



**UNICA**

UNIVERSITÀ  
DEGLI STUDI  
DI CAGLIARI

**Ph.D. DEGREE IN  
PHYSICS**

Cycle XXXVIII

**TITLE OF THE Ph.D. THESIS**

**Search and characterization of pulsars in southern globular  
clusters: new methodologies and results from an S-band survey**

Scientific Disciplinary Sector(s)

PHYSICS(FIS/05)

Ph.D. Student:	Rouhin Nag
Supervisor	Dr.Andrea Possenti
Co-Supervisor	Dr.Marta Burgay

Final exam. Academic Year 2024/2025

Thesis defence session: July 2026

# Acknowledgements

First and foremost, I would like to express my sincere gratitude to the University of Cagliari and the Osservatorio Astronomico di Cagliari (INAF–OAC) for giving me this opportunity and for providing an excellent environment to carry out my research.

I am deeply grateful to my supervisors, Dr. Andrea Possenti and Dr. Marta Burgay, for their constant guidance, support, and encouragement throughout this journey. Their mentorship has been invaluable. I would also like to thank the OAC Pulsar Group for welcoming me so warmly and for their continuous help, discussions, and support over the years.

I would like to extend my sincere thanks to Dr. Michael Kramer, Dr. Vivek Venkatraman Krishnan, Dr. Vishnu Balakrishnan, and Dr. Paulo Freire for giving me the opportunity to visit and work at the Max Planck Institute for Radio Astronomy in Bonn. I am also grateful to all members of the COMPACT group and the Pulsar Group for welcoming me with such enthusiasm and for providing a stimulating and supportive research environment.

I owe everything to my parents, whose sacrifices, unwavering belief in me, and unconditional love have made this journey possible. Their support has been the foundation on which all of this stands.

A very special and heartfelt acknowledgement goes to my late maternal grandfather, Nirmal Dasgupta. He was the one who first introduced me to science and opened my eyes to the wonders of the world. I hope I have made you proud.

A very special and heartfelt thanks goes to Dr. Moulika Hazra, for being my pathfinder, my guide, and my rock. Thank you for standing by me through every high and low and for always being my voice of clarity. This would not have been possible without you. I would also like to thank her parents, Mr. S. S. Hazra and Mrs. Jaba Hazra, for their warmth, love, and support over the years.

A special thanks to Dr. Alessandro Ridolfi, Dr. Federico Abbate, and Dr. Matteo Trudu for being the best office mates one could ask for. Whether it was helping me navigate bureaucracy, giving me rides to the observatory, or patiently entertaining my many (sometimes ridiculous) ideas—scientific or otherwise—you made this journey significantly easier and far more enjoyable.

I would also like to thank Dr. Scott Ransom and Dr. Zichen Pan for their constant willingness to help and for always entertaining my curiosity.

I would also like to express my gratitude to Dr. Debabrata Mitra and Mr. Shubroto Banerjee, my school science teachers, for making science feel exciting and inspiring enough to set me on this path. I also thank Fr. Peter James SJ. and Fr. Vincent

Hansda SJ for their kind words and support whenever it was needed.

I would also like to thank Dr. Swarat Chaudhuri and Dr. Sumit Kumar Roy for their prompt help and support. Without your help, this project would have failed before it lifted off and I am forever grateful.

I would also like to thank everyone who has ever given me a lift to the observatory—your help, often at odd hours, has not gone unnoticed and is deeply appreciated.

Finally, I would like to thank all my friends and relatives who have been with me before and throughout this journey. Your support, encouragement, and presence have meant more than I can express.

This thesis is not solely the result of my own efforts, but the culmination of the roles played by all those mentioned above. Each of you has contributed in your own way—some in profound and defining moments, others in smaller but no less meaningful ways—often stepping in at exactly the right time. Without your guidance, support, and presence throughout this journey, this work may not have come to fruition. In that sense, this thesis is as much yours as it is mine.

# Abstract

This thesis explores the search for radio pulsars in globular clusters and the development and improvement of computational methods to enhance their detection in challenging observational conditions. Pulsars are rapidly rotating neutron stars—remnants of massive stars that have undergone supernova explosions—and they represent some of the most extreme objects in the Universe. Studying them allows us to probe physical conditions that cannot be recreated in laboratories on Earth, including ultra-dense matter, intense magnetic fields, and strong gravitational environments.

Beyond their intrinsic interest, pulsars play a broader role in modern astrophysics. Their remarkably regular emission makes them powerful tools for precision measurements, enabling studies of fundamental physics, stellar evolution, and the structure of our Galaxy. They are used to test theories of gravity under extreme conditions, to trace the distribution of matter in the interstellar medium, and to understand how stars live and die. In this sense, pulsars are not only fascinating objects in their own right, but also versatile probes that connect multiple areas of physics and astronomy.

Globular clusters provide a particularly compelling environment in which to study pulsars. These systems are dense, gravitationally bound collections of stars, often containing hundreds of thousands to millions of members packed into a relatively small volume. The high density leads to frequent interactions between stars, which can significantly alter their evolutionary paths. As a result, globular clusters are exceptionally efficient at producing exotic systems, including rapidly spinning millisecond pulsars in binary systems. These objects offer a unique window into how stellar systems evolve under conditions that are fundamentally different from those in the more isolated regions of the Galaxy.

Despite their importance, detecting pulsars in globular clusters is inherently difficult. As radio signals travel through space, they are affected by the interstellar medium, which disperses and distorts them. In addition, many pulsars reside in binary systems, where orbital motion can smear their signals and make them harder to detect using standard search techniques. These challenges become more severe for compact and rapidly evolving systems, requiring both carefully designed observations and more advanced data analysis methods.

This thesis addresses these challenges through a combination of observational work and methodological development. It is structured as follows.

Chapter 1 introduces the basic properties of neutron stars and pulsars, including how they form, how their emission is observed, and why they are important tools in

astrophysics.

Chapter 2 focuses on pulsars in globular clusters, describing how the dense environments of these systems lead to the formation of unusual and exotic pulsar populations. The role of stellar interactions and binary evolution is discussed.

Chapter 3 presents the techniques used to search for pulsars in radio data. It explains how signals are processed and analysed, and outlines the key challenges involved in detecting both isolated pulsars and those in binary systems.

Chapter 4 presents the first globular cluster pulsar survey conducted with the MeerKAT radio telescope at S-band frequencies (approximately 2.4 GHz central frequency). This chapter describes the observing strategy, data processing pipeline, and search methods used in this work, and places the survey in the broader context of pulsar searches across different observing frequencies.

Chapter 5 introduces an improved computational approach to searching for pulsars, with a focus on high-performance computing and GPU-based methods. This includes the development of a time-domain search technique designed to improve sensitivity to systems that are difficult to detect due to complex orbital motion, along with a discussion of the computational strategies used to make such searches feasible.

The thesis concludes with appendices that provide additional technical details, including methods for estimating pulsar flux densities and considerations relevant to the observational setup.

Overall, this work aims to improve our ability to detect and study pulsars in dense stellar environments, thereby contributing to a broader understanding of fundamental physics, stellar evolution, and the behaviour of matter under extreme conditions.

# Contents

<b>List of Acronyms</b>	<b>11</b>
<b>List of Symbols</b>	<b>13</b>
<b>1 Pulsar generalities</b>	<b>25</b>
1.1 Neutron Stars and the Discovery of Pulsars . . . . .	25
1.2 Pulsar Geometry and the Lighthouse Model . . . . .	26
1.3 Birth of Pulsars . . . . .	29
1.3.1 Core-Collapse Supernovae and Neutron Star Formation . . . . .	29
1.3.2 Initial Spin and Magnetic Field Properties . . . . .	30
1.4 Rotation-Powered Pulsars and Spin-Down Evolution . . . . .	30
1.4.1 Rotational Energy as the Power Source . . . . .	31
1.4.2 Spin-Down Laws and Magnetic Braking . . . . .	31
1.4.3 Characteristic Age and Magnetic Field Estimates . . . . .	32
1.5 Accretion-Powered Pulsars and Binary Evolution . . . . .	33
1.5.1 Binary Systems and Mass Transfer . . . . .	34
1.5.2 Accretion-Powered Pulsars . . . . .	34
1.5.3 Spin-Up and the Recycling Scenario . . . . .	35
1.5.4 Binary Pulsars and Evolutionary Diversity . . . . .	35
1.6 Millisecond Pulsars and Binary Motion . . . . .	36
1.6.1 Defining Properties of Millisecond Pulsars . . . . .	36
1.6.2 Binary Motion and Orbital Modulation . . . . .	37
1.6.3 Binary Millisecond Pulsars and Evolutionary Implications . . . . .	37
1.6.4 The $P-\dot{P}$ Diagram and the Pulsar Fauna . . . . .	38
1.7 Radio Emission Properties of Pulsars . . . . .	42
1.7.1 Broadband Radio Emission from Pulsars . . . . .	42
1.7.2 Flux Density as an Observable Quantity . . . . .	43
1.7.3 Spectral Index and Frequency Dependence . . . . .	43
1.7.4 Polarisation Properties of Pulsar Emission . . . . .	44

1.8	Pulsars as Astrophysical and Physical Laboratories . . . . .	44
1.8.1	Probing Fundamental Physics in Strong-Field Regimes . . . . .	45
1.8.2	Magnetospheric Physics and Plasma Processes . . . . .	45
1.8.3	Stellar Evolution and Binary Interactions . . . . .	46
1.8.4	The Interstellar Medium and Galactic Structure . . . . .	46
1.8.5	Cosmological and Multi-Messenger Applications . . . . .	46
<b>2</b>	<b>Pulsar Science with Globular clusters</b>	<b>47</b>
2.1	From Isolated Evolution to Dynamical Environments . . . . .	47
2.2	Structural Parameters and Dynamical Evolution of Globular Clusters	50
2.2.1	From Structure to Dynamics: the basic ingredients . . . . .	50
2.2.2	Crossing time: the “fast” dynamical clock . . . . .	51
2.2.3	Two-body relaxation: why clusters evolve at all . . . . .	51
2.2.4	Mass segregation and dynamical friction: why remnants col- lect in the core . . . . .	54
2.2.5	Close encounters and gravitational focusing . . . . .	54
2.2.6	The encounter rate parameter $\Gamma$ . . . . .	56
2.2.7	Binary processing, core collapse, and binary heating . . . . .	56
2.3	Neutron Stars and the Dynamical Formation of Recycled Pulsars . .	57
2.3.1	Natal kicks and the neutron-star retention problem . . . . .	58
2.3.2	Mass segregation and the central concentration of compact remnants . . . . .	61
2.3.3	Binary processing and the pathways to recycling in clusters . .	61
2.3.4	Globular Clusters as Laboratories for Exotic Pulsar Systems .	62
2.4	Pulsars as Probes of Globular Cluster Gravitational Potentials . . . .	64
2.4.1	Minimal timing concepts: pulse times of arrival and phase- coherent models . . . . .	65
2.4.2	From Doppler shifts to apparent spin derivatives . . . . .	65
2.4.3	Observed spin derivatives: intrinsic and kinematic contributions	66
2.4.4	Pulsars as accelerometers in the cluster potential . . . . .	67
2.4.5	Projected separations, angular offsets, and acceleration en- velopes . . . . .	68
2.4.6	Central mass concentrations and possible black holes . . . . .	70
<b>3</b>	<b>Observational Techniques for Radio Pulsar Searches</b>	<b>72</b>
3.1	Introduction . . . . .	72
3.2	Radio Telescopes and Pulsar Observing Systems . . . . .	74
3.3	Propagation Effects in the Interstellar Medium . . . . .	75

3.3.1	Dispersion . . . . .	76
3.3.2	Faraday Rotation . . . . .	78
3.3.3	Scattering . . . . .	80
3.3.4	Interstellar Scintillation . . . . .	82
3.4	Radio Telescopes and Pulsar Search Data Acquisition . . . . .	84
3.4.1	Radio telescopes for pulsar observations . . . . .	85
3.4.2	Front-end receivers and signal amplification . . . . .	85
3.4.3	Digital back-ends and channelisation . . . . .	86
3.4.4	Search-mode observations . . . . .	86
3.4.5	Pulsar search data formats . . . . .	87
3.4.6	Pulsar search software: the PRESTO toolkit . . . . .	87
3.5	Pulsar pre-search Processing . . . . .	88
3.5.1	Radio-frequency interference and noise in pulsar search data . . . . .	89
3.5.2	Why RFI can mimic pulsars . . . . .	91
3.5.3	RFI mitigation in a PRESTO-based workflow . . . . .	91
3.5.4	Clipping, masking, and DM= 0 filtering . . . . .	92
3.5.5	Birdies and periodic terrestrial signals . . . . .	92
3.5.6	White noise, red noise, and the limits of mitigation . . . . .	94
3.5.7	De-dispersion . . . . .	95
3.5.8	Red-noise mitigation . . . . .	96
3.6	Periodic Pulsar Search Techniques . . . . .	97
3.6.1	Fourier-domain periodicity searches . . . . .	97
3.6.2	Implementation in PRESTO: the <code>realfft</code> stage . . . . .	101
3.6.3	Harmonic structure and harmonic summing . . . . .	102
3.6.4	Binary motion and Doppler modulation . . . . .	104
3.6.5	Line-of-sight acceleration and jerk . . . . .	105
3.6.6	Signal drift in the Fourier domain . . . . .	107
3.6.7	Candidate generation, sifting, and folding . . . . .	114
3.7	Additional pulsar search techniques . . . . .	119
3.7.1	Fast Folding Algorithm (FFA) searches . . . . .	119
3.7.2	Single-pulse searches . . . . .	120
3.7.3	Stacking and phase-modulation techniques . . . . .	120
3.7.4	Template-bank searches for compact binaries . . . . .	121
<b>4</b>	<b>Results from the first MeerKAT S-band pulsar survey</b>	<b>122</b>
4.1	Introduction . . . . .	122
4.2	Target selection and observations . . . . .	124
4.3	Data Analysis . . . . .	126

