovery of 60 pulsars*, *MNRAS* **450** 2922, arXiv: [1504.08000](#) [[astro-ph.HE](#)] (cit. on pp. 3, 56, 98).
- Nomoto, K. (1982), “The origin of the Crab Nebula and electron capture supernova of 8-10 solar-mass stars”, *Supernovae: A Survey of Current Research*, ed. by M. J. Rees and R. J. Stoneham, vol. 90, NATO Advanced Study Institute (ASI) Series C 205 (cit. on p. 6).
- Nomoto, K. (1984), *Evolution of 8-10 solar mass stars toward electron capture supernovae. I - Formation of electron-degenerate O + NE + MG cores.*, *ApJ* **277** 791 (cit. on p. 6).
- Nomoto, K., F. K. Thielemann and K. Yokoi (1984), *Accreting white dwarf models for type I supern. III. Carbon deflagration supernovae.*, *ApJ* **286** 644 (cit. on pp. 5, 36).
- Nomoto, K. (1987), *Evolution of 8–10  $M_{\text{sun}}$  Stars toward Electron Capture Supernovae. II. Collapse of an O + NE + MG Core*, *ApJ* **322** 206 (cit. on p. 6).
- Oswald, L. S., A. Karastergiou and S. Johnston (2023), *Pulsar polarization: a partial-coherence model*, *MNRAS* **525** 840, arXiv: [2307.14265](#) [[astro-ph.HE](#)] (cit. on p. 13).
- Oswald, L. S. et al. (2021), *The Thousand-Pulsar-Array programme on MeerKAT - V. Scattering analysis of single-component pulsars*, *MNRAS* **504** 1115, arXiv: [2104.01081](#) [[astro-ph.HE](#)] (cit. on pp. 16, 17, 133, 143).
- Özel, F. and P. Freire (2016), *Masses, Radii, and the Equation of State of Neutron Stars*, *ARA&A* **54** 401, arXiv: [1603.02698](#) [[astro-ph.HE](#)] (cit. on pp. 7, 35).



