



UNICA

UNIVERSITÀ
DEGLI STUDI
DI CAGLIARI

**Ph.D. DEGREE IN
Physics**

Cycle XXXVIII

TITLE OF THE Ph.D. THESIS

**Exploiting pulsar timing at low radio frequencies: applications to Pulsar Timing
Arrays, interstellar medium and HII regions**

Scientific Disciplinary Sector(s)

FIS/05 ASTROFISICA E ASTRONOMIA

Ph.D. Student:	Francesco Iraci
Supervisor	Dr. Andrea Possenti
Co-Supervisor	Dr. Caterina Tiburzi, Prof. Alberto Sesana

Final exam. Academic Year 2024/2025
Thesis defence session: February 2026

Contents

Abstract	1
1 Introduction	3
1.1 Pulsars	3
1.1.1 Formation	3
1.1.2 Properties	4
1.1.3 Evolution	4
1.1.4 Emission mechanism	6
1.2 Propagation of radio signal through the IISM	7
1.2.1 Dispersion of radio signals	7
1.2.2 Scattering	10
1.2.3 Faraday rotation	11
1.2.4 Dispersion measure variations	12
1.2.5 Solar wind	13
1.2.6 The turbulence of the IISM	14
1.3 Pulsar timing	14
1.3.1 The pulsar timing model	16
1.3.2 Classification of noise processes in pulsar timing	17
1.3.3 Impact of propagation effects on timing	18
1.3.4 Applications of pulsar timing	19
1.4 Pulsar Timing Arrays	20
1.5 Thesis structure	21
2 Data and methods	23
2.1 Observations	23
2.1.1 LOFAR	23
2.1.2 NenuFAR	26
2.2 Pulsar timing at low-radio frequencies	28
2.2.1 Template creation and timing procedure	28
2.3 Noise in pulsar timing	30
2.3.1 Noise modelling methods	30

3	Datasets: combination of LOFAR and EPTA	35
3.1	Motivations for this combination	35
3.2	The dataset	35
3.2.1	EPTA-InPTA DR2	35
3.2.2	LOFAR	36
3.2.3	NenuFAR	36
3.3	Data combination	37
3.4	Noise analysis	38
3.5	Noise analysis results	41
3.5.1	Observations for the DM variations	41
3.5.2	Observations for the RN	45
3.6	Model selection	47
3.6.1	Comparing the favored models of DR2new+ and DR2low	47
3.6.2	Discussing the results from the noise model selection	49
3.7	SW impact	52
3.8	Conclusions	54
4	Discrete ionised regions along the LoS	57
4.1	HII regions	57
4.1.1	WISE catalogue	58
4.1.2	HH14 catalogue	58
4.2	HII along the LoS	58
4.3	Analysis	59
4.3.1	Power spectrum of the density fluctuations	61
4.3.2	Scintillation	64
4.3.3	Results	66
4.4	Discussion	69
4.4.1	Structure function and turbulence spectrum	69
4.4.2	Scintillation	70
4.5	Conclusions	70
5	Methods to estimate DM variations	75
5.1	Scope of the work	75
5.2	Simulations and software	75
5.3	Methods of analysis	79
5.3.1	Epoch-wise dispersion measure	79
5.3.2	DMX	79
5.3.3	DM GP	79
5.4	Results	82
5.5	Discussion	85
5.5.1	Absorption of red noise in the dispersion measure reconstruction	85
5.5.2	Comparison of the methods based on precision and accuracy	85
5.5.3	Whitening of the dispersion measure residuals	87
5.5.4	Alternative simulation settings	89
5.6	Conclusions	94
6	Summary and Conclusions	97

A	Appendix to Chap. 3	101
A.1	Parameters of the noise models used in this work	101
A.2	Estimation of the ecorr parameters	102
A.3	Solar wind parameters	103
B	Appendix to Chap. 4	105
B.1	Generation of simulated DM time series from the free spectrum	105
B.2	Transformation of Proper Motion and Computation of Effective Velocity	106
B.2.1	Coordinate Conversion	106
B.2.2	Angle Between Coordinate Axes	106
B.2.3	Proper Motion Rotation	106
B.2.4	Transverse Velocity Projection	106
B.2.5	Effective Velocity	107

List of Figures

1.1	P- $\dot{P}$ diagram	5
1.2	Impact of DM errors on pulse profiles of PSR J0358+5413.	9
1.3	Pulse profile of PSR J0614+2229 with the characteristic exponential scattering tail.	11
1.4	Examples of timing residuals of PSR J0751+1807 for different model imperfections.	17
1.5	Post-fit timing residuals for PSR J0358+5413 after fitting for a single DM value.	19
1.6	Timing residuals for PSR J0030+0451 showing the effect of the solar wind.	20
2.1	Distribution of LOFAR stations in Europe	24
2.2	LOFAR antenna types.	25
2.3	NenuFAR antennas (left) and a view of the field with the mini-arrays (right).	27
3.1	Mollweide sky projection of the 25 DR2new+ pulsars in equatorial coordinates.	37
3.2	Time vs frequency plot for DR2low	39
3.3	Timing residuals of DR2low	40
3.4	Spectral index vs. amplitude posterior distributions of the power-law PSD for DMv (left) and RN (right) obtained for twelve pulsars from DR2low (blue) and DR2new+ (orange) using the standard noise models.	42
3.5	Same as Fig. 3.4, but for the remaining six pulsars.	43
3.6	Posterior distributions of spectral indices (γ) and amplitudes (A) of DMv and RN for PSRs J1024–0719 (top) and J1738+0333 (bottom) obtained with DR2low (blue) and DR2new+ (red), using the standard noise models.	48
3.7	Posterior distribution of the chromatic index (blue violins) and their 99.7% credible intervals (errorbars) obtained for the pulsars that favor the presence of CN ₄ in their noise model.	51
3.8	Free spectra for DMv (referred to 1.4GHz) and RN for PSR J0030+0451 (top) and PSR J1022+1001 (bottom).	53

4.1	Distribution of H II region radii from the WISE catalogue for H II that intersect the LoS of pulsars in the full pulsar catalogue. . .	59
4.2	DM time series of PSRs J0358+5413 (top), J0407+1607 (middle), J2225+6535 (bottom).	67
4.3	Structure functions of PSRs J0358+5413 (top), J0407+1607 (middle), J2225+6535 (bottom).	67
4.4	Wavenumber spectra for PSRs J0358+5413 (light-blue dots), J0407+1607 (green squares), and J2225+6535 (coral diamonds).	68
4.5	Normalized signal-to-noise ratio (S/N divided by $\sqrt{T_{\text{int}}}$) and autocorrelation function (ACF) for the three pulsars, displayed from top to bottom: J0358+5413 (blue), J0407+1607 (green), and J2225+6535 (orange).	72
4.6	C_n^2 as a function of β for PSR J0358+5413.	73
5.1	Simulations carried out with injected values of γ_{RN} , A_{RN} , and γ_{DM} of, respectively, 3.7, -13.3 , and $8/3$	78
5.2	DM time series and DM residuals recovered by EW (blue, first and second panels), DMX (red, third and fourth panels), and DM GP (green, fifth and sixth panels).	83
5.3	DM residuals of 100 realizations for the five considered methods: EW, DMX no RN, DM GP no RN, DMX, DM GP.	84
5.4	Histograms of DM residuals normalized by their uncertainty for EW (blue), DMX no RN (magenta), DM GP no RN (orange), DMX (red), and DM GP (green).	86
5.5	As in Fig. 5.4, but the spectral indices are $\gamma_{RN} = 2.5$ and $\gamma_{DM} = 3.2$	86
5.6	DM residuals obtained by applying DM GP (with RN modeling) to an individual simulation over the 100 realizations for values of the noise parameters $A_{RN} = -13.6$, $A_{DM} = -13.3$, $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$, EFAC and EQUAD as in Table 5.1, and $\sigma_{\text{ToA}} = 5\mu\text{s}$	90
5.7	Power spectra of the DM residuals left by the EW (left panel), DMX (central panel), and DM GP (right panel) methods.	90
5.8	DM residuals and their corresponding histogram for L-band simulated data.	92
5.9	Power spectra of the DM residuals obtained using EW (left), DMX (center), and DM GP (right), compared with the injected signal (light blue) in the L band.	93
5.10	DM time series injected in the data with an empirical power spectrum (Donner, 2022).	93
5.11	DM residuals and the corresponding power spectra for the custom DM injection scenario.	94
A.1	Posterior distribution of the ecorr parameters (logarithmic scale) and their 99.7% credible intervals (errorbars) obtained for the LOFAR systems for each pulsars.	102
A.2	Posterior distribution of the SW parameters for PSRs J0030+0451, J0751+1807, J1022+1001 and J1744-1134.	103

List of Tables

3.1	Comparison of the favoured noise models obtained with DR21low and with DR2new+ for the twelve pulsars considered in this work.	47
3.2	Estimated values (medians with 95% confidence intervals) of the amplitude and spectral index of RN, DMv and CN ₄ obtained with DR21low by using the favoured noise models displayed in Table 3.1.	50
4.1	Galactic parameters, distances of the six pulsars and the H II regions intersecting their LoS. The table reports also the pulsar transverse velocity.	60
4.2	Best-fit spectral indices β obtained from the structure function analysis using the power-law model of Equation (4.16).	68
4.3	For each pulsar, the table lists the pulsar name, the fractional distance of the H II region along the line of sight (s), the effective transverse velocity v_{eff} , the best-fit refractive timescale Δt_r (referred at 153.8 MHz), and the derived diffractive timescale Δt_d , computed using Equation (4.22).	69
5.1	Input noise parameter values adopted for our simulations.	77
5.2	Mean uncertainty on the DM measurements for each method and pair of (A_{RN}, A_{DM}) values in the case of $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$	84
5.3	Mean uncertainty on the DM measurements for each method and pair of (A_{RN}, A_{DM}) in the case of $\gamma_{RN} = 2.5$, $\gamma_{DM} = 3.2$	85
5.4	Reduced χ^2 of a Gaussian fit to the histograms reported in Fig. 5.4. $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$	88
5.5	Reduced χ^2 of a Gaussian fit to the histograms reported in Fig. 5.4. $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$	88
5.6	Reduced χ^2 obtained from a Gaussian fit on the histograms reported in Fig. 5.5. $\gamma_{RN} = 2.5$, $\gamma_{DM} = 3.2$	89
5.7	Mean DM uncertainty across methods (EW, DMX, DM GP) for simulated ToAs with fixed uncertainties of 1 μs and 0.1 μs	91
5.8	As in Table 5.7, but for L-band ToAs with $\sigma_{\text{temp}} = 1 \mu\text{s}$ and injected DM noise parameters of $A_{DM} = -13.6$ and $\gamma_{DM} = 3.7$	91

5.9	Reduced χ^2 of the normalized DM residuals. Obtained from a Gaussian fit on histograms of DM residuals normalized by their uncertainty for the EW, DMX, and DM GP methods, over ToAs simulated in the L band with $\sigma_{\text{temp}} = 1 \mu s$ and injected DM noise ($A_{DM}, \gamma_{DM} = -13.3, 8/3$).	92
A.1	Approach for modelling the dispersion caused by solar winds inspired by Susarla et al., 2024.	101

Abstract

High-precision pulsar timing is strongly affected by the interstellar medium (ISM), which introduces variations in dispersion measure (DM) and other propagation effects that impact timing residuals. This thesis focuses on characterising DM variations and understanding how the turbulent ISM influences pulsar signals. Using low-frequency observations, statistical analyses of DM time series, and complementary simulations, I investigate the properties of interstellar turbulence along different lines of sight, including those intersecting dense ionised regions. By combining multiple approaches, this work aims to provide a coherent view of how the ISM affects pulsar timing and to quantify its role as a source of noise in high-precision experiments.

To disentangle intrinsic pulsar spin variations from propagation effects, I combine high-frequency pulsar timing data with low-frequency observations from the LOFAR and NenuFAR telescopes. These low-frequency datasets enhance sensitivity to DM variations, enabling a more precise separation of radio-frequency-dependent (chromatic) effects, such as DM fluctuations, from radio-frequency-independent (achromatic) noise processes, such as intrinsic spin irregularities. This improved characterisation of DM noise components also reveals the high sensitivity of these observations to variations in the Solar Wind, highlighting limitations in the Solar Wind models currently applied in pulsar timing array analyses. Overall, the combined dataset allows for a more accurate modelling of both pulsar-intrinsic and propagation-induced contributions to timing residuals.

I also exploit LOFAR's high sensitivity to study DM variations driven by interstellar turbulence, which influence pulsar signal propagation and timing precision. While turbulence in the diffuse ISM is well described by a Kolmogorov-like spectrum, its properties within dense, structured environments such as HII regions remain poorly constrained. Several pulsars observed with LOFAR have lines of sight intersecting HII regions, offering a rare opportunity to probe turbulence under extreme plasma conditions. I analyse DM time series using statistical tools such as the structure function and examine scintillation observables to extract complementary constraints on plasma structure across multiple scales. This multi-probe approach allows building a coherent picture of turbulence in HII regions and assessing its contribution to propagation-induced noise in high-precision pulsar timing.

Finally, I study DM variations using simulations. Since multiple methods are employed in the pulsar timing array (PTA) community to estimate DM variations, I create mock datasets incorporating all the typical noise components present in PTA observations. By comparing these methods on controlled simulated data, I assess their relative performance and determine which approaches provide the most reliable estimates of DM variations.

Together, these observational and simulation-based studies provide a comprehensive view of how interstellar turbulence affects pulsar timing and help improve methods for DM characterisation in high-precision experiments.

CHAPTER 1

Introduction

This chapter presents an introduction to pulsar astrophysics, focusing on their fundamental properties, the influence of the interstellar medium on their observed signals, and the pulsar timing technique. It concludes with a short overview of the Pulsar Timing Array experiment.

1.1 Pulsars

Spinning, flashing, and ticking like celestial metronomes, pulsars are the remnants of stellar explosions that now serve as some of the most stable clocks in the Universe.

1.1.1 Formation

Neutron stars represent the compact remnants of stars with masses between approximately 9.5 and $25 M_{\odot}$ (Karakas and Lattanzio, 2014), that are expected to undergo a supernova event at the end of their evolution (Kippenhahn, Weigert, and Weiss, 1990). Two main formation channels can lead to the birth of a neutron star: core-collapse supernovae and electron-capture supernovae. Core-collapse supernovae occur when the iron core of a massive star becomes gravitationally unstable and collapses under its own weight, typically involving progenitors with initial masses above $\sim 10 M_{\odot}$. The collapse releases a huge amount of gravitational energy, part of which is transferred to the outer layers, causing the supernova explosion. In contrast, electron-capture supernovae are thought to occur in stars with initial masses roughly between 8 and $10 M_{\odot}$. In these stars, the degenerate core is mainly composed of oxygen, neon, and magnesium, and the capture of electrons by these nuclei triggers the core collapse. The sudden contraction of the core also releases energy that drives the ejection of the stellar envelope. As a consequence, while most of the mass of the progenitor star is expelled during the supernova explosion, the core undergoes a gravitational collapse. This collapse is eventually halted by neutron degeneracy pressure, which provides an outward force that balances gravity. The result is

a compact, self-gravitating object primarily composed of degenerate neutrons. Neutron stars typically have radii of about 10 km, masses around $1.4 M_{\odot}$, and average densities on the order of $10^{14} \text{ g cm}^{-3}$. This equilibrium configuration is what we define as a neutron star.

Shortly after their formation, some neutron stars become observable as pulsars, rapidly rotating and highly magnetized stellar remnants that emit beams of electromagnetic radiation from their magnetic poles. If these beams cross the Earth as the star spins, they appear as periodic pulses, often with remarkable regularity. The first pulsar was discovered in 1967 by Jocelyn Bell (Hewish et al., 1968), marking the beginning of a new era in astrophysics.

1.1.2 Properties

Neutron stars are characterized by two fundamental parameters that govern their observable behaviour: the spin period P , its time derivative $\dot{P}$ (the spin-down). The spin period P can range from milliseconds to several seconds, reflecting the rate at which the neutron star rotates. Its gradual increase over time ($\dot{P} > 0$) indicates a continuous loss of rotational energy, which is commonly attributed to the emission of magneto-dipole radiation. This process can be described as

$$I_{NS} \omega \dot{\omega} = \frac{2}{3} \frac{1}{c^3} \omega^4 B_s^2 R_{NS}^6 \sin^2 \alpha, \quad (1.1)$$

where I_{NS} is the neutron star moment of inertia, $\omega = 2\pi/P$ is the angular velocity, R_{NS} the stellar radius, B_s is the surface magnetic field strength, and α the angle between the magnetic and rotational axes. By assuming $\alpha = 90^\circ$ and adopting standard values $I_{NS} = 10^{45} \text{ g cm}^2$ and $R_{NS} = 10^6 \text{ cm}$, we obtain an estimate of the dipolar magnetic field at the surface B_s in Gauss (G):

$$B_s = 3.2 \times 10^{19} \sqrt{P \dot{P}} \text{ G}. \quad (1.2)$$

Similarly, integrating Equation (1.1) in time leads to the spin-down age (or characteristic age):

$$\tau_c = \frac{P}{2\dot{P}} \left(1 - \frac{P_0^2}{P^2} \right) \sim \frac{P}{2\dot{P}}, \quad (1.3)$$

where P_0 is the initial spin period, usually negligible with respect to the current P . These derived parameters allow us to place neutron stars within the so-called P - $\dot{P}$ diagram, a powerful tool to investigate their evolutionary tracks and to distinguish between different populations, such as normal pulsars and recycled (millisecond) pulsars. In this diagram, lines of constant B_s and constant τ_c can be drawn, providing a clear visualization of their spin and magnetic evolution (Figure 1.1). The P - $\dot{P}$ diagram therefore provides a convenient framework to compare neutron star populations and to trace their evolutionary histories. In the following sections, we will use this diagram as a reference tool to investigate the physical properties and evolutionary scenarios of different classes of pulsars.

1.1.3 Evolution

Newly formed neutron stars, typically observed as young pulsars, are born in the upper-left region of the P - $\dot{P}$ diagram. They are characterised by short




Figure 1.1: Pulsar data from the ATNF pulsar catalogue (Manchester et al., 2005), version 2.6.0. The death line is shown following case C of Chen and Ruderman (1993). Note that we use a single death line for simplicity, although a more realistic representation would involve a “death valley”, i.e. a region where the pulsar emission gradually fades out.

spin periods and relatively high spin-down rates, which reflect strong surface magnetic fields (see Equation 1.2).

Under the assumption of magnetic dipole braking in vacuum and a constant surface magnetic field, a neutron star spins down over time: its rotational period gradually increases while its spin-down rate $\dot{P}$ decreases (as follows from Equation 1.1). This secular evolution leads to a trajectory in the P - $\dot{P}$ diagram along lines of constant B and crossing lines of constant characteristic age. The star ultimately approaches the so-called *death line*, a theoretical boundary below which either the magnetic field is too weak or the spin rate is too slow to sustain radio emission via the normal, and as yet still highly uncertain, emission mechanism. In addition to these normal pulsars, a distinct population is found in the lower-left region of the P - $\dot{P}$ diagram, characterised by very short spin periods ($P < 30$ ms) and extremely small spin-period derivatives ($\dot{P} \sim 10^{-19} - 10^{-21}$). These so-called millisecond pulsars (MSPs) cannot be explained by standard

magnetic dipole spin-down from birth, as the required characteristic ages would exceed the age of the Universe. The widely accepted formation scenario involves a process referred to as *recycling*: an old neutron star is spun up through the accretion of mass and angular momentum from a binary companion to the point that they emit radio pulses despite their relatively low magnetic field strengths $B \sim 10^8 - 10^9$ G. This process typically occurs in low-mass X-ray binaries, interacting binary systems in which a neutron star accretes matter from a low-mass ($\lesssim 1 M_\odot$) companion star via Roche-lobe overflow. The accretion disk formed around the neutron star emits strongly in X-rays, hence the name of this class of systems (e.g. Bhattacharya and van den Heuvel, 1991; Tauris and van den Heuvel, 2006). Sustained mass transfer allows the neutron star to reach millisecond spin periods. Once the accretion ceases, the object re-emerges above the death line as a radio pulsar with stable, rapid rotation, observed today as an MSP. Strong observational support for this scenario comes from the orbital properties of binary MSPs (Tauris and van den Heuvel, 2023). Those with low-mass white dwarf companions ($M \sim 0.1 - 1 M_\odot$) generally exhibit low eccentricities, as a result of tidal circularization during the prolonged accretion phase. In contrast, highly eccentric binary pulsars often host neutron star companions, suggesting formation through a second supernova explosion, which imparted a velocity kick to the system. Moreover, globular clusters host a much higher fraction of recycled pulsars per unit stellar mass than the Galactic field, due to frequent dynamical interactions. These encounters can disrupt original binaries and lead to exotic systems such as MSPs in eccentric orbits or with main-sequence companions – a clear indication that the recycling donor was replaced during an exchange interaction.

1.1.4 Emission mechanism

At first order, the magnetic field of a pulsar can be described as dipolar. The corotation of this field with the rapidly spinning star generates extremely strong electric fields at the stellar surface, powerful enough to extract charged particles despite the intense gravitational pull of the neutron star. These particles fill the surrounding space, forming a plasma known as the magnetosphere. The region of closed magnetic field lines corotates with the star, while the so-called open magnetosphere is threaded by field lines that extend beyond the light cylinder, where rigid corotation would violate causality. It is within this open zone, or at the boundary between open and closed field lines, that most processes responsible for pulsar emission are thought to occur. Despite decades of study, the mechanism underlying the radio emission remains elusive. One reason is that the radio band carries only a tiny fraction of the pulsar’s total radiative output compared to higher-energy bands. At the same time, the radio signal is coherent, with implied brightness temperatures of $10^{26} - 10^{34}$ K, making it fundamentally different from the incoherent processes typically invoked for X-ray and gamma-ray emission.

Any successful theory of pulsar radio emission must therefore account for the extraordinary coherence and polarization properties of the signal, as well as the striking similarity of emission characteristics across pulsars that span four orders of magnitude in spin period, six orders of magnitude in magnetic field strength, and nearly ten octaves in observed radio frequency range (from ~ 100

MHz up to ~ 100 GHz). Several classes of models have been proposed, usually grouped into three broad categories: coherent emission from particle bunches (antenna-type mechanisms), relativistic plasma emission driven by instabilities and mode conversion, and maser-type processes interpreted as negative absorption. While each class has its own challenges and open issues, they share the common assumption that the emission originates in the inner magnetosphere, where pair cascades supply the dense plasma required for coherent radiation. For a more extensive discussion of pulsar emission mechanisms, the reader is referred to Melrose (1995), Lorimer and Kramer (2005), and Philippov, Timokhin, and Spitkovsky (2020). Progress in distinguishing between these scenarios will rely on wideband radio observations, high-precision polarimetry, and increasingly sophisticated simulations of pulsar magnetospheres.

1.2 Propagation of radio signals through the Ionised Interstellar Medium

As the radio signal emitted by a pulsar travels toward Earth, it passes through the ionized interstellar medium (IISM), which introduces frequency-dependent distortions to the observed signal. The IISM is an ionised medium whose turbulent structure gives rise to free-electron density fluctuations across a broad range of spatial scales. When radio waves propagate through this plasma, they are affected by phenomena such as dispersion, scattering, and Faraday rotation, which alter the pulse's arrival time, intensity, and shape. Accurate modelling of these effects is essential for high-precision pulsar analyses and also provides valuable insights into the physical conditions and turbulence within the IISM.

1.2.1 Dispersion of radio signals

Electromagnetic waves propagating through a cold plasma are governed by the following dispersion relation:

$$\omega^2 = \omega_e^2 + c^2 k^2, \quad (1.4)$$

where ω is the angular frequency of the wave (which is related to the electromagnetic frequency ν by $\omega = 2\pi\nu$), k is the wave number and c is the speed of light in vacuum. The quantity ω_e is the plasma frequency, which characterizes the natural oscillation rate of electrons in the plasma (Draine, 2011), and is given by:

$$\omega_e^2 = \frac{4\pi n_e e^2}{m_e} \quad (1.5)$$

where n_e is the electron number density, e is the elementary charge and m_e is the electron mass. The group velocity of the wave v_g is given by:

$$v_g = \frac{\partial \omega}{\partial k} = c \left(1 - \frac{\omega_e^2}{\omega^2} \right)^{1/2}, \quad (1.6)$$

which allows us to compute the arrival time t :

$$t(\nu) = \int_0^L \frac{dl}{v_g} \approx \int_0^L \frac{dl}{c} \left(1 + \frac{1}{2} \frac{\omega_e^2}{\omega^2} \right), \quad (1.7)$$

where L is the distance to the source and in the second step we assume $\omega \gg \omega_e$. Hence we can compute the difference between the arrival times at frequency ν and at the reference infinite frequency:

$$\delta t(\nu) = t(\nu) - t(\infty) = \frac{e^2}{2\pi\nu^2 m_e} \int_0^L n_e dl = \frac{K}{\nu^2} \text{DM}, \quad (1.8)$$

where $K = \frac{e^2}{2\pi\nu^2 m_e}$ is the dispersion constant to which by convention is given the value of $K^{-1} = 2.41 \cdot 10^{-4} \text{cm}^{-3} \text{pc MHz}^{-2} \text{s}^{-1}$ (see Kulkarni (2020) for a detailed discussion on this constant). We also introduce the *dispersion measure*:

$$\text{DM} = \int_0^L n_e dl \quad (1.9)$$

which quantifies the electron column density along the line of sight by integrating the number density of free electrons along the observer-source path.

Another way to approach this problem is by defining the refractive index of an ionised medium:

$$\mu = \sqrt{1 - \frac{\nu_e^2}{\nu^2}} \quad (1.10)$$

where $\nu_e = \omega_e/2\pi$ is the plasma frequency. The group velocity can then be obtained as

$$v_g = \mu c \quad (1.11)$$

which in the end gives the same result for the time delay as Equation (1.8).

Overall, the dispersion relation shows that the time delay induced by the interstellar medium depends on the observing frequency. In particular, since $\delta t \propto \nu^{-2}$, higher-frequency components of the pulse arrive earlier than lower-frequency ones with a quadratic sweep in the arrival times across the observed bandwidth of a single observation. Figure 1.2 illustrates this effect: in the top panels, each plot shows how the pulse intensity varies with observing frequency (vertical axis) and rotational phase (horizontal axis). The left panel displays the result of a correct de-dispersion (i.e., the removal of the dispersive delay by accounting for a specific value for the DM): the pulse remains aligned at the same phase across all frequency channels. In contrast, the right panel shows the effect of de-dispersing with an incorrect DM, where the pulse appears at different phases at different frequencies. The bottom panels show the corresponding pulse profiles obtained by summing the signal over all frequencies. When the correct DM is used (left), the pulse remains sharp and well-defined. When the DM is incorrect (right), the misalignment across frequency channels causes the integrated pulse to broaden and lose structure.

Given the inverse frequency dependence of the dispersive delay (see equation (1.8)), its magnitude increases at lower frequencies, making the effect easier to characterise. The appropriate way to quantify the sensitivity to DM of a given observing setup is through the fractional bandwidth f_B , as defined by Verbiest and Shaifullah (2018):

$$f_B = \frac{\nu_f - \nu_i}{\nu_c}, \quad (1.12)$$

where ν_f and ν_i are the upper and lower edges of the bandwidth, respectively, and ν_c is the central frequency. Hence, the dispersive delay Δt within the band




Figure 1.2: Impact of DM errors on pulse profiles of PSR J0358+5413. (a, b): Frequency vs phase plots with correct (a) and incorrect (b) de-dispersion. (c, d): Corresponding integrated intensity profiles. Using an incorrect DM value ($\delta DM = 7 \times 10^{-2} \text{ pc cm}^{-3}$) leads to smearing and pulse broadening.

can be computed as

$$\Delta t = K \text{ DM} \left(\frac{1}{\nu_1^2} - \frac{1}{\nu_2^2} \right) = \frac{K \text{ DM}}{\nu_c^2} \frac{2f_B}{\left(1 - \frac{1}{4}f_B^2\right)^2} \quad (1.13)$$

This formulation makes it clear that achieving high sensitivity to DM requires not only low observing frequencies, but also a sufficiently large fractional bandwidth. Figure 6 of Verbiest et al. (2008) clearly illustrates the DM-induced delay as a function of the fractional bandwidth.

1.2.2 Scattering

Beyond dispersion, pulsar signals are also affected by scattering, caused by small-scale fluctuations in the free electron density of the ionised IISM. These inhomogeneities act like a weakly refracting medium, deviating portions of the radio wavefront and producing multipath propagation. In pulsar observations, each recorded pulse profile is typically obtained by folding the signal over many rotations of the pulsar. Under the assumption that the intrinsic pulse shape remains constant over time, the observed profile corresponds to the average of many individual pulses. In the presence of interstellar scattering, however, a fraction of the radiation is deflected by small-scale density irregularities in the ionised interstellar medium, causing delayed arrival times with respect to the direct, unscattered rays. When all these pulses are averaged together, the unscattered radiation – which travels along the shortest path and thus arrives first – dominates the peak of the observed profile. Meanwhile, scattered rays travel along longer paths, producing delayed arrivals. In the most simple and widely adopted model for scattering, the thin screen approximation, all the scattering irregularities are assumed to be concentrated in a narrow layer along the line of sight. In this model, the small-angle deviations of the rays θ caused by random inhomogeneities in the electron density are typically assumed to follow a Gaussian distribution. The corresponding probability density function is

$$P(\theta) \propto \exp\left(-\frac{\theta^2}{2\sigma_\theta^2}\right), \quad (1.14)$$

where σ_θ represents the standard deviation of the scattering angles. The additional geometric path length for a scattered ray translates into a time delay approximately proportional to the square of the deviation angle, $t \sim D\theta^2/(2c)$ where D is the distance between the screen and the observer. Under this assumption, the Gaussian distribution of angles transforms into an exponential distribution of delays:


$$I(t) \propto \exp\left(-\frac{t}{\tau_s}\right), \quad (1.15)$$

where $\tau_s \sim D\sigma_\theta^2/(2c)$ is the characteristic scattering timescale. This produces the familiar exponential tail and the resulting asymmetry in the total pulse profile. It is important to note that this exponential form arises only under these specific assumptions. More complex distributions of inhomogeneities along the line of sight can lead to pulse broadening functions that deviate significantly from a simple exponential, making the real scattering behaviour often more complicated than the idealised thin-screen case. This phenomenon, known as

temporal broadening, arises from the multipath propagation caused by scattering in the interstellar medium. It becomes more pronounced at lower frequencies, where scattering is stronger. This is a consequence of the frequency dependence of the scattering timescale τ_s , which roughly follows a power law,

$$\tau_s \propto \nu^{-\alpha}. \quad (1.16)$$

In the thin screen approximation $\alpha = 4$. For an extended three-dimensional turbulent medium following a Kolmogorov spectrum, $\alpha \simeq 4.4$ (Rickett, 1977). As frequency decreases, radio waves undergo stronger multipath propagation effects, leading to more significant temporal broadening of the pulse. An example of such an effect is shown in Figure 1.3.



(a) Pulse intensity profile.

(b) Frequency versus pulse phase.

Figure 1.3: Pulse profile of PSR J0614+2229 (left) showing the characteristic exponential scattering tail. The right panel displays the frequency-phase plot divided into 20 frequency channels, highlighting the progressive broadening of the pulse at lower frequencies due to scattering.

1.2.3 Faraday rotation

In addition to dispersion and scattering, if the medium is not only ionised but also magnetised, electromagnetic waves are also subject to Faraday rotation. A linearly polarized wave, such as those typically observed from pulsars, can be decomposed into left- and right-hand circularly polarized components. These two components, in the presence of a magnetic field, propagate at different speeds, thereby introducing a phase offset between them. Indeed, the refractive index for a cold magnetised plasma becomes:

$$\mu = \sqrt{1 - \frac{\nu_e^2}{\nu^2} \mp \frac{\nu_e \nu_B}{\nu^3}} \quad (1.17)$$

where $\mp$ reflects the different propagation speeds for the left and right hand circular polarisations, respectively; ν_B is the cyclotron frequency defined as

$$\nu_B = \frac{eB_{\parallel}}{2\pi m_e c} \quad (1.18)$$

with $B_{\parallel}$ being the line-of-sight (LoS) component of the magnetic field B measured in Gauss. As a result, the plane of linear polarization rotates as the wave travels through the medium. The amount of rotation, known as the Faraday rotation angle Ψ , is proportional to the square of the observing wavelength λ and to the so-called *rotation measure* (RM), which quantifies the integrated effect of the electron density and the LoS component of the magnetic field:

$$\Delta\Psi = \lambda^2 \cdot RM \quad (1.19)$$

where RM is defined as

$$RM = \frac{e^3}{2\pi m_e^2 c^4} \int_0^L n_e B_{\parallel} dl. \quad (1.20)$$

The measurement of both RM and DM allows us to quantify the average magnetic field strength along the LoS weighted by electron density:

$$\langle B_{\parallel} \rangle = \frac{\int_0^L n_e B_{\parallel} dl}{\int_0^L n_e dl}. \quad (1.21)$$

1.2.4 Dispersion measure variations

As already mentioned before, DM quantifies the integrated column density of free electrons along the LoS to a pulsar. It is a crucial parameter in pulsar timing, as it introduces an observing frequency-dependent delay on the arrival times of radio pulses.