4.3.1	Data Format for the S-band survey	126
4.3.2	Search operation	127
4.3.3	Sifting and folding	132
4.3.4	Flux density measurements	133
4.4	Results	134
4.4.1	Re-detections	134
4.4.2	New Discoveries and other search results from Glimpse-C01	135
4.4.3	Dispersion measure discrepancy and cluster membership	140
4.4.4	S-band and L-band flux density measurements	143
4.4.5	S-band non-detections and reference spectral-index constraints	144
4.5	Discussion and Conclusions from the survey	149
<b>5</b>	<b>PULSEJET_beta: Jerk Search Pipeline</b>	<b>152</b>
5.1	Going Beyond Acceleration Searches	152
5.2	Time-Domain Resampling Formalism	153
5.3	The PULSEJET_beta Pipeline: Motivation and Design	155
5.3.1	Accessing Extreme Regions of Parameter Space	156
5.3.2	Computational Cost of Jerk Searches	162
5.4	Astrophysical Determination of the Acceleration and Jerk Bounds	165
5.4.1	Keplerian description of the line-of-sight motion	165
5.4.2	Line-of-sight velocity	167
5.4.3	Line-of-sight acceleration	167
5.4.4	Line-of-sight jerk	168
5.4.5	Determining the search bounds	168
5.5	Metric Construction and Stochastic Population of the $(a_{\parallel}, j_{\parallel})$ Template Bank	169
5.5.1	Choice of phase model	170
5.5.2	Mismatch and local metric approximation	171
5.5.3	Derivatives of the acceleration–jerk phase model	172
5.5.4	Projection onto the $(a, j)$ subspace	173
5.5.5	Determinant and local template density	174
5.5.6	Proper volume of the target region	174
5.5.7	From proper volume to template count	175
5.5.8	Sampling the bank from the metric density	175
5.5.9	Saving and interpreting the bank	176
5.5.10	Empirical validation of the bank	177
5.5.11	Relation to exact stochastic banks and the role of the metric approximation	177

5.5.12	Summary of the construction . . . . .	178
5.6	Iterative GPU Resampling over the $(a_{\parallel}, j_{\parallel})$ Template Bank . . . . .	178
5.6.1	Why a fixed-step grid in $a$ and $j$ is inefficient . . . . .	179
5.6.2	From the continuous resampling model to the discrete index map . . . . .	180
5.6.3	Nearest-neighbour resampling in the implemented kernel . . . . .	181
5.6.4	The iterative structure of the pipeline . . . . .	183
5.6.5	Why the jerk term changes the computational character of the resampling . . . . .	183
5.6.6	Why GPUs are the natural implementation target . . . . .	184
5.7	Current Status of the Pipeline and Preliminary Template-Bank Per- formance . . . . .	185
5.8	Injection Tests and Preliminary Comparison with Existing Jerk-Search Implementations . . . . .	187
5.8.1	Validation of the resampler with exact trial parameters . . . . .	187
5.8.2	Illustrative comparison with an acceleration-only search . . . . .	188
5.8.3	Blind recovery tests with small sets of injected pulsars . . . . .	188
5.8.4	Larger injection campaigns . . . . .	190
5.8.5	Template-bank trials versus a naive acceleration-jerk scan . . . . .	191
5.8.6	Preliminary runtime comparison with <code>PrestoZL</code> . . . . .	192
5.8.7	Interpretation of the present results . . . . .	193
<b>6</b>	<b>Conclusions and Future Work</b>	<b>195</b>
6.1	Conclusions . . . . .	195
6.2	Future Work . . . . .	196
6.2.1	Follow-up of the Glimpse-C01 discoveries . . . . .	196
6.2.2	Companion identification and multiwavelength follow-up . . . . .	197
6.2.3	Further development of <code>PULSEJET_beta</code> . . . . .	198
6.3	Data Availability and Software Used . . . . .	199
<b>A</b>	<b>Flux Density Estimation from Folded Pulse Profiles</b>	<b>201</b>
A.1	Flux Density Estimation from Folded Pulse Profiles . . . . .	201
<b>B</b>	<b>Effective SEFD and Gain Determination</b>	<b>203</b>
B.1	Effective SEFD and Gain Determination . . . . .	203

# List of Acronyms

ISM	Interstellar Medium
MSP	Millisecond Pulsar
GC	Globular Cluster
SNR	Supernova Remnant
PWN	Pulsar Wind Nebula
WD	White Dwarf
XDINS	X-ray Dim Isolated Neutron Stars
GR	General Relativity
PTA	Pulsar Timing Array
TOA	Time of Arrival
FAST	Five-hundred-meter Aperture Spherical Telescope
RFI	Radio-Frequency Interference
DM	Dispersion Measure
RM	Rotation Measure
FFT	Fast Fourier Transform
FFA	Fast Folding Algorithm
PSRFITS	Pulsar FITS data format
S/N	Signal-to-noise ratio
ML	Machine Learning
CNN	Convolutional Neural Network
FRB	Fast Radio Burst



# List of Symbols

$M_{\odot}$	Solar mass
$G$	Gravitational constant
km	Kilometre
$P$	Pulsar spin period
$\dot{P}$	First time derivative of the spin period
$\ddot{P}$	Second time derivative of the spin period
$\Omega$	Angular spin frequency of the neutron star
$\dot{\Omega}$	First time derivative of the angular spin frequency
$\ddot{\Omega}$	Second time derivative of the angular spin frequency
$I$	Moment of inertia of the neutron star
$E_{\text{rot}}$	Rotational kinetic energy
$\dot{E}_{\text{rot}}$	Rate of rotational energy loss (spin-down power)
$n$	Braking index
$K$	Braking coefficient in the spin-down law
$\tau_c$	Characteristic age of the pulsar
$\Omega_0$	Angular spin frequency at birth
$f_0$	Intrinsic pulsar spin frequency
$f_{\text{obs}}$	Observed pulsar spin frequency
$\dot{f}$	First time derivative of the spin frequency
$\ddot{f}$	Second time derivative of the spin frequency
$v_{\parallel}$	Line-of-sight component of the pulsar orbital velocity
$a_{\parallel}$	Line-of-sight acceleration of the pulsar
$j_{\parallel}$	Line-of-sight jerk of the pulsar
$c$	Speed of light
$P_b$	Binary orbital period
$T_{\text{obs}}$	Duration of the observation
$\delta\nu$	Fourier frequency resolution ( $1/T_{\text{obs}}$ )
$z$	Fourier bin drift parameter due to acceleration
$\phi$	Pulse phase within one rotation
$\Phi(t)$	Signal phase as a function of time
$I_k$	Intensity in phase bin $k$ of a folded pulse profile
$N_{\text{bin}}$	Number of phase bins used when folding data
$\sigma$	One-dimensional velocity dispersion in a globular cluster
$r_c$	Core radius of a globular cluster

$n_\star$	Stellar number density
$\rho_0$	Central mass density of a globular cluster
$\Gamma$	Encounter rate parameter
$\Sigma$	Interaction cross section
$v_\infty$	Relative velocity at infinity
$m_\star$	Typical stellar mass
$m_1, m_2$	Masses of interacting objects
$a$	Binary semi-major axis
$E_b$	Binary binding energy
$v_{\text{esc}}$	Escape velocity from a globular cluster
$a_{\parallel, \text{GC}}$	Acceleration due to the globular cluster potential
$a_{\text{Gal}}$	Differential Galactic gravitational acceleration
$r$	Three-dimensional distance from the cluster centre
$R_\perp$	Projected distance from the cluster centre on the plane of the sky
$z$	Line-of-sight coordinate relative to the cluster centre
$\theta_\perp$	Angular separation from the cluster centre
DM	Dispersion measure, $\int n_e dl$ along the line of sight
$n_e$	Free electron number density in the interstellar medium
$\Delta t$	Frequency-dependent dispersive delay of the pulsar signal
RM	Rotation measure describing Faraday rotation
$\chi$	Polarisation position angle
$\tau_{\text{sc}}$	Pulse broadening timescale due to interstellar scattering
$\Delta\nu_{\text{d}}$	Diffraction scintillation bandwidth
$\Delta t_{\text{d}}$	Diffraction scintillation timescale
$\lambda$	Wavelength of the radio radiation
$B_\parallel$	Line-of-sight component of the interstellar magnetic field
$M_{\text{p}}$	Mass of the neutron star
$M_{\text{c}}$	Mass of the companion star
$f(M)$	Mass function of the binary system
$x$	Projected semi-major axis ( $a \sin i/c$ )
$T_0$	Epoch of periastron passage
$e$	Orbital eccentricity
$q$	Mass ratio ( $M_{\text{c}}/M_{\text{p}}$ )
$i$	Orbital inclination angle

# List of Figures

- 1.1 Schematic illustration of the pulsar lighthouse model. A neutron star rotates steadily about its rotation axis while possessing a strong magnetic field whose axis is misaligned with the direction of rotation. Radio emission is produced along open magnetic field lines near the magnetic poles and is concentrated into narrow beams aligned with the magnetic field geometry. As the neutron star rotates, these beams sweep through space in a regular manner, producing the observed pulsed emission when the beam crosses the observer’s line of sight. The light cylinder defines the boundary beyond which rigid corotation with the neutron star is no longer possible. This geometric interpretation forms the basis of the lighthouse model of pulsars [1–3]. . . . . 27
- 1.2 Spin period–period derivative ( $P-\dot{P}$ ) diagram showing the main populations of known rotation-powered neutron stars. Ordinary radio pulsars from the ATNF catalogue form a dense population centred around  $P \sim 0.1\text{--}1$  s and  $\dot{P} \sim 10^{-15}\text{--}10^{-13}$  s s<sup>-1</sup>. Millisecond pulsars (MSPs) occupy the lower-left region of the diagram, characterised by short spin periods and extremely small spin-down rates, indicative of old, recycled neutron stars. Magnetars populate the upper-right region, exhibiting long spin periods and large period derivatives consistent with ultra-strong magnetic fields. X-ray Dim Isolated Neutron Stars (XDINS) are found at long periods and intermediate spin-down rates. Lines of constant surface dipole magnetic field strength and characteristic age are overlaid for reference. The diagram provides a unified view of neutron star evolution under magnetic braking and highlights the distinct evolutionary pathways that give rise to the observed pulsar populations. . . . . 39

2.1	Number of known radio pulsars hosted by Galactic globular clusters, separated into binary and isolated systems. The statistics are compiled from the online globular cluster pulsar catalogue maintained by P. C. C. Freire as of 18 march 2026. . . . .	49
2.2	Distribution of core ( $t_{rc}$ ) and half-mass ( $t_{rh}$ ) relaxation times for Galactic globular clusters from the Harris (1996; 2010 edition) catalogue. The dashed line marks $10^{10}$ yr for comparison with typical globular cluster ages. Cluster cores generally relax on timescales $\ll 10^{10}$ yr, implying substantial dynamical evolution over a Hubble time. . . . .	53
2.3	Gravitational focusing enhancement factor, $1+2G(m_1+m_2)/(r_{\min}v_{\infty}^2)$ , as a function of encounter velocity at infinity $v_{\infty}$ for representative periapsis scales $r_{\min}$ . Low velocity dispersions typical of globular clusters strongly increase the effective cross section for close encounters relative to the geometric expectation. . . . .	55
2.4	Retained neutron-star fraction as a function of cluster escape speed, $f_{\text{ret}}(v_{\text{esc}}) = \int_0^{v_{\text{esc}}} f(v) dv$ , for two illustrative natal-kick distributions. The high-kick curve adopts a Maxwellian with one-dimensional dispersion $\sigma_k = 265 \text{ km s}^{-1}$ inferred for ordinary field pulsars [4]. The low-kick curve is included for comparison to illustrate how the presence of low-kick formation channels (e.g. electron-capture or ultra-stripped supernovae) can substantially increase retention in globular clusters [5, 6]. Vertical lines indicate representative globular-cluster escape speeds. . . . .	60
2.5	Conceptual illustration of the maximum and minimum line-of-sight accelerations expected for pulsars as a function of angular separation from the cluster centre. The solid curves represent the LOS acceleration envelope produced by a standard cluster potential, while the dotted and dashed curves illustrate the additional inner contribution that could arise from increasingly compact central mass components. The vertical line marks the cluster core radius $r_c$ . Comparisons between timing-derived pulsar accelerations and such envelopes provide a direct dynamical probe of the cluster mass distribution. . . . .	69