- 
- Padmanabh, P. V. et al. (2023), *The MPIfR-MeerKAT Galactic Plane Survey - I. System set-up and early results*, *MNRAS* **524** 1291, arXiv: 2303.09231 [astro-ph.HE] (cit. on pp. 47, 56, 89–93, 105, 106, 115, 149, 163).
- Pallanca, C. et al. (2013), *The Optical Companion to the Intermediate-mass Millisecond Pulsar J1439-5501 in the Galactic Field*, *ApJ* **773**, 127 127, arXiv: 1306.3787 [astro-ph.SR] (cit. on p. 23).
- Parthasarathy, A. et al. (2019), *Timing of young radio pulsars - I. Timing noise, periodic modulation, and proper motion*, *MNRAS* **489** 3810, arXiv: 1908.11709 [astro-ph.HE] (cit. on p. 144).
- Pelisoli, I. et al. (2023), *A 5.3-min-period pulsing white dwarf in a binary detected from radio to X-rays*, *Nature Astronomy* **7** 931, arXiv: 2306.09272 [astro-ph.SR] (cit. on p. 3).
- Peters, P. C. (1964), *Gravitational Radiation and the Motion of Two Point Masses*, *Physical Review* **136** 1224 (cit. on pp. 83, 135).
- Peters, P. C. and J. Mathews (1963), *Gravitational Radiation from Point Masses in a Keplerian Orbit*, *Physical Review* **131** 435 (cit. on p. 39).
- Petroff, E. et al. (2015), *Identifying the source of perytons at the Parkes radio telescope*, *MNRAS* **451** 3933, arXiv: 1504.02165 [astro-ph.IM] (cit. on p. 58).
- Philippov, A. and M. Kramer (2022), *Pulsar Magnetospheres and Their Radiation*, *ARA&A* **60** 495 (cit. on p. 9).
- Planck Collaboration (2020), *Planck 2018 results. VI. Cosmological parameters*, *A&A* **641**, A6 A6, arXiv: 1807.06209 [astro-ph.CO] (cit. on p. 4).
- Pol, N., M. McLaughlin and D. R. Lorimer (2020), *An Updated Galactic Double Neutron Star Merger Rate Based on Radio Pulsar Populations*, *RNAAS* **4**, 22 22, arXiv: 2002.10225 [astro-ph.HE] (cit. on pp. 136, 138, 139).
- Poutanen, J., G. Lipunova, S. Fabrika, A. G. Butkevich and P. Abolmasov (2007), *Supercritically accreting stellar mass black holes as ultraluminous X-ray sources*, *MNRAS* **377** 1187, arXiv: astro-ph/0609274 [astro-ph] (cit. on p. 28).
- PPTA Collaboration (2023), *Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array*, *ApJ* **951**, L6 L6, arXiv: 2306.16215 [astro-ph.HE] (cit. on pp. 38, 46).
- Prayag, V. et al. (2024), *The TRAPUM Large Magellanic Cloud pulsar survey with MeerKAT - I. Survey set-up and first seven pulsar discoveries*, *MNRAS* **533** 2570, arXiv: 2408.04899 [astro-ph.HE] (cit. on pp. 89, 149).
- Press, W. H., B. P. Flannery and S. A. Teukolsky (1986), *Numerical recipes. The art of scientific computing* (cit. on p. 74).
- Pringle, J. E. and R. A. Wade (1985), *Interacting binary stars* (cit. on p. 28).
- Prša, A. et al. (2016), *Nominal Values for Selected Solar and Planetary Quantities: IAU 2015 Resolution B3*, *AJ* **152**, 41 41 (cit. on p. 24).
- Radhakrishnan, V. and D. J. Cooke (1969), *Magnetic Poles and the Polarization Structure of Pulsar Radiation*, *Astrophys. Lett.* **3** 225 (cit. on p. 12).
- Radhakrishnan, V. and G. Srinivasan (1982), *On the origin of the recently discovered ultra-rapid pulsar*, *Current Science* **51** 1096 (cit. on pp. 1, 2, 27).
- Raghavan, D. et al. (2010), *A Survey of Stellar Families: Multiplicity of Solar-type Stars*, *ApJS* **190** 1, arXiv: 1007.0414 [astro-ph.SR] (cit. on p. 22).
- Rankin, J. M. (1993), *Toward an Empirical Theory of Pulsar Emission. VI. The Geometry of the Conal Emission Region: Appendix and Tables*, *ApJS* **85** 145 (cit. on p. 11).