Several physical mechanisms can lead to temporal variations in the DM :

1. Relative motion between the pulsar, the Earth, and the turbulent IISM: as both the pulsar and the Earth move through the Galaxy, the line of sight to the pulsar shifts, intersecting regions with different electron densities due to the turbulent structure of the interstellar medium (Armstrong, Rickett, and Spangler, 1995). This combination of motion and turbulence results in time-variable integrated electron column density.
2. Presence of discrete structures: localised electron density enhancements such as HII regions (Ocker et al., 2024), filaments, or plasma lenses can induce sudden or systematic changes in DM as the line of sight moves through or near them.
3. Propagation effects from the solar system: contributions from the solar wind and the Earth's ionosphere can introduce additional frequency-dependent delays. The solar wind effects are typically accounted for, but may still impact low-frequency observations if not properly modeled, while the contribution from Earth ionosphere is typically negligible.

4. Radial motion of the pulsar: if the pulsar moves toward or away from the observer, the total path length over which the electron density is integrated changes with time, potentially modifying the observed DM (Lam et al., 2016).

In summary, DM variations are a key source of low-frequency noise in pulsar timing data. Their presence complicates the analysis in high-precision timing experiments and can limit the achievable accuracy of pulsar measurements, especially over long timescales. This is particularly relevant for the Pulsar Timing Array (PTA, Romani, 1989; Foster and Backer, 1990) experiments that rely on high-precision timing to search for correlated signals across a network of pulsars which aim to detect gravitational waves (GWs) in the nanohertz regime. In this context, even small, unmodeled DM fluctuations can mimic or obscure the expected signals. A more detailed discussion of PTAs and their sensitivity to such noise processes will be presented in Section 1.4.

1.2.5 Solar wind

In addition to the interstellar medium, the signal from a pulsar propagates through the Solar System plasma, which introduces an additional, time-dependent DM contribution. The primary source of this plasma is the solar wind (SW), that is a continuous, outward flow of charged particles, mainly electrons and protons, emitted by the Sun and that fills the heliosphere (Biermann, 1951). Due to Earth's orbital motion, the LoS from Earth to a pulsar occasionally passes very close to the Sun. These moments correspond to solar conjunctions, when the pulsar appears near the Sun in the sky, and the radio signal encounters regions of increased electron density. This leads to an enhanced dispersion delay, which depends strongly on the angular distance between the pulsar and the Sun, commonly referred to as the solar elongation angle.

Due to Earth's orbital motion, the LoS from Earth to a pulsar passes close to the Sun – particularly during solar conjunctions, when the pulsar is observed near the Sun in the sky – the radio signal encounters regions of increased electron density. This leads to an enhanced dispersion delay, which depends strongly on the angular distance between the pulsar and the Sun, commonly referred to as the solar elongation angle. A common model for the SW treats it as a time-independent, spherically symmetric distribution of free electrons, with density decreasing with radial distance as $n_e(r) \propto r^{-2}$. Under this assumption, the solar wind contribution to the DM follows a simple, analytic expression that depends on the solar elongation angle:

$$DM_{SW} = n_e \frac{\rho}{r_e \sin \rho} [1 \text{ au}]^2 \quad (1.22)$$

where ρ is the pulsar-Sun-observer angle, n_e is the electron density at 1 au and r_e is the distance between the observatory and the Sun (Tiburzi et al., 2019) and $n_e = 7.9$ at 1 au (optimal value found by Madison et al., 2019). However, recent studies (e.g., Susarla et al., 2024) have shown that the average electron density n_e strongly depends on the ecliptic latitude. In practice, the SW is highly dynamic and anisotropic, particularly during periods of elevated solar activity. As a result, the simple spherical model is often inadequate to capture its true variability, and more sophisticated, time-dependent models have been developed. A detailed description of one such model is provided in Section ??.

1.2.6 The turbulence of the IISM

One of the main goals of studying the IISM is to understand the nature and origin of its turbulent structure. Electron density fluctuations are observed over a broad range of spatial scales and are believed to arise from processes such as stellar feedback, magnetic activity, and large-scale Galactic motions. A standard way to describe these fluctuations is through a wavenumber spectrum, which quantifies how their strength varies with spatial scale: smaller physical scales correspond to higher wavenumbers. Determining both the shape and amplitude of this spectrum is essential for identifying the mechanisms responsible for injecting, cascading, and dissipating energy in the turbulent medium. Since the relevant structures span many orders of magnitude in size, direct measurements of the electron density spectrum are not feasible. Observables such as DM variations and scattering delays probe only restricted scale ranges and depend on the LoS geometry. As a result, interstellar turbulence is typically investigated by comparing these indirect measurements with the predictions of theoretical spectra. A widely adopted framework assumes that the fluctuations follow a power-law spectrum, as expected for fully developed turbulence. In this context, the three-dimensional power spectrum of electron density variations is commonly written as (e.g. Rickett, 1990)

$$P_{\delta n_e}(q) = C_n^2 q^{-\beta}, \quad (1.23)$$

where $q = 2\pi/l$ is the wavenumber associated to the spatial scale l , C_n^2 sets the amplitude of the fluctuations, and β is the spectral index. A Kolmogorov cascade (Kolmogorov, 1941), which describes an incompressible and fully developed turbulent flow, corresponds to $\beta = 11/3$, and is broadly consistent with many observations (Armstrong, Rickett, and Spangler, 1995). By estimating the spectral index and the fluctuation amplitude from pulsar observations, it becomes possible to test whether the turbulence in the IISM follows a Kolmogorov-like behaviour or deviates from it due to additional physical processes. This statistical description applies only over a finite range of spatial scales, bounded by two physical cutoffs. The outer scale l_o corresponds to the largest scale at which energy is injected into the turbulent cascade, while the inner scale l_i marks the smallest scale where dissipation becomes effective. For scales larger than l_o , the fluctuations decorrelate and the spectrum flattens. At scales below l_i , dissipation mechanisms suppress the fluctuations, leading to a steepening or cutoff in the spectrum. This implies that the power-law assumption is satisfied as long as $2\pi/l_o \lesssim q \lesssim 2\pi/l_i$. The normalisation factor C_n^2 varies along different LoS and reflects the local level of turbulence in the IISM. It can be estimated from observables such as pulse broadening, scintillation, or DM time series, and is often used to quantify the turbulence strength of a given interstellar path.

A more detailed investigation of interstellar turbulence, particularly within discrete structures such as HII regions, will be presented in Chapter 4. That analysis will be based on DM time series as well as scintillation observables.

1.3 Pulsar timing

Pulsar timing is a well-established technique that allows precise measurements of a pulsar's rotational and orbital parameters, as well as of the properties of the

medium through which its signal propagates. It requires time-stamped observations of the pulses and a predictive model, the timing model (see Section 1.3.1), which accounts for the pulsar’s rotation, signal propagation, and the relative position and motion of the pulsar and Earth. Based on multiple measurements of the pulse arrival times (the times-of-arrival, or ToAs), the parameters of this model can initially be constrained and ultimately measured with high precision, enabling a wide range of astrophysical studies. In practice, pulsar timing is an iterative process. After the discovery of a new pulsar, a preliminary timing model is constructed using initial estimates of the pulsar’s position, spin period, and dispersion measure, along with basic information about the observatory. The quality of the timing solution is progressively refined as more observations are accumulated and the model is updated.

During each rotation, the pulsar’s radio emission induces current variations in the telescope receivers. These signals are digitised and stored in a standard format (e.g., PSRFITS), organised along the dimensions of observing time, radio frequency channels, pulse phase, and intensity. For practical analyses, the data are grouped into sub-integrations (“subints”), each covering a short interval of observing time, which preserves temporal information while keeping the data manageable. The intensity may include polarisation information, but many timing studies consider only the total intensity (Stokes I). Ideally, pulsar observations are recorded with a very large number of narrow frequency channels and sampled at the highest possible time resolution, in order to preserve the intrinsic pulse shape accurately. However, for most pulsar timing applications, the initial high-resolution data are subsequently reduced by averaging over sub-integrations and by combining frequency channels. This reduction balances the need for manageable data volumes with the requirement to maintain sufficient information for precise timing measurements. One of the first steps in processing these data is *de-dispersion*. As described above, since large portions of the interstellar medium are ionised, the group velocity of radio waves depends on frequency, causing the pulse to be smeared in time. De-dispersion corrects for this effect using knowledge of the total column density of free electrons along the LoS, i.e., the DM. High resolution in frequency and time greatly improves the effectiveness of this correction. After de-dispersion, the pulses are folded over the spin period of the pulsar. For the majority of pulsars, individual pulses are too weak to be detected with sufficient signal-to-noise, so many tens to thousands of pulses are averaged together to produce a stable pulse profile. While single-pulse studies are occasionally performed for very bright sources, standard timing analyses rely on these folded profiles to define the ToAs. Folding preserves the pulse phase structure and forms the basis for all subsequent timing steps. Following folding, the data are typically further averaged over sub-integrations, frequency channels, and polarisation, a process commonly referred to as “scrunching”. In the most extreme case, the data are collapsed along all three axes, resulting in a single total-intensity profile as a function of pulse phase per observation. This averaging reduces data volume and improves signal-to-noise, while retaining enough information to extract accurate ToAs. To obtain the ToAs, a template pulse profile is constructed. This template represents an idealised, noise-free version of the pulsar’s intrinsic profile and can, for example, be created by summing many folded observations. Using cross-correlation between the observed profiles and the template, the rel-

ative phase offsets can be measured precisely, yielding the ToAs. Once the ToAs are generated, they are compared with the predictions from the timing model, and the differences, the timing residuals, are used to update and optimise the model. This optimisation, typically performed through a linearised least-squares fit, refines estimates of the pulsar’s rotational and orbital parameters, astrometry, and dispersion measure. In summary, this iterative process – from high-resolution recording, through de-dispersion, folding, averaging, template generation, cross-correlation, and timing model optimisation – results in a precise and self-consistent set of ToAs and a robust timing solution.

Pulsar timing based on this framework has proven extremely powerful not only for studying pulsar properties, their binary systems, and the ionised interstellar medium, but also for investigating the Solar System. Furthermore, by combining timing data from multiple pulsars, PTAs can search for correlated signals across the sky, such as those induced by GWs.

1.3.1 The pulsar timing model

The pulsar timing model provides a parametric description of the observed ToAs, accounting for the contributions that can be modelled through a well-defined set of physical parameters describing the pulsar and its environment. These include the pulsar rotational ephemeris, astrometric parameters such as position, proper motion, and parallax, and, for pulsars in binary systems, the orbital and post-Keplerian elements. The timing model also incorporates the transformation from the observatory to the Solar System barycentre using a Solar System ephemeris, as well as geometric and relativistic propagation effects, including the Roemer, Einstein, and Shapiro delays, and the bulk dispersive delays due to the interstellar medium and the solar wind (see e.g., Lorimer and Kramer, 2005). The relation between the observed ToAs and the timing model can be written as

$$t_{\text{obs}} = t_{\text{model}}(\boldsymbol{\theta}) + \delta t_{\text{noise}}, \quad (1.24)$$

where t_{obs} denotes the measured pulse time of arrival at the observatory and $t_{\text{model}}(\boldsymbol{\theta})$ is the timing model evaluated for a vector of pulsar parameters $\boldsymbol{\theta}$. The term δt_{noise} accounts for contributions that are not captured by the parametric timing model or that exhibit intrinsically stochastic behaviour, such as instrumental noise, intrinsic rotational irregularities of the pulsar, and time-variable dispersion measure fluctuations. When the timing model parameters are accurately inferred from the data, all predictable structures are removed from the ToAs, and the remaining timing residuals provide a direct probe of stochastic processes.

1.3.1.1 Effects of an incorrect timing model

While the timing model is designed to capture all deterministic contributions to the arrival times of pulses, imperfections in the model inevitably lead to residual systematics. These features are typically recognisable and can be traced back to specific parameters being inaccurate or missing.

For example, if the pulse period is not correctly determined, the residuals will exhibit a linear trend over time, reflecting a constant mismatch in the




Figure 1.4: Examples of timing residuals of PSR J0751+1807 for different model imperfections. **Top-left:** Flat residuals obtained with a correct timing model. **Top-right:** Linear trend produced by an incorrect pulse period. **Bottom-left:** One-year sinusoidal pattern resulting from an error in the pulsar’s position. **Bottom-right:** Sinusoid with increasing amplitude caused by an incorrect proper motion.

assumed rotation rate. An error in the pulsar’s sky position will result in sinusoidal residuals with a one-year period, caused by the incorrect projection of Earth’s motion along the line of sight. Similarly, a wrong estimate of the proper motion introduces a sinusoid whose amplitude increases linearly with time. Other effects can arise from unmodelled binary parameters, inaccurate dispersion corrections, or neglected relativistic contributions. In each case, the form of the residuals provides valuable diagnostic power, guiding improvements to the timing model.

Examples of the impact of timing model inaccuracies are shown in Figure ??, including a linear drift due to an incorrect pulse period (panel ??), and geometric timing errors such as an incorrect position or proper motion (panels ?? and ??).

1.3.2 Classification of noise processes in pulsar timing

In pulsar timing data, *noise* refers to any stochastic or unmodelled variations in the observed pulse times of arrival that cannot be accounted for by the deterministic timing model (see Section 1.3.1). Noise arises from a variety of physical and instrumental sources and manifests on different timescales. A useful clas-

sification distinguishes between short-term and long-term noise, depending on their impact on the timing residuals across different temporal baselines.

Short-term noise in pulsar timing data is often classified as *white noise*, as it is largely uncorrelated between observations. It typically includes radiometer noise and pulse jitter. Radiometer noise arises from thermal fluctuations in the receiver system and follows Gaussian statistics, decreasing with the square root of the observing time and bandwidth. Pulse jitter, instead, originates from stochastic variations in the shape and phase of individual pulses. Although the integrated pulse profile obtained after folding thousands of pulses is considered converged (see Helfand, Manchester, and Taylor, 1975, on this topic), a small residual variability remains, introducing a level of phase noise that does not average down as rapidly as radiometer noise. Together, these white-noise components dominate the timing residuals on short timescales and are largely uncorrelated between successive observations.

Long-term noise includes *red noise* processes, which exhibit more power at low Fourier frequencies and less power at high Fourier frequencies, leading to correlations over months to years. Among these, spin noise caused by irregularities in the pulsar’s rotation leads to slowly varying residuals and is more prominent in younger pulsars; it is often referred to as “timing noise”. Moreover, variations in the DM also contribute significantly (see Section 1.2.4). These DM variations are often smooth and can be modelled to some extent, but leave residual signatures if not properly accounted for. An in-depth description of the propagation effects on timing is given in the following section.

A third source is the so-called system noise, which may stem from back-end instabilities or calibration imperfections, and tends to mimic the statistical behaviour of other long-term processes. The overall noise budget must be carefully characterized, especially in the context of PTAs, where unmodelled or mischaracterised noise can bias the detection of correlated signals such as the gravitational-wave background.

A detailed description of the statistical models used to characterise these noise processes, including their treatment in PTA analyses, is presented in Section 2.3.1.

1.3.3 Impact of propagation effects on timing

1.3.3.1 Dispersion measure variations

As introduced in Section 1.2.4, the DM represents the integrated free electron content along the line of sight and causes frequency-dependent delays in pulsar signals. While timing analyses typically apply a fixed DM correction, temporal variations in the electron density lead to residual, unmodeled delays that can degrade timing precision.

Indeed, even after standard DM corrections, residual delays can persist due to unmodeled or time-variable dispersion. These fluctuations significantly impact high-precision timing analyses. An example is shown in Figure 1.5, where the data are fitted using a single DM value. Despite this correction, timing residuals remain offset across different frequency channels (indicated by different colors), suggesting that the DM varies on timescales shorter than the observing span and that a single DM value is insufficient to describe all frequency-dependent delays.




Figure 1.5: Post-fit timing residuals for PSR J0358+5413 after fitting for a single DM value. The time-variable offset between frequency channels indicates residual dispersion not captured by a static DM correction.

1.3.3.2 Solar wind

As discussed earlier (Section 1.2.5), the solar wind contributes a time-variable component to the dispersion measure along the line of sight. Its effect becomes especially pronounced near solar conjunctions, where the pulsar’s apparent proximity to the Sun causes enhanced dispersive delays. An example where the signature of the solar wind is clearly visible in the timing residuals is shown in Figure 1.6. Around solar conjunctions—i.e., when the pulsar lies closest to the Sun in the sky—the residuals exhibit a sudden increase (i.e., an un-modeled delay), caused by the enhanced electron density in the solar wind along the line of sight. The solar contribution to the total DM is variable in time and can significantly affect timing residuals, particularly in low-frequency observations where dispersive delays scale with ν^{-2} . If not properly corrected, this effect can mimic or mask the signatures of DM variations arising from the interstellar medium, or even introduce correlated structures that contaminate high-precision timing analysis and impact the overall sensitivity of PTA experiments (Tiburzi et al., 2016).

1.3.4 Applications of pulsar timing

Pulsar timing has proven to be an extraordinarily powerful tool in astrophysics and fundamental physics. By precisely measuring the times of arrival of pulses and modelling their variations, it allows us to probe neutron star properties, test theories of gravity, and study the interstellar medium (Lorimer and Kramer, 2005; Kramer, 2016).

One of the most important applications involves the use of pulsars as a galactic-scale detector for gravitational waves, which forms the basis of PTA experiments. These projects rely on timing ensembles of millisecond pulsars to search for correlated signals induced by passing gravitational waves, opening a unique window into the low-frequency gravitational wave universe. The next section introduces the principles behind PTAs, their current status, and their scientific potential.




Figure 1.6: Timing residuals for PSR J0030+0451 showing the effect of the solar wind. Residuals are plotted as a function of time (MJD) and colour-coded by observing frequency. The sharp increase in delay around solar conjunctions is caused by the enhanced electron density in the solar wind, which introduces a frequency-dependent dispersive delay not fully captured by standard timing models.

1.4 Pulsar Timing Arrays

Having discussed the fundamentals of pulsar timing and the various propagation effects that influence pulse signals – including DM variations and interstellar medium turbulence – we now turn to a key application of pulsar timing: PTAs (Romani, 1989; Foster and Backer, 1990).

PTAs exploit the exceptional rotational stability of millisecond pulsars to search for low-frequency GWs in the nanohertz regime. Such long-wavelength signals, produced for instance by inspiraling supermassive black hole binaries (Sesana, 2013), are expected to induce tiny deviations in the pulse ToAs measured on Earth. A key signature of a GW background (GWB) is that these timing deviations are not independent, but are spatially correlated between different pulsars, with a correlation pattern that depends on their angular separation on the sky.

In practice, each pulsar acts as a highly stable clock, and timing residuals can reveal the presence of perturbations affecting the spacetime metric. A passing gravitational wave modifies the proper distance between Earth and the pulsar, introducing a quadrupolar signature in the residuals of a network of pulsars spread across the sky. By monitoring an ensemble of such pulsars over many years, PTAs aim to detect a stochastic background of gravitational waves as a distinctive angular correlation pattern in the residuals, known as the Hellings & Downs curve (Hellings and Downs, 1983). A comprehensive overview of the current status of the field can be found in several recent reviews, including Tiburzi (2018), Burke-Spolaor et al. (2019), and Verbiest et al. (2024).

To date, eight PTA collaborations exist in the world: the European PTA (EPTA EPTA Collaboration et al., 2023a), the North-American Nanohertz Ob-

servatory for Gravitational waves (NANOGrav Agazie et al., 2023b), the Parkes PTA (Zic et al., 2023), the Indian PTA (InPTA Rana et al., 2025), the Chinese PTA (CPTA Lee, 2016), the MeerKAT PTA (MeerTIME Miles et al., 2025a), the African Pulsar Timing (APT¹) consortium and the International PTA (IPTA Perera et al., 2019). In 2023, five PTA collaborations presented results supporting the very first evidence for the presence of a GWB in PTA data (EPTA Collaboration et al., 2023c; Agazie et al., 2023a; Reardon et al., 2023; Xu et al., 2023), followed a year later by the sixth collaboration (Miles et al., 2025b). Although none of these studies proved the GWB presence with a significance higher than the required 5σ threshold (Allen et al., 2023), the results will be improved by the IPTA (formed by EPTA, NANOGrav, PPTA, InPTA, MeerTIME and APT), which will combine the data from all of the individual PTAs. The EPTA will not only contribute with its second data release (described in EPTA Collaboration et al., 2023a and EPTA Collaboration et al., 2023b), but also by enhancing the IPTA dataset. The combination of PTA datasets within the IPTA framework improves the overall electromagnetic frequency-domain coverage, increases the sky coverage by including more pulsars, and benefits from a higher observing cadence.

1.5 Thesis structure

This thesis is organised as follows:

Chapter 2 presents the datasets and methods used throughout this work. I describe the properties of the observations, including observing frequency coverage, time span, and instrumental systems, as well as the procedures adopted for pulsar timing and noise analysis

Chapter 3 focuses on the combination of low-frequency observations from LOFAR and NenuFAR with high-precision EPTA/InPTA data (DR2new+). It shows how this extended frequency coverage improves the characterisation of noise processes, while also discussing the limitations of current models in capturing certain propagation effects

Chapter 4 investigates the presence and impact of discrete ionised regions along the line of sight, such as HII regions. It presents ongoing work aimed at studying the turbulence within these regions, and assesses whether their statistical properties differ from those of the diffuse interstellar medium.

Chapter 5 provides a comparison of different methods for estimating variations in DM, based on simulated datasets. Particular focus is given to evaluating three widely used techniques in the context of PTA analysis, highlighting their strengths and weaknesses.

Finally, Chapter 6 summarises the main findings and outlines directions for future research.

¹<https://africanpulsartiming.github.io>

CHAPTER 2

Data and methods

This chapter provides an overview of the observational data and the analysis methods adopted, with a focus on the LOFAR and NenuFAR datasets and the techniques used for pulsar timing and noise characterization.

2.1 Observations

2.1.1 LOFAR

The LOw-Frequency ARray (LOFAR, van Haarlem et al., 2013) is a cutting-edge radio interferometer specifically designed to observe the sky at low radio frequencies. One of its most notable features is the wide field of view, which is particularly advantageous for surveys and studies of transient radio sources, and also enables efficient monitoring of multiple pulsars and regular calibration observations. This makes LOFAR a powerful instrument for pulsar timing, low-frequency pulsar surveys, and large-scale sky surveys. LOFAR consists of a network of modules (referred to as *stations*) distributed across Europe (see Figure 2.1), each made of two radio-antennae fields: the Low Band Antennas (LBA), operating between 10 and 90 MHz, and the High Band Antennas (HBA) (Figure 2.2a), covering the 110 to 240 MHz range (Figure 2.2b). The stations are operated together as an interferometer as the International LOFAR Telescope (ILT): their distribution is given by a dense central region of 24 stations in Exloo, the Netherlands (referred to as the *Core*), other 16 Dutch stations outside Exloo (the *remote stations*), and 13 *international* stations distributed across various partner countries.

The network continues to expand, with new members joining the collaboration and planning to establish their own stations. Figure 2.1 shows the distribution of the LOFAR stations, including the ones that will be part of the array once the upgrade to LOFAR2.0 will be completed¹.

The international LOFAR stations can operate as stand-alone interferometers and are primarily used for pulsar observations, enabling regular monitoring

¹<https://www.lofar.eu>, LOFAR official website.





Figure 2.1: Distribution of the LOFAR stations in Europe, as presented in the LOFAR2.0 White Paper (https://www.lofar.eu/wp-content/uploads/2023/04/LOFAR2_0_White_Paper_v2023.1.pdf).

of multiple sources. Most pulsar observations are carried out using the HBA antennas of eight international stations: Nançay, France (FR606); Onsala, Sweden (SE607); and six stations belonging to the German Long Wavelength (GLOW) consortium, namely Effelsberg (DE601), Unterweilenbach (DE602), Tautenburg (DE603), Potsdam-Bornim (DE604), Jülich (DE605), and Norderstedt (DE609). Observations are typically performed on a weekly basis, with observing sessions lasting between one and three hours. Each international station is equipped with 96 LBA dipoles and 96 HBA tiles. Further information on the GLOW consortium is available on the official website².


The more sensitive LOFAR Core plays a key role in complementary monthly-to-bimonthly observations due to its higher signal quality. The Core consists of multiple stations, each with 48 LBA dipoles and 48 (2×24) HBA tiles, whose signals can be coherently combined to form a tied-array beam, significantly increasing sensitivity and improving resolution. While each dipole has an almost all-sky field of view, the resulting beam narrows as signals are combined. Within each HBA tile, initial beamforming is performed in analog before further digital processing at the station level. Observing modes depend on the clock system in use; for standard pulsar observations, a Rubidium clock is typically employed, enabling bandwidths of 100-200 MHz.

The two LOFAR front-ends, HBA and LBA, exhibit different bandpass responses (see Figure 2.2b). The HBA response is relatively uniform across its frequency range, whereas the LBA bandpass is not flat over the 10-85 MHz band, showing a sensitivity peak at mid-band frequencies.

²<https://www.glowconsortium.de>



(a) HBA tiles



(b) LBA dipoles

Figure 2.2: LOFAR antenna types. **(a)** HBA tiles and **(b)** LBA dipoles. For each panel, the image on the left shows the antenna array, while the image on the right displays the median averaged spectrum of the corresponding elements, as presented in van Haarlem et al. (2013).

2.1.1.1 Pulsar observations

During a standard LOFAR observation with an international station in Germany operating in standalone time, the signals from the individual antennas of that station are coherently summed locally to form a single station beam. In this mode no interferometric combination with the Core or other stations is carried out, and the station effectively behaves as an autonomous telescope. The beamforming is handled by the local beamformer, which produces a tied signal directed toward the target source. Once the acquisition is complete, the beamformed data are transferred via high-speed links to the Max Planck Institute for Radio Astronomy (MPIfR) and Forschungszentrum Jülich (FZJ), where they are handled by the LuMP (LOFAR und MPIfR Pulsare) system. LuMP manages the recording of the data streams to disk and prepares them for near real-time processing. The software package DSPSR (van Straten and Bailes, 2011) is then used to fold the observations, computing phase-resolved averages over time and frequency. This step reduces the data volume from the order of terabytes to a few gigabytes per observation. The folded archives are stored in the timer or PSRFITS format (Hotan, van Straten, and Manchester,

2004) as data cubes with several dimensions: time (10-second sub-integrations), frequency (195.3125 kHz wide), full Stokes polarization, and rotational phase (typically 128-1024 bins, depending on the observed source). For instrumental reasons, the observing band is split into three or four frequency “lanes”, each covering a different part of the total band.

Over the years the observing setup has changed. Early observations (2012) covered 47.7 MHz (244 channels) centered at 138.8 MHz. In 2013 a new mode was adopted, reducing data rates while doubling the frequency band, yielding 488 channels across 95.3 MHz centered at 149.9 MHz. A further adjustment occurred in 2015 when additional stations were incorporated. To accommodate these stations, the observing bandwidth of the German international stations (excluding DE601) was reduced to 71.5 MHz (366 channels), and the center frequency was shifted to 153.8 MHz to remain in the most sensitive part of the band. The same configuration was adopted for station DE601 in 2019. All frequency shifts were always performed by integer multiples of the channel width, ensuring that the absolute frequencies of the channels remained fixed across the dataset.

In parallel to the standalone observations with the international stations, a separate observing programme is carried out with the LOFAR Core. In this case, the signals from many Dutch stations are coherently combined at the correlator to form a tied-array beam. The resulting effective collecting area is much larger than that of any individual international station, which allows for the detection of significantly fainter pulsars. However, because Core time is not allocated by station-owner policies and is in high demand for both pulsar work and other LOFAR science, these observations are less frequent and typically scheduled once every one to two months, with individual scans lasting between 5 and 30 minutes. The raw data are processed using the dedicated pulsar pipeline Pulp (Stappers et al., 2011; Kondratiev et al., 2016), producing PSRFITS archives with the same time, frequency, polarization and phase resolution as the GLOW dataset. Each Core observation covers 78.1 MHz of bandwidth (400 channels) centered at 148.9 MHz.

2.1.1.2 Radio frequency interferences

Astronomical signals at LOFAR frequencies are many orders of magnitude fainter than emissions generated by human technology, so the very first step before processing the data is to identify and remove contaminating radio-frequency interference (RFI). The RFI excision was carried out with a customised version of the SURGICAL algorithm from the COASTGUARD python package (Lazarus et al., 2016) wrapped inside the PSRCHIVE framework (Hotan, van Straten, and Manchester, 2004).

2.1.2 NenuFAR

The New Extension in Nançay upgrading LOFAR (NenuFAR) is a highly-sensitive low-frequency radio interferometer located at the Nançay Radio Observatory, in the Sologne region of central France. The site hosts several radio telescopes operating below 3 GHz and has long been a key location for French radio astronomy.

Optimised for the 10-85 MHz band (corresponding to wavelengths between 4 and 30 meters), hence matching the LBA’s frequency coverage, NenuFAR operates in the lowest spectral window accessible from Earth, and it compensates for the peaked LBA bandpass with a more uniform sensitivity across the bandwidth. As such, NenuFAR is one of the most powerful instruments in the world at these frequencies. Within this range, the telescope offers excellent sensitivity, enabling high-quality observations of pulsars, transients, and cosmic background signals.

NenuFAR is composed of crossed inverted-V dipole antennas (see Figure 2.3a, similar in orientation to those used in LOFAR). Each antenna is equipped with a low-noise preamplifier that provides a stable response across the 10-85 MHz band. Thanks to this setup, the system is sky-noise limited across the entire frequency range.

The NenuFAR array follows a hierarchical layout (Zarka et al., 2012), where the basic elements, the crossed dipole antennas, are grouped into mini-arrays, each consisting of 19 antennas arranged in a hexagonal configuration. These mini-arrays, with a diameter of about 25 meters, are analog-phased to form a broad and steerable beam. A total of 96 mini-arrays are densely packed within a central core region approximately 400 meters in diameter, while an additional 6 mini-arrays are positioned further out, extending the array’s baseline to distances of up to 3 km.




Figure 2.3: NenuFAR antennas (left) and a view of the field with the mini-arrays (right). Image from <https://nenufar.obs-nancay.fr/en/homepage-en/>.

Pulsar observations have been conducted as part of the NenuFAR Pulsar Key Project, an ongoing monitoring program active since 2019. The observing cadence is bi-weekly or monthly (depending on the source), and the program currently monitors around 40 pulsars, including a few millisecond pulsars. For a detailed description of the observing setup and strategy, see Bondonneau et al. (2021).

The RFI cleaning stage at NenuFAR is performed using a dedicated pipeline called `NenuPlot.py`³, a tool for folded pulsar observations built on top of PSRCHIVE. This program processes folded FITS archives and applies RFI mitigation through a modified version of the COASTGUARD software.

³The source code is available on GitHub: <https://github.com/louisbondonneau/NenuPlot>

2.2 Pulsar timing at low-radio frequencies

2.2.1 Template creation and timing procedure

Once the data have been cleaned of RFIs, the next step is to extract the pulse Times of Arrival (ToAs) and refine the pulsar’s timing model. This process relies heavily on the construction of a high signal-to-noise template profile, which serves as a reference for measuring the arrival times of pulses across all observations.

2.2.1.1 Template creation

A pulse template represents the noise-free average profile of a pulsar, providing a stable reference for timing measurements. To construct such a template, multiple observations collected at different epochs are summed together using the PSRCHIVE tool PSRADD. This process ensures a high signal-to-noise ratio and preserves the intrinsic shape of the pulse across frequency. In the case of LOFAR and NenuFAR, observations acquired at solar angles smaller than 40° are excluded from the template construction. At such small angles, the line of sight passes close to the Sun, where the density of free electrons is significantly higher. This excess of electrons adds an extra contribution to the dispersion measure, which can lead to pulse broadening and biased timing. Section 1.2.5 describes how the solar wind is typically modelled in pulsar timing analyses.