3.1	<p>Illustration of dispersive delay in a pulsar signal across a radio observing band. <i>Top left:</i> Simulated dynamic spectrum showing the characteristic frequency-dependent arrival time of a dispersed pulse in the ionised interstellar medium (ISM), where lower-frequency radiation arrives later than higher-frequency radiation following the cold-plasma dispersion law (<math>t \propto f^{-2}</math>). <i>Top right:</i> The same signal after correcting for the dispersive delay using the correct dispersion measure (DM), resulting in the pulse aligning across all frequency channels. <i>Bottom panels:</i> Band-averaged time series before and after dedispersion. Before correction the pulse is significantly smeared due to the frequency dependent delays, while after dedispersion the intrinsic narrow pulse is recovered with substantially higher peak signal-to-noise ratio. . . . .</p>	77
3.2	<p>Illustration of Faraday rotation in a magnetised ionised medium. <i>Left panel:</i> the physical effect of Faraday rotation. As linearly polarised radiation propagates through the interstellar medium (ISM), the plane of polarisation rotates by an angle that increases with the square of the observing wavelength. <i>Right panel:</i> the corresponding observational relation between the polarisation position angle <math>\psi</math> and <math>\lambda^2</math>. According to <math>\psi = \psi_0 + \text{RM} \lambda^2</math>, the slope of this relation yields the rotation measure (RM). The labelled points correspond to the wavelengths illustrated in the left panel, demonstrating how the physical rotation of the polarisation vector translates into the linear <math>\psi</math>-<math>\lambda^2</math> relation measured in pulsar polarimetric observations. . . . .</p>	79
3.3	<p>Simulated frequency evolution of pulse broadening caused by multipath scattering in the ISM. The same intrinsic pulse profile is shown at three observing frequencies assuming a scattering timescale scaling as <math>\tau_{\text{sc}} \propto f^{-4}</math>. At higher observing frequencies the pulse remains narrow, while at lower frequencies the pulse develops an increasingly extended exponential tail. This broadening reduces the peak signal-to-noise ratio and can strongly limit the detectability of pulsars in low-frequency surveys. . . . .</p>	81

3.4	Simulated dynamic spectrum illustrating interstellar scintillation of a pulsar signal. The main panel shows the observed intensity as a function of time and observing frequency. The patchy bright and dim regions (“scintles”) arise from interference between multiple propagation paths through the turbulent ionised interstellar medium. The top and right panels show the mean intensity variations with time and frequency, respectively. Scintillation can strongly modulate the apparent brightness of pulsars and may therefore affect their detectability in individual observations. . . . .	83
3.5	Illustration of common forms of radio-frequency interference (RFI) encountered in pulsar search data. The upper panels show simplified dynamic spectra illustrating individual types of contamination: <i>narrow-band RFI</i> , which appears as persistent horizontal features affecting specific frequency channels; <i>broadband impulsive RFI</i> , which produces vertical features affecting many channels simultaneously for short time intervals; and <i>periodic terrestrial RFI</i> , which appears as regularly repeating bursts in time. For comparison, the lower-right panel shows a simulated dispersed pulsar signal, which produces a frequency-dependent sweep across the observing band due to interstellar dispersion. The bottom panel shows a combined dynamic spectrum containing all of these components simultaneously, illustrating how different forms of interference and astrophysical signals can co-exist in real pulsar search observations. . . . .	90
3.6	Illustration of the effect of zero-DM subtraction on undispersed terrestrial interference in pulsar search data. The upper panels show the dynamic spectrum before and after zero-DM subtraction. Broadband terrestrial RFI, which arrives simultaneously across all frequency channels and therefore peaks at $DM \approx 0$ , appears as vertical stripes in the dynamic spectrum. Subtracting the mean across frequency channels at each time sample strongly suppresses this undispersed component while largely preserving the dispersed pulsar signal. The lower panels show the time series after dedispersion at the pulsar dispersion measure. Prior to mitigation, terrestrial bursts contaminate the dedispersed time series and obscure the astrophysical signal. After zero-DM subtraction, the RFI contribution is significantly reduced and the pulsar peak becomes clearly identifiable. . . . .	93

3.7	Illustration of the power of Fourier-domain periodicity searches. <i>Top panels:</i> a short-duration noisy time series and its Fourier spectrum. Although the periodic signal is difficult to identify by eye in the time domain, the FFT already reveals a localized excess at the signal frequency. <i>Middle panels:</i> the same experiment for a longer observation. The time-domain signal still appears noise dominated, but the Fourier peak becomes much more significant because the periodic component adds coherently over a larger number of cycles. <i>Bottom panels:</i> comparison of the clean periodic component with the noise realization in the time domain, and the corresponding growth of coherent signal power in the Fourier spectrum. This demonstrates why the FFT is such a powerful tool for uncovering repeating signals buried in noise.	100
3.8	Illustration of the line-of-sight orbital acceleration $a_{\parallel}(t)$ and line-of-sight jerk $j_{\parallel}(t)$ over one full binary orbit for a circular system. The upper panel shows the periodic variation of the line-of-sight acceleration, which determines the instantaneous drift rate of the observed pulsar spin frequency through $\dot{f}_{\text{obs}} = (f_0/c) a_{\parallel}$ . The lower panel shows the corresponding line-of-sight jerk, which determines the rate of change of that drift through $\ddot{f}_{\text{obs}} = (f_0/c) j_{\parallel}$ . Together, these quantities provide the physical origin of Doppler smearing in Fourier-domain pulsar searches.	107
3.9	Illustration of the effect of binary acceleration on a pulsar signal in the Fourier domain. In the absence of binary motion (top panels), the periodic signal remains at a fixed frequency throughout the observation and appears as a sharp peak in the FFT. Under non-zero line-of-sight acceleration (central panels), the observed spin frequency drifts during the observation and the signal power is spread across multiple Fourier bins. This Doppler smearing reduces the peak significance in a standard FFT search and motivates the use of acceleration-search techniques (see comparison in bottom panels).	110
3.10	A typical pulsar diagnostic plot prepared by PRESTO's Prepfold routine.	116
131figure.caption.55		
4.2	Time-phase plots and two pulse periods of the integrated pulse profiles for the four newly discovered pulsars in Glimpse-C01 (J1848–0129C, D, E and F), shown in order.	136

4.3	Observed spin period as a function of orbital phase for J1848–0129C (top) and J1848–0129D (bottom). Blue points represent measured spin periods from individual observations, while the solid black curves indicate the predictions obtained from the orbital solutions listed in Table 4.3. . . . .	139
4.4	Intra-cluster dispersion-measure dispersion, $\sigma_{\text{DM}}$ , as a function of the nominal cluster dispersion measure (defined as the median pulsar DM) for Galactic globular clusters hosting at least five known radio pulsars. The solid line shows the best-fitting power-law relation of the form $\sigma_{\text{DM}} = ADM^\alpha$ , with $A = 0.008$ and $\alpha = 1.06$ (bootstrap 16–84% interval: 0.90–1.21). The shaded region represents the root-mean-square scatter in log space (0.27 dex) about the fit. . . . .	142
5.0	Comparison between the parameter-space coverage of the Fourier-domain PRESTO search (blue wedge) and a representative time-domain search band (orange), shown for acceleration (left column) and jerk (right column). The three rows correspond to harmonic sums of $h = 4, 8,$ and $16$ . Increasing $h$ improves sensitivity to narrow pulse profiles through harmonic summing, but simultaneously makes the Fourier-domain search boundary more restrictive, particularly at short spin periods. The vertical dotted lines indicate the crossing period below which the Fourier-domain wedge no longer fully covers the chosen time-domain band. This effect is most pronounced for jerk, where the loss of coverage occurs precisely in the region occupied by rapidly spinning pulsars in extremely compact binaries. . . . .	161
5.1	Evolution of the resampling index shift, $(\text{out\_idx} - \text{in\_idx})$ , for increasing trial jerk for the same value of acceleration. In the pure-acceleration case ( $j = 0$ ), the mapping remains only weakly distorted, whereas progressively larger jerk values introduce an increasingly strong non-linear deformation of the output index as a function of input index. This reflects the quadratic contribution of the jerk term to the time-domain resampling map. As $j$ increases, the resampler must account for progressively larger departures from the original regular sampling pattern, implying a larger correction to the time series before the FFT is performed. . . . .	163

5.2	Local increment of the resampling index, $\Delta \text{out\_idx}$ , for increasing trial jerk keeping the value of acceleration constant. While the $j = 0$ case remains comparatively regular, the inclusion of jerk introduces a position-dependent variation in the local stepping of the output index. This means that neighbouring output samples no longer map back to the input array with a uniform stride. Instead, the access pattern becomes increasingly irregular as the jerk increases, which is precisely the behaviour that drives up the computational cost of the resampling stage. In practice, this leads to less regular memory access and motivates the use of massively parallel GPU implementations when extending time-domain searches beyond constant acceleration. .	164
5.3	Coarse mismatch map of the current stochastic $(a, j)$ template bank. The colour at each point shows the mismatch to the nearest template, evaluated using the projected metric on the acceleration–jerk manifold. Darker regions indicate better local coverage, while brighter patches correspond to under-covered regions where the nearest-template mismatch is higher. The patchy structure, particularly toward the boundaries of the search box, shows that the present Monte Carlo / MCMC bank generator does not yet provide sufficiently uniform coverage for a production-quality search. . . . .	186
5.4	Illustrative recovery of the same injected pulsar using an acceleration-only search (left) and the jerk-enabled <code>PULSEJET_beta</code> search (right). The <code>PULSEJET_beta</code> candidate shows a substantially cleaner and stronger recovery, consistent with the expectation that a signal containing significant jerk is not optimally corrected by an acceleration-only search. .	189
5.5	Recovery fractions for the 10-pulsar blind injection tests using the current stochastic template-bank implementation. The left panel shows the result obtained with the original template bank, for which 7 out of 10 injected pulsars were recovered. The right panel shows the result after enlarging the template bank so that the injected pulsars lay more centrally within the populated $(a, j)$ region, yielding 9 out of 10 recoveries. The improved performance in the enlarged-bank test supports the conclusion that the dominant limitation of the present implementation is not the jerk-resampling kernel itself, but incomplete template-bank coverage, especially near the boundaries of the acceleration–jerk search space. . . . .	190

5.6	Comparison between the number of trials required by a naive rectangular scan in acceleration and jerk and by the metric-informed template-bank approach, shown for several observation lengths. The right-hand panel shows the corresponding reduction factor. Even in its present beta form, the template-bank method reduces the total number of required trials by a factor of about 3.4. . . . .	192
5.7	Preliminary comparison between <code>PrestoZL</code> and <code>PULSEJET_beta</code> in terms of per-run jerk-search runtime and total candidate count. In the tested configuration, <code>PULSEJET_beta</code> completed the run more quickly but produced many more candidates. At the present stage this should be interpreted only as an exploratory computational comparison, since the two pipelines probe the parameter space differently and <code>PULSEJET_beta</code> is still under active development. . . . .	193
B.1	Mean system-equivalent flux density per antenna across the full MeerKAT S-band receiver bandwidth, digitised from SRAO performance documentation. . . . .	205
B.2	Mean system-equivalent flux density per antenna evaluated for individual sub-bands across the observing bandwidth. . . . .	205

# List of Tables

4.1	Summary of globular-cluster pointings in the first MeerKAT S-band pulsar survey. For each target we list the cluster identifier, J2000 right ascension and declination of the telescope pointing, the observation date in MJD, the nominal cluster dispersion measure (DM) used for dedispersing the data, the total integration length, and the sampling interval. The columns “Re-det” and “New-det” indicate whether any known pulsar was re-detected or a new pulsar discovered in that pointing. The DM range searched for each globular cluster was $\pm 20\%$ of the nominal cluster DM (taken as the median for clusters with multiple pulsars and as the DM of pulsar A in clusters with only one known pulsar), with the DM step sizes calculated using PRESTO’s <code>DDplan.py</code> . . . . .	125
4.2	Parameters for the pulsars in the globular cluster Glimpse-C01. Newly discovered pulsars are shown in boldface. Listed are the barycentric spin period $P$ , dispersion measure (DM), and mean flux density $S$ measured at MeerKAT S-band. Uncertainties reported in parentheses correspond to $1\sigma$ errors on the final quoted digit(s). Flux-density uncertainties are derived from the off-pulse noise of the folded profiles, as described in Section A.1. . . . .	136
4.3	Best-fitting orbital parameters for J1848–0129C and J1848–0129D, together with their relative $1\sigma$ uncertainties, obtained using a modified version of the PRESTO routine <code>fit_circular_orbit.py</code> . . . . .	138
4.4	Flux densities of pulsars measured in MeerKAT S-band and L-band observations. The $S_S$ values are from single-epoch detections, whereas $S_L$ values are averages over multiple epochs, where available. Spectral indices are computed assuming $S_\nu \propto \nu^\alpha$ between 1.2 and 2.4 GHz. Uncertainties are $1\sigma$ ; values in parentheses denote the uncertainty on the final quoted digit(s). . . . .	147

4.5	Flux density measurements for the Terzan 5 pulsars detected in our single-epoch MeerKAT S-band observations. Owing to the large number of pulsars detected in this cluster, we present the Terzan 5 measurements separately from the rest of the sample. The $S_{1400}$ and $S_{2000}$ values are taken from Martsen et al. [7], while $S_{2400}$ gives our measured flux densities at a central frequency of $\sim 2406$ MHz. The spectral index $\alpha$ is computed using the literature $S_{1400}$ value together with our measured $S_{2400}$ , assuming a power-law spectrum, $S_\nu \propto \nu^\alpha$ . No literature flux density measurement was available for PSR J1748–2446am.	148
6.1	Principal software packages and pipelines used in this thesis, listed in alphabetical order.	200

# Chapter 1

## Pulsar generalities

### 1.1 Neutron Stars and the Discovery of Pulsars

Neutron stars are among the most extreme astrophysical objects known, representing the compact remnants of massive stars that have ended their lives in core-collapse supernova explosions. The theoretical existence of neutron stars was first proposed shortly after the discovery of the neutron by Baade and Zwicky [8], who suggested that supernovae could result in the formation of extremely dense stellar remnants composed primarily of neutrons. Subsequent theoretical work established neutron stars as objects with masses comparable to that of the Sun but radii of only about 10–12 km, implying mean densities exceeding nuclear saturation density.