- Ransom, S. M. et al. (2014), *A millisecond pulsar in a stellar triple system*, *Nature* **505** 520, arXiv: [1401.0535 \[astro-ph.SR\]](#) (cit. on pp. 2, 39).
- Ransom, S. (2011), *PRESTO: Pulsar Exploration and Search Toolkit*, Astrophysics Source Code Library, record ascl:1107.017 (cit. on pp. 93, 97, 103).
- Ransom, S. M., S. S. Eikenberry and J. Middleditch (2002), *Fourier Techniques for Very Long Astrophysical Time-Series Analysis*, *AJ* **124** 1788, arXiv: [astro-ph/0204349 \[astro-ph\]](#) (cit. on pp. 52, 53, 55, 56).
- Ransom, S. M. et al. (2004), *Green Bank Telescope Discovery of Two Binary Millisecond Pulsars in the Globular Cluster M30*, *ApJ* **604** 328, arXiv: [astro-ph/0310347 \[astro-ph\]](#) (cit. on p. 16).
- Rappaport, S., P. Podsiadlowski, P. C. Joss, R. Di Stefano and Z. Han (1995), *The relation between white dwarf mass and orbital period in wide binary radio pulsars*, *MNRAS* **273** 731 (cit. on pp. 31, 37, 113).
- Rickett, B. J. (1977), *Interstellar scattering and scintillation of radio waves.*, *ARA&A* **15** 479 (cit. on p. 17).
- Ridolfi, A., P. C. C. Freire, Y. Gupta and S. M. Ransom (2019), *Upgraded Giant Metrewave Radio Telescope timing of NGC 1851A: a possible millisecond pulsar - neutron star system*, *MNRAS* **490** 3860, arXiv: [1909.06163 \[astro-ph.HE\]](#) (cit. on p. 77).
- Ridolfi, A. et al. (2021), *Eight new millisecond pulsars from the first MeerKAT globular cluster census*, *MNRAS* **504** 1407, arXiv: [2103.04800 \[astro-ph.HE\]](#) (cit. on pp. 103, 119, 149).
- Ridolfi, A. et al. (2022), *TRAPUM discovery of 13 new pulsars in NGC 1851 using MeerKAT*, *A&A* **664**, A27 A27, arXiv: [2203.12302 \[astro-ph.HE\]](#) (cit. on pp. 47, 89, 149).
- Ridolfi, A. (2020), *PYSOLATOR: Remove orbital modulation from a binary pulsar and/or its companion*, Astrophysics Source Code Library, record ascl:2003.012 (cit. on p. 134).
- Riley, T. E. et al. (2019), *A NICER View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation*, *ApJ* **887**, L21 L21, arXiv: [1912.05702 \[astro-ph.HE\]](#) (cit. on pp. 7, 35).
- Riley, T. E. et al. (2021), *A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy*, *ApJ* **918**, L27 L27, arXiv: [2105.06980 \[astro-ph.HE\]](#) (cit. on pp. 7, 35).
- Roberts, M. S. E. (2013), “Surrounded by spiders! New black widows and redbacks in the Galactic field”, *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, ed. by J. van Leeuwen, vol. 291 127, arXiv: [1210.6903 \[astro-ph.HE\]](#) (cit. on p. 31).
- Romani, R. W., R. Narayan and R. Blandford (1986), *Refractive effects in pulsar scintillation.*, *MNRAS* **220** 19 (cit. on p. 17).
- Romani, R. W., D. Kandel, A. V. Filippenko, T. G. Brink and W. Zheng (2022), *PSR J0952-0607: The Fastest and Heaviest Known Galactic Neutron Star*, *ApJ* **934**, L17 L17, arXiv: [2207.05124 \[astro-ph.HE\]](#) (cit. on p. 34).
- Roy, A. E. (2005), *Orbital motion* (cit. on pp. 23, 28, 159, 161).
- Ruderman, M., J. Shaham and M. Tavani (1989), *Accretion Turnoff and Rapid Evaporation of Very Light Secondaries in Low-Mass X-Ray Binaries*, *ApJ* **336** 507 (cit. on p. 31).
- Scargle, J. D. (1982), *Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data.*, *ApJ* **263** 835 (cit. on p. 64).
- Sengar, R. et al. (2022), *The High Time Resolution Universe Pulsar Survey - XVII. PSR J1325-6253, a low eccentricity double neutron star system from an ultra-stripped supernova*, *MNRAS* **512** 5782, arXiv: [2204.06869 \[astro-ph.HE\]](#) (cit. on pp. 3, 134).