When summing multiple archives, special attention must be paid to whether DM derivatives are present in the timing model. If so, the resulting template may appear artificially broadened, since PSRADD does not account for DM variations over time. To avoid this effect, the headers of each archive are adjusted to correct for the time-dependent DM, ensuring proper phase alignment. This results in a narrower and more accurate template profile.

Once the archives are combined, the template is smoothed to further suppress residual noise. Two PSRCHIVE tools are commonly used for this step: PSRSMOOTH (Demorest et al., 2013), which applies a wavelet-based smoothing algorithm, and PAAS, which fits an analytic model composed of a sum of von Mises functions (Jammalamadaka and Sengupta, 2001). These functions resemble Gaussians in shape but are periodic, making them better suited for modelling data defined on a circular domain such as pulse phase. Each component is defined as:

$$f(x) = A \cdot e^{k(\cos(2\pi(x-\mu))-1)} \quad (2.1)$$

where A is the amplitude, k controls the width of the component, and μ is the central phase. The phase variable x spans the interval from 0 to 1. This approach is especially useful when dealing with noisy data, as it can produce a smooth, analytic template even in low signal-to-noise conditions. However, since the model is constrained by the functional form, the resulting profile may not perfectly match the true shape of the pulsar signal. The final result is a clean, high-fidelity standard profile that can be reliably used for ToA extraction.

It is important to note that, in both LOFAR and NenuFAR, the final templates are not fully averaged (or *scrunched*, in data-analysis jargon) in frequency, in order to preserve the frequency-dependent structure of the pulse profile. At such low observing frequencies, the shape of the pulse changes more notice-

ably across the band, and retaining sufficient frequency resolution allows these variations to be properly tracked when they occur (Lorimer and Kramer, 2005).

For LOFAR, the number of frequency channels in the template ranges from 8 to 20 depending on the pulsar’s signal-to-noise ratio, while for NenuFAR a fixed configuration of 8 channels is adopted. This setup is motivated by the strong profile evolution seen at low radio frequencies: the pulse shape can vary significantly across the band due to both intrinsic and propagation effects. Preserving frequency resolution prevents the mixing of profiles with different morphology, which would otherwise bias the timing. This is particularly important in narrow-band timing, where the data are divided into many narrow frequency channels so that the pulse profile can be treated as approximately constant within each channel. This assumption only holds if the channel width is small enough, which makes the choice of frequency resolution particularly important in our case. Avoiding full frequency scrunching thus ensures that the extracted ToAs remain accurate and that DM-related delays or shape changes do not degrade the timing precision.

2.2.1.2 ToA extraction and timing refinement

The most fundamental measurement in pulsar timing is the ToA, which marks the arrival time of a pulse at the telescope, referenced to a standard clock. Several methods exist to compute ToAs, but the most widely used approach relies on cross-correlating observed profiles to the high S/N template, that is created as described in the previous section. The tool used for this purpose is PAT, part of the PSRCHIVE suite, which offers various algorithms for extracting ToAs. Following Taylor (1992), the cross-correlation is done in the Fourier domain, where a time shift corresponds to a continuous phase rotation, enabling ToA measurements with precision finer than a single profile bin. Each ToA corresponds to a given frequency channel and time stamp, and is associated with metadata such as observing backend, frequency resolution. Once the phase offset between the observed profile and the template has been measured, it is converted into a time shift Δt . The corresponding ToA is then obtained by adding this shift to the reference timestamp of the observation, t_0 , which typically corresponds to the start of the observation. Once all measured ToAs are collected, they are compared with the expected arrival times from the timing model to compute the timing residuals (see Section 1.3.1). To minimize these residuals, the timing model can be fitted, yielding an updated set of parameters. The resulting ephemeris includes improved estimates of rotational parameters (e.g., spin period and its derivative), astrometric information (e.g., position and proper motion), and DM. The updated ephemeris can then be applied to the original archives, allowing the extraction of more precise ToAs. A new template is subsequently built using the same method as before, but now with improved phase alignment across the frequency band. This process – extracting ToAs and refitting the timing model – is repeated iteratively, refining both the template and the timing parameters, until convergence is reached and a consistent set of ToAs and a final timing model are obtained.

For comparison, at higher frequencies such as L-band (1400 MHz), pulse profiles are observed to be generally more stable across the observing band, and the effects of profile evolution and dispersion are reduced. As a result,

fewer frequency channels are typically sufficient, and the iterative procedure for refining the template and ToAs is less critical than at low radio frequencies.

2.3 Noise in pulsar timing

In pulsar timing, one aims to measure the ToA of pulses with the highest possible precision and to minimize the weighted root mean square of the timing residuals. However, various noise processes affect the latter, introducing deviations between the observed ToAs and those predicted by the deterministic timing model. Accurately characterising these noise sources is therefore essential not only to improve the timing precision, but also to avoid biased estimates of astrophysical parameters or false detections of physical signals, such as gravitational waves.

Noise in timing data is traditionally classified into white and red components. In the low-radio-frequency regime, such as that probed by LOFAR and NenuFAR, red noise is often dominated by propagation effects through the ionised interstellar medium, particularly DM variations induced by turbulent plasma along the line of sight.

In the following sections, we describe the specific models and tools used to carry out this analysis.

2.3.1 Noise modelling methods

The analysis and characterization of noise in pulsar timing data employ a variety of approaches, reflecting the complexity of the underlying physical processes and the observational strategies. While Bayesian inference frameworks, such as those implemented in software packages like ENTERPRISE (Ellis et al., 2020) and TEMPONEST (Lentati et al., 2014), have become widely used for simultaneously estimating timing parameters and stochastic noise components, alternative methods remain important, especially for certain noise contributions. In particular, variations in the DM, which are especially significant at low radio frequencies, can be modelled either through stochastic noise processes incorporated in a global Bayesian framework, or by fitting DM values on an observation-by-observation basis using traditional timing techniques. These complementary approaches allow flexibility in handling DM variations depending on data quality and analysis goals.

In Bayesian approaches, the analysis is typically performed by including the noise components directly in the time-domain likelihood function (van Haasteren and Vallisneri, 2014; Chalumeau et al., 2022). This function is generally expressed as a multivariate Gaussian, assuming the noise-reduced timing residuals follow a Gaussian distribution. The residuals, denoted as δt , represent the differences between the observed and modelled ToAs, after subtraction of a deterministic timing model. The likelihood L is written as:

$$\log L(\tilde{\delta t}|\tilde{\theta}, \tilde{\theta}') = -\frac{1}{2} \left(\left[\tilde{\delta t} - \tilde{y}(\tilde{\theta}') \right]^T C^{-1}(\tilde{\theta}) \left[\tilde{\delta t} - \tilde{y}(\tilde{\theta}') \right] \right) + \text{cst.}, \quad (2.2)$$

where the timing residuals $\vec{\delta t}$ are reduced by any deterministic time-domain waveform $\vec{y}$ with parameters $\vec{\theta}'$, and where the covariance matrix is composed

of the sum of all the covariance matrices from all the stochastic components λ with parameters θ_λ such that $C(\vec{\theta}) = \sum_\lambda C^\lambda(\vec{\theta}_\lambda)$. This formulation allows for a coherent treatment of multiple noise sources and their covariances, and it forms the basis for the use of Markov Chain Monte Carlo (MCMC) methods to explore the parameter space.

In the following sections, we detail the characterization of the main noise components that impact our observations.

2.3.1.1 White noise

In an ideal scenario where the timing model perfectly describes the observed ToAs, the resulting timing residuals, once normalized by their uncertainties, would follow a normal distribution. In practice, however, the measurement of ToAs and their associated uncertainties, typically attributed to radiometer noise, are affected by various sources of stochastic error. These errors are usually stochastic and uncorrelated in time, and are therefore referred to as white noise, as they correspond to a flat power spectral density (PSD) across all sampling frequencies.

White noise does not introduce long-timescale trends in the residuals and manifests as a diagonal covariance matrix in the time domain. While it lacks temporal correlations, its accurate characterization remains crucial, as it directly impacts the weighting of individual observations and, consequently, the estimation of all other model parameters. There are three main sources of white noise in pulsar timing analysis:

Radiometer noise and EFAC: The formal ToA uncertainties reflect the precision of the measurement process and are primarily limited by radiometer noise. These uncertainties are estimated during the standard ToA computation via template matching, but they may not capture all sources of small, uncorrelated error such as minor instrumental instabilities. To account for this, an empirical correction factor called EFAC (Extra FACTor) is applied as a multiplicative rescaling of the reported uncertainties.

EQUAD: Additional uncorrelated white noise contributions may arise from independent sources such as digitization effects or small instrumental errors not captured by EFAC. These are modelled via an additive term called EQUAD (Extra QUADrature), added in quadrature to the formal uncertainties.

Pulse phase jitter and ECORR: As mentioned in Section 1.3.2, when folding a finite number of pulses during a single observing epoch, small statistical fluctuations in pulse shape may occur due to the limited number of averaged rotations. This phenomenon, known as pulse phase jitter, introduces a white noise component in the timing residuals, independent from radiometer and instrumental noise. Since the number of pulses per epoch is typically stable across the dataset, this noise component is usually characterized per receiver-backend system. Moreover, during a single epoch, it is common to extract multiple simultaneous ToAs across adjacent frequency channels. Because these ToAs are derived from the same folded pulse train, the jitter-induced noise will be fully correlated across channels

within the epoch, but uncorrelated between epochs. This motivates the inclusion of an Extra CORRelated (ECORR) noise term, which models this specific structure: white noise that is epoch-uncorrelated but frequency-correlated within each epoch (Cordes and Shannon, 2010; Shannon and Cordes, 2012; Liu et al., 2012).

The white noise terms are incorporated into the likelihood covariance matrix C as follows:

$$C_{\alpha\beta}^{\text{WN}} = \begin{cases} (\text{EF}_\alpha^2 \times \sigma_{\text{init},i}^2 + \text{EQ}_\alpha^2) \delta_{ij} \delta_{\alpha\beta}, & \text{if } \alpha \text{ does not include EC term,} \\ [(\text{EF}_\alpha^2 \times \sigma_{\text{init},i}^2 + \text{EQ}_\alpha^2) \delta_{ij} + \text{EC}_\alpha^2 \delta_{\text{E}(i)\text{E}(j)}] \delta_{\alpha\beta}, & \text{else,} \end{cases} \quad (2.3)$$

where EF, EQ and EC stand for EFAC, EQUAD and ECORR respectively, the indices i, j and α, β range over the ToAs and the observing systems respectively and $\text{E}(i)$ indexes the epoch of the ToA i . The symbol δ refers to the Kronecker delta, as in the following parts of this thesis.

2.3.1.2 Time-correlated stochastic signals

Time-correlated stochastic signals in pulsar timing, such as red noise or DM variations, are modeled as Gaussian processes (GPs). Each component, denoted by λ , is described by a covariance matrix that captures the correlations between timing residuals at different epochs.

Following van Haasteren and Vallisneri, 2014 and Rasmussen and Williams, 2006, the covariance matrix for a given component λ can be expressed using a set of basis functions $F_\mu^\lambda(t)$ and a covariance matrix in the weight space $\Phi_{\mu\nu}^\lambda$:

$$C^{\lambda, \text{GP}}(t_i, t_j) = \sum_{\mu, \nu} F_\mu^\lambda(t_i) \Phi_{\mu\nu}^\lambda F_\nu^\lambda(t_j), \quad (2.4)$$

where t_i and t_j are the observing epochs, and the indices μ and ν run over the number of basis functions used to describe the process. This formulation allows a flexible and computationally efficient representation of the GP, while preserving the correlations in time. The resulting covariance matrix $C^{\lambda, \text{GP}}$ enters the time-domain likelihood (Equation 2.2), enabling a coherent treatment of multiple time-correlated noise components simultaneously. By construction, different components λ are additive, and the total covariance matrix of the stochastic signals is

$$C_{\text{total}} = \sum_{\lambda} C^{\lambda, \text{GP}}. \quad (2.5)$$

The choice of basis functions and the number of components is guided by the expected spectral properties of each noise type, allowing the GP to accurately capture the long-term correlations observed in the timing residuals.

This framework is used to model both achromatic (frequency-independent) red noise and time-dependent chromatic (frequency-dependent) signals arising from the complex properties of the IISM. Chromatic effects include DM variations that are not captured by low-order DM derivatives, as well as additional

frequency-dependent delays such as scattering. Under the thin-screen approximation with inhomogeneities of comparable size, scattering can produce a ν^{-4} delay (Cordes and Shannon, 2010)⁴.

These noise components, labeled **RN**, **DM ν** , and **CN $_{\chi}$** (where χ is the chromatic index), are modelled in the PSD domain using an incomplete Fourier basis approach. In this framework, the basis functions are constructed as sine and cosine pairs for each Fourier frequency f , and the corresponding Gaussian process covariance matrices are defined as:

$$\Phi_{\mu\nu}^{\lambda} = S^{\lambda}(f_{\mu}) \Delta f \delta_{\mu\nu}, \quad (2.6)$$

where $S(f_{\mu})$ is the PSDs amplitude at frequency f_{μ} and Δf is the spectral resolution, equal to the inverse of the total time span of the time series. Unless it is specified, the PSD of these components are modelled as power-laws:

$$S^{\lambda}(f) = \frac{A_{\lambda}^2}{12\pi^2} \left(\frac{f}{f_{\text{yr}}} \right)^{-\gamma_{\lambda}} f_{\text{yr}}^{-3}, \quad (2.7)$$

where f_{yr} refers to the frequency of 1yr^{-1} and A_{λ} , γ_{λ} are respectively the amplitude referred at f_{yr} and the spectral index of the noise component λ . The different noise components are $\lambda = \{\text{RN}, \text{DM}\nu, \text{CN}_{\chi}\}$, which differ from their chromaticity (i.e., dependence on the observed radio-frequency ν) that is added in the basis functions as $F_{\mu}^{\lambda}(t) \propto (\nu/\nu_{\text{ref}})^{-\chi}$, where $\chi = 0, 2, n$ for **RN**, **DM ν** and **CN $_n$** respectively, and ν_{ref} being a reference frequency set at 1.4 GHz in this work. While the chromatic index χ is fixed for these components, it is also possible to fit for this parameter simultaneously with other parameters and hyper-parameters of the model.

The timing model parameter uncertainties are as well modelled as GPs also by using the covariance-function expansion (cf. Equation 2.4), but they are analytically marginalized over. The basis functions are the timing model parameter design matrix and the Gaussian process covariance matrix is a diagonal matrix with very large values (Chalumeau et al., 2022).

2.3.1.3 Solar wind

The SW is an outflow of charged particles from the Sun which can significantly affect the DM especially for observations of pulsars near the ecliptic plane. Properly modelling the impact of SW is crucial for precise timing analysis, especially in presence of very low frequency data. A detailed description of the latest model developed within the PTA community to account for the SW is presented in Susarla et al. (2024), where they include a deterministic and a stochastic component to the model. The deterministic contribution follows the standard formulation of Equation (1.22), in which the SW-induced DM depends on the electron density at 1 au parametrised by n_e (denoted as **NE_SW** in **TEMPO2**) and the pulsar-Sun-observer angle. The stochastic term accounts for secular variations in the SW by incorporating a stochastic SW signature

⁴The chromatic scaling of IISM-induced delays strongly depends on the detailed structure of the electron density distribution and may vary between pulsars. For instance, a $\nu^{-4.4}$ dependence is predicted for a turbulent extended medium described by Kolmogorov theory, while refractive propagation effects can lead to a steeper scaling, up to $\nu^{-6.4}$ (Shannon and Cordes, 2016).

modelled with GPs. This approach, implemented as SWGP, is available in the chromatic noise module of ENTERPRISE_EXTENSIONS (Taylor et al., 2021). It uses a power spectral density described by the following power-law formula:

$$S_{\text{SW}} = A_{\text{SW}}^2 \left(\frac{f}{\text{yr}^{-1}} \right)^{-\gamma_{\text{SW}}} \text{yr}^3, \quad (2.8)$$

where S_{SW} is the power spectral density for SWGP, A_{SW} is the SW amplitude at a frequency of $1/\text{yr}$, γ_{SW} is the spectral index, and f is the Fourier frequency.

2.3.1.4 *Ad hoc* noise component: exponential dips

It is known a priori for PSR J1713+0747 that timing residuals contain multiple exponential dips: a brief radio-frequency dependent drop (where the ToAs are measured before the time predicted by the timing model) followed by an exponential recovery for several weeks/months (Coles et al., 2015; Lam et al., 2018; Xu et al., 2021). Depending on the corresponding chromatic index χ^E , the origin of these events are either considered to be related with a drop in the IISM content along the line-of-sight (Lam et al., 2018, if $\chi^E = 2$), or related to magnetospheric effects if $\chi^E < 2$ and if profile changes are observed at the time of the event. Following (EPTA Collaboration et al., 2023b), this effect is modelled as

$$y^E(A_E, \tau_E, t_{0,E}, \chi_E) = \begin{cases} 0, & \text{if } t < t_{0,E} \\ A_E (\nu/1.4 \text{ GHz})^{-\chi_E} \exp\left(-\frac{t-t_{0,E}}{\tau_E}\right), & \text{if } t \geq t_{0,E} \end{cases} \quad (2.9)$$

where A^E , χ^E , t_0^E and τ^E are respectively the amplitude (in residual unit), the chromatic index, the reference epoch (in MJD) and the relaxation time (in days) of the exponential dip modelled for a time t and a radio-frequency ν .

CHAPTER 3

Datasets: combination of LOFAR and EPTA

This chapter presents the combination of low-frequency data from LOFAR with the second data release of the European Pulsar Timing Array. It describes the procedure adopted to merge the datasets and the methods used for noise modelling. Particular attention is given to the sensitivity of the combined dataset to propagation effects in the interstellar medium, especially those related to the solar wind (based on Iraci et al., 2025).

3.1 Motivations for this combination

The integration of very low-frequency radio observations, specifically from LOFAR (100-200 MHz) and NenuFAR (30-90 MHz), into a PTA analysis framework represents a key step of this work. By combining these data with the latest release from the European and Indian PTAs (restricted to the DR2new+ subset, which includes only new backend systems EPTA Collaboration et al., 2023b), we construct a new dataset, DR2low, which spans 12 pulsars, ~ 11 years of observations, and an unprecedented frequency range from 30 to 2500 MHz. The primary aim of this work is to explore the role of very low frequency data in improving the characterization of noise processes in PTA datasets – particularly dispersion measure variations – which are critical for enhancing the sensitivity of PTA experiments to gravitational wave signals. This chapter demonstrates the feasibility and scientific value of including very low frequency data in PTA analyses and lays the foundation for their systematic incorporation into future international PTA data releases, including the upcoming IPTA DR3, which targets a 5σ detection of the stochastic gravitational wave background.

3.2 The dataset

3.2.1 EPTA-InPTA DR2

The data combination and noise analysis for the EPTA/InPTA are presented in EPTA Collaboration et al. (2023a) and EPTA Collaboration et al. (2023b)

respectively. In this data release, EPTA and InPTA combined data collected from seven main European and Indian radio telescopes, i.e. the Nançay Radio Telescope (NRT), the 100-m Telescope at Effelsberg, the Lovell Telescope, the upgraded Giant Meterwave Radio Telescope (GMRT), the Westerbork Synthesis Radio Telescope, the Sardinia Radio Telescope and “Large European Array for Pulsars” (LEAP), from mid-1990s to 2021. Instead of analyzing the 42 pulsars presented in the first European data release (Desvignes et al., 2016), the authors selected 25 pulsars only, following an approach to maximize the sensitivity to the GWB and continuous GWs with a minimum number of pulsars (Speri et al., 2022). These pulsars and their characteristics are reported in Tables B1-B7 of EPTA Collaboration et al. (2023a).

The data were combined using time offsets known as **JUMPs** with respect to the reference dataset, chosen to be the L-band NUPPI backend from the NRT. After that, initial timing solutions were obtained using the TEMPO2 timing software, and refined using the TEMPONEST Bayesian toolkit (Lentati et al., 2014) to introduce white noise (WN) parameters, red noise (RN; e.g., originating from rotational instabilities), and chromatic noise models (i.e., long-period noise that has some dependence on the observing frequency).

By considering only the new EPTA – more precise and more uniformly distributed across the frequency bands – combined with the InPTA dataset, the so called **DR2new+** dataset was constructed. This combination showed evidence for a GWB with a statistical significance exceeding 3.5σ (EPTA Collaboration et al., 2023c). The importance of radio-frequency coverage in the EPTA data was further investigated in Ferranti et al. (2025), which demonstrated that poor radio frequency coverage can significantly reduce the sensitivity of PTAs to a GWB signal. As discussed in Section 3.3, this work focuses exclusively on the **DR2new+** dataset.

3.2.2 LOFAR

LOFAR data acquisition, processing, and timing are described in detail in Chapter 2.1.1.

Since 2013 a large pulsar monitoring campaign has been conducted using the HBA band of the LOFAR core and the six German international stations (DE601, DE602, DE603, DE604, DE605 and DE609) – some details are given by Kondratiev et al. (2016), Bilous et al. (2016), Donner et al. (2020), and Donner (2022). More than 100 pulsars were observed with a monthly to bimonthly cadence with the core, and with a weekly cadence with the German stations, and similar observing programs started in many other countries hosting the international stations. Among the observed sources, 12 millisecond pulsars (MSPs) are in common with the twenty-five presented in EPTA DR2.

LOFAR ToAs were generated following the same procedure outlined in Chapter 2.2. In the end, LOFAR ToAs are frequency resolved with a number of frequency channels that depends on the pulsar S/N.

3.2.3 NenuFAR

An overview of the NenuFAR data and of the monitoring program, which has been running since 2019 under the Pulsar Key Project, is given in Chapter 2.1.2.




Figure 3.1: Mollweide sky projection of the 25 DR2new+ pulsars (blue stars) in equatorial coordinates. Red circles and green squares display sources observed by respectively LOFAR and NenuFAR. The names of the pulsars studied in this work are written next to their position. The orange line (and area) are representing the Sun’s position (± 20 degrees angle) along the year, denoting the influence zone of the Solar wind.

The program follows a biweekly to monthly cadence and targets approximately 40 pulsars, including two MSPs also present in the EPTA DR2 list: PSRs J0030+0451 and J1022+1001.

ToAs were generated using the procedure described in Chapter 2.2. NenuFAR ToAs are used with an 8-channel frequency resolution.

3.3 Data combination

In the context of PTA studies, combining data from different observatories is essential to increase the cadence, the frequency coverage and the time span, thereby improving sensitivity to the stochastic GWB (Siemens et al., 2013). This section outlines the procedures used to combine the low-radio frequency datasets of LOFAR and NenuFAR with DR2new+ and describes the basic timing analysis performed to prepare the data for the noise analysis. In Figure 3.1 we display the sky positions of the 25 MSPs from the DR2new+, emphasizing the names of the 12 sources studied in this work. For sake of clarity, the data set composed of the combination of DR2new+, LOFAR and NenuFAR is referred as DR2low in the rest of the document. To address systematic time delays between observing systems, we fit for time offsets (JUMPS) for each observing station. Specifically, we include one JUMP for each international LOFAR station, one for the LOFAR core, and one for NenuFAR when present. Another important aspect to address when combining frequency-resolved data (one ToA per frequency channel, timed against the relevant frequency channels of a frequency-resolved template) is the difference between the DMs that were used to de-disperse the template profiles. More precisely put, when timing frequency-resolved data against a frequency-resolved template, the ToAs of each frequency channel are measured relative to a phase reference that is defined by the relevant frequency

channel of the template. Through dedispersion of the template profile, the relative phases of the different channels are changed with respect to one another, introducing a frequency dependence of the ToA’s phase reference. This frequency-dependent phase reference is different for each template used and is defined by the combination of the DM used to dedisperse the template and the actual interstellar DM affecting the data that were used to construct the template. Consequently, this effect is fully covariant (and indistinguishable from) a different DM affecting the data sets derived from a given template profile. In order to account for this effect, we therefore apply a constant DM offset to all but one of the frequency-resolved datasets. Specifically, in our analysis, frequency-resolved ToAs were available for the LOFAR, NenuFAR, and InPTA data, the latter of which was used as a reference for the other two to compute the DM offset.

Once the datasets are prepared, we start from the EPTA/InPTA timing model and include the additional JUMPs. After fitting the combined dataset, we observe that only a few key parameters exhibit significant changes—namely, the DM and its derivative (if present), the rotational period and its derivative, and the electron density at 1 AU to model the solar wind contribution. All other timing model parameters remain consistent with their initial values, likely due to the larger ToA uncertainties in the low-frequency data.

In Figures 3.2 and 3.3, we show the frequency coverage and timing residuals for each of the 12 combined pulsars. The ToAs, timing models, and low-frequency templates for all the pulsars are available on Zenodo¹.

3.4 Noise analysis: Parameter estimation and model selection

Our noise modeling approach follows the framework outlined in Section 2.3.1, which accounts for WN, time-correlated stochastic signals (DM, RN), SW, and exponential dips.

In this work, we perform several analyses per pulsar, including parameter estimation and model selection methods, in order to (1) get information on the properties of the chosen noise components described below, (2) analyze the correlations between them, and (3) obtain the most favored combination among the noise components, which could be compared against the results published in EPTA Collaboration et al. (2023b). This enables us to understand the implications of including LOFAR and NenuFAR data with DR2new+.

For the sake of clarity, since we consider multiple noise models per pulsar, we describe the noise models from their Gaussian process components (apart from the timing model uncertainties), and the number of PSD frequencies used. As an example, the noise model commonly adopted in PTA analyses includes marginalized timing model parameter uncertainties, white noise, an achromatic RN component, and DM variations. The RN and DM components were both modelled as power-law processes, with 30 and 100 frequency bins respectively. This noise model will be referred to as RN30+DMv100 throughout this document.

For the model selection, we applied a similar approach as the one described in EPTA Collaboration et al. (2023b), where we defined a set of candidate

¹<https://doi.org/10.5281/zenodo.16794085>




Figure 3.2: Time in year and MJD (resp. for the top and bottom x-axes) vs. observing frequency (in MHz) shown in logarithmic scale for the observations of the twelve pulsars considered in this work displayed in green, orange and blue for DR2new+, LOFAR and NenuFAR data, respectively. The grey vertical areas emphasize the periods when the pulsar-sun angles are equal or lower than 20 degrees.




Figure 3.3: Same as in Figure 3.2 but displaying the post-fit timing residuals (in μs) for the y-axis.

noise models, fit for the number of frequency bins (i.e., ν in Equation 2.4) used to model the power-law PSD GPs, and computed a Bayes factor among the different chosen models with their most favored number of frequency bins. The selected models are WN (white noise only), WN+RN, WN+DM ν , WN+RN+DM ν , WN+DM ν +CN₄, WN+RN+DM ν +CN₄, where CN₄ denotes a chromatic noise component with a fixed ν^{-4} frequency dependence, commonly used to model scattering-induced delays. These models are identical to the models in EPTA Collaboration et al. (2023b).

In all runs, we included the SW model described in Table A.1, which consists of a deterministic component, `n_earth`, and a stochastic term, SWGP for a specific set of pulsars. Previous analyses based on LOFAR-only data indicated strong evidence in favor of this choice, with significantly improved model fit when the SWGP component is included. While this model may not fully capture all aspects of the solar wind signal – particularly around annual harmonics – it generally leads to a more accurate description of the data, with notably reduced white noise levels in several systems.

3.5 Noise analysis - Results with standard models

We performed a parameter estimation analysis following the approach described in Section 3.4 on two data sets, DR2new+ and DR2low, by using `standard noise models`, composed of the timing model TM, WN, RN30, DM ν 100 and the SW following the description of Section 2.3.1.3. In Figures 3.4 and 3.5, we show the credible intervals (at confidence levels equivalent to 1, 2 and 3- σ in a Gaussian distribution) of the noise parameter posterior distributions for DR2new+ (orange) and DR2low (blue). For a better clarity, the color is lighter for the 3- σ contour levels. The left-hand column presents the DM variations posterior distributions, with $\log_{10} A_{DM}$ on the y-axis and γ_{DM} on the x-axis, while the right-hand column displays the corresponding parameters for the RN. For reference, we display the two principal directions² of the posteriors distributions for the same parameters obtained in Agazie et al. (2024) on the NANOGrav 15-year dataset (Agazie et al., 2023b) and PPTA DR3 (Zic et al., 2023) with solid and dashed lines respectively. The green vertical lines indicate the spectral index values expected from Kolmogorov turbulence (for the DM variations; left column), and the one predicted from a population of GW-driven supermassive black hole binaries with circular orbits (for the RN; right column).

3.5.1 Observations for the DM variations

3.5.1.1 Comparing precision between DR2new+ and DR2low

As expected, adding very-low frequency data from LOFAR and NenuFAR allowed to provide stronger constraints of the DM ν parameters (i.e., DR2low having narrower posterior distributions than DR2new+) for most of the twelve pulsars, apart for PSRs J1022+1001, J1713+0747 and J1918–0642 for which they remain similar. In particular, we observed a very significant improvement for PSRs J0030+0451, J1024–0719, J1730–2304 and J1738+0333, which remained

²The principal directions are evaluated by projecting the distributions onto the eigenvectors, with asymmetric scaling factors corresponding to 3- σ in a Gaussian distribution.




Figure 3.4: Spectral index vs. amplitude posterior distributions of the power-law PSD for DMv (left) and RN (right) obtained for twelve pulsars from DR2low (blue) and DR2new+ (orange) using the standard noise models. The green vertical lines show the predicted spectral index from a Kolmogorov turbulence in the IISM for DMv (left), and from circular and GW-driven SMBHBs (right). The black crosses display the 99.7% credible intervals of posterior distributions obtained with NANOGrav-15yr and PPTA DR3, showing the principal directions. This panel shows the first six pulsars.

largely unconstrained in the DR2new+ analysis.




Figure 3.5: Same as Fig. 3.4, but for the remaining six pulsars.

3.5.1.2 Consistency between DR2new+ and DR2low

The new constraints remain fully consistent with DR2new+ for seven of the twelve pulsars: PSRs J0030+0451, J0751+1807, J1024-0719, J1730-2304, J1738+0333, J1857+0943 and J1918-0642, while a slight shift toward lower amplitudes is observed for PSR J1713+0747, with a tension of only $\sim 0.77\sigma$ according to the description of a tension metric described in EPTA Collaboration et al., 2023b based on Raveri and Doux, 2021. For PSR J1918-0642, we noticed a slight increase of the wide posterior distribution at the high spectral index/low amplitudes in the parameter space. However, the constraints are significantly different for PSRs J1012+5307, J1022+1001, J1640+2224, and J1744-1134, where the posterior distributions are shifted toward higher spec-

tral indices (from about 1 to 2) and lower amplitudes (logarithmic values ~ 0.3 to 0.5). These discrepancies may arise from several factors. First, the **DR2new+** dataset mainly includes observations at high frequencies, which may limit its ability to disentangle chromatic (DMv) and achromatic (RN) noise components. For PSR J1640+2224, **DR2new+** contains only data at L and S bands, making it inherently unreliable for DM characterization. Furthermore, PSRs J1022+1001, J1744–1134, and even J1012+5307 – despite its relatively high ELAT of 38.76° – are affected by the SW signal, which, as discussed in Section 3.7, biases the DM recovery. Another, although less-likely, explanation could be found in DM chromaticity (Cordes, Shannon, and Stinebring, 2016).

3.5.1.3 Comparison with predicted spectral index

The spectral indices inferred from both data sets are consistent with the expected Kolmogorov value of $8/3$ (indicated by green vertical lines) for six pulsars: PSRs J0751+1807, J1024–0719, J1730–2304, J1738+0333, J1857+0943, and J1918–0642. Interestingly, while this was not the case with **DR2new+**, the spectral index becomes consistent with $8/3$ for PSRs J1012+5307 and J1640+2224, and shows a mild agreement for J1713+0747. Although the constraints for PSRs J0030+0451, J1022+1001, and J1744–1134 remain below the expected value, the distributions for the latter two exhibit an increasing trend. While the increase in spectral index from **DR2new+** to **DR2low** could be attributed to frequency-dependent ray path averaging, we highlight the significant impact of solar wind effects on these three pulsars. Such effects may bias the estimation of DM variation parameters due to limitations in current solar wind models (cf. Section 3.7), potentially resulting in spectral index constraints below the expected $8/3$ value. In the context of pulsar timing data, a spectral index lower than the true value typically indicates excess power at higher frequencies in the PSD, which may arise from unmodeled contributions to dispersion variations, such as those induced by the solar wind.