The formation of a neutron star occurs when a massive star, typically with an initial mass greater than  $\sim 8 M_{\odot}$ , exhausts its nuclear fuel and can no longer support itself against gravitational collapse. During the collapse of the stellar core, protons and electrons combine via inverse beta decay to form neutrons, while the outer layers of the star are expelled in a supernova explosion. Conservation of angular momentum during this rapid collapse leads to a dramatic increase in the rotation rate of the remnant, while magnetic flux conservation amplifies the magnetic field to strengths typically in the range  $10^{11}$ – $10^{13}$  G [9, 10].

The observational manifestation of neutron stars as pulsars was discovered serendipitously in 1967 by Jocelyn Bell Burnell and Antony Hewish during a radio survey designed to study interplanetary scintillation. The discovery consisted of highly regular, narrow radio pulses with a period of approximately 1.3 s, originating from a fixed position on the sky [11]. The remarkable stability of the signal initially led to speculation about an artificial origin, but it was soon recognised that the source, later identified as CP 1919 (now PSR B1919+21), was astrophysical in nature. This discovery marked the first observational confirmation of rotating neutron stars and

inaugurated the field of pulsar astronomy.

Pulsars are now understood as rotating, magnetised neutron stars that emit beams of electromagnetic radiation from regions near their magnetic poles. When the magnetic axis is inclined with respect to the rotational axis, these beams sweep through space as the star rotates, producing a lighthouse-like effect. An observer detects a pulse of radiation each time a beam crosses the line of sight, with the observed pulse period corresponding directly to the neutron star's rotation period [1].

The radio emission from pulsars is believed to originate within the neutron star magnetosphere, where strong electric fields accelerate charged particles along open magnetic field lines. Although the detailed emission mechanism remains an active area of research, the resulting radiation is highly coherent and narrowly beamed, leading to the characteristic pulsed signals observed at Earth [2]. Pulsar rotation periods span a wide range, from milliseconds to several seconds, reflecting both their birth properties and subsequent evolutionary history.

Over time, isolated pulsars lose rotational energy through electromagnetic radiation and particle winds, leading to a gradual increase in their spin periods. This process, commonly referred to as spin-down, provides a direct probe of the neutron star's magnetic field and energy loss mechanisms. Despite this slow evolution, pulsars exhibit extraordinary rotational stability, making them precise celestial clocks and powerful tools for a wide range of astrophysical applications.

It is important to note that pulsar emission is intrinsically directional, and only a fraction of the Galactic neutron star population is observable as pulsars. The detected population is therefore shaped by both intrinsic properties, such as luminosity and beam geometry, and observational selection effects. This foundational understanding of pulsars forms the basis for more complex systems, including millisecond pulsars, binary pulsars, and pulsars in dense stellar environments.

## 1.2 Pulsar Geometry and the Lighthouse Model

The characteristic pulsed emission observed from pulsars arises primarily from geometric effects associated with rotation rather than from an intrinsically intermittent radiation process. This interpretation, commonly known as the *lighthouse model*, provides a simple and intuitive explanation for why pulsars appear as periodically flashing sources and why only a subset of neutron stars are detectable as pulsars [1, 3, 12].

Figure 1.1 presents a schematic illustration of this geometric picture. At the




Figure 1.1: Schematic illustration of the pulsar lighthouse model. A neutron star rotates steadily about its rotation axis while possessing a strong magnetic field whose axis is misaligned with the direction of rotation. Radio emission is produced along open magnetic field lines near the magnetic poles and is concentrated into narrow beams aligned with the magnetic field geometry. As the neutron star rotates, these beams sweep through space in a regular manner, producing the observed pulsed emission when the beam crosses the observer's line of sight. The light cylinder defines the boundary beyond which rigid corotation with the neutron star is no longer possible. This geometric interpretation forms the basis of the lighthouse model of pulsars [1-3].

centre of the system lies the neutron star, an extremely compact remnant formed during a core-collapse supernova. The neutron star rotates steadily about a fixed rotation axis, shown in the figure as a vertical direction. Surrounding the star is an intense magnetic field whose large-scale structure is well approximated by a dipole configuration, similar in form to the magnetic field of a bar magnet [2, 9].

A key element of the lighthouse model is that the magnetic field axis is generally misaligned with the rotation axis. This misalignment causes regions associated with magnetic activity to sweep through space as the neutron star rotates. The magnetic field divides the surrounding magnetosphere into regions of closed magnetic field lines, which loop back onto the neutron star, and open magnetic field lines, which extend outward from the magnetic poles. The boundary between these regions is associated with the light cylinder, which marks the outer limit within which magnetic field lines can corotate rigidly with the neutron star [13–15].

In the conventional picture, radio emission is usually associated with coherent plasma processes operating along open magnetic field lines above the magnetic poles of the neutron star. However, the emission location remains an active area of research, especially for millisecond pulsars. Recent work has suggested that some radio components in MSPs may instead arise much farther out, possibly in or near the current sheet beyond the light cylinder, colocated with high-energy emission [16]. Therefore, the polar-cap/open-field-line picture should be regarded as the standard baseline interpretation rather than a complete description of all pulsar radio emission. This emission is not radiated uniformly in all directions but is instead confined to relatively narrow beams aligned with the magnetic field geometry. As the neutron star rotates, these beams sweep across space in a regular and predictable manner, analogous to the rotating beam of a lighthouse [1, 2, 17, 18].

An observer detects a pulsar only if the sweeping radio beam intersects their line of sight during each rotation of the neutron star. When this condition is satisfied, the observer perceives a brief pulse of radiation once per rotation period, giving rise to the highly regular pulse trains that characterise pulsar signals. If the beam does not cross the observer’s line of sight, the neutron star remains undetected as a pulsar despite continuing to emit radiation. This geometric selection effect strongly influences the observed pulsar population and implies that the known sample represents only a fraction of the total neutron star population in the Galaxy [3, 19, 20].

The lighthouse model successfully accounts for several fundamental observational properties of pulsars, including the stability of their pulse periods, the repeatability of pulse profiles, and the strong dependence of detectability on viewing geometry. Although the detailed physical mechanisms responsible for particle acceleration and

coherent radio emission within the magnetosphere remain an active area of research, the geometric framework illustrated in Fig. 1.1 forms the foundation of modern pulsar astronomy and underpins much of the observational interpretation of pulsar data [2, 12, 15].

## 1.3 Birth of Pulsars

Pulsars are formed as a natural consequence of the evolution and death of massive stars. Their observable properties are largely determined by physical conditions established at birth and by the subsequent evolution driven by energy loss and environmental interactions. An understanding of pulsar birth therefore provides the foundation for interpreting the wide range of pulsar behaviours and populations observed today.

### 1.3.1 Core-Collapse Supernovae and Neutron Star Formation

The formation of a pulsar begins with the gravitational collapse of a massive star at the end of its nuclear-burning lifetime. Stars with initial masses exceeding approximately eight solar masses develop iron cores that cannot undergo further fusion. Once nuclear support is exhausted, the core collapses on a dynamical timescale, while the outer layers of the star are expelled in a core-collapse supernova explosion [8, 9, 21].

During collapse, densities rise to values comparable to or exceeding nuclear saturation density. Under these conditions, electrons and protons combine via inverse beta decay, producing a neutron-rich core and a burst of neutrinos that carry away a significant fraction of the gravitational binding energy. The collapse is halted by neutron degeneracy pressure and strong nuclear forces, leaving behind a neutron star with a mass typically between one and two solar masses and a radius of order ten kilometres [10, 22].

Observational evidence directly linking pulsars to supernova explosions is provided by the association of young pulsars with supernova remnants. Classic examples include the Crab and Vela pulsars, which are located near the centres of expanding remnants and exhibit spin-down ages consistent with the independently estimated ages of their associated supernovae [3, 12, 23]. These systems provide compelling confirmation that pulsars represent an early evolutionary phase of neutron stars formed in core-collapse supernovae.

### 1.3.2 Initial Spin and Magnetic Field Properties

The defining characteristics of a newly formed pulsar—rapid rotation and an intense magnetic field—arise naturally from conservation laws acting during stellar collapse. As the stellar core contracts by several orders of magnitude in radius, conservation of angular momentum leads to a substantial increase in rotational frequency. Even a progenitor core rotating relatively slowly prior to collapse can therefore give rise to a neutron star with an initial spin period ranging from milliseconds to seconds [2, 9].

This process can be expressed schematically by angular momentum conservation,

$$I_{\text{prog}} \Omega_{\text{prog}} \approx I_{\text{NS}} \Omega_{\text{NS}}, \quad (1.1)$$

where the dramatic reduction in the moment of inertia during collapse results in a corresponding increase in the neutron star’s rotation rate. This simple relation captures the essential physical origin of pulsars as rapidly rotating objects, without requiring detailed modelling of the collapse itself.

In parallel, conservation of magnetic flux during collapse leads to a strong amplification of the magnetic field. As the cross-sectional area of the collapsing core decreases, the magnetic field strength increases accordingly, producing surface magnetic fields that are many orders of magnitude larger than those found in ordinary stars. These strong magnetic fields dominate the structure of the neutron star magnetosphere and are essential for the generation of pulsar emission [9, 15, 24].

Newborn pulsars therefore emerge as rapidly rotating, strongly magnetised neutron stars with substantial reservoirs of rotational energy. Young pulsars are consequently highly energetic systems, often powering bright radio emission, high-energy radiation, and relativistic particle winds. As pulsars evolve, their rotation gradually slows and their emission properties change, leading to the diverse population observed today. The mechanisms responsible for this long-term evolution define the class of rotation-powered pulsars and will be examined in detail in the following section, where the processes governing pulsar spin-down and energy loss are introduced.

## 1.4 Rotation-Powered Pulsars and Spin-Down Evolution

The majority of known radio pulsars are powered by the gradual loss of their rotational kinetic energy. These objects are commonly referred to as *rotation-powered*

*pulsars* and include both young, energetic pulsars and older systems whose emission remains detectable despite substantial spin-down. The evolution of a pulsar’s rotation rate provides a natural framework for understanding its age, energetics, and long-term behaviour.

### 1.4.1 Rotational Energy as the Power Source

A rotating neutron star possesses a large reservoir of kinetic energy associated with its spin. To first order, the total rotational energy of a neutron star can be expressed as

$$E_{\text{rot}} = \frac{1}{2}I\Omega^2, \quad (1.2)$$

where  $I$  is the moment of inertia of the neutron star and  $\Omega$  is its angular rotation frequency. For a canonical moment of inertia  $I \simeq 10^{45} \text{ g cm}^2$ , this corresponds to  $E_{\text{rot}} \simeq 2 \times 10^{46} \text{ erg}$  for a pulsar with  $P = 1 \text{ s}$ , and rises to  $\simeq 2 \times 10^{50} \text{ erg}$  for  $P = 10 \text{ ms}$ . Expressed in terms of the energy radiated by the Sun in one year, these values correspond to roughly  $1.6 \times 10^5$  and  $1.6 \times 10^9 L_{\odot} \text{ yr}$ , respectively. The rotational reservoir is therefore especially large for young or rapidly rotating neutron stars. [2, 9].

As pulsars emit electromagnetic radiation and drive relativistic particle winds, they lose rotational energy over time. The rate of energy loss, often referred to as the spin-down luminosity, is given by

$$\dot{E}_{\text{rot}} = I\Omega\dot{\Omega}, \quad (1.3)$$

where  $\dot{\Omega}$  denotes the time derivative of the angular frequency. This quantity provides a direct measure of the power available to drive pulsar emission and associated phenomena such as pulsar wind nebulae [1, 23].