- 
- Shakura, N. I. and R. A. Sunyaev (1973), *Black holes in binary systems. Observational appearance.*, *A&A* **24** 337 (cit. on p. 28).
- Shamohammadi, M. et al. (2023), *Searches for Shapiro delay in seven binary pulsars using the MeerKAT telescope*, *MNRAS* **520** 1789, arXiv: 2212.04051 [astro-ph.HE] (cit. on pp. 32, 37, 112).
- Shapiro, I. I. (1964), *Fourth Test of General Relativity*, *Phys. Rev. Lett.* **13** 789 (cit. on p. 71).
- Shklovskii, I. S. (1970), *Possible Causes of the Secular Increase in Pulsar Periods.*, *Soviet Ast.* **13** 562 (cit. on p. 86).
- Siebert, T. (2019), *Vertical position of the Sun with  $\gamma$ -rays*, *A&A* **632**, L1 L1, arXiv: 1910.09575 [astro-ph.HE] (cit. on p. 6).
- Singh, S. et al. (2022), *The GMRT High Resolution Southern Sky Survey for Pulsars and Transients. III. Searching for Long-period Pulsars*, *ApJ* **934**, 138 138, arXiv: 2206.00427 [astro-ph.HE] (cit. on p. 54).
- Smith, N. (2013), *The Crab nebula and the class of Type II<sub>n</sub>-P supernovae caused by sub-energetic electron-capture explosions*, *MNRAS* **434** 102, arXiv: 1304.0689 [astro-ph.HE] (cit. on p. 6).
- Splaver, E. M. et al. (2002), *Probing the Masses of the PSR J0621+1002 Binary System through Relativistic Apsidal Motion*, *ApJ* **581** 509, arXiv: astro-ph/0208281 [astro-ph] (cit. on p. 88).
- Spruit, H. C. (2009), “The source of magnetic fields in (neutron-) stars”, *Cosmic Magnetic Fields: From Planets, to Stars and Galaxies*, ed. by K. G. Strassmeier, A. G. Kosovichev and J. E. Beckman, vol. 259 61 (cit. on p. 8).
- Staelin, D. H. (1969), *Fast folding algorithm for detection of periodic pulse trains.*, *IEEE Proceedings* **57** 724 (cit. on p. 52).
- Stappers, B. and M. Kramer (2016), “An Update on TRAPUM”, *MeerKAT Science: On the Pathway to the SKA 9 9* (cit. on pp. 89, 149).
- Staveley-Smith, L. et al. (1996), *The Parkes 21 CM multibeam receiver*, *PASA* **13** 243 (cit. on p. 46).
- Stovall, K. et al. (2014), *The Green Bank Northern Celestial Cap Pulsar Survey. I. Survey Description, Data Analysis, and Initial Results*, *ApJ* **791**, 67 67, arXiv: 1406.5214 [astro-ph.HE] (cit. on p. 138).
- Sturrock, P. A. (1971), *A Model of Pulsars*, *ApJ* **164** 529 (cit. on p. 9).
- Sushch, I., B. Hnatyk and A. Neronov (2011), *Modeling of the Vela complex including the Vela supernova remnant, the binary system  $\gamma^2$  Velorum, and the Gum nebula*, *A&A* **525**, A154 A154, arXiv: 1011.1177 [astro-ph.GA] (cit. on p. 101).
- Swart, G. P., P. E. Dewdney and A. Cremonini (2022), *Highlights of the SKA1-Mid telescope architecture*, *Journal of Astronomical Telescopes, Instruments, and Systems* **8**, 011021 011021 (cit. on p. 153).
- Takahashi, K., T. Yoshida and H. Umeda (2013), *Evolution of Progenitors for Electron Capture Supernovae*, *ApJ* **771**, 28 28, arXiv: 1302.6402 [astro-ph.SR] (cit. on pp. 5, 6, 36).
- Tauris, T. M. (2011), “Five and a Half Roads to Form a Millisecond Pulsar”, *Evolution of Compact Binaries*, ed. by L. Schmidtbreick, M. R. Schreiber and C. Tappert, vol. 447, Astronomical Society of the Pacific Conference Series 285, arXiv: 1106.0897 [astro-ph.HE] (cit. on pp. 22, 28–30).
- Tauris, T. M., N. Langer and M. Kramer (2011), *Formation of millisecond pulsars with CO white dwarf companions - I. PSR J1614-2230: evidence for a neutron star born massive*, *MNRAS* **416** 2130, arXiv: 1103.4996 [astro-ph.SR] (cit. on pp. 28, 32, 112, 147).