3.5.1.4 Comparison with NANOGrav 15-yr

We notice a clear consistency of the posterior distributions from **DR2low** with the constraints obtained with NANOGrav 15-yr dataset (black solid lines) for seven of the eleven considered pulsars: PSRs J0030+0451, J1024–0719, J1640+2224, J1713+0747, J1738+0333, J1857+0943 and J1918–0642, providing an improvement in the robustness of our results. The constraints for PSR J1744–1144 are consistent for the spectral index, but **DR2low** has a lower amplitude. We also observe a strong difference for PSR J1012+5307, where NANOGrav 15-yr appears to be more consistent with **DR2new+**, with higher amplitude and lower spectral index. Such property could be explained by the fact that NANOGrav 15-yr shares the same observing radio frequencies with **DR2new+** for this pulsar. Although, the NANOGrav data set does not fully overlap in time with **DR2low** since it contains ToAs from 2004 for this pulsar (starting almost 7 years before **DR2low**). Thus, any effects implying non-stationary DM variations would yield to different spectral properties as well. We attribute the large differences for PSRs J1022+1001 and J1730–2304 to the short time span of NANOGrav for

these two pulsars, with about 5 and 4 years respectively.

3.5.1.5 Comparison with PPTA DR3

The posterior distributions from PPTA DR3 (dashed lines) are fully consistent with DR2low for four of the seven overlapping pulsars: PSRs J1024–0719, J1713+0747, J1730–2304 and J1857+0943, while mildly consistent with PSR J1744–1134, such as for NANOGrav. However, the posterior distribution of PPTA DR3 for PSR J1022+1001 appears to be more consistent with DR2new+, with a lower spectral index than DR2low. The very poor constraints for PSR J0030+0451 might be due to the short time span for PPTA data, with about 3 years.

3.5.2 Observations for the RN

3.5.2.1 Comparing precision between DR2new+ and DR2low

We note a large impact on the RN constraints after adding data from LOFAR and NenuFAR. The posterior distributions become significantly narrower for PSRs J1022+1001, J1713+0747, J1730–2304, J1744–1134 and J1918–0642, where even the $3\text{-}\sigma$ contour levels appear now constrained, while it was only the case for PSR J1713+0747 with DR2new+. We also observe an improvement of precision for PSRs J0030+0451, J0751+1807, J1640+2224 and J1857+0943, where a peak appears more prominent in contour levels $\leq 2\sigma$. Furthermore, we note an improvement in the upper limits for PSRs J1024–0719 and J1738+0333, for which including LOFAR data helps reducing the covariances between RN and DMv (see the following paragraph for more details). We note a similar precision between both data sets only for PSR J1012+5307.

3.5.2.2 Consistency between DR2new+ and DR2low

The RN distributions from DR2low are fully consistent with DR2new+ for PSRs J0030+0451 and J1713+0747, where DR2low results in a more prominent peak. However, we obtain a constrained RN for PSRs J0751+1807, J1640+2224, J1730–2304, J1744–1134, J1857+0943 and J1918–0642, while it appears fully unconstrained for DR2new+. This improved constraint on the RN component is a new result, but it remains highly consistent with DR2new+ for PSRs J0751+1807, J1640+2224, J1730–2304, and J1857+0943, with tensions of only 0.35σ , 0.06σ , 0.04σ , and 0.18σ , respectively. It is in mild tension for PSR J1918–0642 (1.02σ) and PSR J1744–1134 (1.32σ), where DR2low yields a RN highly constrained at low spectral index ($\gamma_{\text{RN}} \sim 1.1$) and high amplitudes ($\log_{10} A_{\text{RN}} \sim -13.3$). Conversely, while RN was slightly constrained for PSRs J1024–0719 and J1738+0333 with DR2new+, it appears significantly reduced with DR2low. This behaviour can be observed in Figure 3.6 (top and bottom plots for PSRs J1024–0719 and J1738+0333 respectively), where we observe strong “L-shaped” correlations between RN and DMv, as described in Ferranti et al., 2025 in the case of poor radio frequency coverage. Adding lower frequencies and thus improving the precision on DM variations has removed these correlations, avoiding

a false positive for the RN in these cases. For PSR J1012+5307, the distribution is slightly shifted compared to DR2new+, mainly toward a higher spectral index ($\sim$ from 1 to 2), with a mild tension of about 1.46σ . Interestingly, we observe for the first time a bimodal posterior distribution for the RN of PSR J1022+1001, which lies inside the poorly constrained distribution for DR2new+ (tension at only 0.48σ). We observe a peak constrained with a low spectral index ($\gamma_{\text{RN}} \sim 0$), while the second one favours a steeper power-law ($\gamma_{\text{RN}} \sim 5$). This feature could reveal a disentanglement between short-term signals with low spectral index, that could be related with pulse shape variations (Liu et al., 2015; Padmanabh et al., 2020), and longer term processes that could be from different origin (e.g., nHz gravitational waves or spin noise). It also means that a careful treatment may be required to properly model the RN of this pulsar, as was already pointed out in EPTA Collaboration et al., 2023b.

3.5.2.3 Comparison with a predicted value

As described in Section 3.4, the achromatic red noise might have multiple origins (e.g., pulsar’s spin noise, nanoHertz gravitational waves), which implies some caveats about comparing it against predicted values from single pulsar noise analyses. Nevertheless, it is in practice interesting to apply such comparisons with predicted value for the GWB (see the recent summary by Verbiest et al., 2024), in particular regarding the slight tensions from the most recent results published by PTAs. In this case, we compare the constraints on the spectral index of the RN against the predicted value of $13/3$ from the simplistic case of a GWB caused by a population of GW-driven and circular SMBHBs (green vertical line). In the case of DR2low, the spectral index value of $13/3$ is within the 1σ contour levels for PSRs J1024–0719, J1640+2224, J1738+0333, J1857+0943 and J1918–0642, and withing the 2σ contour levels for PSRs J0030+0451 J0751+1807, J1022+1001 (second peak), J1713+0747 and J1730–2304. It is however strongly inconsistent with the lower spectral indices of PSRs J1012+5307 and J1744–1134 (with 1D marginalized medians resp. equal to 1.75 and 1.10), for which the major contribution to the RN might be from another origin (e.g., pulsar-related effects).

3.5.2.4 Comparison with other PTAs

As mentioned in Section 3.5.1, the short time spans for PSRs J1730–2304 and J1022+1001 with NANOGrav 15-yr (about 4 and 5 years respectively), and for PSR J0030+0451 (~ 3 years) with PPTA DR3 are the cause of their large uncertainties. We find high consistency in the RN properties between DR2low and other PTAs for PSRs J0030+0451, J1012+5307, J1022+1001 (where only the peak at low spectral index shows up for PPTA), J1024–0719, J1640+2224, J1730–2304, J1738+0333, J1857+0943 and J1918–0642, with tensions lower than 0.6σ . We observe mild tensions for PSR J1713+0747 (1.85σ and 1.50σ resp. for NANOGrav and PPTA), where the constraints of DR2low imply larger amplitudes than the other PTAs. We emphasize that the shorter time spans for DR2low (half of the other’s data sets) could be the cause of such inconsistencies. Given the differences in the timing and noise models employed, we recommend

a thorough investigation of the combined datasets to evaluate their consistency for these pulsars in future studies. For PSR J1744–1134, we observe consistent results with PPTA ($\sim 0.43\sigma$) but find a moderate tension with NANOGrav ($\sim 2.18\sigma$) which favors a higher spectral index and a lower amplitude as obtained in DR2new+. Inconsistencies between PTAs were already found in Agazie et al., 2024 for this pulsar, hence a combination of the datasets in an IPTA context might help to better understand these features.

3.6 Noise analysis - Results from model selection

We performed an analysis aimed at finding the noise model that is most favored by the DR2low data, following a similar approach as EPTA Collaboration et al., 2023b. The posterior distributions of ECORR parameters for every pulsars/systems are shown in Figure A.1, where only five pulsars have unconstrained distributions (PSRs J0571+1807, J1024–0719, J1738+0333, J1857+0943 and J1918–0642). The SW parameters of PSRs J0030+0451, J0751+1807, J1022+1001 and J1744–1134 are presented in Figure A.2.

In this section, we first intend to provide a comparison with the results that were obtained with DR2new+, then interpret the results for DR2low.

Pulsar	DR2low			DR2new+		
J0030+0451	RN13	DMv100	CN ₄ 77	RN10	X	X
J0751+1807	X	DMv92	CN ₄ 38	X	DMv50	X
J1012+5307	RN87	DMv20	CN ₄ 100	RN100	DMv47	X
J1022+1001	RN98	DMv99	CN ₄ 99	RN30	DMv100	X
J1024–0719	X	DMv18	X	X	DMv15	X
J1640+2224	RN54	DMv40	CN ₄ 40	X	DMv47	X
J1713+0747	RN73	DMv74	CN ₄ 12	RN73	DMv72	X
J1730–2304	RN18	DMv52	X	X	DMv11	X
J1738+0333	RN29	DMv28	X	X	DMv24	X
J1744–1134	RN22	DMv100	CN ₄ 31	RN19	DMv98	X
J1857+0943	RN87	DMv99	X	X	DMv92	X
J1918–0642	RN36	DMv24	CN ₄ 42	X	DMv75	X

Table 3.1: Comparison of the favoured noise models obtained with DR2low and with DR2new+ for the twelve pulsars considered in this work.

3.6.1 Comparing the favored models of DR2new+ and DR2low

Table 3.1 presents the final noise models obtained with DR2low (left) from this work, and DR2new+ (right) from the work in EPTA Collaboration et al., 2023b. It explicitly describes the noise components (RN, DMv and CN₄) modeled as GPs with a power-law power spectral density, along with the corresponding favored number of frequency bins. For instance, for PSR J0030+0451 the favoured




Figure 3.6: Posterior distributions of spectral indices (γ) and amplitudes (A) of DMv and RN for PSRs J1024-0719 (top) and J1738+0333 (bottom) obtained with DR2low (blue) and DR2new+ (red), using the standard noise models.

DR2low model includes a RN component described using 13 frequency bins, DMv modelled with 100 bins, and the additional chromatic noise CN₄ modelled with 77 bins, while the DR2new+ analysis favours only a RN component with 10 bins. It is important to note that model selection was not applied to PSR J1022+1001 for the EPTA DR2 data sets (including DR2new+), as it was found that a single power-law model was insufficient to adequately describe the data across both low and high frequencies. Consequently, we exclude this pulsar from the comparison with DR2new+, as the RN30+DMv100 model was retained for this data set.

We observe a significant change in the final noise models between the two data sets, where the favored models are the same only for PSR J1024–0719, with just DMv. Firstly, we observe that the RN is favored for ten of the twelve pulsars with DR2low, while it was only the case for five of them with DR2new+. This behavior supports the notion that incorporating low-frequency data can enhance PTA sensitivity. As demonstrated in EPTA Collaboration et al. (2023c), the inclusion of InPTA DR1 – featuring observations in the 300-500 MHz range – alongside EPTA DR2 led to a modest increase in the evidence for a GWB signal. In the case of LOFAR and NenuFAR, we anticipate a stronger improvement since the frequencies are lower (< 200 MHz), and also because it contains a longer time span and higher cadence, allowing to better sample signals in the power spectral domain. Additionally, we observe the presence of higher order chromatic noise for eight pulsars, while none were observed with DR2new+. Lastly, while the DM variations are favored for all pulsars apart from PSR J0030+0451 with DR2new+, they are now present for all of them with DR2low. This property is not surprising thanks to the high sensitivity of low frequency data toward DM variations.

For the selected number of frequency bins, we observe comparable results for the RN component, with relatively low values (~ 10 -22) for PSRs J0030+0451 and J1744–1134, and significantly higher values (~ 87 -100) for PSRs J1012+5307 and J1713+0747 – indicating the presence of constrained power at higher Fourier frequencies. The outcomes are also similar for the DM variations with PSRs J1024–0719, J1640+2224, J1713+0747, J1738+0333, J1744–1134 and J1857+0943, while significant differences appear for PSRs J0751+1807, J1012+5307, J1730–2304 and J1918+0642.

3.6.2 Discussing the results from the noise model selection

Table 3.2 provides the estimated amplitude (set to a reference frequency of 1 yr⁻¹, and 1.4 GHz for DMv and CN₄) and slope (medians with 95% confidence intervals) for the selected components. The last column displays the Bayes factors in favour of the selected model, against a model composed of RN+DMv, with favored number of frequency bins specifically obtained for this model.

As described in the previous subsection, the selected models include all the noise components (RN, DMv and CN₄) for seven pulsars: PSRs J0030+0451, J1012+5307, J1022+1001, J1640+2224, J1713+5307, J1744–1134 and J1918–0642. For these, we observe very high Bayes factors (> 100) in favor of including CN₄ over RN+DMv, apart for PSRs J1713+0747 and J1918–0642, with only $\mathcal{B}_{\text{RN+DMv}}^{\text{RN+DMv+CN}_4} = 16.0$ and 1.7 respectively. The additional chromatic noise is also included in the favored model for PSR J0751+1807 (DMv + CN₄), where

the apparent RN in Figure 3.4 becomes highly unconstrained once CN_4 is included. We still notice low Bayes factors $\mathcal{B}_{\text{RN+DMv}}^{\text{DMv+CN}_4} = 9.4$ and $\mathcal{B}_{\text{RN+DMv+CN}_4}^{\text{DMv+CN}_4} = 1.2$, indicating only mild support for the final model. It is similar for PSR J1024–0719, where the favored model DMv is poorly favored against RN+DMv: $\mathcal{B}_{\text{RN+DMv+CN}_4}^{\text{DMv}} = 1.4$. This shows a lack of support for the presence of RN in the data, also being consistent for this pulsar with a similar analysis applied on the PPTA DR2 (Goncharov et al., 2020). For the three remaining pulsars, PSRs J1730–2304, J1738+0333 and J1857+0943, the selected models remain composed of RN+DMv.

Pulsar	Red noise		DM variations		Chromatic noise		$\mathcal{B}_{\text{RN+DMv}}^{\text{Fav}}$
	$\log_{10} A$	γ	$\log_{10} A$	γ	$\log_{10} A$	γ	
J0030+0451	$-14.79^{+1.24}_{-2.02}$	$4.61^{+2.26}_{-3.56}$	$-13.49^{+0.10}_{-0.09}$	$0.96^{+0.30}_{-0.27}$	$-17.15^{+0.48}_{-0.77}$	$3.12^{+1.78}_{-1.16}$	$\geq 10^3$
J0751+1807	X	X	$-13.37^{+0.24}_{-0.64}$	$2.69^{+1.72}_{-0.77}$	$-15.15^{+0.17}_{-0.47}$	$0.74^{+1.57}_{-0.68}$	9.4
J1012+5307	$-12.92^{+0.14}_{-0.13}$	$1.41^{+0.50}_{-0.48}$	$-13.73^{+0.14}_{-0.12}$	$1.72^{+0.72}_{-0.65}$	$-16.30^{+0.16}_{-0.22}$	$1.21^{+0.53}_{-0.72}$	$\geq 10^3$
J1022+1001	$-13.04^{+0.13}_{-0.12}$	$0.15^{+0.39}_{-0.15}$	$-13.27^{+0.10}_{-0.09}$	$1.59^{+0.34}_{-0.32}$	$-16.52^{+0.20}_{-0.54}$	$1.30^{+1.47}_{-0.67}$	$\geq 10^3$
J1024–0719	X	X	$-13.74^{+0.33}_{-0.41}$	$3.38^{+1.51}_{-1.10}$	X	X	1.4
J1640+2224	$-13.26^{+0.18}_{-2.27}$	$0.28^{+1.49}_{-0.27}$	$-14.26^{+0.55}_{-1.77}$	$3.19^{+3.71}_{-1.83}$	$-15.93^{+0.19}_{-1.78}$	$1.79^{+4.13}_{-0.86}$	$\geq 10^3$
J1713+0747	$-13.87^{+0.16}_{-0.43}$	$2.07^{+1.83}_{-0.63}$	$-13.80^{+0.11}_{-0.12}$	$1.72^{+0.62}_{-0.50}$	$-15.88^{+0.39}_{-1.23}$	$2.30^{+2.53}_{-1.89}$	16.0
J1730–2304	$-13.83^{+0.63}_{-1.63}$	$3.20^{+3.46}_{-2.43}$	$-13.64^{+0.30}_{-0.78}$	$2.19^{+2.77}_{-1.21}$	X	X	X
J1738+0333	$-15.10^{+2.21}_{-4.60}$	$3.72^{+3.12}_{-3.52}$	$-13.05^{+0.28}_{-0.51}$	$2.45^{+1.52}_{-0.89}$	X	X	X
J1744–1134	$-13.34^{+0.22}_{-0.30}$	$1.05^{+1.20}_{-0.86}$	$-13.68^{+0.10}_{-0.10}$	$1.69^{+0.38}_{-0.32}$	$-15.90^{+0.17}_{-0.25}$	$0.53^{+0.92}_{-0.49}$	124.0
J1857+0943	$-15.62^{+1.93}_{-4.15}$	$3.72^{+3.09}_{-3.48}$	$-13.31^{+0.14}_{-0.26}$	$1.88^{+1.11}_{-0.53}$	X	X	X
J1918–0642	$-15.26^{+1.90}_{-4.49}$	$3.45^{+3.30}_{-3.21}$	$-14.48^{+1.21}_{-1.15}$	$5.03^{+1.86}_{-3.15}$	$-16.33^{+1.28}_{-3.29}$	$3.34^{+3.43}_{-2.96}$	1.7

Table 3.2: Estimated values (medians with 95% confidence intervals) of the amplitude and spectral index of RN, DMv and CN_4 obtained with DR2low by using the favoured noise models displayed in Table 3.1.

3.6.2.1 Comparison with standard noise models

Here, we compare our results with those obtained using the standard noise models – namely TM, WN, RN30, and DM100, as defined in Section 3.5. The estimation of RN (if present) and DMv parameters are highly consistent with the ones obtained with the standard noise models (cf. Figures 3.4 and 3.5) for PSRs J0030+0451, J0751+1807, J1024–0719, J1713+0747, J1730–2304, J1738+0333, J1744–1134 and J1857+0943. Mild differences are observed for PSR J1012+5307, with DMv being slightly shifted toward a lower spectral index, and PSR J1918–0642, where the RN and DMv are respectively less and more constrained. For the case of PSR J1022+1001, while DMv remains fully consistent, the RN becomes highly constrained to the peak at the low spectral index (cf. Figure 3.4), due to the inclusion of many frequency bins (RN98). However, we emphasize that such a noise model is expected to be insufficient to properly include the power of the RN at low PSD frequencies. A dedicated study for this pulsar is still required to ensure a proper characterization of the noise. Lastly, we observe significant changes on both RN and DMv parameter posterior distributions for PSR J1640+2224, where the RN gets highly constrained at a low spectral index and high amplitude, while the DMv displays a bimodal shape, with a main peak consistent with the standard noise model and a secondary peak favoring high spectral indices ($\gamma > 7$) and low amplitudes

($\log_{10} A < -15$).

3.6.2.2 Examining the value of the chromatic index

To validate the choice of the chromatic index of CN_4 , we performed an additional parameter estimation using the same model as in Table 3.1, but instead of fixing its chromatic index, we treat it as a free parameter χ . Figure 3.7 displays the posterior distributions of the chromatic indices obtained for all of the pulsars that include CN_4 in the noise model and the error bars display the 99.7% credible interval of the distribution. We observe a strong consistency




Figure 3.7: Posterior distribution of the chromatic index (blue violins) and their 99.7% credible intervals (errorbars) obtained for the pulsars that favor the presence of CN_4 in their noise model. Here we use the favored noise model and replace the CN_4 noise component by a Gaussian process power-law PSD signal with the same number of frequency bins, but a free chromatic index, with a prior defined as $\mathcal{U}(0, 10)$. The black horizontal line displays the fixed value for the case of CN_4 .

with the predicted value of 4 from the thin screen approximation (shown as a black horizontal line in the figure) for PSR J1012+5307, with $\chi = 4.00^{+0.8}_{-0.6}$ (median with 68% credible intervals). The results are also consistent with this predicted value for PSRs J0751+1807, J1713+0747 and J1744–1134, that all display constraints favoring higher values, and also for PSR J1640+2224, where the posterior distribution is favoring slightly lower values. The chromatic index of PSR J1918–0642 appears constrained at lower values, with a main peak centered at $\chi = 1.6^{+0.4}_{-0.4}$. However, we also observe a secondary peak centered at 4, leading to a 99.7% credible interval consistent with the predicted value. Thus, we suggest conducting a dedicated analysis of this pulsar, particularly when incorporating low-radio-frequency data. The two remaining pulsars, PSRs J0030+0451 and J1022+1001 provide chromatic indices that are highly inconsistent with thin screen turbulence theory, being respectively $\chi = 3.0^{+0.1}_{-0.1}$ and $0.92^{+0.09}_{-0.08}$. We emphasize the presence of strong dispersion variations caused by

the solar wind for both pulsars, which would induce a reduction of the measured value in case of imperfect modeling (cf. next section) since DMv has a chromaticity of 2. Furthermore, we could also expect a lower value for PSR J1022+1001 as we pointed a bimodal distribution of RN parameters that is poorly modeled with a single power-law. Thus, an imperfect treatment of RN, for which the chromaticity is equal to zero, would lead to reduction of the measured chromatic index.

3.7 Impact of solar wind

In this section, we study in more detail the impact of the SW model on PSRs J0030+0451 and J1022+1001. Specifically, we aim to investigate the differences in the noise properties when the SW signature is absent by inspecting the free spectra of DM variations and RN, where the power spectral density amplitude in each frequency bin is treated as an independent free parameter, rather than being constrained by a power-law model. For this reason, we analyzed two truncated datasets per pulsar by removing the DR2low data taken closer than 45° and 75° to the Sun, respectively. The 45° threshold follows Tiburzi et al. (2019), where it is shown that the SW effects are well modelled beyond this angular distance from the Sun (see their Figure 12). The more conservative 75° cut is introduced because NenuFAR – whose data were absent in Tiburzi et al. (2019) – offers an even higher sensitivity to DM effects than LOFAR. This stricter cut ensures that any residual and unmodelled SW contribution is removed, even in the presence of these highly sensitive datasets.

The adopted noise model is the same as in Table 3.1, however while in the previous analysis the SW was mitigated following Table A.1, in this case it is not anymore part of model. In Figure 3.8 we show the resulting free-spectra. By inspecting the DM spectra of the DR2low we observed that there is a noticeable increase in power at the harmonics of 1 yr^{-1} compared to the two cases with solar angle cuts. This suggests that the adopted SW model does not fully capture the entirety of the signal dynamics leading to a periodic leakage into the DM spectrum. Hence, the simple power-law model compensates by flattening its spectral index and increasing the number of Fourier components as shown in Section 3.5 and 3.6 for the solar pulsars such as PSRs J0030+0451, J1022+1001 and J1744–1134. This compensatory behaviour highlights the limitations of the current SW modelling in accurately characterizing its induced noise. In future works one could explore some improvement of the current SW model by changing the number of Fourier bins and the power-law assumption of the PSD. We also noticed that the 45° - and 75° -cut datasets yield consistent results across almost all bins in the spectrum. In contrast, the RN does not appear to be affected by the SW since no differences in the free-spectra are detected among the three datasets.




Figure 3.8: Free spectra for DMv (referred to 1.4GHz) and RN for PSR J0030+0451 (top) and PSR J1022+1001 (bottom). Each violin represents the posterior distribution of the RMS amplitude in a given Fourier bin. Blue violins correspond to the full DR2low dataset; green and orange violins show the results obtained after applying solar angle cuts of 45° and 75°, respectively. Vertical dashed lines indicate the harmonics of the 1 yr⁻¹ frequency, which is representative of the SW effect. The number of Fourier bins displayed matches that of the preferred noise model listed in Table 3.1.

3.8 Conclusions

In this work, we combined high-precision timing data from the EPTA DR2new+ with low-frequency observations from LOFAR and NenuFAR. Among the 25 MSPs included in the DR2new+, we identified respectively 12 and 2 sources in common with the LOFAR and NenuFAR datasets. The time span of the common pulsars almost perfectly overlaps with the one of DR2new+. The addition of low-frequency radio data significantly extends the frequency coverage of the dataset, enhancing our sensitivity to propagation effects such as DM variations and higher order chromatic delays. This new dataset is referred to as DR2low. We performed multiple iterations of the noise analysis, following the approach outlined in EPTA Collaboration et al. (2023b), and we compared the results obtained with DR2low against those from DR2new+. As a first step, we adopted a **standard noise model** consisting of the TM, WN, RN30, and DMv100 components. In addition, we applied SW corrections on those pulsars with low ecliptic latitudes and sufficient observing cadence around solar conjunctions (see Table A.1). The corner plots for the **standard noise model** are presented in Figures 3.4 and 3.5.

Overall, the results show a significant improvement in constraining the DM variations, as expected with the increased low-frequency coverage of the dataset. Moreover, this enhanced sensitivity allows to better distinguish between DM variations and RN, or at least to reduce the correlation between the two processes. This improvement is evident in Figure 3.6, where the typical L-shaped correlation between RN and DMv has been mitigated by the inclusion of low-frequency data. This result is consistent with the simulations presented by Ferranti et al. (2025). Furthermore, we showed improved consistency with the results from NANOGrav 15y and PPTA DR3.

We then performed a more in depth analysis to find the favoured noise model for DR2low. We notice a significant change in the final noise model with respect to DR2new+ (see Table 3.1). The extended time span and high cadence of the low-frequency data (approximately 10 years) now allow us to better sample signals across the power spectral domain. This improvement enables to accurately redistribute power into the correct noise components. As a result, we find that ten pulsars have significant RN in the final model in contrast of the five of DR2new+. Moreover, DMv are present in all the twelve pulsars while for DR2new+ this component was missing for PSR J0030+0451. We also find an additional noise component with a fixed chromatic index of 4 in the final models of 8 out of 12 pulsars, whereas none were detected in DR2new+. To validate the chosen chromatic index, we fit for this parameter (χ) in an additional analysis. Five out of eight pulsars show consistency with the expected value from a thin screen with same-size inhomogeneities, $\chi = 4$. A more detailed analysis is required for PSR J1918–0642, which displays a main peak at $\chi < 2$, but also a secondary peak that is consistent with $\chi = 4$. For PSRs J0030+0451 and J1022+1001, the chromatic index is inconsistent with 4, possibly due to the strong DM variations caused by the solar wind.

Finally, we conducted a detailed study on the impact of the SW model for PSR J0030+0451 and PSR J1022+1001. We applied two thresholds to remove epochs from DR2low with solar angles smaller than 45° and 75° , respectively. We analyzed the datasets using the same final noise model shown in Table 3.1,

while excluding the SW model in the truncated datasets. Figure 3.8 shows the free spectra of DMv and RN. We observe clear excess power in the DMv around the harmonics of the 1 yr^{-1} frequency. This suggests that the current SW model does not fully capture all features of this process, causing the excess power to be absorbed into the DMv modeling.

The inclusion of low-frequency data has proven to significantly improve the noise modeling in PTAs. These data help to better disentangle noise components and offer enhanced sensitivity to DM variations and other chromatic signals. Moreover, the addition of low-frequency observations has been shown to improve the accuracy of astrometric parameters, as demonstrated in Liu et al. (2025). This initial integration of LOFAR and NenuFAR data is a crucial step, as both datasets will be included in the upcoming IPTA DR3 (Good and International Pulsar Timing Array Team, 2023), which is set to become the most extensive PTA dataset ever created, with the widest frequency coverage. Furthermore, LOFAR and NenuFAR will also be part of the next EPTA Data Release 3, which will feature a larger sample of pulsars in common between the datasets. In this context, the present analysis lays important groundwork for future studies. In particular, it will be especially interesting to assess the effect of these very low-frequency data on the detection and characterization of a GWB in upcoming data releases. The extended frequency coverage provided by LOFAR and NenuFAR has the potential to improve sensitivity to the GWB and will be an important aspect to explore in future work.

Finally, it is important to highlight the significance of future low-frequency observations, particularly with SKA-low. The expected improvements in sensitivity and time coverage provided by SKA-low will be indispensable for advancing pulsar timing array studies and for ensuring the continued success of gravitational wave detection efforts.

CHAPTER 4

The effect of discrete ionised regions along the line of sight

This chapter investigates the influence of discrete ionised regions, such as HII regions, on the propagation of pulsar signals in the LOFAR sample. It further explores how turbulence properties can be inferred from DM time series and investigates the corresponding scintillation behaviour (Iraci et al., in prep).

4.1 Discrete structures in the turbulent IISM: H II regions

As mentioned in the previous chapter, an important contribution to DM variations arise because the IISM is a turbulent medium (see Section 1.2.6), where fluctuations in the electron density are expected over a wide range of spatial scales. Each observation samples a slightly different LoS through the IISM, and changes in DM naturally occur. The IISM in our Galaxy is structured into multiple phases, ranging from the diffuse ionized medium that extends across and beyond the Galactic plane, to compact and very dense regions. These dense environments significantly influence pulsar timing through their enhanced electron densities. A prominent example is offered by H II regions, clouds of ionized hydrogen formed under the ultraviolet radiation emitted by young, massive O- and B-type stars. They can be described in terms of Strömgren spheres, i.e. ionized volumes in equilibrium between photoionization and recombination (Strömgren, 1939). These regions are fundamental tracers of recent star formation and are typically associated with spiral arms and other sites of active stellar birth. They were first identified through their optical emission (Gum, 1955; Sharpless, 1959), but modern surveys have greatly expanded their census using radio and infrared observations, revealing thousands of regions that are invisible in optical wavelengths due to large distances and/or strong extinction. Two of the most widely used catalogues for H II regions are the WISE catalogue (v2.3; Anderson et al., 2014) and HH14 (Hou and Han, 2014). The WISE catalogue is based on infrared data (IR) and spans all Galactic longitudes and latitudes $|b| \leq 8^\circ$, including cross-matches with previously known H II regions at higher and lower latitudes. On the other hand, HH14 logs H II regions from the published literature, covering multiple wavelengths, sensitivities, and sky

regions, and providing spectrophotometric distance estimates for many sources.

4.1.1 WISE catalogue

Thanks to the high sensitivity and low extinction of IR observations, WISE can detect H II regions throughout the entire Galaxy. The catalogue provides centroid positions and angular sizes derived from their IR morphology, which typically shows warm dust emission at $20\ \mu\text{m}$ surrounded by $10\ \mu\text{m}$ emission coincident with the radio continuum. The latter traces free-free emission from ionized gas, while the mid-IR arises from dust heated by the same UV radiation, making their spatial correlation a clear signature of H II regions.

The sample is considered nearly complete for regions ionized by O-type stars and largely complete for early B-type stars (Anderson et al., 2014; Armentrout et al., 2021). Possible contaminants such as planetary nebulae or background galaxies are excluded through cross-matching with existing surveys and radio recombination line detections, ensuring that the fraction of misclassified sources is below 5% (Anderson et al., 2011).

4.1.2 HH14 catalogue

The HH14 catalogue (Hou and Han, 2014) compiles data for about 2500 H II regions, combining optical and radio detections collected from a wide range of surveys in the literature. In total, it provides roughly 4500 spectral line measurements, with distance estimates available for nearly 1900 sources. Unlike the WISE catalogue, HH14 does not include most of the IR-only candidates, but it offers a diverse set of distance determinations, including kinematic, parallax, and spectrophotometric estimates. The latter are derived by associating each region with its main ionizing star and estimating the distance thanks to the stellar spectral type and photometry, introducing uncertainties related to stellar classification, extinction, and luminosity dispersion. HH14 also includes several extended or high-latitude H II regions for which WISE does not report distances. Since all of the sources in HH14 are associated with the detection of ionized gas lines, the catalogue is expected to contain an almost pure sample of genuine H II regions.

4.2 Identification of H II regions along the LoS

Our aim is to analyze pulsars from the LOFAR sample whose LoS intersects known H II regions from the WISE and HH14 catalogs, building on the work by Ocker et al. (2024) where 600 such pulsars (observed in other monitoring campaigns) were identified. The advantage of using data from the LOFAR telescope instead of other observatories is the very high sensitivity of low-frequency observations to plasma processes, that enables us to investigate interstellar medium effects along these LoS with unprecedented detail. While similar studies could in principle be carried out with other low-frequency instruments, such as NenuFAR, the analysis presented here is based exclusively on LOFAR observations.

We adopted the same approach as Ocker et al. (2024) to identify H II regions along the pulsar LoS. Given the pulsar distance, the H II region distance, and its angular radius θ_{HII} , we first require that the pulsar distance is larger than

the minimum distance of the H II region, and then determine whether the LoS intersects the region.

A valid cross-match with the WISE catalogue is defined when the pulsar’s projected position at the H II distance lies within $2\theta_{\text{H II}}$. The choice of a maximum impact parameter – defined as the projected separation between the pulsar LoS and the centre of the region – twice the angular size accounts for the fact that the IR-derived radius may underestimate the true extent of the region, as ionizing radiation can leak beyond the nominal boundary.