Rotation-powered pulsars are therefore characterised by emission whose ultimate energy source is the star’s rotational kinetic energy. As the pulsar spins down, the available energy decreases, leading to an evolutionary sequence in which young pulsars are typically more luminous and energetic than their older counterparts.

### 1.4.2 Spin-Down Laws and Magnetic Braking

The gradual slowdown of pulsar rotation is driven by external torques acting on the neutron star and its magnetosphere. Although the detailed torque depends on magnetospheric physics, a widely used phenomenological description assumes that

the spin-down can be expressed as a power law in the rotation rate,

$$\dot{\Omega} = -K \Omega^n, \quad (1.4)$$

where  $K$  is a constant that encapsulates the strength of the braking torque (and may depend on magnetic field strength, geometry, and magnetospheric state), and  $n$  is the *braking index* [2, 12, 15].

A convenient and observable definition of the braking index can be obtained directly from Eq. 1.4. Differentiating with respect to time gives

$$\ddot{\Omega} = -Kn \Omega^{n-1} \dot{\Omega}. \quad (1.5)$$

Rearranging and eliminating  $K$  using Eq. 1.4 yields

$$n = \frac{\Omega \ddot{\Omega}}{\dot{\Omega}^2}. \quad (1.6)$$

In other words, the braking index is a measurable quantity that characterises how the spin-down rate itself evolves with time. For the idealised case of pure magnetic dipole radiation in vacuum, one expects  $n = 3$ , while observed values for young pulsars frequently deviate from this, indicating that additional physics such as magnetospheric currents, particle winds, or evolving magnetic fields contribute to the braking torque [3, 24, 25].

### 1.4.3 Characteristic Age and Magnetic Field Estimates

The power-law spin-down model in Eq. 1.4 can be integrated to relate a pulsar's current spin state to the time elapsed since birth. Separating variables gives

$$\frac{d\Omega}{\Omega^n} = -K dt. \quad (1.7)$$

Integrating from the birth rotation rate  $\Omega_0$  to the current rotation rate  $\Omega$  over a time interval  $t$  yields, for  $n \neq 1$ ,

$$t = \frac{1}{(n-1)K} (\Omega^{1-n} - \Omega_0^{1-n}). \quad (1.8)$$

Using Eq. 1.4 to eliminate  $K$  (i.e.  $K = -\dot{\Omega}/\Omega^n$ ) gives an equivalent expression in terms of the currently observed spin parameters:

$$t = \frac{\Omega}{(n-1)|\dot{\Omega}|} \left[ 1 - \left( \frac{\Omega}{\Omega_0} \right)^{n-1} \right]. \quad (1.9)$$

This is the general age estimate under the assumptions of a constant braking index and a time-independent braking coefficient.

A commonly used special case is obtained by assuming dipole-like braking with  $n = 3$  and that the pulsar was born spinning much faster than it rotates today (so that  $\Omega_0 \gg \Omega$ ). Under these assumptions, the bracketed term in Eq. 1.9 approaches unity and the age estimate reduces to the *characteristic age*:

$$\tau_c \equiv \frac{\Omega}{2|\dot{\Omega}|}. \quad (1.10)$$

Expressed in terms of the spin period  $P = 2\pi/\Omega$  and its derivative  $\dot{P}$ , this becomes the familiar form

$$\tau_c = \frac{P}{2\dot{P}}. \quad (1.11)$$

The characteristic age is widely used in pulsar population studies because it can be computed directly from timing measurements [2, 3, 12]. However, it is not a direct measurement of true age: deviations from  $n = 3$ , significant initial spin periods, or time-variable torques can all lead to large differences between  $\tau_c$  and the actual time since birth [24, 26].

For completeness, the braking index in Eq. 1.6 may also be expressed in terms of period derivatives. Using  $\Omega = 2\pi/P$  one obtains

$$n = 2 - \frac{P\ddot{P}}{\dot{P}^2}, \quad (1.12)$$

which is sometimes convenient when period derivatives are tabulated directly [2].

## 1.5 Accretion-Powered Pulsars and Binary Evolution

While the majority of radio pulsars are powered by the loss of rotational kinetic energy, a significant and astrophysically important subset of neutron stars derive their observable emission from interactions with a binary companion. In such systems, the presence of a companion star fundamentally alters the evolutionary pathway of

the neutron star, giving rise to accretion-powered pulsars and, in many cases, recycled systems. Binary evolution therefore plays a central role in shaping the observed pulsar population.

### 1.5.1 Binary Systems and Mass Transfer

Binary systems are common outcomes of massive star formation, and many neutron stars are born with stellar companions. The subsequent evolution of such systems depends sensitively on the orbital separation, the mass of the companion, and the evolutionary state of both stars. As the companion evolves, it may expand and transfer mass onto the neutron star through stellar winds or Roche-lobe overflow, initiating an accretion phase [27, 28].

Accretion proceeds through the formation of an accretion disk when the transferred material possesses sufficient angular momentum. Within the disk, viscous processes transport angular momentum outward, allowing matter to spiral inward toward the neutron star. The release of gravitational potential energy during accretion can produce intense high-energy emission, often dominating the observed luminosity of the system [29].

The onset of accretion introduces torques that can strongly influence the neutron star's rotation. Depending on the accretion rate and magnetospheric interaction, these torques may either spin the neutron star up or down, in stark contrast to the monotonic spin-down experienced by isolated, rotation-powered pulsars. Binary interaction thus provides an external mechanism capable of reshaping a neutron star's spin evolution.

### 1.5.2 Accretion-Powered Pulsars

Accretion-powered pulsars are neutron stars whose observable emission is primarily driven by the release of gravitational energy from infalling matter rather than by rotational energy loss. These systems are most commonly observed as X-ray pulsars, where accreted material is channelled by the neutron star's magnetic field onto the magnetic poles, producing hot spots that rotate with the star and give rise to pulsed X-ray emission [29, 30].

In accretion-powered systems, the spin evolution of the neutron star is governed by the balance between accretion torques and magnetic stresses at the magnetospheric boundary. Sustained accretion can transfer substantial angular momentum to the neutron star, leading to a progressive increase in its rotation rate. This process can dramatically alter the spin period and magnetic field configuration over

timescales much shorter than those associated with isolated spin-down evolution [27, 28].

Accretion-powered pulsars therefore represent a distinct evolutionary phase in which the neutron star’s observable properties are dominated by its binary environment. These systems are particularly important as progenitors of recycled pulsars, providing the physical mechanism required to spin neutron stars up to very short periods.

### 1.5.3 Spin-Up and the Recycling Scenario

The long-term outcome of sustained accretion in a binary system is encapsulated in the *recycling* scenario. In this framework, an old, slowly rotating neutron star is spun up through prolonged mass and angular momentum transfer from a companion star. As accretion proceeds, the neutron star’s spin period decreases while its magnetic field strength is reduced, either through field decay or magnetic burial beneath accreted material [27, 31].

Once accretion ceases—typically after the companion evolves into a white dwarf—the neutron star re-emerges as a rapidly rotating, weakly magnetised pulsar observable primarily at radio wavelengths. These objects are known as millisecond pulsars and exhibit exceptional rotational stability, reflecting both their advanced age and their modified magnetic field structure [2].

Binary interaction and recycling therefore provide a natural explanation for the existence of millisecond pulsars and their strong association with binary systems. This evolutionary pathway contrasts sharply with that of isolated pulsars and highlights the importance of external torques and environmental effects in shaping pulsar populations.

### 1.5.4 Binary Pulsars and Evolutionary Diversity

Binary pulsars exhibit a wide range of orbital configurations and companion types, including white dwarfs, neutron stars, and, in rare cases, non-degenerate stellar companions. The diversity of observed systems reflects the complex interplay between stellar evolution, mass transfer, and dynamical interactions [3, 28].

Some binary pulsars retain ongoing interactions with their companions, while others represent the end products of past accretion episodes. In dense stellar environments such as globular clusters, close encounters and exchange interactions further enrich the binary pulsar population, producing exotic systems that are rarely found in the Galactic field. These systems provide valuable laboratories for studying stel-

lar dynamics, binary evolution, and the long-term consequences of accretion-driven spin-up.

The central role of binary evolution and recycling in producing millisecond pulsars motivates a detailed examination of such systems. In the following section, the properties of millisecond pulsars and their significance for modern astrophysics are discussed in greater detail, forming a key foundation for the subsequent chapters of this thesis.

## 1.6 Millisecond Pulsars and Binary Motion

Millisecond pulsars (MSPs) form a distinct and astrophysically important subclass of the pulsar population. They are characterised by exceptionally short spin periods, typically of only a few milliseconds, weak inferred magnetic fields, and remarkable rotational stability. These properties distinguish MSPs sharply from the population of normal, isolated pulsars and point to a fundamentally different evolutionary history.

### 1.6.1 Defining Properties of Millisecond Pulsars

Millisecond pulsars are understood to be old neutron stars that have undergone significant spin-up through prolonged interaction with a binary companion. As a result of this evolutionary pathway, MSPs exhibit spin periods that are one to two orders of magnitude shorter than those of canonical pulsars, while their magnetic fields are correspondingly weaker. This combination leads to much slower intrinsic spin-down rates and exceptionally stable rotation over long timescales [2, 27, 31].

The rotational stability of MSPs rivals that of atomic clocks, making them uniquely valuable for high-precision timing experiments. Their long-term stability reflects both their advanced age and the reduced influence of magnetospheric torque fluctuations compared to young, strongly magnetised pulsars. Consequently, MSPs play a central role in a wide range of astrophysical applications, from tests of gravitational theories to studies of stellar evolution [3, 32].

A defining observational characteristic of MSPs is their strong association with binary systems. The majority of known MSPs are found in binaries, typically with white dwarf companions, providing direct evidence for their formation through accretion-driven recycling. This binary nature has profound implications for both their observed spin properties and their detectability.

### 1.6.2 Binary Motion and Orbital Modulation

The presence of a binary companion introduces additional complexity into the observed signal from a millisecond pulsar. As the pulsar orbits the system's centre of mass, its line-of-sight velocity varies periodically, leading to Doppler modulation of the observed spin frequency. This orbital motion causes the apparent pulsar period to vary as a function of orbital phase, even when the intrinsic spin period of the neutron star remains constant [2, 33].

To first order, the observed spin frequency can be expressed as

$$f_{\text{obs}}(t) = f_0 \left( 1 + \frac{v_{\parallel}(t)}{c} \right), \quad (1.13)$$

where  $f_0$  is the intrinsic spin frequency of the pulsar,  $v_{\parallel}(t)$  is the line-of-sight component of the pulsar's orbital velocity, and  $c$  is the speed of light. This Doppler modulation leads to time-dependent variations in the observed frequency and phase of the pulsar signal.

For millisecond pulsars, whose intrinsic spin frequencies are very high, even modest orbital velocities can produce significant frequency modulation over the course of an observation. As a result, binary motion must be carefully accounted for in both pulsar searches and timing analyses. Failure to correct for orbital Doppler effects can smear the pulsar signal in frequency space, reducing sensitivity and potentially rendering fast-spinning systems undetectable [34, 35].

### 1.6.3 Binary Millisecond Pulsars and Evolutionary Implications

The prevalence of binary systems among MSPs highlights the central role of binary evolution in shaping this population. In most cases, the companion star has evolved into a low-mass white dwarf following the accretion phase that spun up the neutron star. The resulting systems often exhibit nearly circular orbits, reflecting the dissipative nature of mass transfer through an accretion disk [28, 36].

Binary millisecond pulsars thus represent the end products of a long evolutionary sequence that begins with the formation of a neutron star in a supernova and proceeds through an extended phase of mass transfer and angular momentum exchange. Their observed properties encode information about both their accretion history and their subsequent dynamical evolution.

The strong connection between millisecond pulsars, binary evolution, and orbital dynamics makes them particularly sensitive probes of stellar evolution and gravi-

tational physics. Understanding their formation and observational characteristics is therefore essential for interpreting pulsar populations and motivates dedicated searches and timing studies, which form a central focus of this thesis.

#### 1.6.4 The $P-\dot{P}$ Diagram and the Pulsar Fauna

One of the most powerful tools for understanding the global properties and evolutionary pathways of neutron stars is the spin period–period derivative ( $P-\dot{P}$ ) diagram. By placing pulsars according to their measured rotation period  $P$  and its first time derivative  $\dot{P}$ , this diagram provides a compact, observationally driven summary of how neutron stars are born, evolve, and populate distinct physical regimes. For this reason, the  $P-\dot{P}$  diagram is often described as a “map” of the neutron star population, or colloquially as a view of the “pulsar fauna” [24, 26, 37, 38].

Figure 1.2 shows the  $P-\dot{P}$  diagram used throughout this thesis. It includes all rotation-powered neutron stars compiled from the Australia Telescope National Facility (ATNF) Pulsar Catalogue<sup>1</sup>[37] and magnetars from the McGill Online Magnetar Catalog<sup>2</sup>[39], updated to the current catalog releases as of January 2026. Additional populations, such as X-ray Dim Isolated Neutron Stars (XDINS) and candidate white dwarf pulsars, are overlaid for comparison.