- Tauris, T. M., N. Langer and M. Kramer (2012), *Formation of millisecond pulsars with CO white dwarf companions - II. Accretion, spin-up, true ages and comparison to MSPs with He white dwarf companions*, *MNRAS* **425** 1601, arXiv: [1206.1862 \[astro-ph.SR\]](#) (cit. on pp. [29](#), [30](#), [32](#), [35](#)).
- Tauris, T. M. and T. Sennels (2000), *Formation of the binary pulsars PSR B2303+46 and PSR J1141-6545. Young neutron stars with old white dwarf companions*, *A&A* **355** 236, arXiv: [astro-ph/9909149 \[astro-ph\]](#) (cit. on pp. [30](#), [31](#)).
- Tauris, T. M. and E. P. J. van den Heuvel (2006), “Formation and evolution of compact stellar X-ray sources”, *Compact stellar X-ray sources*, vol. 39 623 (cit. on pp. [28](#), [29](#)).
- Tauris, T. M. et al. (2017), *Formation of Double Neutron Star Systems*, *ApJ* **846**, 170 170, arXiv: [1706.09438 \[astro-ph.HE\]](#) (cit. on pp. [6](#), [29](#), [30](#), [33](#), [35](#), [133](#), [134](#)).
- Tauris, T. M., N. Langer and P. Podsiadlowski (2015), *Ultra-stripped supernovae: progenitors and fate*, *MNRAS* **451** 2123, arXiv: [1505.00270 \[astro-ph.SR\]](#) (cit. on pp. [33](#), [36](#), [134](#)).
- Tauris, T. M. and G. J. Savonije (1999), *Formation of millisecond pulsars. I. Evolution of low-mass X-ray binaries with  $P_{\text{orb}} > 2$  days*, *A&A* **350** 928, arXiv: [astro-ph/9909147 \[astro-ph\]](#) (cit. on pp. [30](#), [31](#), [37](#), [113](#)).
- Tauris, T. M. and E. P. J. van den Heuvel (2023), *Physics of Binary Star Evolution. From Stars to X-ray Binaries and Gravitational Wave Sources* (cit. on pp. [28](#), [29](#)).
- Taylor, J. H. (1992), *Pulsar Timing and Relativistic Gravity*, *Philosophical Transactions of the Royal Society of London Series A* **341** 117 (cit. on p. [68](#)).
- Taylor, J. H., L. A. Fowler and P. M. McCulloch (1979), *Measurements of general relativistic effects in the binary pulsar PSR1913 + 16*, *Nature* **277** 437 (cit. on pp. [1](#), [22](#), [38](#), [60](#), [83](#)).
- Taylor, J. H. and G. R. Huguenin (1969), *Two New Pulsating Radio Sources*, *Nature* **221** 816 (cit. on p. [53](#)).
- Taylor, J. H. and J. M. Weisberg (1982), *A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16*, *ApJ* **253** 908 (cit. on pp. [1](#), [22](#), [38](#), [83](#)).
- Taylor, J. H. and J. M. Weisberg (1989), *Further Experimental Tests of Relativistic Gravity Using the Binary Pulsar PSR 1913+16*, *ApJ* **345** 434 (cit. on pp. [71](#), [72](#), [88](#)).
- Thorsett, S. E., Z. Arzoumanian, M. M. McKinnon and J. H. Taylor (1993), *The Masses of Two Binary Neutron Star Systems*, *ApJ* **405** L29, arXiv: [astro-ph/9303002 \[astro-ph\]](#) (cit. on p. [31](#)).
- Thorsett, S. E. and D. Chakrabarty (1999), *Neutron Star Mass Measurements. I. Radio Pulsars*, *ApJ* **512** 288, arXiv: [astro-ph/9803260 \[astro-ph\]](#) (cit. on p. [31](#)).
- Truran, J. W. and M. Livio (1986), *On the Frequency of Occurrence of Oxygen-Neon-Magnesium White Dwarfs in Classical Nova Systems*, *ApJ* **308** 721 (cit. on p. [5](#)).
- Turner, J. D. et al. (2024), *TRAPUM search for pulsars in supernova remnants and pulsar wind nebulae - I. Survey description and initial discoveries*, *MNRAS* **531** 3579, arXiv: [2405.11899 \[astro-ph.HE\]](#) (cit. on pp. [47](#), [89](#), [149](#)).
- United Kingdom Hydrographic Office (2011), *The Astronomical Almanac*, US Government Printing Office (cit. on pp. [157](#), [158](#)).
- van den Heuvel, E. P. J. (2019), *High-Mass X-ray Binaries: progenitors of double compact objects*, *IAU Symposium* **346** 1, arXiv: [1901.06939 \[astro-ph.HE\]](#) (cit. on p. [29](#)).
- van der Plas, J. T. (2018), *Understanding the Lomb-Scargle Periodogram*, *ApJS* **236**, 16 16, arXiv: [1703.09824 \[astro-ph.IM\]](#) (cit. on p. [65](#)).
- van Kerkwijk, M. H. and S. R. Kulkarni (1999), *A Massive White Dwarf Companion to the Eccentric Binary Pulsar System PSR B2303+46*, *ApJ* **516** L25, arXiv: [astro-ph/9901149 \[astro-ph\]](#) (cit. on p. [31](#)).