In contrast, since the HH14 catalogue does not provide angular sizes, we adopt a fixed maximum impact parameter of 20 pc. This is a reasonable assumption given that the distribution of H II radii from the WISE catalogue, obtained by cross-matching all H II regions with the full pulsar catalogue, shows that most regions have radii smaller than 10 pc (half the maximum impact parameter; see Figure 4.1).




Figure 4.1: Distribution of H II region radii from the WISE catalogue for H II that intersect the LoS of pulsars in the full pulsar catalogue.

4.3 Analysis

We found 6 pulsars from the LOFAR sample with H II along the LoS, only from the WISE catalogue. Table 4.1 reports the 6 pulsars, the corresponding H II and the impact parameter for each intersection. We found that three pulsars, namely PSRs J0358+5413, J0407+1607, and J2225+6535, have multiple intersections with different H II regions.

Pulsar				H II				
Name	l [deg]	b [deg]	d_{so} [kpc]	H II l [deg]	H II b [deg]	H II d_{HII} [kpc]	r_{impact} [pc]	$v_{\text{psr},\perp}$ [km/s]
J0358+5413	148.19	0.81	$1.09^{+0.23}_{-0.16}$	148.540	-0.240	0.850 ± 0.150	16.4	62^{+12}_{-9}
				150.050	-1.110	0.320	14.9	
J0407+1607	176.62	-25.66	1.337	176.240	-20.880	0.150 ± 0.050	12.536	107 ± 27
				178.910	-20.120	0.150 ± 0.050	15.510	
J0609+2130	189.2	1.0	1.2	189.032	0.784	0.210	1.106	87 ± 125
J0614+2229	188.8	2.4	$3.55^{+0.44}_{-0.26}$	189.032	0.784	0.210	5.965	$21.0^{+4.3}_{-2.3}$
J1955+2908	65.8	0.4	6.3 ± 1.3	65.914	0.597	3.9 ± 0.5	11.627	128 ± 27
J2225+6535	108.6	6.8	$0.832^{+0.169}_{-0.102}$	107.920	5.780	0.764 ± 0.027	17.098	766^{+154}_{-93}
				108.860	6.150	0.764 ± 0.027	9.733	

Table 4.1: Galactic parameters, distances, and transverse velocities of the pulsars. Distances: J0358+5413 from Chatterjee et al. (2004); J0609+2130 and J0407+1607 from the electron-density model Cordes and Lazio (2002); J0614+2229 and J2225+6535 from Deller et al. (2019); J1955+2908 from Yao, Manchester, and Wang (2017). Transverse velocities of J0407+1607 and J0609+2130 are derived from LOFAR timing. For H II regions, the table reports the galactic coordinates, the distance, and the impact parameter.

To process the LOFAR data we followed the same procedure described in Section 2.2. The template’s bandwidth was divided into 10 frequency channels, hence we obtained 10 ToAs per observation, and we computed the DM time series for each pulsar using the ENTERPRISE software suite (Ellis et al., 2020). Within a Bayesian framework, along the lines of what was done in Section 3.7, we derived the free power spectra of the DM time series using 80 Fourier components, where amplitudes of the frequency bins are fitted independently. The only exception was PSR J2225+6535, for which this approach resulted in an unreliable reconstruction of the DM time series. In the free-spectrum approach implemented in ENTERPRISE, the DM variations are described by a set of Fourier components whose lowest sampled temporal frequency is fixed at $1/T_{\text{span}}$. For this pulsar, part of the DM variability is likely occurring on timescales longer than the available data span, and therefore cannot be captured by the adopted Fourier basis. As a consequence, the reconstructed DM time series does not provide a reliable representation of the underlying long-term DM behaviour, despite the inclusion of low-order DM derivatives (DM1 and DM2) in the timing model. For this pulsar, we used TEMPO2 to calculate an epoch-by-epoch DM value, and we smoothed the resulting time series with a running average. The width of the smoothing window was matched to the highest Fourier bins included in the ENTERPRISE analysis.

To study the impact of the H II regions along each pulsar’s LoS, we followed two different approaches: the first focuses on the power spectrum of the density fluctuations $P_{\delta n_e}$ to search for deviations from Kolmogorov turbulence, the second is focused on scintillation.

4.3.1 Power spectrum of the density fluctuations

Following the first approach, we aim to build the $P_{\delta n_e} = C_n^2 q^{-\beta}$ spectrum (see Section 1.2.6) from pulsar observations. Our main goal is to estimate the amplitude C_n^2 , the wavenumber q , and the spectral index β . Observed DM variations over time, $DM(t)$, can trace stochastic electron density fluctuations along the pulsar LoS and hence provide constraints on the wavenumber spectrum. Following Ocker et al. (2021) and Lam et al. (2016), the DM structure function

$$D_{DM}(\tau) = \langle [DM(t + \tau) - DM(t)]^2 \rangle \quad (4.1)$$

can be related to the amplitude of the density wavenumber spectrum through

$$\hat{C}_n^2 = \frac{1}{K_{DM}(\beta)} \left[\frac{D_{DM}(\tau)}{d_{so}(v_{\text{eff},\perp} \tau)^{\beta-2}} \right], \quad q \approx \frac{2\pi}{v_{\text{eff},\perp} \tau}. \quad (4.2)$$

In this expression, $D_{DM}(\tau)$ is the DM structure function at lag τ , measuring the typical squared DM difference over that timescale; $v_{\text{eff},\perp}$ is the effective transverse velocity of the pulsar-Earth-ISM system; d_{so} is the distance to the pulsar along the line of sight; q represents the characteristic wavenumber corresponding to the timescale τ ; and $K_{DM}(\beta)$ is a numerical factor that depends on the spectral index β and it is given by (Lam et al., 2016):

$$K_{DM}(\beta) = \frac{8\pi^2}{(\beta - 2) 2^{\beta-2}} \frac{\Gamma(2 - \beta/2)}{\Gamma(\beta/2)}, \quad (4.3)$$

where Γ denotes the Gamma-function. In the Kolmogorov turbulence assumption $K_{\text{DM}}(11/3) \approx 88.3$.

We now discuss each of the quantities entering Equation (4.2), along with the approximations adopted in our analysis. First, we assume the thin screen approximation, placing the scattering screen at the distance of the H II region, d_{HII} . Given that the typical physical extent of an H II region along the line of sight is of the order of a few to tens of au (i.e., $\lesssim 10^{-7}$ kpc), which is negligible compared to the uncertainty on its distance, the exact placement of the screen within the region (centre or edges) has an insignificant effect on the inferred scattering geometry. Given the pulsar distance d_{so} , we define the fractional distance of the screen as

$$s = \frac{d_{\text{so}} - d_{\text{HII}}}{d_{\text{so}}}. \quad (4.4)$$

The effective transverse velocity of the pulsar image on the screen, $v_{\text{eff},\perp}$, neglecting the motion of the medium, is then written as (Cordes and Rickett, 1998)

$$v_{\text{eff},\perp} = (1 - s)v_{\text{psr},\perp} + s v_{\oplus,\perp}, \quad (4.5)$$

where $v_{\text{psr},\perp}$ and $v_{\oplus,\perp}$ are the transverse velocities of the pulsar and the Earth relative to the LOS. A detailed derivation of the coordinate transformation used to properly compute these velocities is given in Appendix B.2. Finally, the DM structure function $D_{\text{DM}}(\tau)$ directly links the observed time series with the underlying density wavenumber spectrum, completing the set of quantities needed to compute $P_{\delta n_e}$.

4.3.1.1 Structure function

To characterise variability in a time series, the structure function D is one a powerful statistical tool, especially in studies of the IISM. It is defined as the mean squared difference between data points separated by a given time lag, and it quantifies how the signal changes over different timescales. This makes it particularly effective for identifying red noise: long-term, time-correlated variations manifest as a characteristic rise in the structure function at large lags, because points separated by longer timescales are more likely to differ systematically when the signal is correlated over time. Moreover, the structure function is closely connected to the PSD of the analysed quantity (Keith et al., 2013).

The structure function $D_x(\tau)$ of a time series $x(t)$ quantifies the typical squared difference between values separated by a time lag τ :

$$D_x(\tau) = \langle [x(t + \tau) - x(t)]^2 \rangle, \quad (4.6)$$

where $\langle \cdot \rangle$ denotes an ensemble or time average. This definition assumes that the increments of the time series, $x(t + \tau) - x(t)$, are *stationary*, meaning that their statistical properties depend only on the lag τ and not on the absolute time t . This condition is typically satisfied for stochastic processes such as red noise. If the time series is *wide-sense stationary*, i.e., it has a constant mean and an autocovariance that depends only on the lag τ , the structure function can be related to the autocovariance function

$$C(\tau) = \langle (x(t) - \mu)(x(t + \tau) - \mu) \rangle, \quad (4.7)$$

where $\mu = \langle x(t) \rangle$ is the mean of the process, as

$$\begin{aligned} D_x(\tau) &= \langle x(t)^2 \rangle + \langle x(t+\tau)^2 \rangle - 2\langle x(t)x(t+\tau) \rangle \\ &= 2[C(0) - C(\tau)]. \end{aligned} \quad (4.8)$$

Here, $C(\tau)$ is the autocovariance function (sometimes normalized to form the autocorrelation function). According to the Wiener-Khinchin theorem (Cohen, 1998), the autocovariance function of a real-valued process $x(t)$ is related to its one-sided PSD $P(f)$ by:

$$C(\tau) = \int_0^\infty \cos(2\pi f\tau) P(f) df. \quad (4.9)$$

Thus, the structure function can also be written in terms of the PSD as

$$D_x(\tau) = 2 \int_0^\infty [1 - \cos(2\pi f\tau)] P(f) df. \quad (4.10)$$

For a stochastic process such as $DM(t)$, a power-law wavenumber spectrum of electron-density fluctuations with spectral index β ,

$$P_\delta(q) \propto q^{-\beta}, \quad (4.11)$$

corresponds to a red-noise time series with a PSD

$$P(f) \propto f^{-\gamma}, \quad \gamma = \beta - 1, \quad (4.12)$$

where A is a normalization constant. By substituting $a = f\tau$, the structure function becomes

$$D_x(\tau) = A \tau^{\gamma-1} \int_0^\infty 2(1 - \cos(2\pi a)) a^{-\gamma} da, \quad (4.13)$$

which converges for $1 < \gamma < 3$. Therefore, for Kolmogorov turbulence ($\beta = 11/3$), $\gamma = \beta - 1 = 8/3$, giving

$$D_x(\tau) \propto \tau^{5/3}. \quad (4.14)$$

In practice, the structure function is often modeled as a simple power law,

$$D_x(\tau) = A \tau^\alpha, \quad (4.15)$$

where α is the measured slope. In logarithmic space, this becomes linear,

$$\log D_x(\tau) = \log A + \alpha \log \tau, \quad (4.16)$$

allowing α to be estimated via a linear fit. The slope α is related to the spatial wavenumber spectral index β by

$$\alpha = \beta - 2. \quad (4.17)$$

Hence, measuring α provides a direct estimate of the underlying turbulence spectrum.

4.3.1.2 Structure function and uncertainty estimation

For each pulsar, we computed the structure function of the DM time series following the standard definition of Equation (4.6). We first calculated all possible pairs of DM measurements and their time separations, and then averaged the squared DM differences within logarithmically spaced time-lag bins. The first bin was set to 7 days, consistent with the shortest observing cadence of the ILT monitoring, and each subsequent bin was increased by a logarithmic factor of 1.2. The central lag τ of each bin was taken as the geometric mean of its boundaries. To avoid poor statistics at large lags, we restricted the calculation to $\tau \leq T_{span}/3$, where T_{span} is the total time span of the data.

To interpret the structure function reliably, it is essential to quantify its uncertainties. Two main sources contribute to the error budget of $\mathcal{D}_{DM}(\tau)$:

- white noise: it arises from uncertainties in the individual DM measurements. Even in the absence of intrinsic DM variability, these errors propagate into fluctuations in the structure function, especially at small lags where the white noise dominates. The variance of the squared DM differences within each bin provides an empirical estimate of the white-noise contribution to the structure function uncertainty at that lag.
- red noise: it reflects the stochastic nature of the underlying red signal. To estimate this contribution, we use the free spectrum of the DM time series obtained with the Bayesian analysis tool ENTERPRISE (Ellis et al., 2020). We then generate numerous, different realizations of synthetic DM time series with the same power spectral properties and we compute the structure function for each of these simulated series, producing a statistical ensemble that defines a confidence envelope around the original, measured structure function. This envelope reflects the range of expected fluctuations due to the red noise alone, assuming perfect DM measurements.

These two sources of uncertainty describe different aspects of the variability in the structure function: the white noise contribution captures the impact of measurement errors and of the finite number of data pairs contributing to each lag, while the red noise envelope quantifies the expected spread in the structure function due to the stochastic nature of the underlying process.

4.3.2 Scintillation

As a pulsar signal propagates through the IISM, it encounters fluctuations in the electron density. These irregularities slightly deflect the radio waves, so that radiation reaching the observer has travelled along multiple paths. This corresponds to the scattering phenomenon described in Section 1.2.2. The interference between the light rays following these different paths, that can be constructive and destructive, produces scintillation, observed as a modulation of the signal intensity in time and frequency.

The turbulence of the IISM determines the strength, characteristic scales, and regime of scintillation. Two principal regimes are usually distinguished: *diffractive interstellar scintillation* (DISS) arises from small-scale fluctuations in

electron density (Rickett, 1977; Cordes and Rickett, 1998). It produces a characteristic pattern of narrowband and short-timescale intensity variations. The decorrelation bandwidth and scintillation timescale depend on frequency and distance, and can be directly measured from dynamic spectra; *refractive interstellar scintillation* (RISS) is caused by larger-scale inhomogeneities in the electron density (Rickett, 1977; Rickett, 1986; Cordes, Pidwerbetsky, and Lovelace, 1986; Coles et al., 1987). It leads to broader-band and slower variations in intensity, typically on timescales of days to months. RISS modulates both the mean flux density and the properties of DISS, and provides information about the large-scale structure of the scattering medium. Overall, scintillation provides a direct probe of the amplitude and scales of electron-density fluctuations along the LoS.

Pulsar observations, especially at low radio frequencies, typically occur in the strong scattering regime, where the turbulent interstellar medium induces large phase variations in the wavefront. In this regime, diffractive and refractive scintillation patterns are not independent, but are statistically connected through the Fresnel scale r_F . Following Rickett (1990), the diffractive scale l_d corresponds to the small transverse separation over which the wave phase changes by roughly one radian, producing the rapid, narrowband fluctuations observed in DISS. In contrast, the refractive scale l_r characterizes larger transverse separations where the cumulative effect of phase gradients leads to slower, broader-band modulations, as seen in RISS. By definition, in the strong scattering regime $l_d \ll r_F \ll l_r$ and these physical scales satisfy

$$r_F^2 = l_d l_r, \quad (4.18)$$

while the Fresnel scale itself is set by the geometry and the observing wavelength

$$r_F^2 = \frac{\lambda d_{\text{eff}}}{2\pi}, \quad d_{\text{eff}} = \frac{d_{sl} d_{lo}}{d_{so}}, \quad (4.19)$$

where λ is the observing wavelength, d_{sl} the pulsar-screen distance, d_{lo} is the screen-observer distance (which in our case coincides with d_{HII}), and as before d_{so} is the pulsar distance. Physically, r_F marks the boundary between the diffractive and refractive regimes, linking the rapid, small-scale fluctuations of DISS to the slower, large-scale modulations of RISS.

In the strong scattering regime, the diffractive and refractive scintillation timescales are related to the transverse scales via the effective velocity v_{eff} of the pulsar–Earth–scattering screen system:

$$\Delta t_d = \frac{l_d}{v_{\text{eff}}}, \quad \Delta t_r = \frac{l_r}{v_{\text{eff}}}. \quad (4.20)$$

The effective velocity accounts for the relative transverse motion of the pulsar, the Earth, and the scattering screen (see Appendix B.2). Using the fractional screen distance s of Equation (4.4), the effective distance becomes $d_{\text{eff}} = d_{so} s(1 - s)$, and the Fresnel scale can be expressed as

$$r_F^2 = \frac{\lambda d_{so}}{2\pi} s(1 - s). \quad (4.21)$$

From Equation (4.18) the diffractive and refractive timescales satisfy

$$\Delta t_d \Delta t_r = \frac{r_F^2}{v_{\text{eff}}^2} = \frac{\lambda d_{so}}{2\pi v_{\text{eff}}^2} s(1 - s). \quad (4.22)$$

In deriving this relation, we assume a *thin-screen geometry*, in which all scattering is concentrated in a single plane. In our case, this plane corresponds to the H II region intersecting the line of sight, which we identify as the dominant scattering screen. Under this approximation, both diffractive and refractive effects arise from the *same screen*. This allows us to directly link the observed scintillation timescales to the position and properties of the H II region along the LoS.

To investigate the amplitude of refractive scintillation, we use the time series of the S/N measured across individual observing epochs, normalized by the square root of the integration time of each observation, as expected from the radiometer equation. While LOFAR data are not flux-calibrated, the S/N serves as a practical proxy for the pulsar intensity. Potential variations due to array configuration, elevation-dependent gain, or other instrumental effects are expected to be small for our dataset, given the relatively short duration and observing setup of each epoch; a quantitative assessment of these effects has not been performed. From this S/N time series, we compute the autocorrelation function (ACF) (Romani, Narayan, and Blandford, 1986) as a function of the time lag τ , which quantifies how correlated the signal is with itself at different temporal offsets. We then fit the central part of the ACF with a Lorentzian profile,

$$ACF(\tau) = \frac{A}{\Delta t_r^2 + \tau^2}, \quad (4.23)$$

where A is the amplitude and Δt_r is the characteristic decay timescale¹. The fit is performed over the range of lags up to the first zero crossing of the ACF, thus excluding the zero-lag region where the correlation is dominated by noise (see Figure 4.5). The resulting timescale Δt_r is taken as the refractive scintillation timescale, reflecting the broad, slowly varying intensity modulations induced by large-scale electron-density inhomogeneities in the scattering screen.

Because direct measurements of the diffractive scintillation timescale Δt_d are not feasible with our data (the frequency resolution is insufficient to resolve the narrowband, rapid DISS structure), we inferred Δt_d by combining the measured Δt_r with the fractional screen distance s through Equation (4.22). Under the assumption that the turbulent plasma responsible for both diffractive and refractive scintillation is confined to the H II region along the LoS, the predicted Δt_d quantifies the characteristic scale of the small-scale fluctuations within that region.

4.3.3 Results

4.3.3.1 Turbulence spectrum

The analysis we are presenting focuses on only three of the six pulsars of Table 4.1: PSRs J0358+5413, J0407+1607, J2225+6535, because of the low S/N of the remaining three pulsars in our list. Figure 4.2, 4.3, 4.4 show, respectively, the DM time series for the three pulsars analysed, the corresponding structure functions and the wavenumber spectra calculated as reported in the previous sections.

¹see also an alternative approach using the flux structure function; Daszuta, Lewandowski, and Kijak (2013).




Figure 4.2: DM time series of PSRs J0358+5413 (top), J0407+1607 (middle), J2225+6535 (bottom). Shaded regions indicate the 99.7% confidence intervals (3σ uncertainties).




Figure 4.3: Structure functions of PSRs J0358+5413 (top), J0407+1607 (middle), J2225+6535 (bottom). Shaded regions indicate the 99.7% confidence intervals (3σ uncertainties).

PSR	β
J0358+5413	3.2 ± 0.6
J0407+1607	3.3 ± 0.6
J2225+6535	3.2 ± 0.4

Table 4.2: Best-fit spectral indices β obtained from the structure function analysis using the power-law model of Equation (4.16).




Figure 4.4: Wavenumber spectra for PSRs J0358+5413 (light-blue dots), J0407+1607 (green squares), and J2225+6535 (coral diamonds). Grey crosses represent the spectra of 18 pulsars from the NANOGrav 12-year dataset (labelled as NG12y), from the work of Ocker et al. (2021).

The results of the linear fit (in log-log space) for the spectral index β are presented in Table 4.2. As a comparison, Figure 4.4 also displays the wavenumber spectra of 18 pulsars from the NANOGrav 12-year dataset (NG12y, Alam et al., 2021), as derived by Ocker et al. (2021) from their DM time-series analysis. These sources do not have any H II region along their LoS.

4.3.3.2 Scintillation

Figure 4.5 shows the time series of the S/N normalised by the integration time for the three pulsars and the corresponding ACFs. To build these time series we used uniquely data collected by the international LOFAR stations because of their higher observing cadence.

Table 4.3 reports the measured values of the refractive timescales, and the corresponding diffractive timescales associated to the HII, as from Equation 4.22. In the cases of PSRs J0407+1607 and J2225+6535, we calculate the diffractive timescale only once, even though they intersect multiple H II regions. This is because Δt_d depends solely on d_{HII} , which has the same value at each intersection for both pulsars (see Equation (4.22)).

PSR	s	$v_{\text{eff},\perp}$ (km/s)	Δt_r (days)	Δt_d (sec)
J0358+5413	$0.234^{+0.174}_{-0.210}$	54^{+11}_{-9}	78 ± 2	83^{+67}_{-71}
	$0.712^{+0.070}_{-0.084}$	39^{+4}_{-3}	78 ± 2	105^{+57}_{-36}
J0407+1607*	0.888 ± 0.037	38 ± 4	235 ± 10	$3.83^{+1.6}_{-1.4}$
J2225+6535*	$0.089^{+0.151}_{-0.089}$	693^{+153}_{-130}	13 ± 6	$0.01^{+0.02}_{-0.01}$

Table 4.3: For each pulsar, the table lists the pulsar name, the fractional distance of the H II region along the line of sight (s), the effective transverse velocity v_{eff} , the best-fit refractive timescale Δt_r , and the derived diffractive timescale Δt_d , computed using Equation (4.22). All the scintillation timescales are quoted at the central observing frequency of LOFAR data ($\nu_c = 153.8$ MHz). For PSRs J0407+1607 and J2225+6535 (marked with *), the estimated Δt_d is reported only once, as both H II regions are located at the same distance d_{HII} , resulting in identical Δt_d values (see Table 4.1).

4.4 Discussion

4.4.1 Structure function and turbulence spectrum

To calculate the log-log linear fit of the DM structure functions we did not use all of their available bins. This is because of how the DM time series were reconstructed with ENTERPRISE, where we adopted a free-spectrum model with a finite number of Fourier bins. In our case, the free-spectra were computed using 80 frequency bins, and the highest sampled frequency corresponds to a period of roughly 45 days. Because the Gaussian-process reconstruction relies on a sum of sinusoidal components, any fluctuation at frequencies higher than that highest bin cannot be captured. As a consequence, the fit is only meaningful for lags exceeding that threshold. This also explains the absence of the typical flattening of the structure function at short lags, where the white noise starts to dominate, as high-frequency fluctuations are not included in the model. Although this limitation affects the shortest lags of the DM structure functions, the underlying spectrum of electron-density fluctuations inferred from the measured slopes remains consistent with a Kolmogorov wavenumber spectrum. This is shown in Table 4.2, where the power-law indices for all three pulsars are compatible with the expected $\beta = 11/3$, suggesting that the turbulence within the H II regions follows the same statistical properties as in the diffuse IISM.

Moreover, Figure 4.4 we also note that the three pulsars lie near the upper edge of the distribution of the NG12y sample, which does not include any pulsar whose LoS crosses an HII region. Here, we plot the power spectrum of electron-density fluctuation $P_{\delta n_e}$ which is directly related to C_n^2 (see Equation (1.23)). The trend of higher $P_{\delta n_e}$ along LoSs intersecting HII regions could be an increase in C_n^2 although the sample is still too small to draw a robust conclusion.

We also stress that a direct comparison of the calculated C_n^2 values is not straightforward. The NG12y analysis by Ocker et al. (2021) assumes a fixed Kolmogorov spectral index, while here we explicitly fit for β . Moreover, none of the pulsars in their sample have an H II region crossing the LoS; therefore, they assume the scattering screen to be halfway between the Earth and the

pulsar (i.e., $s = 0.5$). Since C_n^2 depends on β (see Equation 4.2), adopting a Kolmogorov slope leads to typical values of $C_n^2 \sim 10^{-3.5} \text{ m}^{-20/3}$ in the NG12y sample, albeit with considerable spread (Ocker et al., 2021). In our case, the inferred values of β are compatible with the Kolmogorov value within the associated uncertainties, the best-fit values of β differ slightly from the Kolmogorov one, and this results in C_n^2 values that can vary by several orders of magnitude. As an example, Figure 4.6 shows that, for PSR J0358+5413, even a change of a few 10^{-1} in β can shift C_n^2 by several orders of magnitude, and this spread includes the aforementioned average NG12 value.

Finally, PSR J2225+6535 reaches the lowest wavenumbers among our targets, thanks to its very high transverse velocity.

4.4.2 Scintillation

From Table 4.3, we note that the best-fit refractive timescales Δt_r vary significantly among the three pulsars. PSR J0358+5413 and PSR J0407+1607 exhibit relatively long timescales (~ 78 days and ~ 235 days, respectively), while PSR J2225+6535 shows a much shorter timescale (~ 13 days), which is reasonable given its very high effective velocity. For the latter, the noisy S/N time series introduces additional uncertainty, and the derived parameters should be treated with caution. Overall, these differences reflect the combination of varying LoS velocities and H II region locations along the pulsar sightlines. From the refractive timescales reported in Table 4.3, we can derive the corresponding refractive spatial scales at 153.8 MHz $l_r = v_{\text{eff}} \Delta t_r$. As expected, these are found to be significantly smaller than the typical size of H II regions (around 20 pc) and also smaller than the impact parameters along the LoS (see Table 4.1). Specifically, we obtain $l_r = 2.4_{-0.4}^{+0.5}$ au for PSR J0358+5413, $l_r = 17_{-2}^{+3}$ au for PSR J0407+1607, and $l_r = 51 \pm 24$ au for PSR J2225+6535.

The diffractive timescale, Δt_d , also reported in Table 4.3, provides an additional check on our results. To validate these estimates, they should be compared with direct observational measurements. In future, we aim to better explore the relation between scintillation and DM fluctuations by considering available Δt_d measurements at LOFAR frequencies. Moreover, by definition $D_{\text{DM}}(\Delta t_d) = 1$, so Δt_d can serve as a direct observational anchor for the structure function. This would allow us to verify whether the power-law model derived from DM variations remains consistent across the full range of temporal scales.

4.5 Conclusions

In this work we studied the impact of discrete H II regions along the line of sight of pulsars observed with the LOFAR telescope. We demonstrated how pulsar observations can be used to probe the turbulence of the IISM.

Using the low observing frequencies reached by LOFAR and high-quality observations and short cadence, we calculated refined DM timeseries and their associated turbulence spectra with an extensive baseline of wavenumbers. Our analysis of the spectral index was carried out rigorously, without assuming an a-priori Kolmogorov spectrum, and we did not find significant deviations from the expected Kolmogorov behaviour.

We also carried out a scintillation analysis. While direct measurements of the diffractive timescale were not feasible with the available data, we were able to study the refractive scintillation using a proxy for the flux (S/N normalised by the square root of the integration time). The resulting refractive timescales correspond to spatial scales smaller than the typical H II region size, consistent with the possibility that refractive scintillation originates within the H II structures. Future direct measurements of the diffractive timescale, achievable with high-resolution low-frequency observations using LOFAR or NenuFAR (see e.g., Bondonneau et al., 2021; Wu et al., 2022), will be important to link these two approaches – the turbulence spectrum analysis and the scintillation study – providing a quantitative comparison and allowing us to assess whether the observed scintillation is indeed dominated by turbulence within the H II region.




Figure 4.5: Normalized signal-to-noise ratio (S/N divided by $\sqrt{T_{\text{int}}}$) and auto-correlation function (ACF) for the three pulsars, displayed from top to bottom: J0358+5413 (blue), J0407+1607 (green), and J2225+6535 (orange). Within each pulsar panel, the upper part shows the normalized S/N measured in individual observing epochs, while the lower part displays the corresponding ACF. Different markers indicate different observing stations: circle markers for DE601, diamond for DE602, square for DE603, cross for DE604, triangle for DE605, and plus for DE609.




Figure 4.6: C_n^2 as a function of β for PSR J0358+5413. The best-fit value is $\beta_{\text{fit}} = 3.2$ (red circle), while the Kolmogorov expectation is shown for comparison (green diamond). The corresponding C_n^2 values differ by nearly five orders of magnitude.

Methods to estimate Dispersion Measure variations

This chapter presents a comparison of three methods for estimating DM variations from pulsar timing data. It describes the principles behind each approach, evaluates their performance using simulated datasets, and assesses their effectiveness through various comparison metrics (based on Iraci et al. (2024)).

5.1 Scope of the work

In this work we compare the functionality and reliability of three commonly used schemes to measure temporal variations in interstellar propagation effects in pulsar timing data. The aim is to understand how the choice of method influences the recovery of DM variations, both in terms of accuracy and potential biases. This comparison provides insight into the limitations and strengths of each approach, offering guidance for current analyses. Future developments will extend this investigation to assess how these methods may affect the recovery of gravitational wave signals in PTA experiments.

5.2 Simulations and software

To compare the effectiveness of the various PTA methods of calculating and accounting for DM variations, we evaluated their application to simulated data, which were produced using the `libstempo` software package (Vallisneri, 2020).

5.2.0.1 Simulations with `libstempo`

`libstempo` is a Python wrapper of the `TEMP02` software package (Hobbs, Edwards, and Manchester, 2006) that allows the use of all of the `TEMP02` functionalities within a Python-based environment. In particular, we exploited the `libstempo.toasim.fake_pulsar` library to simulate 10 narrow-band ToAs per observing epoch in the most commonly used LOFAR frequency band for pulsar observations, between 100 and 200 MHz, with 5 μ s fixed ToA error bars. The

simulated observing epochs have a regular fortnightly cadence and cover a time span of $T_{span} = 3000$ days. To build simulations as close as possible to the typical PTA dataset, we injected stochastic noise in the form of white noise (WN), achromatic red noise (RN), and DM variations.

The WN reflects instrumental errors and instrumental sensitivity, as well as intrinsic pulse jitter. We modeled it by considering two parameters: EFAC, which is a multiplicative factor that takes into account ToA measurement errors; and EQUAD, which is added in quadrature and accounts for any other WN that may be given by profile variations and possible systematic errors.

$$\sigma_{\text{ToA}} = \sqrt{(\text{EFAC} \cdot \sigma_{\text{temp}})^2 + \text{EQUAD}^2}, \quad (5.1)$$

where σ_{temp} is the template-fitting error, and was given as an input to our simulations¹.

Achromatic RN, arising from intrinsic pulsar spin irregularities, and DM variations, are both time-correlated noise processes that are modeled with a Fourier basis of N_f coefficients and a power-law power spectrum (EPTA Collaboration et al., 2023b). We distinguished between chromatic and achromatic noise processes using a chromatic index, α , equal, respectively, to 2 and 0. The time delay induced on a ToA with frequency, ν , at an epoch, t , is then:

$$t_{\text{delay}}(t) = \sum_{i=1}^{N_f} \sqrt{\mathcal{P}_i} F_i \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{-\alpha}. \quad (5.2)$$

$\mathcal{P}_i$ is the power-law power spectral density with hyperparameters $\hat{A}$, the normalized amplitude at the frequency of 1 yr^{-1} , and γ , the spectral index. It is written as follows:

$$\mathcal{P}_i = \frac{\hat{A}_n^2}{12\pi^2 T_{span}} \frac{\text{yr}^3}{\left(\frac{f_i}{1\text{yr}^{-1}} \right)^{-\gamma_n}}, \quad (5.3)$$

where T_{span} is the total time span, $n = \text{RN or DM}$, and we refer to the amplitude as $A_n = \log_{10}(\hat{A}_n)$. F is the matrix of the cosine functions at each Fourier frequency, $f_i = i/T_{span}$, and time, t :

$$F_i(t) = \sqrt{2} \cos(2\pi f_i t + \phi), \quad (5.4)$$

with ϕ being a random phase drawn from $\mathcal{U}(0, 2\pi)$. The radio frequency is ν and the reference frequency, ν_{ref} , was set to 1.4 GHz. The time series of the injected noise signal is then a sum over a finite number, N_f , of cosine functions (see the second and third rows of Figure 5.1 for an example). Strictly speaking, the time series generated here is not a GP, as we do not draw the Fourier amplitudes from a normal distribution. Instead, we construct a single realization of the noise with the correct power-law spectrum and units consistent with ENTERPRISE conventions. This approach is sufficient for our purposes, since the subsequent recovery with ENTERPRISE treats the noise as a GP. Our goal is therefore to produce a time series that reflects the desired spectral properties, rather than to simulate the full statistical ensemble of the noise. For both RN and DM noise, we considered $N_f = 30$ components. This corresponds to a frequency

¹We used the same convention as EPTA Collaboration et al. (2023b) for σ_{ToA} ; however, some PTA collaborations use the TEMPO2 definition (see Verbiest et al., 2016)

EFAC = 1.2 EQUAD = $2 \cdot 10^{-6}$ s		
$(\gamma_{RN} = 3.7; \gamma_{DM} = 8/3)$	or	$(\gamma_{RN} = 2.5; \gamma_{DM} = 3.2)$
$A_{RN} = -12.3$	$A_{RN} = -12.3$	$A_{RN} = -12.3$
$A_{DM} = -12.6$	$A_{DM} = -13.6$	$A_{DM} = -14.6$
$A_{RN} = -13.3$	$A_{RN} = -13.3$	$A_{RN} = -13.3$
$A_{DM} = -12.6$	$A_{DM} = -13.6$	$A_{DM} = -14.6$
$A_{RN} = -14.3$	$A_{RN} = -14.3$	$A_{RN} = -14.3$
$A_{DM} = -12.6$	$A_{DM} = -13.6$	$A_{DM} = -14.6$

Table 5.1: Input noise parameter values adopted for our simulations.