**Spin period and spin-down rate.** The horizontal axis of the diagram represents the spin period  $P$  of the neutron star. This quantity spans more than five orders of magnitude, from  $\sim 1$  ms for the fastest known millisecond pulsars to tens of seconds for the slowest magnetars. The vertical axis shows the period derivative  $\dot{P}$ , which measures how rapidly the pulsar’s rotation is slowing down.

Physically,  $\dot{P}$  encodes the rate at which the neutron star is losing rotational kinetic energy due to external torques acting on its magnetosphere. For isolated, rotation-powered pulsars, this spin-down is primarily driven by electromagnetic braking associated with the star’s magnetic field and particle outflows [1, 13, 25]. A larger value of  $\dot{P}$  therefore corresponds to stronger braking and more rapid rotational evolution.

**Derived physical quantities.** Although the  $P-\dot{P}$  diagram displays only two directly measured observables, it implicitly encodes several fundamental physical properties of neutron stars. Under the assumption of magnetic dipole braking, the

<sup>1</sup><https://www.atnf.csiro.au/research/pulsar/psrcat/>

<sup>2</sup><http://www.physics.mcgill.ca/~pulsar/magnetar/main.html>




Figure 1.2: Spin period–period derivative ( $P-\dot{P}$ ) diagram showing the main populations of known rotation-powered neutron stars. Ordinary radio pulsars from the ATNF catalogue form a dense population centred around  $P \sim 0.1\text{--}1$  s and  $\dot{P} \sim 10^{-15}\text{--}10^{-13}$  s s $^{-1}$ . Millisecond pulsars (MSPs) occupy the lower-left region of the diagram, characterised by short spin periods and extremely small spin-down rates, indicative of old, recycled neutron stars. Magnetars populate the upper-right region, exhibiting long spin periods and large period derivatives consistent with ultra-strong magnetic fields. X-ray Dim Isolated Neutron Stars (XDINS) are found at long periods and intermediate spin-down rates. Lines of constant surface dipole magnetic field strength and characteristic age are overlaid for reference. The diagram provides a unified view of neutron star evolution under magnetic braking and highlights the distinct evolutionary pathways that give rise to the observed pulsar populations.

surface dipole magnetic field strength can be estimated as

$$B \simeq 3.2 \times 10^{19} \sqrt{P\dot{P}} \text{ G}, \quad (1.14)$$

while the characteristic age is given by

$$\tau_c = \frac{P}{2\dot{P}}. \quad (1.15)$$

Lines of constant magnetic field strength and constant characteristic age are overlaid in Fig. 1.2 to guide interpretation. While these quantities should be treated as order-of-magnitude estimates rather than precise measurements, they provide a valuable framework for understanding population-level trends [26, 40].

**Ordinary radio pulsars.** The bulk of the pulsar population occupies a dense region centred around  $P \sim 0.1\text{--}1$  s and  $\dot{P} \sim 10^{-15}\text{--}10^{-13}$  s s<sup>-1</sup>. These “ordinary” or “canonical” radio pulsars are typically young to middle-aged neutron stars formed in core-collapse supernovae. Their inferred magnetic field strengths cluster around  $10^{11}\text{--}10^{13}$  G, and their characteristic ages range from  $\sim 10^4$  to  $\sim 10^7$  yr.

The diagonal distribution of this population reflects long-term spin-down evolution: pulsars are born at short periods and large  $\dot{P}$  and gradually migrate toward longer periods and smaller spin-down rates as they lose rotational energy [24, 38].

**Millisecond pulsars.** Millisecond pulsars (MSPs) form a clearly distinct population in the lower-left corner of the diagram, characterised by spin periods of a few milliseconds and extremely small period derivatives ( $\dot{P} \lesssim 10^{-19}$  s s<sup>-1</sup>). These properties imply weak inferred magnetic fields ( $\sim 10^8\text{--}10^9$  G) and characteristic ages that often exceed a gigayear.

The location of MSPs in the  $P\text{--}\dot{P}$  diagram provides compelling evidence for the recycling scenario discussed in the previous sections: these neutron stars have been spun up through prolonged accretion in binary systems, which simultaneously reduces their effective magnetic field strength and suppresses subsequent spin-down [27, 31, 36]. Their extreme rotational stability makes MSPs invaluable tools for high-precision timing experiments and gravitational-wave detection efforts [32].

**Magnetars.** At the opposite extreme lie magnetars, which populate the upper-right region of the diagram. These objects exhibit long spin periods ( $P \sim 2\text{--}12$  s) combined with very large period derivatives ( $\dot{P} \sim 10^{-11}\text{--}10^{-10}$  s s<sup>-1</sup>), implying surface magnetic fields in excess of  $10^{14}$  G. Such fields are several orders of magni-

tude stronger than those of ordinary pulsars and dominate the star’s energetics and emission properties [26, 39].

The position of magnetars in the  $P-\dot{P}$  diagram highlights the role of magnetic field strength as a key parameter in neutron star evolution. Their rapid spin-down and relatively short characteristic ages are consistent with a picture in which magnetic energy, rather than rotational energy, powers much of their observed high-energy activity.

**X-ray Dim Isolated Neutron Stars.** X-ray Dim Isolated Neutron Stars (XDINS) occupy a region characterised by long spin periods ( $P \sim 3\text{--}12$  s) and intermediate period derivatives. Unlike ordinary radio pulsars, XDINS are predominantly detected through their thermal X-ray emission and are radio-quiet or extremely radio-faint. Their inferred magnetic fields are typically of order  $10^{13}$  G, placing them between canonical pulsars and magnetars in magnetic field strength.

The nature and evolutionary status of XDINS remain active areas of research. They are often interpreted as nearby, cooling neutron stars whose radio emission is either intrinsically weak or geometrically unfavourable for detection [? ]. Their placement in the  $P-\dot{P}$  diagram suggests a possible evolutionary link with high-magnetic-field radio pulsars and magnetars, although the details of this connection are still debated.

**White dwarf pulsars and other outliers.** A small number of objects have been proposed as “white dwarf pulsars”, occupying regions of the diagram at very long spin periods but comparatively small  $\dot{P}$ . While these sources do not follow the standard neutron star magnetic braking framework, their inclusion serves to illustrate the broader landscape of compact-object rotation and highlights the physical assumptions underlying the interpretation of the  $P-\dot{P}$  diagram [24].

**A unified view of neutron star evolution.** Taken as a whole, the  $P-\dot{P}$  diagram provides a unifying framework for understanding the diversity of neutron star manifestations. Distinct observational classes—ordinary pulsars, millisecond pulsars, magnetars, and XDINS—occupy well-defined regions of the diagram, reflecting differences in magnetic field strength, age, and evolutionary history. At the same time, the continuous distribution of sources suggests that these classes are not entirely disjoint, but instead represent different evolutionary pathways through a shared physical parameter space.

This population-level perspective will be particularly important in the context of this thesis, which focuses on millisecond pulsars in dense stellar environments.

The location of MSPs in the  $P-\dot{P}$  diagram encapsulates the cumulative effects of binary evolution, accretion, and long-term spin-down, and provides essential context for interpreting both new discoveries and non-detections in globular cluster pulsar surveys.

## 1.7 Radio Emission Properties of Pulsars

While the rotational and evolutionary properties of pulsars can be understood through their spin parameters, observational studies ultimately rely on the characteristics of the radiation emitted by these objects. Pulsars are most commonly detected and studied at radio wavelengths, where their emission is highly coherent, strongly polarised, and broadband in nature. Understanding the basic properties of pulsar radio emission is therefore essential for interpreting both detections and non-detections, as well as for placing observational results in a broader physical context.

This section introduces the fundamental concepts related to pulsar radio emission, including the spectral behaviour of the radiation, the definition of flux density, and the polarisation properties of pulsar signals. These topics are discussed at a qualitative and phenomenological level, providing the necessary background for later chapters without assuming prior familiarity with radio astronomy techniques.

### 1.7.1 Broadband Radio Emission from Pulsars

Pulsar radio emission is characterised by its broadband nature, extending over several orders of magnitude in frequency, from tens of megahertz to several gigahertz. Unlike thermal radiation, which has a characteristic spectral shape determined by temperature, pulsar radio emission is non-thermal and is widely believed to arise from coherent radiation processes within the neutron star magnetosphere [1, 40, 41].

Coherent emission implies that the radiated power scales with the square of the number of participating charges, allowing pulsars to produce extremely high brightness temperatures that cannot be explained by incoherent mechanisms such as synchrotron or thermal emission. Although the precise physical mechanism responsible for pulsar radio emission remains an open problem, a wide range of models invoke plasma processes operating along open magnetic field lines near the magnetic poles of the neutron star [42].

From an observational perspective, the broadband nature of pulsar emission means that pulsars can, in principle, be detected across a wide range of radio frequencies. However, the emitted power is not distributed uniformly with frequency,

and the observed brightness of a pulsar depends sensitively on the observing band.

### 1.7.2 Flux Density as an Observable Quantity

The fundamental observable used to quantify the brightness of a pulsar at radio wavelengths is the flux density. The flux density, denoted  $S_\nu$ , is defined as the received power per unit collecting area per unit bandwidth at a given observing frequency  $\nu$ , and is typically expressed in units of janskys ( $1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1}$ ).

For a pulsar observed at a frequency  $\nu$ , the measured flux density depends on several factors, including the intrinsic emission strength of the source, its distance from the observer, and propagation effects introduced by the ISM. Unlike intrinsic quantities such as spin period or magnetic field strength, the measured flux density is therefore not an inherent property of the neutron star alone, but an observational quantity affected by the intrinsic luminosity, distance, beaming geometry, and propagation effects such as scintillation. The ability to detect or accurately measure this flux density, however, depends on instrumental sensitivity and calibration.[38, 40].

Pulsar flux densities are also known to vary significantly with time. These variations can arise from changes in the emission process, geometric effects related to beam structure, or propagation effects such as interstellar scintillation, which can modulate the observed signal strength on timescales ranging from seconds to days [43, 44]. As a result, flux density measurements for pulsars are often subject to substantial uncertainties and should be interpreted with appropriate caution.

### 1.7.3 Spectral Index and Frequency Dependence

The frequency dependence of pulsar radio emission is commonly described using a power-law model for the flux density,

$$S_\nu \propto \nu^\alpha, \tag{1.16}$$

where  $\alpha$  is the spectral index. For most radio pulsars, the spectral index is negative, indicating that pulsars are generally brighter at lower radio frequencies. Typical values of  $\alpha$  lie in the range  $\alpha \sim -1.4$  to  $-1.8$ , although significant scatter exists within the population [45–47].

The spectral index provides a compact, phenomenological description of how pulsar brightness changes with observing frequency. Steep spectra imply that pulsars may be readily detectable at low frequencies but become increasingly faint at higher frequencies, while flatter spectra allow emission to remain detectable across a broader frequency range. Differences in spectral index have been reported between

various pulsar populations, including canonical pulsars and millisecond pulsars, although the underlying physical causes of these differences remain an active topic of investigation [48, 49].

It is important to emphasise that the power-law description is an approximation. Many pulsars exhibit spectral turnovers at low frequencies, spectral flattening or breaks at higher frequencies, or complex spectral shapes that cannot be captured by a single power-law index [47, 50]. Nevertheless, the spectral index remains a widely used and informative parameter in population studies and survey interpretation.

#### 1.7.4 Polarisation Properties of Pulsar Emission

One of the most distinctive features of pulsar radio emission is its high degree of polarisation. Pulsar signals are often strongly linearly polarised, with a significant circularly polarised component also commonly observed. The polarisation properties of pulsars provide important insights into the geometry of the emission region and the structure of the neutron star magnetosphere [17, 51].

The observed polarisation is typically described in terms of the Stokes parameters, which quantify the total intensity, linear polarisation, and circular polarisation of the radiation. Systematic variations in the polarisation position angle across the pulse profile are commonly observed and are often interpreted within the framework of geometrical models that relate the emission beam to the rotating magnetic field of the neutron star [17, 38].

Propagation effects within the magnetosphere and the ISM can further modify the observed polarisation signal. In particular, Faraday rotation, caused by the interaction of linearly polarised radiation with magnetised plasma along the line of sight, can rotate the plane of polarisation in a frequency-dependent manner [52]. While such effects can complicate the interpretation of polarisation measurements, they also provide valuable probes of magnetic fields in both the pulsar environment and the intervening ISM.

### 1.8 Pulsars as Astrophysical and Physical Laboratories

Beyond their intrinsic interest as exotic stellar remnants, pulsars occupy a unique position in modern astrophysics as precision tools for studying a wide range of physical phenomena. Their extreme properties—rapid rotation, intense magnetic fields, enormous densities, and exceptional rotational stability—allow pulsars to

probe regimes of physics that are inaccessible to terrestrial experiments or to other classes of astronomical objects.

This section provides a broad overview of the key scientific areas in which pulsars play a central role, highlighting why they are widely regarded as natural laboratories for fundamental physics, astrophysics, and cosmology.

### 1.8.1 Probing Fundamental Physics in Strong-Field Regimes

Neutron stars represent the densest stable form of matter known in the Universe, with central densities exceeding those found in atomic nuclei. As a result, pulsars provide a unique window into the behaviour of matter under extreme conditions, where strong nuclear forces, relativistic gravity, and quantum effects coexist [10, 22? ].