- 
- van Straten, W., P. Demorest and S. Osłowski (2012), *Pulsar Data Analysis with PSRCHIVE*, *Astronomical Research and Technology* **9** 237, arXiv: 1205.6276 [astro-ph.IM] (cit. on p. 94).
- van Veen, B. D. and K. M. Buckley (1988), *Beamforming: A versatile approach to spatial filtering*, *IEEE ASSP Magazine* **5** 4 (cit. on p. 46).
- Venkatraman Krishnan, V. et al. (2020), *Lense-Thirring frame dragging induced by a fast-rotating white dwarf in a binary pulsar system*, *Science* **367** 577, arXiv: 2001.11405 [astro-ph.HE] (cit. on p. 2).
- Verbunt, F. and C. Zwaan (1981), *Magnetic braking in low-mass X-ray binaries.*, *A&A* **100** L7 (cit. on p. 31).
- Viganò, D. et al. (2013), *Unifying the observational diversity of isolated neutron stars via magneto-thermal evolution models*, *MNRAS* **434** 123, arXiv: 1306.2156 [astro-ph.SR] (cit. on p. 19).
- Voisin, G. et al. (2020), *An improved test of the strong equivalence principle with the pulsar in a triple star system*, *A&A* **638**, A24 A24, arXiv: 2005.01388 [gr-qc] (cit. on pp. 2, 39).
- Wang, J. et al. (2022a), *A comparative analysis of pulse time-of-arrival creation methods*, *A&A* **658**, A181 A181, arXiv: 2111.13482 [astro-ph.IM] (cit. on p. 68).
- Wang, N., R. N. Manchester and S. Johnston (2007), *Pulsar nulling and mode changing*, *MNRAS* **377** 1383, arXiv: astro-ph/0703241 [astro-ph] (cit. on pp. 108, 109).
- Wang, Y. et al. (2022b), *Discovery of PSR J0523-7125 as a Circularly Polarized Variable Radio Source in the Large Magellanic Cloud*, *ApJ* **930**, 38 38, arXiv: 2205.00622 [astro-ph.HE] (cit. on p. 13).
- Watson, D. et al. (2019), *Identification of strontium in the merger of two neutron stars*, *Nature* **574** 497, arXiv: 1910.10510 [astro-ph.HE] (cit. on p. 38).
- Weisberg, J. M. and Y. Huang (2016), *Relativistic Measurements from Timing the Binary Pulsar PSR B1913+16*, *ApJ* **829**, 55 55, arXiv: 1606.02744 [astro-ph.HE] (cit. on pp. 24, 83, 85, 87).
- Weisberg, J. M. and J. H. Taylor (2002), *General Relativistic Geodetic Spin Precession in Binary Pulsar B1913+16: Mapping the Emission Beam in Two Dimensions*, *ApJ* **576** 942, arXiv: astro-ph/0205280 [astro-ph] (cit. on p. 11).
- Wette, K. (2023), *Searches for continuous gravitational waves from neutron stars: A twenty-year retrospective*, *Astroparticle Physics* **153**, 102880 102880, arXiv: 2305.07106 [gr-qc] (cit. on p. 38).
- Whiteoak, J. B. Z. and A. J. Green (1996), *The MOST supernova remnant catalogue (MSC).*, *A&AS* **118** 329 (cit. on p. 129).
- Wolszczan, A. (1991), *A nearby 37.9-ms radio pulsar in a relativistic binary system*, *Nature* **350** 688 (cit. on p. 138).
- Wolszczan, A. and D. A. Frail (1992), *A planetary system around the millisecond pulsar PSR1257 + 12*, *Nature* **355** 145 (cit. on p. 60).
- Wongphechauxsorn, J. et al. (2024), *The High Time Resolution Universe Pulsar survey - XVIII. The reprocessing of the HTRU-S Low Lat survey around the Galactic Centre using a Fast Folding Algorithm pipeline for accelerated pulsars*, *MNRAS* **527** 3208, arXiv: 2310.14008 [astro-ph.HE] (cit. on pp. 3, 56).
- Yang, H. and R. A. Chevalier (2015), *Evolution of the Crab Nebula in a Low Energy Supernova*, *ApJ* **806**, 153 153, arXiv: 1505.03211 [astro-ph.HE] (cit. on p. 6).
- Yao, J. M., R. N. Manchester and N. Wang (2017), *A New Electron-density Model for Estimation of Pulsar and FRB Distances*, *ApJ* **835**, 29 29, arXiv: 1610.09448 [astro-ph.GA] (cit. on pp. 49, 158).