The WN parameters EFAC and EQUAD were fixed in all the realisations, respectively, to 1.2 and $2 \mu\text{s}$ and $\sigma_{\text{temp}} = 5 \mu\text{s}$. We used two different pairs of spectral indices: the first and more realistic with a steeper RN ($\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$); and the second and opposite case, which was produced with a flatter RN spectrum ($\gamma_{RN} = 2.5$, $\gamma_{DM} = 3.2$). For each pair, we used nine different combinations of (A_{RN}, A_{DM}) , as reported in each cell.

of approximately 1.15×10^{-7} Hz, comparable to twice the typical sampling frequency of the observations, which is one measurement every 14 days (i.e., about 2×10^{-7} Hz).

5.2.0.2 Simulation parameters

Table 5.1 shows all the parameters adopted for our simulations, with σ_{temp} , EFAC, and EQUAD being fixed, respectively, to $5 \mu\text{s}$, 1.2, and $2 \mu\text{s}$ for all the cases. Concerning the noise spectral indices, we first assumed a realistic spectral index for DM variations from the Kolmogorov turbulence theory (Armstrong, Rickett, and Spangler, 1995) ($\gamma_{DM} = 8/3$) and a steeper one for the RN ($\gamma_{RN} = 3.7$). Secondly, we considered a different pair of spectral indices: $\gamma_{DM} = 3.2$ and $\gamma_{RN} = 2.5$. For each pair of spectral indices, we have studied three fixed values of amplitude of the power spectrum for both of the noise processes, producing a total of nine datasets per pair of spectral indices. A_{RN} and A_{DM} have been chosen within the range of values reported by EPTA Collaboration et al. (2023b).

For each specific set of A_{RN} and A_{DM} and pair of spectral indices, we generated a total of 100 timing residuals realizations. In each realization, the seed for the injected RN was kept constant to maintain the same RN signature across all timing residuals, while the DM noise seed varies with each iteration, resulting in distinct DM time series for all 100 realizations.

Figure 5.1 clarifies our simulation process by reporting three different scenarios, each of them with a different A_{DM} that increases going from left to right. The first row reports the overall timing residuals; the second row shows the delay due to the injected RN; and the third row illustrates the DM variations introduced in our dataset. The last row represents the power spectrum of the noise processes involved in the three scenarios. We note that the first and third rows of each column each report only 1 of the 100 realizations that have been created for that set of parameters.




Figure 5.1: Simulations carried out with injected values of γ_{RN} , A_{RN} , and γ_{DM} of, respectively, 3.7, -13.3, and 8/3. The columns are associated with increasing values of A_{DM} , left to right (as is indicated at the bottom). First row: 1 out of the 100 simulated timing residuals with a color map associated with the observing frequency (as is reported in the top right corner). Second row: Injected RN signal. Third row: 1 out of the 100 different injections of the DM time series. On the right y axis, we report the corresponding time delay at the reference frequency of 150 MHz. Fourth row: Power spectra of the injected noise processes; WN (black), RN (red), and DM variations at LOFAR central frequency of the band of ~ 150 MHz (light blue). The dashed vertical gray lines are the 30 Fourier frequency bins.

5.3 Methods of analysis

Each generated series of timing residuals was analyzed with the following methods.

5.3.1 Epoch-wise dispersion measure

The epoch-wise (EW) DM (EW, e.g. Donner et al., 2019; Tiburzi et al., 2019) models a DM value and a constant time offset, C , via TEMPO2 over the ToAs for each simulated observing epoch. In particular, the fit functional form is

$$\Delta t = \frac{\text{DM}}{\nu^2} + C, \quad (5.5)$$

where the parameters are the dispersion measure, DM, and C , a constant offset that is supposed to absorb any unaccounted-for achromatic process. The observing frequency, ν , corresponds in our simulations to an array of 10 values in the LOFAR frequency range 110-190 MHz and Δt is the time delay contained in the measured ToAs. This EW analysis returns a DM and its corresponding 1σ uncertainty given by the fit for each observation. The advantage of EW is that it simultaneously accounts for the DM and RN effects on the set of parsed ToAs that corresponds to the ToA set for a single observing epoch. Hence, we did not need to separately model the RN.

5.3.2 DMX

DMX is a method developed by the NANOGrav collaboration (e.g., Agazie et al., 2023b, sec 4.1) that performs a piece-wise constant fit of a DM value across an a priori specified temporal window. In our simulations, each window is such that it contains the 10 ToAs of a single observing epoch so that, at the end of the procedure, we obtain as many DMX parameters as the number of observing epochs. In this way, DMX is extremely similar to EW, with two important differences: it uses the PINT software (Luo et al., 2021) suite to perform the fit, and it does not model the constant offset, C . This means that fitting for the DMX parameters does not account for any RN process in the data and, in order to do so, the procedure to calculate an unbiased DM time series is divided into three steps: a first fit for the DMX parameters; an RN modeling round (while correcting for the initial DMX measures) with the ENTERPRISE software suite (Ellis et al., 2020); then a second DMX fit. In the end, DMX will also provide a DM value, with the corresponding uncertainty, for each observation.

5.3.3 DM GP

It is the fully Bayesian-based noise-analysis approach carried out with the ENTERPRISE software suite, able to describe the WN, and the achromatic RN and DM variations, which are modeled as stationary GPs (see details in Section 5.3.3.1 and van Haasteren and Vallisneri, 2014), while the Bayesian inference is performed via the Markov chain Monte Carlo sampler PTMCMCSampler (Ellis and Van Haasteren, 2017). In this approach, the time delay induced on

a ToA with radio frequency, ν , at an epoch, t , is written as

$$\delta t(t) = \sum_{i=1}^{N_f} [a_i \sin(2\pi t f_i) + b_i \cos(2\pi t f_i)] \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{-\alpha}, \quad (5.6)$$

where ν_{ref} , α , and f_i are the same as in Eq. (5.2). The number of frequency components we used was the same as for the injection: $N_f = 30$ for both RN and DM. The weights, a_i and b_i , followed a multivariate Gaussian distribution with zero mean and a covariance matrix, Σ_{jk} , that in the frequency domain is given by

$$\Sigma_{jk} = P_L(f_j) \frac{\delta_{jk}}{T_{\text{span}}}, \quad (5.7)$$

with $j, k = 1 \dots N_f$, the Kronecker delta, δ_{jk} , and the power spectrum, P_L :

$$P_L(f) = \frac{A^2 \text{yr}^3}{12\pi^2} \left(\frac{f}{1\text{yr}^{-1}} \right)^{-\gamma}. \quad (5.8)$$

By marginalizing over the a_i and b_i parameters, this analysis yields the posterior distributions for EFAC and EQUAD, and for the hyperparameters, A and γ , of RN and DM power spectra. The posteriors describing the DM power spectrum allow one to reconstruct the DM noise via the LAFORGE software suite (Hazboun, 2020), as a sum over a set of N_f Fourier components, as is described in Section 5.3.3.1. In particular, we gave as input to LAFORGE a large number, N , of random posterior samples of the noise parameters, and we obtained as output N different time-domain realizations of the time delays corresponding to the same noise process. These delays were then converted into DM values by inverting Equation (1.8), and hence yielded a set of N DM time series (i.e., a DM probability distribution for each epoch). In order to compare the DM GP output with the results of the other two methods, we computed the mean DM and standard deviation, σ_{DM} , at each epoch, and hence obtained one reference value per observation.

5.3.3.1 Evaluating time-domain realizations of stochastic time-correlated signals with LAFORGE

The stochastic and time-correlated signals in PTA data such as DM variations are mainly modelled as GPs. GPs are defined as a collection of (infinite) random variables representing the function values (e.g., time delay of the modelled signal in the timing residuals) for all input locations (i.e., observing epochs). Any finite set of these random variables has a joint Gaussian distribution. Thus, GPs allows us to define the mean value and the variance of the signal at any input location.

GPs can be fully specified with one of each following approaches (Rasmussen and Williams, 2006):

- Weight-space view, where the process is defined as a linear sum of deterministic basis functions $T_\mu(t)$ multiplied by weights a_μ as

$$f(t) \sim \sum_{\mu}^m a_{\mu} T_{\mu}(t), \quad (5.9)$$

with $\mu = 1, \dots, m$ and where the weights are Gaussian random variables as $a_\mu \sim \mathcal{N}(a_\mu^0, \phi_{\mu\nu})$, with a_μ^0 assumed to be a zero vector in our case, and $\phi_{\mu\nu}$ corresponding to the weights covariance matrix.

The PTA likelihood including a GP described in the weight-space view can be written as

$$p(\delta t | a_\mu, GP) = \frac{\exp \left[-\frac{1}{2} \sum_{i,j} \left(\delta t_i - \sum_\mu a_\mu T_\mu(t_i) \right) N_{ij}^{-1} \left(\delta t_j - \sum_\mu a_\mu T_\mu(t_j) \right) \right]}{\sqrt{(2\pi)^n \det(N)}} \times \frac{\exp \left[-\frac{1}{2} \sum_{\mu,\nu} a_\mu \phi_{\mu\nu}^{-1} a_\nu \right]}{\sqrt{(2\pi)^m \det(\phi)}}, \quad (5.10)$$

with δt and N respectively corresponding to the timing residuals and the covariance matrix, $i, j = 1, \dots, n$ and $\mu, \nu = 1, \dots, m$.

- Function-space view, where the GP directly describes the targeted signal as

$$f(t) \sim GP(g(t), k(t, t')), \quad (5.11)$$

including the mean function $g(t)$ and the covariance function $k(t, t')$, also referred to as kernel.

The PTA likelihood in that form is marginalized over the basis weights mentioned above and can be expressed as

$$p(\delta t | GP) = \frac{\exp \left[-\frac{1}{2} \sum_{i,j} \delta t_i (N_{ij} + K_{ij})^{-1} \delta t_j \right]}{\sqrt{(2\pi)^n \det(N + K)}}, \quad (5.12)$$

where $K = k(t, t')$.

In PTA analysis, we usually employ Equation (5.12) and express $k(t, t')$ from the correspondence between both views as (Rasmussen and Williams, 2006)

$$k(t, t') = \sum_{\mu,\nu} T_\mu(t) \phi_{\mu\nu} T_\nu(t'). \quad (5.13)$$

In practice, we set the basis functions T as a finite set of sine/cosine functions and

$$\phi_{\mu\nu} = P(f_\mu) \delta_{\mu\nu} / t_{\text{span}}, \quad (5.14)$$

with the right-hand side terms respectively corresponding to the power spectral density (PSD), the Kronecker delta and the time span of the data set. This approach corresponds to the Fourier-sum approach described in van Haasteren and Vallisneri (2014).

To summarize, we use a time-domain likelihood function with a parameterized covariance matrix that contains a PSD that is most often defined as a

simple power-law described with two parameters: the amplitude A , usually set at the reference frequency of 1 yr^{-1} and the spectral slope γ .

The Gaussian process signal in the time-domain $f(t)$ can be evaluated from Equation (5.9) after estimating the weights a with Equation (5.10). Let us now describe a method to estimate realizations of a , where we perform a maximum likelihood estimation on Equation (5.10) (Lentati et al., 2013). Let us rewrite the logarithm of this equation in a matrix notation,

$$\begin{aligned} \log [p(\delta t|a, GP)] = & -\frac{1}{2}(\delta t - aT)^T N^{-1}(\delta t - aT) - \frac{1}{2}a^T \phi^{-1}a \\ & - \frac{1}{2} [(n+m) \log(2\pi) + \log(\det(N) \det(\phi))], \end{aligned} \quad (5.15)$$

where the third term is just a constant number "cst," and thus

$$\begin{aligned} \log [p(\delta t|a, GP)] = & -\frac{1}{2} [\delta t^T N^{-1} \delta t + (aT)^T N^{-1}(aT) + a^T \phi^{-1}a] \\ & + (T^T N^{-1} \delta t)^T a + \text{cst} \\ = & -\frac{1}{2} [\delta t^T N^{-1} \delta t + a^T (T^T N^{-1} T + \phi^{-1})a] \\ & + (T^T N^{-1} \delta t)^T a + \text{cst}. \end{aligned} \quad (5.16)$$

Let us now write its derivative over the weights,

$$\frac{\partial \log [p(\delta t|a, GP)]}{\partial a} = (T^T N^{-1} T + \phi^{-1})a - (T^T N^{-1} \delta t)^T, \quad (5.17)$$

and finally obtain the maximum likelihood vector of weights $\hat{a}$,

$$\hat{a} = (T^T N^{-1} T + \phi^{-1})^{-1} (T^T N^{-1} \delta t)^T, \quad (5.18)$$

where we evaluate ϕ using Equation (5.14), for which $P(f)$ is computed by using posterior distributions of our model parameters (e.g., A and γ if defined as a power-law) evaluated from a Bayesian analysis. Furthermore, N is computed after applying WN parameters (e.g., EFAC, EQUAD) with values taken from the same posterior distributions. This way, we evaluate time-domain realizations that are marginalized over other processes included during the Bayesian analysis.

5.4 Results

We present here the results of the analysis with each method. A total of three main methods have been considered (EW, DMX, and DM GP), along with two additional ones in which the RN is not modeled (DMX no RN and DM GP no RN).² Figure 5.2 shows the DM time series recovered with the three main methods on the same noise realization, accompanied by the corresponding DM residuals obtained by subtracting the recovered and injected DMs. An example of the DM residuals that can be obtained by combining all the 100 realizations generated for each simulation is shown in Figure 5.3, with the bottom panel representing the RN signal injected in all realizations. In Tables 5.2 and 5.3,




Figure 5.2: DM time series and DM residuals recovered by EW (blue, first and second panels), DMX (red, third and fourth panels), and DM GP (green, fifth and sixth panels). The black line is the injected DM time series. In the last two panels, we overplot, respectively, the DM time series and the DM residuals for the three methods. The DM and RN injected parameters for this realization are: $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$, $A_{RN} = -12.6$, and $A_{DM} = -13.3$




Figure 5.3: Top panel: DM residuals of the 100 realizations for the five considered methods – EW(blue), DMX no RN (magenta), DM GP no RN (orange), DMX (red), and DM GP (green) – for the following set of input parameters: $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$, $A_{RN} = -12.6$, and $A_{DM} = -13.3$. Bottom panel: RN signal injected (the same for all of the 100 realizations; see Section 5.2)

	$\langle \sigma_{DM} \rangle \cdot 10^{-5} \text{ (pc cm}^{-3})$						
A_{RN}	EW	DMX	DM GP	DMX no RN	DM GP no RN		A_{DM}
-12.6	5.0	1.3	1.0	0.8	0.6		-12.3
-13.6	5.0	1.1	0.7	0.8	0.5		-12.3
-14.6	5.0	1.0	0.6	0.8	0.5		-12.3
-12.6	5.0	1.3	1.0	0.8	0.6		-13.3
-13.6	5.0	1.1	0.6	0.8	0.5		-13.3
-14.6	5.0	1.0	0.7	0.8	0.5		-13.3
-12.6	5.0	1.3	0.8	0.8	0.4		-14.3
-13.6	5.0	1.1	0.5	0.8	0.4		-14.3
-14.6	5.0	1.1	0.4	0.8	0.4		-14.3

Table 5.2: Mean uncertainty on the DM measurements for each method and pair of (A_{RN}, A_{DM}) values in the case of $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$.

we report the mean error on the DM estimates for each method and for each set of parameters.

Figures 5.4 and 5.5 report DM residual histograms summarizing the results obtained for all the simulations per pair of spectral indices. In particular, the values reported in the histograms were obtained by normalizing the DM residuals by their error bars, in order to have insights about their Gaussianity by fitting them via a Gaussian function. The χ^2 of these fits are given in Tables 5.5 and 5.6.

²RN modeling is an inherent part of EW, so no separate analysis is possible.

	$\langle \sigma_{DM} \rangle \cdot 10^{-5} \text{ (pc cm}^{-3}\text{)}$						
A_{RN}	EW	DMX	DM GP	DMX no RN	DM GP no RN		A_{DM}
-12.6	5.0	1.3	0.9	0.8	0.6		-12.3
-13.6	5.0	1.1	0.6	0.8	0.5		-12.3
-14.6	5.0	1.0	0.6	0.8	0.5		-12.3
-12.6	5.0	1.3	0.9	0.8	0.5		-13.3
-13.6	5.0	1.0	0.6	0.8	0.5		-13.3
-14.6	5.0	1.0	0.6	0.8	0.5		-13.3
-12.6	5.0	1.3	0.7	0.8	0.4		-14.3
-13.6	5.0	1.0	0.4	0.8	0.4		-14.3
-14.6	5.0	1.0	0.4	0.8	0.4		-14.3

Table 5.3: Mean uncertainty on the DM measurements for each method and pair of (A_{RN}, A_{DM}) in the case of $\gamma_{RN} = 2.5$, $\gamma_{DM} = 3.2$.

5.5 Discussion

5.5.1 Absorption of red noise in the dispersion measure reconstruction

The presence of both RN and DM noise in our initial dataset allows us to compare the ability of each method to correctly identify the chromatic noise processes when achromatic signals are also present. The DM residuals reported in the top panel of Figure 5.3 show that none of the methods that model RN display any evident structure. On the other hand, DM GP no RN and DMX no RN show evident corruptions in the DM recovery, whose structure matches the injected RN signal (bottom panel). The same conclusion can be drawn by looking at the first column of Figures 5.4 and 5.5, in which the distributions of the normalized DM residuals calculated for those two methods are not symmetric and not centered around zero. This indicates that modelling by DM GP and DMX is not inherently restricted to modelling frequency-dependent signals, but is susceptible to absorbing frequency-independent signals as well, especially if those are not being independently modelled in the analysis. This is true as long as the RN signal is sufficiently "loud" to rise above the WN level. When this condition is not satisfied, or when the power spectrum of the RN exceeds the WN power by only a few frequency bins, then the DM recovery is not affected as much by the absence of RN modelling. This can easily be seen in the second and third columns of Figures 5.4 and 5.5, where the distributions are very similar to each other because of the low amplitude of the RN (-13.3 and -14.3 , respectively).

5.5.2 Comparison of the methods based on precision and accuracy

We have based our comparison of the methods on both the yielded precision and accuracy. We define precision as the uncertainty on the measured DM. The method that gives measurements with the lowest uncertainties, and that hence is the most precise, is DM GP, followed by DMX and then EW. This can be inferred from Tables 5.2 and 5.3, where we report the mean uncertainty on the DM recoveries for each method and set of parameters. The mean uncertainties of DM GP are around one order of magnitude smaller than the EW ones,




Figure 5.4: Histograms of DM residuals normalized by their uncertainty for EW (blue), DMX no RN (magenta), DM GP no RN (orange), DMX (red), and DM GP (green). A_{RN} decreases going from left to right (as is indicated at the bottom of the columns); A_{DM} increases going upward (as is indicated on the right of each row). The spectral indices are $\gamma_{RN} = 3.7$ and $\gamma_{DM} = 8/3$.




Figure 5.5: As in Fig. 5.4, but the spectral indices are $\gamma_{RN} = 2.5$ and $\gamma_{DM} = 3.2$.

while DMX sits in the middle of the two. When the noise model lacks the RN component (DMX no RN, DM GP no RN), and hence we fit the residuals only with the DM and WN, the mean DM uncertainty is systematically lower than for the full noise model that includes WN, RN, and DM. This is expected because without fitting for the RN we are reducing the number of parameters in the model. The behavior of the DM residuals when we consider all the 100 simulated realization together is shown in Figure 5.3 for each of the five methods. The least precise method, EW, is also associated with the largest spread in the residuals, while DM GP (with RN modelled) produces a much smaller deviation of the residuals around zero. While precision informs about the uncertainty level, it does not report how far the data points are from the true value. We define accuracy as the level of agreement between a test result and the true value³. The accuracy test was conducted by studying the histograms in Figures 5.4 and 5.5. These were obtained by normalizing the DM residuals by their uncertainties, and hence offering a measure that serves as a proxy for the deviation of each bin in the histogram from the true value in terms of sigmas. At first, we checked the fraction of residuals that lies within the 3σ threshold for each method, which for a Gaussian distribution is expected to be 99.7%. For EW, more than 99.9% of the values fall within 3σ , and this happens mostly because the uncertainties on the DM measurements are the highest among the methods. DMX is the second most accurate method, with a percentage of data points within 3σ greater than 99.5%, followed by DM GP with 98.5%. The remaining methods converge on higher percentages, comparable with the three just mentioned, only when the RN amplitude is at the lowest simulated level (-14.3).

5.5.3 Whitening of the dispersion measure residuals

Other than precision and accuracy, we also investigated the ability of each method to whiten the DM residuals, meaning that we checked whether they are random and independent, without exhibiting correlated structures. This is extremely important for noise modeling, since any structure left in the DM residuals is associated with either the absorption of power in other noise components or with artifacts introduced by the method itself. We have conducted a Gaussianity test on the normalized DM residuals, and the associated reduced chi-squared values χ_{red}^2 are reported in Tables 5.5 and 5.6. For EW, $0.87 \leq \chi_{red}^2 \leq 1.35$, which implies that the data are well described by a Gaussian distribution and that there are no, or very few, remaining time-correlated structures in the DM residuals. For DMX, $0.71 \leq \chi_{red}^2 \leq 1.49$, showing that EW does not appreciably outperform DMX in whitening the DM residuals. For DMX no RN and DM GP no RN, the results are as expected: as long as A_{RN} is high, they are not able to properly recover the DM values, and hence the χ_{red}^2 associated with the histograms of the normalized residuals is very high. If $A_{RN} = -13.6$, DMX no RN can reach χ_{red}^2 values that are closer to 1 with respect to DM GP no RN, which does so only when $A_{RN} = -14.6$. Lastly, DM GP, even though it is the most precise method, is not the most accurate and also not the best at whitening the DM residuals. The distributions in Figures 5.4 and

³ISO 5725-1:2023 in International Organization for Standardization (<https://www.iso.org/home.html>)

Table 5.4: Reduced χ^2 of a Gaussian fit to the histograms reported in Fig. 5.4. $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$.

	χ_{red}^2						
A_{RN}	EW	DMX	DM GP	DMX no RN	DM GP no RN		A_{DM}
-12.6	0.95	1.37	1.53	1108.49	2691.50		-12.3
-13.6	0.91	1.04	1.73	1.34	18.75		-12.3
-14.6	0.91	1.10	1.21	0.80	1.25		-12.3
-12.6	0.87	0.78	1.67	6154.89	6538.24		-13.3
-13.6	0.88	1.01	2.45	1.33	26.07		-13.3
-14.6	0.90	0.84	1.35	0.76	1.00		-13.3
-12.6	0.89	0.86	1.63	1104.90	1579.96		-14.3
-13.6	0.89	0.71	5.25	1.53	71.90		-14.3
-14.6	0.87	0.83	1.26	0.73	0.93		-14.3

Table 5.5: Reduced χ^2 of a Gaussian fit to the histograms reported in Fig. 5.4. $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$.

5.5 are most of the time similar to the DMX ones, but the corresponding χ_{red}^2 are not as close to 1. Time-correlated structures, both short- and long-term, are still present in DM GP DM residuals, emphasized by the aforementioned reduced error bars. The short-term time correlated structures are caused by the high-Fourier-frequency components of the injected signal (the highest of which is $f_{30} = 30/T_{\text{span}}$), while long-term ones are mainly due to the presence of RN in the data. A more explicit representation of this feature is shown in Figure 5.6, where the first two rows show DM residuals for a single realization and 100 realizations, respectively, using DM GP (similar results to DM GP can be obtained with DMX). The third row reports the comparison between the structure left in the DM residuals and the injected RN, showing that the DM modeling is indeed absorbing power from the RN process, as it displays the same signature as the injected RN signal. However, we also need to consider whether these structures have a significant impact on the timing residuals by comparing their power spectrum with the WN level. To do so, we computed a discrete Fourier transform and subsequently the power spectrum of the DM residuals at different A_{RN} and compared it with the level of the WN injected for each of the three methods. The result is shown in Figure 5.7. On average, the WN dominates over the DM residuals, which confirms that the signature remaining in the DM residuals cannot significantly be detected, and therefore will not affect the timing residuals. On the other hand, we can say that EW is not affected at all by the strength of the RN, while it significantly influences DM GP and DMX. Importantly, Figure 5.7 shows that the remaining power associated with the DM residuals is the lowest for DM GP (and DMX). This hints that DM GP (and DMX) diminishes the dispersive noise level in the ToAs to the minimum among the techniques tested, a result that would be beneficial in studies that need the most optimal noise mitigation possible. Hence, while a full study of the DM GP analysis and its interaction with various types of RN (particularly GWs) is beyond the scope of this paper, this finding indicates

	χ_{red}^2						
A_{RN}	EW	DMX	DM GP	DMX no RN	DM GP no RN		A_{DM}
-12.6	1.35	0.98	1.96	365.57	218.71		-12.3
-13.6	1.35	0.82	1.45	0.90	1.99		-12.3
-14.6	1.35	1.03	1.09	0.94	0.98		-12.3
-12.6	1.17	1.49	3.35	168.97	194.42		-13.3
-13.6	1.16	1.32	1.58	1.02	1.20		-13.3
-14.6	1.17	0.76	1.13	0.97	1.12		-13.3
-12.6	0.98	1.29	2.20	33.12	141.30		-14.3
-13.6	0.98	0.96	0.95	0.74	1.79		-14.3
-14.6	0.98	0.95	0.77	0.77	1.45		-14.3

Table 5.6: Reduced χ^2 obtained from a Gaussian fit on the histograms reported in Fig. 5.5. $\gamma_{RN} = 2.5$, $\gamma_{DM} = 3.2$.

that DM GP (and DMX) might work best for whitening timing residuals; for example, in a context such as the search for low-frequency GWs.

5.5.4 Alternative simulation settings

Here, we test the methods on datasets that we simulated by varying specific settings. First, we reduced the uncertainties of the simulated ToAs from our initial value (from $5\mu s$ to $1\mu s$ and $0.1\mu s$) to improve our sensitivity to the RN process and DM variations, and then we simulated the ToAs at L-band frequencies. Finally, we changed the injected DM time series.

5.5.4.1 Testing different time-of-arrival uncertainties

We used $\sigma_{temp} = 5\mu s$, since it represents the typical mean uncertainties on ToAs obtained from LOFAR observations (Donner et al., 2020). However, there might be pulsars with higher ToA precision, and we expect that future facilities such as the low-frequency part of the Square Kilometer Array (Janssen et al., 2015) will improve the ToA sensitivity. For this reason, we ran two new simulation sets, with decreased values for the σ_{temp} , firstly at $1\mu s$ and then at $0.1\mu s$. Furthermore, as this reduces the WN floor level, we enhanced the methods' capability to identify RN and properly disentangle it from DM variations. The WN parameters injected in the simulations were also changed to: EFAC=1, EQUAD= $10^{-8} s$. Chromatic and achromatic red processes were both set to 30 Fourier-frequency components, with $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$, $A_{RN} = -13.6$, and $A_{DM} = -13.3$.

In Table 5.7, we report the DM uncertainty levels for each method and each uncertainty. Given the smaller error bars of the ToAs, the precision on the DM measurements increases, as was expected. As we did previously, we can rank DM GP as the most precise method, followed by DMX and EW.

We recall that an important finding from our simulations (see Section 5.5.3) was about the RN signal left in the DM residuals. In fact, Figure 5.6 (bottom panel) shows that the signal left in the DM residuals matches exactly the one of the injected RN. The corresponding delay, which is thus left on average




Figure 5.6: The top panel shows the DM residuals obtained by applying DM GP (with RN modeling) to an individual simulation over the 100 realizations for values of the noise parameters $A_{RN} = -13.6$, $A_{DM} = -13.3$, $\gamma_{RN} = 3.7$, $\gamma_{DM} = 8/3$, EFAC and EQUAD as in Table 5.1, and $\sigma_{\text{ToA}} = 5\mu\text{s}$. The second panel reports an average of the DM residuals left by DM GP over the 100 simulations performed with the same parameters. The bottom panel shows the injected RN signal (red dots) compared to the time delay associated with the mean DM GP (shaded green). Note that we have a shaded region because a single DM value causes a frequency-dependent delay. Hence, the edges of the green region correspond to the maximum and minimum frequencies of the band.




Figure 5.7: Power spectra of the DM residuals left by the EW (left panel), DMX (central panel), and DM GP (right panel) methods. We also show the WN level (black) and the power spectra for different RN amplitudes: $A_{RN} = -14.6$ (orange with vertical dashes); $A_{RN} = -13.6$ (green with dots); and $A_{RN} = -12.6$ (blue with oblique dashes). With DM GP, we modeled the noise up to a Fourier-frequency of $30/T_{\text{span}} \sim 1.15 \times 10^{-7}$ Hz, the same as the injected signal. That is why in the right-hand panel there is a drop at that specific frequency in the power spectrum of the DM GP residuals.

Table 5.7: Mean DM uncertainty across methods (EW, DMX, DM GP) for simulated ToAs with fixed uncertainties of $1 \mu\text{s}$ and $0.1 \mu\text{s}$

σ_{temp}	$\langle\sigma_{DM}\rangle \cdot 10^{-6} \quad (\text{pc cm}^{-3})$		
	EW	DMX	DM GP
$1 \mu\text{s}$	6.3	2.0	1.6
$0.1 \mu\text{s}$	0.6	0.5	0.2

Table 5.8: As in Table 5.7, but for L-band ToAs with $\sigma_{\text{temp}} = 1 \mu\text{s}$ and injected DM noise parameters of $A_{DM} = -13.6$ and $\gamma_{DM} = 3.7$.

$\langle\sigma_{DM}\rangle \cdot 10^{-4} \quad (\text{pc cm}^{-3})$		
EW	DMX	DM GP
7.8	2.1	0.7

in the timing residuals, is about $\pm 2 \mu\text{s}$. This changes in the two new cases characterized by low ToA uncertainties: when $\sigma_{\text{temp}} = 1 \mu\text{s}$, the mean DM signal that is left due to an erroneous RN absorption contributes to the timing residuals at the level of $\pm 150 \text{ ns}$; when $\sigma_{\text{temp}} = 0.1 \mu\text{s}$, this becomes $\pm 20 \text{ ns}$. Since the injected RN signal was the same, we can confirm that higher precision on the ToAs allows to increase the sensitivity to the RN and improve the accuracy of the measurements of DM variations.

5.5.4.2 Testing times of arrival in the L band

In this part, we report the results obtained by applying the DM computation methods to L-band ToAs. To simulate high-frequency ToAs, we refer to the data collected by the Nançay decimeter radiotelescope (NRT), which has a relatively large bandwidth (512 MHz) centered at 1.4 GHz divided into four subbands (EPTA Collaboration et al., 2023a). The input parameters were: $\text{EFAC} = 1$, $\text{EQUAD} = 10^{-8}\text{s}$, $A_{DM} = -13.3$, and $\gamma_{DM} = 8/3$. We used a conservative $1 \mu\text{s}$ ToA template-fitting error. Concerning the precision of the methods, the results are similar to the ones in the LOFAR frequency case: DM GP is the most precise, followed by DMX and EW. However, the size of these errors is now larger by approximately two orders of magnitude (see Table 5.8). This is due to two main reasons: the first is the increased central frequency of the analysed ToAs (which is 1.4 GHz against the 153 MHz of the previous LOFAR-like case), because DM is inversely proportional to the squared of the observing frequency. The second is the fractional bandwidth, f_B (see Equation 1.13): the higher f_B , the better the precision in the DM measurements (Verbiest and Shaifullah 2018). For the NRT-like case, $f_B \simeq 0.34$, while for the LOFAR-like case, $f_B \simeq 0.41$.