Binary pulsars, particularly those in compact or relativistic systems, enable stringent tests of gravitational theories in the strong-field regime. Measurements of orbital dynamics in such systems have provided some of the most precise confirmations of general relativity, including the indirect detection of gravitational radiation through orbital decay [53, 54]. Unlike Solar System tests, which probe weak gravitational fields, pulsars allow gravity to be tested in environments where relativistic effects are dominant.

In addition to gravity, pulsars offer constraints on the equation of state of ultra-dense matter. Observations of neutron star masses and radii place direct limits on the possible composition and stiffness of matter at supranuclear densities, thereby informing nuclear physics and particle physics models [22? ].

### 1.8.2 Magnetospheric Physics and Plasma Processes

Pulsars host some of the most extreme magnetic environments known, with surface magnetic field strengths ranging from  $10^8$  G in recycled millisecond pulsars to above  $10^{14}$  G in magnetars. These fields dominate the structure and dynamics of the surrounding magnetosphere, giving rise to relativistic plasma processes that remain only partially understood [15, 25? ].

The coherent radio emission observed from pulsars is a direct manifestation of these plasma processes and provides an observational handle on particle acceleration, pair production, and wave–particle interactions in strong magnetic fields. Studies of pulsar emission therefore contribute to a broader understanding of plasma physics under extreme conditions, with implications that extend beyond neutron stars to other magnetised astrophysical systems.

### 1.8.3 Stellar Evolution and Binary Interactions

Pulsars serve as key tracers of stellar evolution, particularly in binary systems. The existence of millisecond pulsars provides compelling evidence for long-term mass transfer and angular momentum exchange in close binaries, offering insights into accretion physics and binary stellar evolution [27, 28].

In dense stellar environments such as globular clusters, pulsars also act as probes of dynamical interactions, including exchange encounters, tidal capture, and binary disruption. The overabundance of millisecond pulsars in globular clusters relative to the Galactic field reflects the importance of stellar dynamics in shaping compact-object populations [55, 56]. These systems therefore link pulsar astrophysics to the study of stellar dynamics and cluster evolution.

### 1.8.4 The Interstellar Medium and Galactic Structure

As pulsar signals propagate from their source to the observer, they interact with the ionised interstellar medium (ISM). Frequency-dependent propagation effects such as dispersion, scattering, and scintillation encode information about the distribution and turbulence of free electrons along the line of sight [43, 57].

By studying large samples of pulsars distributed throughout the Galaxy, astronomers can reconstruct models of the Galactic electron density and probe the structure of the ISM on a wide range of spatial scales. Pulsars thus function as effective beacons that illuminate the otherwise invisible ionised component of the Milky Way.

### 1.8.5 Cosmological and Multi-Messenger Applications

The exceptional rotational stability of millisecond pulsars enables their use as ultra-precise celestial clocks. Ensembles of such pulsars form the basis of pulsar timing arrays, which are designed to detect low-frequency gravitational waves produced by supermassive black hole binaries and other cosmological sources [32, 58]. In this context, pulsars extend gravitational-wave astronomy into a frequency regime inaccessible to ground-based detectors.

Pulsars also play a growing role in multi-messenger astrophysics. Associations between pulsars, supernova remnants, high-energy neutrinos, and gamma-ray sources provide complementary views of energetic astrophysical processes and particle acceleration mechanisms across the electromagnetic spectrum [23, 24].

# Chapter 2

## Pulsar Science with Globular clusters

### 2.1 From Isolated Evolution to Dynamical Environments

Chapter 1.1 developed three ideas that are essential for understanding why globular star clusters are such productive environments for pulsar science. (i) Pulsars evolve in spin as they lose rotational energy, so their observed period  $P$  and period derivative  $\dot{P}$  encode physical evolution. (ii) Millisecond pulsars (MSPs) are typically interpreted as *recycled* neutron stars: they reach millisecond periods by gaining angular momentum during a past episode of mass transfer in a binary system. (iii) Binary motion imprints itself directly on observables, because orbital Doppler shifts modulate the apparent spin period and must be accounted for in searches and follow up studies.

Taken together, these points imply a simple but powerful conclusion: producing recycled pulsars is not only a question of how neutron stars evolve, but also a question of how efficiently an environment can *create, modify, and sustain the binary systems* that enable recycling.

In the Galactic field (a low stellar-density environment), binary systems are formed primarily during star formation and subsequently evolve largely in isolation. The probability that an old neutron star will acquire a new companion or that an existing binary will be repeatedly reconfigured by close stellar encounters is negligibly small over gigayear timescales. Consequently, the dominant pathways for producing recycled pulsars in the field are those available to primordial binaries, governed mainly by internal stellar evolution and mass transfer rather than by

external dynamical interactions [27].

Globular clusters (GCs) represent the opposite regime. They are old ( $\sim 10\text{--}13$  Gyr), gravitationally bound stellar systems containing  $\sim 10^4\text{--}10^6$  stars within characteristic radii of only a few parsecs [59]. Their defining feature is not merely their stellar content, but their extreme crowding in the central regions. Central stellar densities can reach values many orders of magnitude higher than in the Solar neighbourhood, leading to frequent gravitational encounters between stars and binaries over the cluster lifetime [60].

A useful concept in this context is the *relaxation time*: the timescale over which many weak gravitational encounters cumulatively redistribute orbital energies and drive processes such as mass segregation 2.2.4 . Unlike the Galactic field, many globular clusters are dynamically old systems in which relaxation has had sufficient time to strongly modify both stellar and binary populations.

This dynamical evolution changes pulsar demographics in two distinct ways. First, the cluster environment can *assemble* and *reconfigure* binaries through close encounters. A neutron star can acquire a new companion through an *exchange interaction*, in which it replaces one member of an existing binary. Encounters can also *harden* binaries (reduce their orbital separations), increasing the likelihood of mass transfer episodes, or disrupt wide “soft” binaries [61, 62]. Clusters therefore do not merely perturb binaries that already exist; they can actively create and re-wire the binary systems that enable recycling. These dynamical formation channels are strongly suppressed in the Galactic field because close encounters are exceedingly rare.

The impact of these processes is reflected directly in the observed pulsar populations of globular clusters. Figure 2.1 shows the current number of known radio pulsars hosted by individual Galactic globular clusters, separated into binary and isolated systems. These population statistics are compiled from the continuously updated online globular cluster pulsar catalogue maintained by Freire<sup>1</sup>, and therefore represent the census at the time of writing.

A comparison with the Galactic field highlights the remarkable efficiency of globular clusters as factories of millisecond pulsars (MSPs): despite containing less than  $\sim 1\%$  of the stellar mass of the Milky Way, they host a disproportionately large fraction of the known MSP population. In addition, the relative fractions of binary and isolated systems differ significantly between the two environments. While MSPs in the Galactic field are predominantly found in binary systems, globular clusters exhibit a substantially higher fraction of isolated pulsars. This trend, also visible in

---

<sup>1</sup><https://www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html>

359 pulsars in 46 clusters




Figure 2.1: Number of known radio pulsars hosted by Galactic globular clusters, separated into binary and isolated systems. The statistics are compiled from the online globular cluster pulsar catalogue maintained by P. C. C. Freire as of 18 march 2026.

Fig. 2.1, is naturally explained by the high rate of dynamical interactions in cluster cores, where close encounters can both form binaries and subsequently disrupt them, leaving behind isolated recycled pulsars.

The strong cluster-to-cluster variation in pulsar yield emphasizes that pulsar production is linked to underlying cluster properties such as central density and encounter activity [63, 64].

Globular clusters provide an observationally advantageous target class for pulsar surveys. A cluster occupies a compact, well-defined sky position, enabling deep and repeated integrations on the same stellar system. This facilitates the detection of faint pulsars and improves sensitivity to certain faint systems which require longer observation times to be detected.

Finally, the cluster environment affects how pulsar parameters are interpreted. A pulsar embedded in a globular cluster resides within the cluster gravitational potential and can experience a line-of-sight acceleration that contributes to the *observed*  $\dot{P}$  [65]. Separating intrinsic spin evolution from environmental contributions is therefore essential when inferring physical parameters, and conversely allows pulsars to be used as probes of cluster mass distributions.

This chapter develops the physical framework required to connect globular cluster properties to pulsar formation pathways, population trends, and observational

detectability, providing the environmental basis for the subsequent sections and chapters.

## 2.2 Structural Parameters and Dynamical Evolution of Globular Clusters

Section 2.1 motivated globular clusters as environments in which gravity-driven encounters can repeatedly reshape stellar systems. The purpose of this section is to define the structural parameters and dynamical timescales that quantify *how fast* clusters evolve and *how frequently* close encounters occur. These quantities will recur throughout the chapter when connecting cluster properties to compact-binary formation and pulsar populations.

### 2.2.1 From Structure to Dynamics: the basic ingredients

A globular cluster can be treated, to first approximation, as a self-gravitating system in approximate dynamical equilibrium. A widely used phenomenological description of the stellar distribution is provided by *lowered isothermal* (King) models [66]. The detailed model is not required for most of what follows, but it motivates a standard set of structural parameters used throughout the literature and tabulated for Galactic clusters [59].

The key quantities are: (i) a characteristic *size* (core radius  $r_c$  and half-light radius  $r_h$ ), (ii) a characteristic *density* (central mass density  $\rho_0$  or central number density  $n_0$ ), and (iii) a characteristic *speed* scale (velocity dispersion  $\sigma$ ). Together these set the relevant timescales and interaction rates.

Here, the core radius  $r_c$  is the radius within which the stellar density remains approximately constant before declining, while the half-light radius  $r_h$  is the radius enclosing half of the total cluster luminosity. The central mass density  $\rho_0$  and number density  $n_0$  describe the stellar density at the cluster centre in terms of mass and number of stars, respectively.

The velocity dispersion  $\sigma$  can be understood as the typical random speed of stars in the cluster. It is set by the depth of the cluster potential: a deeper potential well requires larger random motions to remain in equilibrium. This connection can be made quantitative using the virial theorem, which for a bound system implies

$$2T + U \simeq 0, \tag{2.1}$$

where  $T$  is the total kinetic energy and  $U$  is the total gravitational potential energy. For order-of-magnitude reasoning one may write  $T \sim \frac{1}{2}M\sigma^2$  and  $U \sim -\alpha GM^2/r$ , where  $M$  is the cluster mass,  $r$  is a characteristic radius (often comparable to  $r_h$ ), and  $\alpha$  is a structure-dependent factor of order unity. This yields the useful scaling

$$\sigma^2 \sim \frac{GM}{r}. \quad (2.2)$$

Thus, even without detailed modelling, one expects more massive and more compact clusters to have larger  $\sigma$ .

### 2.2.2 Crossing time: the “fast” dynamical clock

A first timescale is the *crossing time*, the time it takes a typical star to travel across a region of size  $r$  at speed  $\sigma$ :

$$t_{\text{cross}} \sim \frac{r}{\sigma}. \quad (2.3)$$

For  $r \sim 1$  pc and  $\sigma \sim 10$  km s<sup>-1</sup>,  $t_{\text{cross}} \sim 10^5$  yr. This is extremely short compared to GC ages ( $\sim 10^{10}$  yr), meaning that clusters have undergone an enormous number of crossings. Crossing time therefore measures how quickly the system responds dynamically to perturbations, but it does *not* by itself describe long-term evolution of the stellar distribution.

### 2.2.3 Two-body relaxation: why clusters evolve at all

The long-term evolution of a globular cluster is driven primarily by *two-body relaxation*: the cumulative effect of many weak gravitational deflections during close passages between stars [60, 67]. Each encounter typically changes a star’s velocity by a small amount, but the changes add up as a random walk (diffusion) in velocity space. The relaxation time is the timescale on which a star’s velocity changes by an amount comparable to itself, i.e. the time required for stars in the system to lose memory of their initial velocities.

A standard approximate expression for the relaxation time at radius  $r$  is

$$t_r \simeq \frac{0.34 \sigma^3}{G^2 m_\star^2 n \ln \Lambda}, \quad (2.4)$$

where  $m_\star$  is the typical stellar mass,  $n$  is the local number density, and  $\ln \Lambda$  is the Coulomb logarithm, which encodes the contribution of encounters over a wide range of impact parameters [60, 67].

This expression arises by considering the cumulative effect of many weak, long-

range two-body encounters. Individual encounters produce only small deflections in stellar velocities, but their combined effect can be treated as a diffusion process in velocity space. By integrating over the range of possible impact parameters (from close to distant encounters), one obtains the characteristic rate at which velocities change, leading to Eq. 2.4 (see also Mapelli 68).

Equation 2.4 makes the physics transparent:

- relaxation is *faster* at higher densities ( $t_r \propto 1/n$ ),
- relaxation is *slower* in high-velocity systems ( $t_r \propto \sigma^3$ ),
- smaller-mass perturbers are less effective ( $t_r \propto 1/m_\star^2$ ).