- Yoon, S. C. and N. Langer (2004), *Presupernova evolution of accreting white dwarfs with rotation*, *A&A* **419** 623, arXiv: [astro-ph/0402287](#) [[astro-ph](#)] (cit. on pp. 5, 36).
- Young, M. D., R. N. Manchester and S. Johnston (1999), *A radio pulsar with an 8.5-second period that challenges emission models*, *Nature* **400** 848 (cit. on p. 21).
- Yuan, H. et al. (2015), *Stellar Loci II. A Model-free Estimate of the Binary Fraction for Field FGK Stars*, *ApJ* **799**, 135 135, arXiv: [1412.1233](#) [[astro-ph.SR](#)] (cit. on p. 22).
- Zahn, J. P. (2008), “Tidal dissipation in binary systems”, *EAS Publications Series*, ed. by M. J. Goupil and J. P. Zahn, vol. 29, EAS Publications Series 67, arXiv: [0807.4870](#) [[astro-ph](#)] (cit. on p. 28).
- Zhang, S. B. et al. (2024), *RRAT J1913+1330: An Extremely Variable and Puzzling Pulsar*, *ApJ* **972**, 59 59, arXiv: [2306.02855](#) [[astro-ph.HE](#)] (cit. on p. 109).
- Zhao, J., P. C. C. Freire, M. Kramer, L. Shao and N. Wex (2022), *Closing a spontaneous-scalarization window with binary pulsars*, *Classical and Quantum Gravity* **39**, 11LT01 11LT01, arXiv: [2201.03771](#) [[astro-ph.HE](#)] (cit. on p. 39).
- Zhu, W. W. et al. (2014), *Searching for Pulsars Using Image Pattern Recognition*, *ApJ* **781**, 117 117, arXiv: [1309.0776](#) [[astro-ph.IM](#)] (cit. on p. 94).
- Zhu, W. W. et al. (2019), *Tests of gravitational symmetries with pulsar binary J1713+0747*, *MNRAS* **482** 3249, arXiv: [1802.09206](#) [[astro-ph.HE](#)] (cit. on p. 39).



---

## List of Figures

---

1.1	Isolated stellar evolution and the formation of compact remnants . . . . .	5
1.2	Spin and magnetic evolution of a collapsing star . . . . .	8
1.3	The magnetic dipole model . . . . .	10
1.4	Pulse profiles of <b>PSR J2222–0137</b> and <b>PSR J0509+0856</b> . . . . .	11
1.5	The rotating vector model and some examples . . . . .	12
1.6	The effects of the interstellar medium on radio light . . . . .	14
1.7	Scintillation in <b>PSR J1708–4843</b> . . . . .	15
1.8	Scattering on the pulse profile of <b>PSR J1818–1422</b> . . . . .	16
1.9	The general $P - \dot{P}$ diagram . . . . .	20
1.10	The orbits of pulsars and their companions . . . . .	24
1.11	Distribution of orbital and spin parameters of binary pulsars . . . . .	26
1.12	Diagrams of low-mass and high-mass X-ray binaries . . . . .	27
1.13	Evolutionary diagram of binary pulsars . . . . .	30
1.14	Mass distribution of neutron stars in binary systems . . . . .	34
1.15	Orbital parameters and masses in DNS systems . . . . .	36
1.16	Masses of white dwarf systems . . . . .	37
2.1	Diagram of a radio telescope: receiving and recording . . . . .	43
2.2	Angular resolution and signal response of single dishes and interferometers . . . . .	45
2.3	MeerKAT and Parkes/Murriyang sky beams . . . . .	47
2.4	De-dispersing a pulsar signal . . . . .	49
2.5	Folding a pulsar signal . . . . .	51
2.6	Depiction of the discrete Fourier transform . . . . .	53
2.7	Effect of acceleration on pulsar periodicity . . . . .	54
2.8	Examples of pulsar candidate . . . . .	57
2.9	Examples of radio-frequency interference . . . . .	59
2.10	Spin and acceleration evolution of binary pulsars. . . . .	63
2.11	Example fits to the acceleration ellipse . . . . .	64
2.12	Example of a roughness fit to barycentric spin period . . . . .	65
2.13	Making of a timing template . . . . .	66
2.14	Measuring ToAs from a pulsar archive . . . . .	67
2.15	Examples of ToA residuals . . . . .	73
2.16	Shapiro delay amplitude against eccentric anomaly. . . . .	76
2.17	Lost rotation counts in an orbital solution. . . . .	79