Concerning accuracy, we also report similar conclusions as before: EW has more than 99.76% of the values within 3σ and it is the most accurate method, followed by DMX with more than 99.44% and finally DM GP with 99.2%. In addition, the Gaussianity test provides similar results to the LOFAR-like case. Table 5.9 lists the reduced χ^2 values, which confirm that on average EW and DMX are marginally better at whitening the residuals due to their larger error

Table 5.9: Reduced χ^2 of the normalized DM residuals. Obtained from a Gaussian fit on histograms of DM residuals normalized by their uncertainty for the EW, DMX, and DM GP methods, over ToAs simulated in the L band with $\sigma_{\text{temp}} = 1 \mu\text{s}$ and injected DM noise ($A_{DM}, \gamma_{DM} = -13.3, 8/3$).

χ^2_{red}		
EW	DMX	DM GP
1.07	1.10	1.54




Figure 5.8: Top panel: DM residuals left by the EW (blue), DMX (red), and DM GP (green) methods when run over the 100 realizations of L-band ToAs. Bottom panel: Histograms of the DM residuals normalized by their uncertainty. The input noise parameters are $A_{DM} = -13.3$ and $\gamma_{DM} = 8/3$

bars (and thus lower precision), on single DM recoveries. In the top panel of Figure 5.8 are presented the DM residuals for each of the three methods, while the corresponding normalized histograms are reported in the bottom plot.

Figure 5.9 shows the power spectra of the DM residuals compared with the injected signal. Despite the fact that the WN is dominating in most of the frequency bins, all of the three methods can model the long-term structures present in the data. At the highest frequencies in the spectra, EW and DMX cannot match the injected signal anymore due to the limits imposed by the WN. Instead, DM GP is the only method that produces a power spectrum that does not cross the injection line, meaning that it is not adding any more power at those frequencies. We note that DM GP uses a power-law model to describe the noise process, which is identical to the functional form of the injected signal. Therefore, it is possible that the reported results show a biased result for DM GP, since in real-life datasets we may not have pure power-law noise processes.

To evaluate a more realistic case, in the following section we test the methods on an injected DM time series that does not come from a pure power law.




Figure 5.9: Power spectra of the DM residuals obtained using EW (left), DMX (center), and DM GP (right), compared with the injected signal (light blue) in the L band. The input noise parameters are $A_{DM} = -13.3$ and $\gamma_{DM} = 8/3$. The injected signal power (light-blue) drops at a frequency of $30/T_{span} \sim 1.15 \times 10^{-7}$ Hz. We used the same number of frequency components in the modelling with DM GP.




Figure 5.10: DM time series injected in the data with an empirical power spectrum (Donner, 2022).

5.5.4.3 Custom dispersion measure injection

Here, we test the performance of the three considered techniques over an individual, injected DM time-series with an empirical power-spectrum shape, derived in Donner, 2022 from the LOFAR dataset of PSR J0139+5814 (see Fig. 5.10). The ToAs are still evenly spaced, with an uncertainty set to $5\mu\text{s}$, and the injected WN is defined by $\text{EFAC} = 1.2$ and $\text{EQUAD} = 2\mu\text{s}$. Given the short timescale fluctuations in the injected signal, we used 107 Fourier components to define the power-spectrum used in DM GP, which corresponds to the limit imposed by the Nyquist theorem given the 215 data points. The result of the analysis are reported in Figure 5.11, where we show the DM residuals in the top panel and their power spectra in the bottom panel. As previously noted, EW and DMX provide independent DM measurements (although DMX might introduce some degree of correlations for close-by points; see Lam et al., 2017), while DM GP assumes an underlying model for the power spectrum to describe the DM variations. The DM GP error bars are now larger than the DMX ones, and hence the recoveries are extremely accurate. The bottom panel of Figure 5.11 shows that even with the injected signal resembling a broken power law, DM GP is able to absorb the signal, and hence successfully recover the injected DM




Figure 5.11: DM residuals for EW (blue), DM GP (green), and DMX (red) obtained by taking the difference of the DMs recovered by each method and the injected DM time series of Figure 5.10 with an unknown power spectrum. The WN parameters are EFAC = 1.2 and EQUAD $2\mu\text{s}$ and the ToA uncertainty was set at $5\mu\text{s}$. Bottom panel: Power spectra of the injected signal (black) and of the DM residuals.

time series.

5.6 Conclusions

In this work, we carried out a series of simulations in order to test the performance of the three main methods available in the PTA community to calculate DM variations: EW, DM GP, and DMX. The aim is to provide information on the calculation of the DM time series for, for example, studies of the IISM and the solar wind.

The main conclusion is that, while they all perform well, the most accurate method of calculating DM variations is EW, but DM GP and DMX seem to lower the level of the DM noise in the timing residuals the most.

In more detail, we conclude that all of the methods perform satisfactorily (as long as the RN and DM models are selected correctly for DMX and DM GP). Figure 5.3 demonstrates that, when the RN power is significant, including it in the noise model is mandatory; otherwise, the achromatic, time-correlated structures will be absorbed in the DM modeling in the cases of DMX and DM GP.

With respect to the precision, the mean uncertainties $\langle\sigma_{DM}\rangle$ on the DM recoveries for EW are independent of the amplitudes of the RN processes. This means that the method is only affected by the ToA precision and the frequency coverage of the data. For DMX and DM GP, instead, $\langle\sigma_{DM}\rangle$ depends on the amplitude of the injected RN process (see Tables 5.2 and 5.3). Overall, the method that gives the most precise DM measurements is DM GP.

We also studied the accuracy, defined as the level of agreement between a

test result and the true value while taking into account the uncertainties, by analyzing the distributions of the normalized DM residuals (Figures 5.4 and 5.5). Because of the larger error bars, EW is the most accurate method, having in most of the cases $>99.9\%$ of the data within 3σ . The second most accurate method, according to this criterion, is DMX, while DM GP is less accurate because of its high DM precision, which results in more data points in the tail of the distribution. This is true not only at LOFAR frequencies, but also in the L-band trials, where we simulated an NRT-like frequency coverage.

The accuracy test, though, as is described above, does not imply that the DM residuals are white. Hence, we have conducted a Gaussianity test on each of the normalized DM residuals distributions, which informs one about the ability of each method to remove time-correlated structures from the ToAs. The χ^2_{red} of the Gaussian fits are reported in Tables 5.5 and 5.6, and they show that EW is the method with the narrowest range of χ^2_{red} around 1. DMX shows similar results, but only when the RN is modeled. However, the χ^2_{red} range shown by DMX fluctuates more than for EW. DM GP, which uses a linear combination of Fourier components, produces a time series where close data points are not completely independent. This affects the overall residual distribution, resulting in the widest χ^2_{red} range spanned among the three methods.

To better understand the reasons behind the DMX and DM GP distributions, we studied the effect of the RN on the DM residuals. In Fig. 5.6, we show that even when the RN is modeled, a signature can still be found in the DM residuals yielded by DM GP and DMX. This happens when the WN level exceeds the RN, especially at the highest Fourier frequencies. Nonetheless, Fig. 5.7 shows that, at LOFAR frequencies, the unmodeled RN signal remains below the WN level. In order to properly disentangle the RN and DM noise signal, it is necessary to have high-precision ToAs. We investigated this scenario by simulating datasets with reduced uncertainties, $\sigma_{ToAs} = 1 \mu s, 0.1 \mu s$, and, as was expected, DMX and DM GP no longer show any significant structures related to the injected RN signal in the DM residuals. The aforementioned precision-accuracy hierarchy holds in this simulation scenario as well. Similar conclusions are reached in (albeit limited) simulations carried out in the L band and with an injected DM time series with an empirical spectral shape.

However, we stress that these results are based on simulations that do not take into account uneven observational cadence, variable ToA uncertainty, poor, if not absent, frequency coverage of the data, and also combination of datasets; in other words, the most common characteristics of real-life data. When these additional complexities are present, it is possible that DM GP performs better than other methods thanks to its flexibility and the fact that its characteristics allow a more optimal blending of data with diverse properties. Moreover, the performance of EW and DMX might be affected by more realistic datasets because of their windowing constraints. A firmer conclusion will be obtained once these scenarios are taken into consideration. Also, the impact of these DM recovery schemes on timing model parameters has not been assessed in our work, but Kramer et al. (2021) has already shown that astrometric parameters can be corrupted by some of the schemes tested. Consequently, future work on this topic would require a complete assessment of that aspect, as well.

Last, we reach the important conclusion that DM GP and DMX seem to have the capacity to reduce the DM noise level to the minimum (see Fig. 5.7). Future

works should assess the effectiveness of the analyzed methods in modelling DM variations as a source of noise when a GW background is also present in the data.

Summary and Conclusions

This thesis has explored how low-frequency pulsar timing can improve our understanding of DM variations and, more broadly, of the IISM. Pulsars act as remarkably stable cosmic clocks, and their precise timing makes them powerful tools for both astrophysical studies and fundamental physics. However, the achievable timing precision is limited by propagation effects introduced by the IISM, which produce time-varying chromatic delays.

The central aim of this work was to use low-frequency observations to characterise these dispersive effects, link them to the physical conditions of the intervening plasma, and assess their impact on the noise modelling adopted in PTA analyses. The thesis combined observational and simulation-based approaches, focusing on three main aspects: the synergy between LOFAR and EPTA datasets, the identification of discrete ionized regions along the line-of-sight, and the evaluation of different techniques for modelling DM variations. Together, these studies provide a clearer and more quantitative view of how low-frequency timing enhances both our understanding of the IISM and the reliability of PTA noise models.

Low-frequency telescopes such as LOFAR and NenuFAR have transformed the study of DM variations. Operating below 200 MHz, they are far more sensitive to dispersive effects than traditional high-frequency PTA telescopes. Even small changes in the column density of free electrons along the line of sight result in measurable variations in the timing residuals at these frequencies. This sensitivity makes low-frequency data a unique probe of the IISM. The amplitude and timescale of DM fluctuations observed with LOFAR and NenuFAR provide direct insights into the turbulence spectrum, the distribution of plasma irregularities, and the presence of possible localized structures. Moreover, these observations enable more accurate corrections of chromatic delays, which in turn reduce the contamination of timing models by DM-related noise. Throughout this work, it has become evident that low-frequency data are not only a complementary tool to high-frequency timing, but also an essential ingredient for accurate noise characterisation. Their inclusion allows a clearer distinction between achromatic red noise, often attributed to intrinsic spin irregularities, and chromatic variations induced by the interstellar medium.

One of the main outcomes of this thesis was the successful combination of LOFAR and NenuFAR low-frequency data with high-frequency EPTA observations for common pulsars. This marked the first time that such low-frequency data became part of a PTA dataset. The main challenge lay in preparing the data and developing a consistent procedure for the combination, given that both LOFAR and NenuFAR provide frequency-resolved ToAs and therefore require additional care during integration. The major advantage of this approach was the extremely wide frequency coverage achieved with the combined dataset, which enabled unprecedented precision in characterising both propagation effects and achromatic red noise in the timing residuals. The joint analysis demonstrated that, with low-frequency observations, we could place tighter constraints on the amplitude and spectral slope of DM variations, which are often poorly determined when relying solely on high-frequency data, and remove their correlation with the achromatic red noise. These findings supported the idea that low-frequency timing should become a standard component of PTA noise modelling, as it offers a powerful means of addressing one of the most significant sources of systematic uncertainty in long-term timing analyses. Another important aspect concerned the solar wind (SW) model. Using such low-frequency data allowed us to identify issues in the current SW modelling scheme, which fails to capture the full dynamics of the signal. As a consequence, the residuals left unmodelled by the SW component are partially absorbed by the DM variation model. This represents a crucial insight, especially considering that these low-frequency data will be included in the upcoming data release of the IPTA.

Building on the insights gained from the combined datasets, this thesis also explored how low-frequency pulsar timing can be used to probe the turbulent nature of the interstellar plasma. Variations in DM arise from the turbulent fluctuations of the IISM, which are commonly described by a Kolmogorov power-law spectrum. Searching for possible deviations from this model, the work focused on studying the turbulence through low-frequency observations—owing to their strong sensitivity to propagation effects—and in particular on pulsars whose lines of sight intersect discrete, densely ionised regions (H II regions). By cross-matching pulsar coordinates with Galactic catalogues of H II regions and analysing their distances and velocities, potential associations were identified. A total of six pulsars in the LOFAR catalogue were found to have one or more H II regions along their lines of sight. Using the DM time series of these pulsars, it was possible to recover the power spectrum of density fluctuations associated with the H II regions, under the assumption that the observed variations are dominated by turbulence within the regions themselves. No significant deviations from the expected Kolmogorov spectrum were found, although these are indications of stronger turbulence in pulsars intersecting H II regions. Additionally, a scintillation analysis was performed to test whether the assumption of turbulence dominated by discrete structures was consistent with the observed variability. With LOFAR data, it was possible to study only the refractive scintillation, associated with long-timescale flux variations, while for the diffractive scintillation only estimates could be obtained. Future work should include diffractive scintillation measurements at higher frequencies to compare with the results presented here and refine the picture of the underlying turbulence. Overall, although not all cases were conclusive, the results emphasise that discrete structures in the IISM can play a non-negligible role

in shaping DM variations. Understanding and modelling these contributions is crucial for correctly interpreting low-frequency timing data, as the presence of such structures could help explain some of the deviations from Kolmogorov-like turbulence observed in structure functions and DM power spectra.

Beyond the observational studies, the final analytical component of this thesis consisted of extensive simulations designed to test different methods for estimating DM variations from timing data. Three approaches were compared: the simple epoch-wise (EW) method, the DMX technique, and the Bayesian Gaussian Process framework (DM GP). By controlling the injected signals, it was possible to evaluate the recovery performance of each method and to quantify, through specific statistical metrics, their relative effectiveness. The simulated datasets included not only DM variations but also WN and achromatic RN, in order to reproduce realistic low-frequency timing conditions. These tests provided a clear assessment of the strengths and limitations of each method, offering practical guidance for future analyses. For PTA applications, both DMX and DM GP emerged as the most effective trade-off solutions, achieving significant noise suppression while simultaneously modelling both the chromatic and achromatic components of the signal.

The results presented in this thesis open several perspectives for future research on the ionised interstellar medium. The combined use of DM time series, scintillation measurements, and SW modelling will be essential to build a more comprehensive understanding of electron density fluctuations across multiple spatial and temporal scales. An important next step will be to incorporate SW contributions into DM models in a more dynamic and physically motivated way. As PTAs extend their observational baselines, solar activity cycles will increasingly influence timing precision, particularly for pulsars located near the ecliptic. The development of free-spectrum approaches, where the DM power spectrum is modelled without assuming a fixed slope, will help determine how the power is distributed among different Fourier frequencies and noise components.

Low-frequency pulsar timing is emerging as a key component of precision astrophysics. The results of this thesis demonstrate how the inclusion of data from LOFAR and NenuFAR enhances our understanding of DM variability, noise characterisation, and interstellar turbulence. These findings highlight the enormous potential of future facilities such as the SKA, whose low-frequency capabilities will push this field to a new level of sensitivity and precision. Lessons learned from LOFAR and NenuFAR will be directly applicable to the design of SKA data analysis pipelines, particularly in the areas of DM correction and noise modelling.

Ultimately, low-frequency observations will remain indispensable for both fundamental PTA science and the study of the IISM itself. They not only extend our view of the turbulent plasma between stars but also ensure that the search for nanohertz GWs rests on a solid understanding of every source of timing noise along the way.

APPENDIX A

Appendix to Chap. 3

A.1 Parameters of the noise models used in this work

Pulsar	ELAT (deg.)	SW modelling – DR2low		SW modelling – DR2new+	
		Timing - NE_SW	Noise	Timing - NE_SW	Noise
J1022+1001	−0.06	9.51 ± 0.03 (fitted)	<code>n_earth + SWv</code>	7.2 ± 0.5 (fitted)	marg.
J1730−2304	0.19	9.9 ± 0.4 (fitted)	<code>n_earth</code>	7.9 (fixed)	marg.
J0030+0451	1.45	8.20 ± 0.03 (fitted)	<code>n_earth + SWv</code>	7.9 (fixed)	marg.
J0751+1807	−2.81	7.9 (fixed)	<code>n_earth + SWv</code>	7.9 (fixed)	marg.
J1744−1134	11.81	3.62 ± 0.08 (fitted)	<code>n_earth + SWv</code>	7.9 (fixed)	marg.
J1918−0642	15.35	4 (fixed)	marg.	7.9 (fixed)	marg.
J1024−0719	−16.04	4 (fixed)	marg.	7.9 (fixed)	marg.
J1738+0333	26.88	4 (fixed)	marg.	7.9 (fixed)	marg.
J1713+0747	30.70	4 (fixed)	marg.	7.9 (fixed)	marg.
J1857+0943	32.32	4 (fixed)	marg.	7.9 (fixed)	marg.
J1012+5307	38.76	4 (fixed)	marg.	7.9 (fixed)	marg.
J1640+2224	44.06	4 (fixed)	marg.	7.9 (fixed)	marg.

Table A.1: Approach for modelling the dispersion caused by solar winds inspired by Susarla et al., 2024. The pulsar names and ecliptic latitudes (in degree) are displayed respectively in the first and second columns. The methods are described for DR2low (third major column) and DR2new+ (last major column), where the timing analysis include the constant term only (either fixed or fitted with Tempo2), and where the noise analysis includes the constant term, by marginalizing over its error ("marg."; cf. Section 3.4) or fitting for it ("n_{earth}") and eventually adds its time-varying component ("SWv"), where the fluctuations of the electron density at one astronomical unit are modelled as a power-law PSD parameterised with its amplitude A_{SW} referred at $f = 1 \text{ yr}^{-1}$ and spectral index γ_{SW} .

A.2 Estimation of the ecorr parameters




Figure A.1: Posterior distribution of the ecorr parameters (logarithmic scale) and their 99.7% credible intervals (errorbars) obtained for the LOFAR systems for each pulsars.

A.3 Solar wind parameters




Figure A.2: Posterior distribution of the SW parameters for PSRs J0030+0451, J0751+1807, J1022+1001 and J1744-1134.

APPENDIX B

Appendix to Chap. 4

B.1 Generation of simulated DM time series from the free spectrum

The PSD $S(f)$ obtained from the free-spectrum analysis is defined as:

$$S(f_i) = \rho_i^2 \cdot T_{\text{span}}, \quad (\text{B.1})$$

where each ρ_i is sampled from the posterior using enterprise, and T_{span} is the total time span of the observations. The corresponding Fourier basis functions are:

$$F_{\text{even},i}(t) = \sin(2\pi f_i t), \quad (\text{B.2})$$

$$F_{\text{odd},i}(t) = \cos(2\pi f_i t), \quad (\text{B.3})$$

where t is the array of observing epochs and f_i are the Fourier frequencies. The red noise signal in timing residuals is reconstructed by summing over sine and cosine components:

$$\delta t(t) = \sum_i [a_i \sin(2\pi f_i t) + b_i \cos(2\pi f_i t)] \cdot D \left(\frac{\nu_{\text{ref}}}{\nu} \right)^2, \quad (\text{B.4})$$

where D is the dispersion constant, ν is the observing frequency, and $\nu_{\text{ref}} = 1400$ MHz is a fixed reference frequency.

Assuming that the power is equally distributed between the sine and cosine terms, the Fourier amplitudes a_i and b_i are sampled independently from a normal distribution:

$$a_i, b_i \sim \mathcal{N} \left(0, \sqrt{\frac{S(f_i)}{2}} \right). \quad (\text{B.5})$$

To convert the reconstructed signal from residual units (time) to dispersion measure units, we use the inverse of the frequency scaling:

$$DM_{\text{red}}^{(n)}(t) = \frac{\delta t(t)}{D} \cdot \left(\frac{\nu}{\nu_{\text{ref}}} \right)^2, \quad (\text{B.6})$$

where $n = 1, 2, \dots, N$ denotes the realization index. This procedure generates an ensemble of DM time series with the same power spectral characteristics as the real data, allowing for statistical comparisons of their structure functions.

B.2 Transformation of Proper Motion and Computation of Effective Velocity

Given a pulsar with known equatorial coordinates (α, δ) , proper motion components (μ_α, μ_δ) , and total transverse velocity V_T , we compute the components of the transverse velocity in the *ecliptic coordinate system*, i.e., along ecliptic latitude (β) and longitude (λ) .

B.2.1 Coordinate Conversion

Equatorial coordinates (α, δ) are first converted to degrees and used to define an ICRS (equatorial) position. A rotation governed by the obliquity of the ecliptic $\varepsilon \simeq 23.44^\circ$ transforms these into ecliptic coordinates (λ, β) .

B.2.2 Angle Between Coordinate Axes

At the pulsar's position, the local directions of the ecliptic basis vectors $(\hat{\lambda}, \hat{\beta})$ are rotated with respect to the equatorial basis $(\hat{\alpha}, \hat{\delta})$ by an angle θ , computed via:

$$\sin \beta = \sin \delta \cos \varepsilon - \cos \delta \sin \varepsilon \sin \alpha \quad (\text{B.7})$$

$$\tan \lambda = \frac{\sin \alpha \cos \varepsilon + \tan \delta \sin \varepsilon}{\cos \delta} \quad (\text{B.8})$$

$$\theta = \arccos \left(\frac{\cos \varepsilon - \sin \beta \sin \delta}{\cos \beta \cos \delta} \right) \quad (\text{B.9})$$

To preserve orientation, we apply $\theta \rightarrow -\theta$ if $90^\circ < \alpha < 270^\circ$.

B.2.3 Proper Motion Rotation

We rotate the proper motion vector to the ecliptic frame using:

$$\begin{bmatrix} \mu_\lambda \\ \mu_\beta \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \mu_\alpha \\ \mu_\delta \end{bmatrix} \quad (\text{B.10})$$

This yields the proper motion components along λ and β .

B.2.4 Transverse Velocity Projection

Assuming that the transverse velocity vector shares the direction of proper motion, we compute:

$$\mu_{\text{tot}} = \sqrt{\mu_\lambda^2 + \mu_\beta^2} \quad (\text{B.11})$$

$$v_\lambda = V_T \cdot \frac{\mu_\lambda}{\mu_{\text{tot}}} \quad (\text{B.12})$$

$$v_\beta = V_T \cdot \frac{\mu_\beta}{\mu_{\text{tot}}} \quad (\text{B.13})$$

These are the ecliptic components of the pulsar's transverse velocity vector $\vec{v}_{\text{psr}}$.

B.2.5 Effective Velocity

The *effective velocity* $\vec{v}_{\text{eff}}$ is the apparent transverse velocity of the pulsar image on the scattering screen as seen by the observer. It is given by a weighted combination of the pulsar $\vec{v}_{\text{psr}}$ and Earth's $\vec{v}_{\oplus}$ transverse velocities:

$$\vec{v}_{\text{eff}}(t) = (1 - s) \vec{v}_{\text{psr}} + s \vec{v}_{\oplus}(t), \quad (\text{B.14})$$

where $s = (d_{\text{psr}} - d_{\text{screen}})/d_{\text{screen}}$ is the fractional distance of the screen given the pulsar distance d_{psr} and the screen distance d_{screen} . We approximate the Earth's velocity components in the ecliptic frame as:

$$v_{\oplus,\lambda}(t) = V_0 \sin \beta \cos(2\pi t) \quad (\text{B.15})$$

$$v_{\oplus,\beta}(t) = -V_0 \sin(2\pi t) \quad (\text{B.16})$$

with $V_0 = 2\pi \text{ AU}/1 \text{ yr} \approx 30 \text{ km s}^{-1}$, t in years, and β the pulsar's ecliptic latitude. Finally, the effective velocity becomes:

$$\vec{v}_{\text{eff}}(t) = (1 - s) \begin{bmatrix} v_{\lambda} \\ v_{\beta} \end{bmatrix} + s \begin{bmatrix} v_{\oplus,\lambda}(t) \\ v_{\oplus,\beta}(t) \end{bmatrix} \quad (\text{B.17})$$

Bibliography

- [1] G. Agazie et al. “Comparing Recent Pulsar Timing Array Results on the Nanohertz Stochastic Gravitational-wave Background”. In: *The Astrophysical Journal* 966.1 (2024), p. 105.
- [2] G. Agazie et al. “The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background”. In: *The Astrophysical Journal Letters* 951.1, L8 (July 2023), p. L8. DOI: [10.3847/2041-8213/acdac6](https://doi.org/10.3847/2041-8213/acdac6). arXiv: [2306.16213](https://arxiv.org/abs/2306.16213) [[astro-ph.HE](#)].
- [3] G. Agazie et al. “The NANOGrav 15 yr Data Set: Observations and Timing of 68 Millisecond Pulsars”. In: *The Astrophysical Journal Letters* 951.1, L9 (July 2023), p. L9. DOI: [10.3847/2041-8213/acda9a](https://doi.org/10.3847/2041-8213/acda9a). arXiv: [2306.16217](https://arxiv.org/abs/2306.16217) [[astro-ph.HE](#)].
- [4] M. F. Alam et al. “The NANOGrav 12.5 yr Data Set: Observations and Narrowband Timing of 47 Millisecond Pulsars”. In: *ApJS* 252.1, 4 (Jan. 2021), p. 4. DOI: [10.3847/1538-4365/abc6a0](https://doi.org/10.3847/1538-4365/abc6a0). arXiv: [2005.06490](https://arxiv.org/abs/2005.06490) [[astro-ph.HE](#)].
- [5] B. Allen et al. “The International Pulsar Timing Array checklist for the detection of nanohertz gravitational waves”. In: *arXiv e-prints*, arXiv:2304.04767 (Apr. 2023), arXiv:2304.04767. DOI: [10.48550/arXiv.2304.04767](https://doi.org/10.48550/arXiv.2304.04767). arXiv: [2304.04767](https://arxiv.org/abs/2304.04767) [[astro-ph.IM](#)].
- [6] L. D. Anderson et al. “The Green Bank Telescope H II Region Discovery Survey. II. The Source Catalog”. In: *ApJS* 194.2, 32 (June 2011), p. 32. DOI: [10.1088/0067-0049/194/2/32](https://doi.org/10.1088/0067-0049/194/2/32). arXiv: [1103.5085](https://arxiv.org/abs/1103.5085) [[astro-ph.GA](#)].
- [7] L. D. Anderson et al. “The WISE Catalog of Galactic H II Regions”. In: *ApJS* 212.1, 1 (May 2014), p. 1. DOI: [10.1088/0067-0049/212/1/1](https://doi.org/10.1088/0067-0049/212/1/1). arXiv: [1312.6202](https://arxiv.org/abs/1312.6202) [[astro-ph.GA](#)].
- [8] W. P. Armentrout et al. “A VLA Census of the Galactic H II Region Population”. In: *ApJS* 253.1, 23 (Mar. 2021), p. 23. DOI: [10.3847/1538-4365/abd5c0](https://doi.org/10.3847/1538-4365/abd5c0). arXiv: [2103.02049](https://arxiv.org/abs/2103.02049) [[astro-ph.GA](#)].
- [9] J. W. Armstrong, B. J. Rickett, and S. R. Spangler. “Electron Density Power Spectrum in the Local Interstellar Medium”. In: *ApJ* 443 (Apr. 1995), p. 209. DOI: [10.1086/175515](https://doi.org/10.1086/175515).

- [10] D. Bhattacharya and E. P. J. van den Heuvel. “Formation and evolution of binary and millisecond radio pulsars”. In: *Physics Reports* 203.1-2 (Jan. 1991), pp. 1–124. DOI: [10.1016/0370-1573\(91\)90064-S](https://doi.org/10.1016/0370-1573(91)90064-S).
- [11] L. Biermann. “Kometenschweife und solare Korpuskularstrahlung”. In: *ZAp* 29 (Jan. 1951), p. 274.
- [12] A. V. Bilous et al. “A LOFAR census of non-recycled pulsars: average profiles, dispersion measures, flux densities, and spectra”. In: *A&A* 591, A134 (June 2016), A134. DOI: [10.1051/0004-6361/201527702](https://doi.org/10.1051/0004-6361/201527702). arXiv: [1511.01767](https://arxiv.org/abs/1511.01767) [[astro-ph.SR](#)].
- [13] L. Bondonneau et al. “Pulsars with NenuFAR: Backend and pipelines”. In: *A&A* 652, A34 (Aug. 2021), A34. DOI: [10.1051/0004-6361/202039339](https://doi.org/10.1051/0004-6361/202039339). arXiv: [2009.02076](https://arxiv.org/abs/2009.02076) [[astro-ph.IM](#)].
- [14] S. Burke-Spolaor et al. “The astrophysics of nanohertz gravitational waves”. In: *Astronomy and Astrophysics Review* 27.1, 5 (June 2019), p. 5. DOI: [10.1007/s00159-019-0115-7](https://doi.org/10.1007/s00159-019-0115-7). arXiv: [1811.08826](https://arxiv.org/abs/1811.08826) [[astro-ph.HE](#)].
- [15] A. Chalumeau et al. “Noise analysis in the European Pulsar Timing Array data release 2 and its implications on the gravitational-wave background search”. In: *MNRAS* 509.4 (Feb. 2022), pp. 5538–5558. DOI: [10.1093/mnras/stab3283](https://doi.org/10.1093/mnras/stab3283). arXiv: [2111.05186](https://arxiv.org/abs/2111.05186) [[astro-ph.HE](#)].
- [16] S. Chatterjee et al. “Pulsar Parallaxes at 5 GHz with the Very Long Baseline Array”. In: *ApJ* 604.1 (Mar. 2004), pp. 339–345. DOI: [10.1086/381748](https://doi.org/10.1086/381748). arXiv: [astro-ph/0312044](https://arxiv.org/abs/astro-ph/0312044) [[astro-ph](#)].
- [17] K. Chen and M. Ruderman. “Pulsar Death Lines and Death Valley”. In: *ApJ* 402 (Jan. 1993), p. 264. DOI: [10.1086/172129](https://doi.org/10.1086/172129).
- [18] L. Cohen. “The generalization of the Wiener-Khinchin theorem”. In: *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP '98 (Cat. No.98CH36181)*. Vol. 3. 1998, 1577–1580 vol.3. DOI: [10.1109/ICASSP.1998.681753](https://doi.org/10.1109/ICASSP.1998.681753).
- [19] W. A. Coles et al. “Pulsar Observations of Extreme Scattering Events”. In: *ApJ* 808.2, 113 (Aug. 2015), p. 113. DOI: [10.1088/0004-637X/808/2/113](https://doi.org/10.1088/0004-637X/808/2/113). arXiv: [1506.07948](https://arxiv.org/abs/1506.07948) [[astro-ph.SR](#)].
- [20] W. A. Coles et al. “Refractive Scintillation in the Interstellar Medium”. In: *ApJ* 315 (Apr. 1987), p. 666. DOI: [10.1086/165168](https://doi.org/10.1086/165168).
- [21] J. M. Cordes and T. J. W. Lazio. “NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations”. In: *arXiv e-prints*, astro-ph/0207156 (July 2002), astro-ph/0207156. DOI: [10.48550/arXiv.astro-ph/0207156](https://doi.org/10.48550/arXiv.astro-ph/0207156). arXiv: [astro-ph/0207156](https://arxiv.org/abs/astro-ph/0207156) [[astro-ph](#)].
- [22] J. M. Cordes, A. Pidwerbetsky, and R. V. E. Lovelace. “Refractive and Diffractive Scattering in the Interstellar Medium”. In: *ApJ* 310 (Nov. 1986), p. 737. DOI: [10.1086/164728](https://doi.org/10.1086/164728).
- [23] J. M. Cordes and B. J. Rickett. “Diffractive Interstellar Scintillation Timescales and Velocities”. In: *ApJ* 507.2 (Nov. 1998), pp. 846–860. DOI: [10.1086/306358](https://doi.org/10.1086/306358).

- [24] J. M. Cordes and R. M. Shannon. “A Measurement Model for Precision Pulsar Timing”. In: *arXiv e-prints*, arXiv:1010.3785 (Oct. 2010), arXiv:1010.3785. DOI: [10.48550/arXiv.1010.3785](https://doi.org/10.48550/arXiv.1010.3785). arXiv: [1010.3785](https://arxiv.org/abs/1010.3785) [[astro-ph.IM](#)].
- [25] J. M. Cordes, R. M. Shannon, and D. R. Stinebring. “FREQUENCY-DEPENDENT DISPERSION MEASURES AND IMPLICATIONS FOR PULSAR TIMING”. In: *The Astrophysical Journal* 817.1 (2016), p. 16. DOI: [10.3847/0004-637X/817/1/16](https://doi.org/10.3847/0004-637X/817/1/16). URL: <https://dx.doi.org/10.3847/0004-637X/817/1/16>.
- [26] M. Daszuta, W. Lewandowski, and J. Kijak. “Scintillation observations of PSR B0823+26”. In: *MNRAS* 436.3 (Dec. 2013), pp. 2492–2499. DOI: [10.1093/mnras/stt1750](https://doi.org/10.1093/mnras/stt1750). arXiv: [1310.8076](https://arxiv.org/abs/1310.8076) [[astro-ph.HE](#)].
- [27] A. T. Deller et al. “Microarcsecond VLBI Pulsar Astrometry with PSR π II. Parallax Distances for 57 Pulsars”. In: *ApJ* 875.2, 100 (Apr. 2019), p. 100. DOI: [10.3847/1538-4357/ab11c7](https://doi.org/10.3847/1538-4357/ab11c7). arXiv: [1808.09046](https://arxiv.org/abs/1808.09046) [[astro-ph.IM](#)].
- [28] P. B. Demorest et al. “Limits on the Stochastic Gravitational Wave Background from the North American Nanohertz Observatory for Gravitational Waves”. In: *ApJ* 762.2, 94 (Jan. 2013), p. 94. DOI: [10.1088/0004-637X/762/2/94](https://doi.org/10.1088/0004-637X/762/2/94). arXiv: [1201.6641](https://arxiv.org/abs/1201.6641) [[astro-ph.CO](#)].
- [29] G. Desvignes et al. “High-precision timing of 42 millisecond pulsars with the European Pulsar Timing Array”. In: *Monthly Notices of the Royal Astronomical Society* 458.3 (Mar. 2016), pp. 3341–3380. ISSN: 0035-8711. DOI: [10.1093/mnras/stw483](https://doi.org/10.1093/mnras/stw483). eprint: <https://academic.oup.com/mnras/article-pdf/458/3/3341/8006796/stw483.pdf>. URL: <https://doi.org/10.1093/mnras/stw483>.
- [30] J. Y. Donner et al. “Dispersion measure variability for 36 millisecond pulsars at 150 MHz with LOFAR”. In: *A&A* 644, A153 (Dec. 2020), A153. DOI: [10.1051/0004-6361/202039517](https://doi.org/10.1051/0004-6361/202039517). arXiv: [2011.13742](https://arxiv.org/abs/2011.13742) [[astro-ph.HE](#)].
- [31] J. Y. Donner et al. “First detection of frequency-dependent, time-variable dispersion measures”. In: *A&A* 624, A22 (Apr. 2019), A22. DOI: [10.1051/0004-6361/201834059](https://doi.org/10.1051/0004-6361/201834059). arXiv: [1902.03814](https://arxiv.org/abs/1902.03814) [[astro-ph.GA](#)].
- [32] J. Y. Donner. “Studies of the Ionised Interstellar Medium at Low Radio Frequencies using LOFAR”. PhD thesis. Universität Bielefeld, Bielefeld (Germany), 2022. URL: https://pub.uni-bielefeld.de/download/2963304/2963305/Dissertation_Julian_Donner_2022.pdf.
- [33] B. T. Draine. *Physics of the Interstellar and Intergalactic Medium*. 2011.
- [34] J. Ellis and R. Van Haasteren. *jellis18/PTMCMCSampler: Official Release*. Version 1.0.0. Oct. 2017. DOI: [10.5281/zenodo.1037579](https://doi.org/10.5281/zenodo.1037579).
- [35] J. A. Ellis et al. *ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite*. Version v3.0.0. Sept. 2020. DOI: [10.5281/zenodo.4059815](https://doi.org/10.5281/zenodo.4059815).
- [36] EPTA Collaboration et al. “The second data release from the European Pulsar Timing Array. I. The dataset and timing analysis”. In: *A&A* 678, A48 (Oct. 2023), A48. DOI: [10.1051/0004-6361/202346841](https://doi.org/10.1051/0004-6361/202346841). arXiv: [2306.16224](https://arxiv.org/abs/2306.16224) [[astro-ph.HE](#)].

- [37] EPTA Collaboration et al. “The second data release from the European Pulsar Timing Array. II. Customised pulsar noise models for spatially correlated gravitational waves”. In: *A&A* 678, A49 (Oct. 2023), A49. DOI: [10.1051/0004-6361/202346842](https://doi.org/10.1051/0004-6361/202346842). arXiv: [2306.16225](https://arxiv.org/abs/2306.16225) [[astro-ph.HE](#)].
- [38] EPTA Collaboration et al. “The second data release from the European Pulsar Timing Array. III. Search for gravitational wave signals”. In: *A&A* 678, A50 (Oct. 2023), A50. DOI: [10.1051/0004-6361/202346844](https://doi.org/10.1051/0004-6361/202346844). arXiv: [2306.16214](https://arxiv.org/abs/2306.16214) [[astro-ph.HE](#)].
- [39] I. Ferranti et al. “Impact of the observation frequency coverage on the significance of a gravitational wave background detection in pulsar timing array data”. In: *A&A* 694, A38 (Feb. 2025), A38. DOI: [10.1051/0004-6361/202452805](https://doi.org/10.1051/0004-6361/202452805). arXiv: [2410.18179](https://arxiv.org/abs/2410.18179) [[astro-ph.HE](#)].
- [40] R. S. Foster and D. C. Backer. “Constructing a Pulsar Timing Array”. In: *ApJ* 361 (Sept. 1990), p. 300. DOI: [10.1086/169195](https://doi.org/10.1086/169195).
- [41] B. Goncharov et al. “Identifying and mitigating noise sources in precision pulsar timing data sets”. In: *Monthly Notices of the Royal Astronomical Society* 502.1 (Nov. 2020), pp. 478–493. ISSN: 0035-8711. DOI: [10.1093/mnras/staa3411](https://doi.org/10.1093/mnras/staa3411). eprint: <https://academic.oup.com/mnras/article-pdf/502/1/478/36143958/staa3411.pdf>. URL: <https://doi.org/10.1093/mnras/staa3411>.
- [42] D. Good and International Pulsar Timing Array Team. “Creating a New International Pulsar Timing Array Dataset”. In: *American Astronomical Society Meeting Abstracts*. Vol. 241. American Astronomical Society Meeting Abstracts. Jan. 2023, 438.02, p. 438.02.
- [43] C. S. Gum. “A Survey of Southern HII Regions”. In: *MNRAS* 67 (Jan. 1955), p. 155.
- [44] J. S. Hazboun. *La Forge*. Version 0.3.0. Oct. 2020. DOI: [10.5281/zenodo.4152550](https://doi.org/10.5281/zenodo.4152550). URL: <https://doi.org/10.5281/zenodo.4152550>.
- [45] D. J. Helfand, R. N. Manchester, and J. H. Taylor. “Observations of pulsar radio emission. III. Stability of integrated profiles.” In: *ApJ* 198 (June 1975), pp. 661–670. DOI: [10.1086/153644](https://doi.org/10.1086/153644).
- [46] R. W. Hellings and G. S. Downs. “Upper limits on the isotropic gravitational radiation background from pulsar timing analysis.” In: *ApJL* 265 (Feb. 1983), pp. L39–L42. DOI: [10.1086/183954](https://doi.org/10.1086/183954).
- [47] A. Hewish et al. “Observation of a Rapidly Pulsating Radio Source”. In: *Nature* 217.5130 (Feb. 1968), pp. 709–713. DOI: [10.1038/217709a0](https://doi.org/10.1038/217709a0).
- [48] G. B. Hobbs, R. T. Edwards, and R. N. Manchester. “TEMPO2, a new pulsar-timing package - I. An overview”. In: *MNRAS* 369.2 (June 2006), pp. 655–672. DOI: [10.1111/j.1365-2966.2006.10302.x](https://doi.org/10.1111/j.1365-2966.2006.10302.x). arXiv: [astro-ph/0603381](https://arxiv.org/abs/astro-ph/0603381) [[astro-ph](#)].
- [49] A. W. Hotan, W. van Straten, and R. N. Manchester. “PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis”. In: *PASA* 21.3 (Jan. 2004), pp. 302–309. DOI: [10.1071/AS04022](https://doi.org/10.1071/AS04022). arXiv: [astro-ph/0404549](https://arxiv.org/abs/astro-ph/0404549) [[astro-ph](#)].

- [50] L. G. Hou and J. L. Han. “The observed spiral structure of the Milky Way”. In: *A&A* 569, A125 (Sept. 2014), A125. DOI: [10.1051/0004-6361/201424039](https://doi.org/10.1051/0004-6361/201424039). arXiv: [1407.7331](https://arxiv.org/abs/1407.7331) [[astro-ph.GA](#)].
- [51] F. Iraci et al. “Combining the second data release of the European Pulsar Timing Array with low-frequency pulsar data”. In: *A&A* 704, A109 (Dec. 2025), A109. DOI: [10.1051/0004-6361/202555516](https://doi.org/10.1051/0004-6361/202555516). arXiv: [2510.04639](https://arxiv.org/abs/2510.04639) [[astro-ph.HE](#)].
- [52] F. Iraci et al. “Pulsar timing methods for evaluating dispersion measure time series”. In: *A&A* 692, A170 (Dec. 2024), A170. DOI: [10.1051/0004-6361/202450740](https://doi.org/10.1051/0004-6361/202450740). arXiv: [2410.22170](https://arxiv.org/abs/2410.22170) [[astro-ph.HE](#)].
- [53] S. R. Jammalamadaka and A. Sengupta. *Topics in circular statistics*. Vol. 5. world scientific, 2001.
- [54] G. Janssen et al. “Gravitational Wave Astronomy with the SKA”. In: *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*. Apr. 2015, 37, p. 37. DOI: [10.22323/1.215.0037](https://doi.org/10.22323/1.215.0037). arXiv: [1501.00127](https://arxiv.org/abs/1501.00127) [[astro-ph.IM](#)].
- [55] A. I. Karakas and J. C. Lattanzio. “The Dawes review 2: nucleosynthesis and stellar yields of low-and intermediate-mass single stars”. In: *Publications of the Astronomical Society of Australia* 31 (2014), e030.
- [56] M. J. Keith et al. “Measurement and correction of variations in interstellar dispersion in high-precision pulsar timing”. In: *MNRAS* 429.3 (Mar. 2013), pp. 2161–2174. DOI: [10.1093/mnras/sts486](https://doi.org/10.1093/mnras/sts486). arXiv: [1211.5887](https://arxiv.org/abs/1211.5887) [[astro-ph.GA](#)].
- [57] R. Kippenhahn, A. Weigert, and A. Weiss. *Stellar structure and evolution*. Vol. 192. Springer, 1990.
- [58] A. Kolmogorov. “The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds’ Numbers”. In: *Akademiia Nauk SSSR Doklady* 30 (Jan. 1941), pp. 301–305.
- [59] V. I. Kondratiev et al. “A LOFAR census of millisecond pulsars”. In: *A&A* 585, A128 (Jan. 2016), A128. DOI: [10.1051/0004-6361/201527178](https://doi.org/10.1051/0004-6361/201527178). arXiv: [1508.02948](https://arxiv.org/abs/1508.02948) [[astro-ph.HE](#)].
- [60] M. Kramer et al. “Strong-Field Gravity Tests with the Double Pulsar”. In: *Phys. Rev. X* 11 (4 2021), p. 041050. DOI: [10.1103/PhysRevX.11.041050](https://doi.org/10.1103/PhysRevX.11.041050). URL: <https://link.aps.org/doi/10.1103/PhysRevX.11.041050>.
- [61] M. Kramer. “Pulsars as probes of gravity and fundamental physics”. In: *International Journal of Modern Physics D* 25.14, 1630029-61 (July 2016), pp. 1630029–61. DOI: [10.1142/S0218271816300299](https://doi.org/10.1142/S0218271816300299). arXiv: [1606.03843](https://arxiv.org/abs/1606.03843) [[astro-ph.HE](#)].
- [62] S. R. Kulkarni. “Dispersion measure: Confusion, Constants & Clarity”. In: *arXiv e-prints*, arXiv:2007.02886 (July 2020), arXiv:2007.02886. DOI: [10.48550/arXiv.2007.02886](https://doi.org/10.48550/arXiv.2007.02886). arXiv: [2007.02886](https://arxiv.org/abs/2007.02886) [[astro-ph.HE](#)].
- [63] M. T. Lam et al. “A Second Chromatic Timing Event of Interstellar Origin toward PSR J1713+0747”. In: *ApJ* 861.2, 132 (July 2018), p. 132. DOI: [10.3847/1538-4357/aac770](https://doi.org/10.3847/1538-4357/aac770). arXiv: [1712.03651](https://arxiv.org/abs/1712.03651) [[astro-ph.HE](#)].

- [64] M. T. Lam et al. “Systematic and Stochastic Variations in Pulsar Dispersion Measures”. In: *ApJ* 821.1, 66 (Apr. 2016), p. 66. DOI: [10.3847/0004-637X/821/1/66](#). arXiv: [1512.02203 \[astro-ph.HE\]](#).
- [65] M. T. Lam et al. “The NANOGrav Nine-year Data Set: Excess Noise in Millisecond Pulsar Arrival Times”. In: *ApJ* 834.1, 35 (Jan. 2017), p. 35. DOI: [10.3847/1538-4357/834/1/35](#). arXiv: [1610.01731 \[astro-ph.HE\]](#).
- [66] P. Lazarus et al. “Prospects for high-precision pulsar timing with the new Effelsberg PSRIX backend”. In: *MNRAS* 458.1 (May 2016), pp. 868–880. DOI: [10.1093/mnras/stw189](#). arXiv: [1601.06194 \[astro-ph.IM\]](#).
- [67] K. J. Lee. “Prospects of Gravitational Wave Detection Using Pulsar Timing Array for Chinese Future Telescopes”. In: *Frontiers in Radio Astronomy and FAST Early Sciences Symposium 2015*. Ed. by L. Qain and D. Li. Vol. 502. Astronomical Society of the Pacific Conference Series. Feb. 2016, p. 19.
- [68] L. Lentati et al. “TEMPONEST: a Bayesian approach to pulsar timing analysis”. In: *MNRAS* 437.3 (Jan. 2014), pp. 3004–3023. DOI: [10.1093/mnras/stt2122](#). arXiv: [1310.2120 \[astro-ph.IM\]](#).
- [69] L. Lentati et al. “Hyper-efficient model-independent Bayesian method for the analysis of pulsar timing data”. In: *Phys. Rev. D* 87.10, 104021 (May 2013), p. 104021. DOI: [10.1103/PhysRevD.87.104021](#). arXiv: [1210.3578 \[astro-ph.IM\]](#).
- [70] K. Liu et al. “Profile-shape stability and phase-jitter analyses of millisecond pulsars”. In: *MNRAS* 420.1 (Feb. 2012), pp. 361–368. DOI: [10.1111/j.1365-2966.2011.20041.x](#). arXiv: [1110.4759 \[astro-ph.HE\]](#).
- [71] K. Liu et al. “Single-pulse and profile-variability study of PSR J1022+1001”. In: *Monthly Notices of the Royal Astronomical Society* 449.1 (Mar. 2015), pp. 1158–1169. ISSN: 0035-8711. DOI: [10.1093/mnras/stv397](#). eprint: [https://academic.oup.com/mnras/article-pdf/449/1/1158/4155665/stv397.pdf](#). URL: [https://doi.org/10.1093/mnras/stv397](#).
- [72] K. Liu et al. “The impact on astrometry by solar-wind effect in pulsar timing”. In: *MNRAS* 536.3 (Jan. 2025), pp. 2603–2617. DOI: [10.1093/mnras/stae2727](#). arXiv: [2412.07103 \[astro-ph.HE\]](#).
- [73] D. R. Lorimer and M. Kramer. *Handbook of pulsar astronomy*. Vol. 4. Cambridge university press, 2005.
- [74] J. Luo et al. “PINT: A Modern Software Package for Pulsar Timing”. In: *ApJ* 911.1, 45 (Apr. 2021), p. 45. DOI: [10.3847/1538-4357/abe62f](#). arXiv: [2012.00074 \[astro-ph.IM\]](#).
- [75] D. R. Madison et al. “The NANOGrav 11 yr Data Set: Solar Wind Sounding through Pulsar Timing”. In: *ApJ* 872.2, 150 (Feb. 2019), p. 150. DOI: [10.3847/1538-4357/ab01fd](#). arXiv: [1808.07078 \[astro-ph.SR\]](#).
- [76] R. N. Manchester et al. “The Australia Telescope National Facility Pulsar Catalogue”. In: *The Astronomical Journal* 129.4 (Apr. 2005), pp. 1993–2006. DOI: [10.1086/428488](#). arXiv: [astro-ph/0412641 \[astro-ph\]](#).

- [77] D. B. Melrose. “The models for radio emission from pulsars—The outstanding issues”. In: *Journal of Astrophysics and Astronomy* 16 (June 1995), p. 137. DOI: [10.1007/BF02714830](https://doi.org/10.1007/BF02714830).
- [78] M. T. Miles et al. “The MeerKAT Pulsar Timing Array: the 4.5-yr data release and the noise and stochastic signals of the millisecond pulsar population”. In: *MNRAS* 536.2 (Jan. 2025), pp. 1467–1488. DOI: [10.1093/mnras/stae2572](https://doi.org/10.1093/mnras/stae2572). arXiv: [2412.01148](https://arxiv.org/abs/2412.01148) [[astro-ph.HE](#)].
- [79] M. T. Miles et al. “The MeerKAT Pulsar Timing Array: the first search for gravitational waves with the MeerKAT radio telescope”. In: *MNRAS* 536.2 (Jan. 2025), pp. 1489–1500. DOI: [10.1093/mnras/stae2571](https://doi.org/10.1093/mnras/stae2571). arXiv: [2412.01153](https://arxiv.org/abs/2412.01153) [[astro-ph.HE](#)].
- [80] S. K. Ocker et al. “An In Situ Study of Turbulence near Stellar Bow Shocks”. In: *ApJ* 922.2, 233 (Dec. 2021), p. 233. DOI: [10.3847/1538-4357/ac2b28](https://doi.org/10.3847/1538-4357/ac2b28). arXiv: [2107.10371](https://arxiv.org/abs/2107.10371) [[astro-ph.GA](#)].
- [81] S. K. Ocker et al. “Implications for Galactic Electron Density Structure from Pulsar Sightlines Intersecting H II Regions”. In: *ApJ* 974.1, 10 (Oct. 2024), p. 10. DOI: [10.3847/1538-4357/ad6a51](https://doi.org/10.3847/1538-4357/ad6a51). arXiv: [2406.07664](https://arxiv.org/abs/2406.07664) [[astro-ph.GA](#)].
- [82] P. V. Padmanabh et al. “Revisiting profile instability of PSR J1022+1001”. In: *Monthly Notices of the Royal Astronomical Society* 500.1 (Oct. 2020), pp. 1178–1187. ISSN: 0035-8711. DOI: [10.1093/mnras/staa3174](https://doi.org/10.1093/mnras/staa3174). eprint: <https://academic.oup.com/mnras/article-pdf/500/1/1178/34370222/staa3174.pdf>. URL: <https://doi.org/10.1093/mnras/staa3174>.
- [83] B. B. P. Perera et al. “The International Pulsar Timing Array: second data release”. In: *MNRAS* 490.4 (Dec. 2019), pp. 4666–4687. DOI: [10.1093/mnras/stz2857](https://doi.org/10.1093/mnras/stz2857). arXiv: [1909.04534](https://arxiv.org/abs/1909.04534) [[astro-ph.HE](#)].
- [84] A. Philippov, A. Timokhin, and A. Spitkovsky. “Origin of Pulsar Radio Emission”. In: *Phys. Rev. Lett.* 124 (24 2020), p. 245101. DOI: [10.1103/PhysRevLett.124.245101](https://doi.org/10.1103/PhysRevLett.124.245101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.245101>.
- [85] P. Rana et al. “The Indian Pulsar Timing Array data release 2: I. Dataset and timing analysis”. In: *PASA* 42, e108 (July 2025), e108. DOI: [10.1017/pasa.2025.10066](https://doi.org/10.1017/pasa.2025.10066). arXiv: [2506.16769](https://arxiv.org/abs/2506.16769) [[astro-ph.IM](#)].
- [86] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [87] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. 2006.
- [88] M. Raveri and C. Doux. “Non-Gaussian estimates of tensions in cosmological parameters”. In: *Physical Review D* 104.4 (2021), p. 043504.
- [89] D. J. Reardon et al. “Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array”. In: *ApJL* 951.1, L6 (July 2023), p. L6. DOI: [10.3847/2041-8213/acdd02](https://doi.org/10.3847/2041-8213/acdd02). arXiv: [2306.16215](https://arxiv.org/abs/2306.16215) [[astro-ph.HE](#)].

- [90] B. J. Rickett. “Interstellar scattering and scintillation of radio waves.” In: *Annual Review of Astronomy and Astrophysics* 15 (Jan. 1977), pp. 479–504. DOI: [10.1146/annurev.aa.15.090177.002403](https://doi.org/10.1146/annurev.aa.15.090177.002403).
- [91] B. J. Rickett. “Radio propagation through the turbulent interstellar plasma.” In: *Annual Review of Astronomy and Astrophysics* 28 (Jan. 1990), pp. 561–605. DOI: [10.1146/annurev.aa.28.090190.003021](https://doi.org/10.1146/annurev.aa.28.090190.003021).
- [92] B. J. Rickett. “Refractive Interstellar Scintillations of Radio Sources”. In: *ApJ* 307 (Aug. 1986), p. 564. DOI: [10.1086/164444](https://doi.org/10.1086/164444).
- [93] R. W. Romani, R. Narayan, and R. Blandford. “Refractive effects in pulsar scintillation.” In: *MNRAS* 220 (May 1986), pp. 19–49. DOI: [10.1093/mnras/220.1.19](https://doi.org/10.1093/mnras/220.1.19).
- [94] R. W. Romani. “Timing a Millisecond Pulsar Array”. In: *Timing Neutron Stars*. Ed. by H. Ögelman and E. P. J. van den Heuvel. Dordrecht: Springer Netherlands, 1989, pp. 113–117. ISBN: 978-94-009-2273-0. DOI: [10.1007/978-94-009-2273-0_6](https://doi.org/10.1007/978-94-009-2273-0_6). URL: https://doi.org/10.1007/978-94-009-2273-0_6.
- [95] A. Sesana. “Systematic investigation of the expected gravitational wave signal from supermassive black hole binaries in the pulsar timing band.” In: *MNRAS* 433 (June 2013), pp. L1–L5. DOI: [10.1093/mnrasl/slt034](https://doi.org/10.1093/mnrasl/slt034). arXiv: [1211.5375](https://arxiv.org/abs/1211.5375) [astro-ph.CO].
- [96] R. M. Shannon and J. M. Cordes. “Modelling and mitigating refractive propagation effects in precision pulsar timing observations”. In: *MNRAS* 464.2 (Oct. 2016), pp. 2075–2089. ISSN: 0035-8711. DOI: [10.1093/mnras/stw2449](https://doi.org/10.1093/mnras/stw2449). eprint: <https://academic.oup.com/mnras/article-pdf/464/2/2075/18517118/stw2449.pdf>. URL: <https://doi.org/10.1093/mnras/stw2449>.
- [97] R. M. Shannon and J. M. Cordes. “Pulse Intensity Modulation and the Timing Stability of Millisecond Pulsars: A Case Study of PSR J1713+0747”. In: *ApJ* 761.1, 64 (Dec. 2012), p. 64. DOI: [10.1088/0004-637X/761/1/64](https://doi.org/10.1088/0004-637X/761/1/64). arXiv: [1210.7021](https://arxiv.org/abs/1210.7021) [astro-ph.SR].
- [98] S. Sharpless. “A Catalogue of H II Regions.” In: *ApJS* 4 (Dec. 1959), p. 257. DOI: [10.1086/190049](https://doi.org/10.1086/190049).
- [99] X. Siemens et al. “The stochastic background: scaling laws and time to detection for pulsar timing arrays”. In: *Classical and Quantum Gravity* 30.22, 224015 (Nov. 2013), p. 224015. DOI: [10.1088/0264-9381/30/22/224015](https://doi.org/10.1088/0264-9381/30/22/224015). arXiv: [1305.3196](https://arxiv.org/abs/1305.3196) [astro-ph.IM].
- [100] L. Speri et al. “Quality over quantity: Optimizing pulsar timing array analysis for stochastic and continuous gravitational wave signals”. In: *Monthly Notices of the Royal Astronomical Society* 518.2 (Nov. 2022), pp. 1802–1817. ISSN: 0035-8711. DOI: [10.1093/mnras/stac3237](https://doi.org/10.1093/mnras/stac3237). eprint: <https://academic.oup.com/mnras/article-pdf/518/2/1802/47224477/stac3237.pdf>. URL: <https://doi.org/10.1093/mnras/stac3237>.
- [101] B. W. Stappers et al. “Observing pulsars and fast transients with LOFAR”. In: *A&A* 530, A80 (June 2011), A80. DOI: [10.1051/0004-6361/201116681](https://doi.org/10.1051/0004-6361/201116681). arXiv: [1104.1577](https://arxiv.org/abs/1104.1577) [astro-ph.IM].

- [102] B. Strömgren. “The Physical State of Interstellar Hydrogen.” In: *ApJ* 89 (May 1939), p. 526. DOI: [10.1086/144074](https://doi.org/10.1086/144074).
- [103] S. C. Susarla et al. “Exploring the time variability of the solar wind using LOFAR pulsar data”. In: *A&A* 692, A18 (Dec. 2024), A18. DOI: [10.1051/0004-6361/202450680](https://doi.org/10.1051/0004-6361/202450680). arXiv: [2409.09838](https://arxiv.org/abs/2409.09838) [[astro-ph.SR](#)].
- [104] T. M. Tauris and E. P. J. van den Heuvel. “Formation and evolution of compact stellar X-ray sources”. In: *Compact stellar X-ray sources*. Ed. by W. H. G. Lewin and M. van der Klis. Vol. 39. 2006, pp. 623–665. DOI: [10.48550/arXiv.astro-ph/0303456](https://doi.org/10.48550/arXiv.astro-ph/0303456).
- [105] T. M. Tauris and E. P. J. van den Heuvel. *Physics of Binary Star Evolution. From Stars to X-ray Binaries and Gravitational Wave Sources*. 2023. DOI: [10.48550/arXiv.2305.09388](https://doi.org/10.48550/arXiv.2305.09388).
- [106] J. H. Taylor. “Pulsar Timing and Relativistic Gravity”. In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660 (Oct. 1992), pp. 117–134. DOI: [10.1098/rsta.1992.0088](https://doi.org/10.1098/rsta.1992.0088).
- [107] S. R. Taylor et al. *enterprise_extensions*. v2.4.3. 2021. URL: https://github.com/nanograv/enterprise_extensions.
- [108] C. Tiburzi et al. “A study of spatial correlations in pulsar timing array data”. In: *MNRAS* 455.4 (Feb. 2016), pp. 4339–4350. DOI: [10.1093/mnras/stv2143](https://doi.org/10.1093/mnras/stv2143). arXiv: [1510.02363](https://arxiv.org/abs/1510.02363) [[astro-ph.IM](#)].
- [109] C. Tiburzi et al. “On the usefulness of existing solar wind models for pulsar timing corrections”. In: *MNRAS* 487.1 (July 2019), pp. 394–408. DOI: [10.1093/mnras/stz1278](https://doi.org/10.1093/mnras/stz1278). arXiv: [1905.02989](https://arxiv.org/abs/1905.02989) [[astro-ph.HE](#)].
- [110] C. Tiburzi. “Pulsars Probe the Low-Frequency Gravitational Sky: Pulsar Timing Arrays Basics and Recent Results”. In: *PASA* 35, e013 (Mar. 2018), e013. DOI: [10.1017/pasa.2018.7](https://doi.org/10.1017/pasa.2018.7). arXiv: [1802.05076](https://arxiv.org/abs/1802.05076) [[astro-ph.IM](#)].
- [111] M. Vallisneri. *libstempo: Python wrapper for Tempo2*. Astrophysics Source Code Library, record ascl:2002.017. Feb. 2020.
- [112] M. P. van Haarlem et al. “LOFAR: The LOW-Frequency ARray”. In: *A&A* 556, A2 (Aug. 2013), A2. DOI: [10.1051/0004-6361/201220873](https://doi.org/10.1051/0004-6361/201220873). arXiv: [1305.3550](https://arxiv.org/abs/1305.3550) [[astro-ph.IM](#)].
- [113] R. van Haasteren and M. Vallisneri. “New advances in the Gaussian-process approach to pulsar-timing data analysis”. In: *Phys. Rev. D* 90.10, 104012 (Nov. 2014), p. 104012. DOI: [10.1103/PhysRevD.90.104012](https://doi.org/10.1103/PhysRevD.90.104012). arXiv: [1407.1838](https://arxiv.org/abs/1407.1838) [[gr-qc](#)].
- [114] W. van Straten and M. Bailes. “DSPSR: Digital Signal Processing Software for Pulsar Astronomy”. In: *PASA* 28.1 (Jan. 2011), pp. 1–14. DOI: [10.1071/AS10021](https://doi.org/10.1071/AS10021). arXiv: [1008.3973](https://arxiv.org/abs/1008.3973) [[astro-ph.IM](#)].
- [115] J. P. W. Verbiest et al. “Precision Timing of PSR J0437-4715: An Accurate Pulsar Distance, a High Pulsar Mass, and a Limit on the Variation of Newton’s Gravitational Constant”. In: *ApJ* 679.1 (May 2008), pp. 675–680. DOI: [10.1086/529576](https://doi.org/10.1086/529576). arXiv: [0801.2589](https://arxiv.org/abs/0801.2589) [[astro-ph](#)].
- [116] J. P. W. Verbiest et al. “The International Pulsar Timing Array: First data release”. In: *MNRAS* 458.2 (May 2016), pp. 1267–1288. DOI: [10.1093/mnras/stw347](https://doi.org/10.1093/mnras/stw347). arXiv: [1602.03640](https://arxiv.org/abs/1602.03640) [[astro-ph.IM](#)].

- [117] J. P. W. Verbiest and G. M. Shaifullah. “Measurement uncertainty in pulsar timing array experiments”. In: *Classical and Quantum Gravity* 35.13, 133001 (July 2018), p. 133001. DOI: [10.1088/1361-6382/aac412](https://doi.org/10.1088/1361-6382/aac412).
- [118] J. P. W. Verbiest et al. “Status report on global pulsar-timing-array efforts to detect gravitational waves”. In: *Results in Physics* 61, 107719 (June 2024), p. 107719. DOI: [10.1016/j.rinp.2024.107719](https://doi.org/10.1016/j.rinp.2024.107719). arXiv: [2404.19529](https://arxiv.org/abs/2404.19529) [[astro-ph.HE](#)].
- [119] Z. Wu et al. “Pulsar scintillation studies with LOFAR. I. The census”. In: *A&A* 663, A116 (July 2022), A116. DOI: [10.1051/0004-6361/202142980](https://doi.org/10.1051/0004-6361/202142980). arXiv: [2203.10409](https://arxiv.org/abs/2203.10409) [[astro-ph.HE](#)].
- [120] H. Xu et al. “A sustained pulse shape change in PSR J1713+0747 possibly associated with timing and DM events”. In: *The Astronomer’s Telegram* 14642 (May 2021), p. 1.
- [121] H. Xu et al. “Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I”. In: *Research in Astronomy and Astrophysics* 23.7, 075024 (July 2023), p. 075024. DOI: [10.1088/1674-4527/acdfa5](https://doi.org/10.1088/1674-4527/acdfa5). arXiv: [2306.16216](https://arxiv.org/abs/2306.16216) [[astro-ph.HE](#)].
- [122] J. M. Yao, R. N. Manchester, and N. Wang. “A New Electron-density Model for Estimation of Pulsar and FRB Distances”. In: *ApJ* 835.1, 29 (Jan. 2017), p. 29. DOI: [10.3847/1538-4357/835/1/29](https://doi.org/10.3847/1538-4357/835/1/29). arXiv: [1610.09448](https://arxiv.org/abs/1610.09448) [[astro-ph.GA](#)].
- [123] P. Zarka et al. “LSS/NenuFAR: The LOFAR Super Station project in NanÅşay”. In: Dec. 2012.
- [124] A. Zic et al. “The Parkes Pulsar Timing Array third data release”. In: *PASA* 40, e049 (Dec. 2023), e049. DOI: [10.1017/pasa.2023.36](https://doi.org/10.1017/pasa.2023.36). arXiv: [2306.16230](https://arxiv.org/abs/2306.16230) [[astro-ph.HE](#)].