2.18	The dracula algorithm. . . . .	81
2.19	Visual depiction of the post-Keplerian orbital phenomena . . . . .	82
2.20	Orbital decay and Shapiro delay in the double pulsar . . . . .	84
2.21	$M_p$ – $M_c$ diagrams from PK parameters. . . . .	87
3.1	Sky coverage of the MMPGS surveys . . . . .	90
3.2	Tiling of primary and coherent beams in the MMGPS-L . . . . .	92
3.3	MMGPS-L pulsar candidates with CandyJar . . . . .	96
3.4	Multibeam frequency masks with multiTRAPUM . . . . .	98
3.5	Fourier frequency birdies with multiTRAPUM . . . . .	99
3.6	Discovery plot of <b>PSR J1306–6043</b> . . . . .	101
3.7	Harmonics of <b>PSR J1208–5935</b> and <b>PSR J1015–5368</b> . . . . .	102
3.8	Localisation maps of MMGPS pulsars . . . . .	104
3.9	Pulse profiles of MMGPS-L discoveries . . . . .	106
3.10	Pulse profiles of MMGPS-S discoveries . . . . .	107
3.11	Pulse profiles of some isolated MMGPS-L discoveries . . . . .	108
3.12	Intermittency of <b>PSR J1353–6341</b> . . . . .	109
3.13	MMGPS-L binaries in the $P_b - x$ and $P_b - P_s$ diagrams . . . . .	111
3.14	MMGPS-L white dwarf binaries in the $P_b - M_c$ and $P_s - M_c$ diagrams . . . . .	113
3.15	Orbital spin modulation of <b>PSR J1306–6043</b> . . . . .	115
3.16	Orbital spin modulation of <b>PSR J1208–5936</b> . . . . .	117
3.17	Scintillation in <b>PSR J1708–4843</b> . . . . .	120
3.18	Orbital fit of <b>PSR J1708–4843</b> . . . . .	121
3.19	<b>PSR J1108–5936</b> and <b>PSR J1543–5439</b> : profiles and orbital solutions . . . . .	123
3.20	Scintillation and measurements in <b>PSR J1015–5358</b> . . . . .	126
3.21	<b>SNR G340.6+0.3</b> observations . . . . .	128
3.22	Pulsar candidates in <b>SNR G340.6+0.3</b> . . . . .	130
4.1	Orbital campaign of <b>PSR J1208–5936</b> . . . . .	132
4.2	Orbital decay of <b>PSR J1208–5936</b> . . . . .	136
4.3	Sky coverage of radio pulsar surveys. . . . .	138
5.1	Timing noise in <b>PSR J1227–6208</b> . . . . .	145
6.1	Predicted PK measurements in <b>PSR J1208–5936</b> and <b>PSR J1227–6208</b> . . . . .	152
A.1	Angles of pulsar binaries . . . . .	159

---

## List of Tables

---

3.2 Binary systems discovered by the MMGPS-L . . . . .	112
--	-----