A commonly used global form relates relaxation to the crossing time:

$$t_r \sim \frac{0.1 N}{\ln N} t_{\text{cross}}, \quad (2.5)$$

where  $N$  is the number of stars [60]. This expression can be understood by noting that a star experiences a small deflection during each crossing of the cluster, and that many such crossings (of order  $N/\ln N$ ) are required for these deflections to accumulate into a significant change in its velocity. Thus, the relaxation time can be interpreted as the number of crossings needed for the system to lose memory of its initial conditions.

Physically,  $t_{\text{cross}}$  measures how rapidly stars move through the system, while  $t_r$  measures how long it takes for many weak encounters to significantly alter stellar velocities. The fact that  $t_r \gg t_{\text{cross}}$  in most stellar systems reflects that relaxation is a cumulative, diffusive process rather than the result of a single strong encounter.

To illustrate the dynamical significance of relaxation for real globular clusters, Fig. 2.2 shows the distribution of relaxation times tabulated in the Harris catalogue <sup>2</sup>. The blue histogram corresponds to the *core relaxation time*  $t_{rc}$ , evaluated at the cluster centre, while the orange histogram shows the *half-mass (evaluated at half-light radius) relaxation time*  $t_{rh}$  [59]. The dashed vertical line marks  $10^{10}$  yr, a representative globular cluster age. The key point is that  $t_{rc}$  is typically far shorter than the cluster age, implying that most cluster cores have undergone many relaxation cycles and are therefore dynamically evolved. The longer  $t_{rh}$  values reflect the lower densities at larger radii; clusters may be dynamically old in their cores even when their outer regions evolve more slowly. This separation of timescales is central to understanding why compact objects become centrally concentrated and why encounter-driven processes are most efficient in cluster cores.

<sup>2</sup><https://physics.mcmaster.ca/~harris/mwgc.dat>

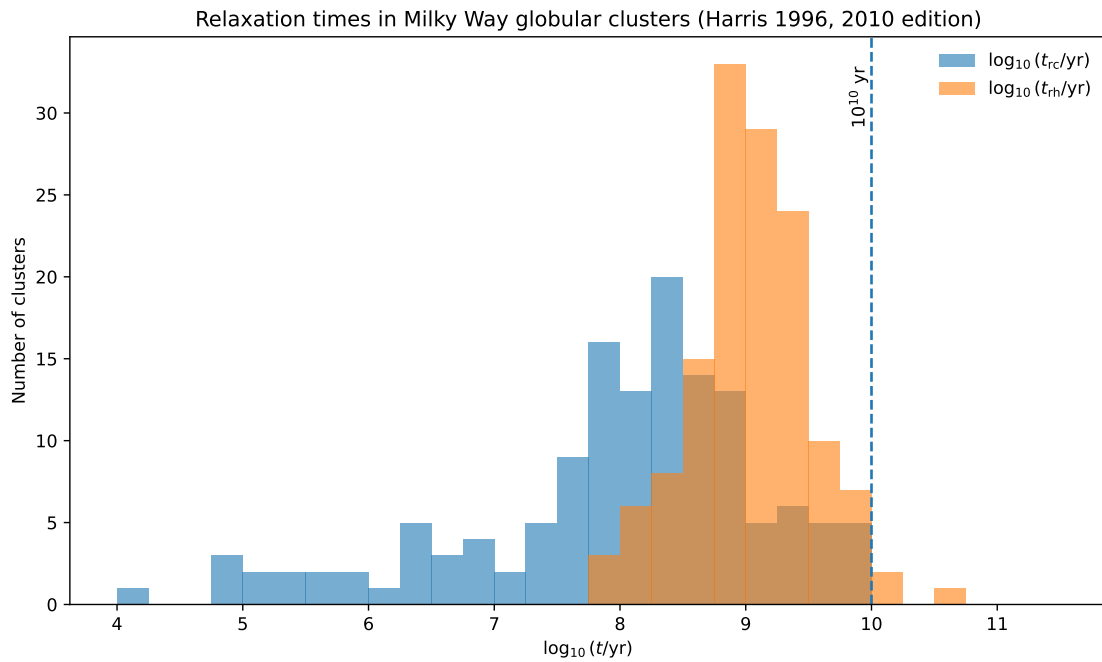


Figure 2.2: Distribution of core ( $t_{rc}$ ) and half-mass ( $t_{rh}$ ) relaxation times for Galactic globular clusters from the Harris (1996; 2010 edition) catalogue. The dashed line marks  $10^{10}$  yr for comparison with typical globular cluster ages. Cluster cores generally relax on timescales  $\ll 10^{10}$  yr, implying substantial dynamical evolution over a Hubble time.

## 2.2.4 Mass segregation and dynamical friction: why remnants collect in the core

Relaxation tends to redistribute kinetic energy between stars. In the limit of complete energy equipartition, one expects  $m\sigma_m^2 \approx \text{const}$ , where  $m$  is the stellar mass and  $\sigma_m$  is the velocity dispersion of stars of mass  $m$ , so heavier stars have smaller velocity dispersions. Even though clusters do not achieve perfect equipartition in general, the qualitative outcome is robust: heavier objects systematically lose kinetic energy and drift inward, while lighter stars preferentially populate larger radii. This process is known as *mass segregation* [60, 67].

For compact remnants (white dwarfs, neutron stars, and stellar-mass black holes), mass segregation implies an enhanced central concentration. This is one of the most important dynamical facts for pulsar science: the objects most relevant for compact binary formation are preferentially delivered into the high-density core where encounter rates are maximized.

## 2.2.5 Close encounters and gravitational focusing

To connect structure to interaction rates, we need a quantitative notion of what it means for two stars to have a “close encounter”. Consider two objects approaching with relative speed  $v_\infty$  at large separation and with an impact parameter  $b$ . Pure geometry would suggest an interaction cross section  $\Sigma \sim \pi b^2$ . In gravitational systems, however, trajectories are bent: gravity focuses passing stars toward smaller separations.

For an encounter to reach a periapsis distance  $r_{\min}$ , one can show that the effective cross section is approximately

$$\Sigma(r_{\min}) \simeq \pi r_{\min}^2 \left( 1 + \frac{2G(m_1 + m_2)}{r_{\min} v_\infty^2} \right), \quad (2.6)$$

where  $m_1$  and  $m_2$  are the masses of the interacting objects and  $v_\infty$  is their relative velocity at large separation (i.e. before the encounter) [60].

The periapsis distance  $r_{\min}$  is the relevant quantity because the strength and outcome of an encounter are determined by the closest approach between the two objects. Processes such as tidal interactions, physical collisions, or binary perturbations depend sensitively on this minimum separation rather than on the initial impact parameter.

The second term represents gravitational focusing. When  $v_\infty$  is small, gravitational attraction significantly deflects trajectories, increasing the likelihood of close

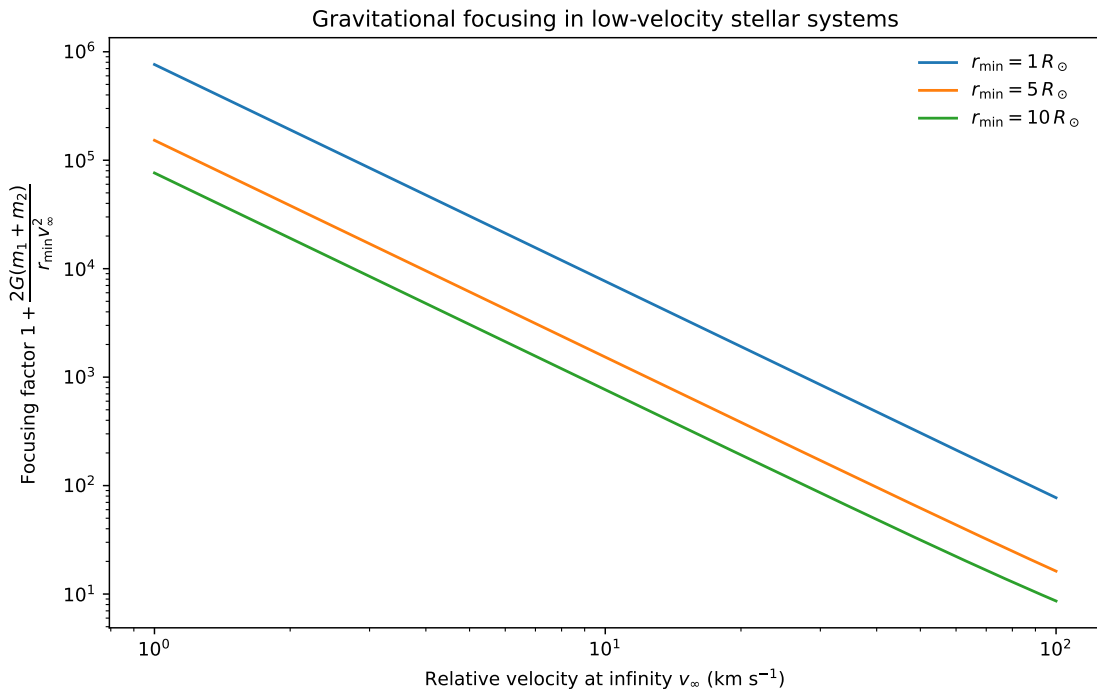


Figure 2.3: Gravitational focusing enhancement factor,  $1 + 2G(m_1 + m_2)/(r_{\min} v_\infty^2)$ , as a function of encounter velocity at infinity  $v_\infty$  for representative periapsis scales  $r_{\min}$ . Low velocity dispersions typical of globular clusters strongly increase the effective cross section for close encounters relative to the geometric expectation.

encounters and making the effective cross section much larger than the purely geometric one.

The importance of gravitational focusing for globular clusters is illustrated in Fig. 2.3, which shows the focusing enhancement factor appearing in Eq. 2.6 as a function of encounter velocity at infinity  $v_\infty$  for representative periapsis scales  $r_{\min}$ . Because the enhancement scales as  $v_\infty^{-2}$ , it increases rapidly in the low-velocity regime characteristic of globular clusters ( $\sigma \sim$  a few to  $\sim 10$  km s<sup>-1</sup>). In this regime, close encounters are not governed purely by geometric cross sections; instead, the effective interaction probability is substantially amplified by gravitational focusing.

The encounter *rate* per target object is then roughly

$$\mathcal{R} \sim n \Sigma v_\infty, \quad (2.7)$$

with  $n$  the number density of potential perturbers. Equations 2.6–2.7 summarize the core reason globular clusters are encounter-rich environments: high  $n$  increases the number of potential encounters, and low  $v_\infty$  enhances gravitational focusing.

## 2.2.6 The encounter rate parameter $\Gamma$

For population studies it is convenient to define a single quantity that scales with the global rate of close encounters in the cluster core. A widely used proxy is the *encounter rate parameter*  $\Gamma$  [62, 63]. Its standard scaling can be motivated in a few steps.

The total number of encounters per unit time in a volume  $V$  is schematically

$$\frac{dN_{\text{enc}}}{dt} \sim \int_V n^2 \langle \Sigma v \rangle dV, \quad (2.8)$$

because encounters involve *pairs*, giving the  $n^2$  dependence. If we approximate the core as having roughly uniform density  $n_0$  and characteristic speed  $\sigma$  within a volume  $V_c \sim r_c^3$ , then

$$\frac{dN_{\text{enc}}}{dt} \propto n_0^2 r_c^3 \langle \Sigma v \rangle. \quad (2.9)$$

In regimes where gravitational focusing dominates, Eq. 2.6 implies  $\Sigma \propto 1/v^2$ , so  $\Sigma v \propto 1/v$ . Identifying  $v$  with the velocity dispersion  $\sigma$  gives the commonly quoted scaling

$$\Gamma \propto \frac{\rho_0^2 r_c^3}{\sigma}, \quad (2.10)$$

where  $\rho_0$  is used instead of  $n_0$  (they differ by a factor of the typical stellar mass). The purpose of Eq. 2.10 is not precision, but clarity: clusters with higher central densities and larger cores have much larger encounter rates, while higher velocity dispersions reduce focusing and therefore reduce the interaction rate.

Empirically,  $\Gamma$  correlates with the abundance of dynamically formed compact binaries (notably X-ray binaries), supporting the interpretation that close encounters are a dominant formation channel in dense clusters [63].

## 2.2.7 Binary processing, core collapse, and binary heating

Two-body relaxation drives an outward transport of energy: the cluster core tends to lose energy and contract, while the outer regions gain energy and expand. In the absence of an internal energy source, this process can lead to a runaway increase in central density, historically described as the gravothermal catastrophe and observationally associated with core-collapsed clusters [60, 67].

Observationally, core collapse is identified through the surface brightness profile of a cluster, which shows a steep central cusp rather than a flat core. Such clusters lack a well-defined core radius and exhibit a continuous rise in brightness toward the centre, indicating a very high central stellar density.

An important internal energy source is provided by binaries. Consider a binary with component masses  $m_1$  and  $m_2$  and semi-major axis  $a$ . Its binding energy is

$$E_b = \frac{Gm_1m_2}{2a}. \quad (2.11)$$

This can be compared to the characteristic kinetic energy of a typical cluster star,

$$E_{\text{kin}} \sim \frac{1}{2}m_\star\sigma^2, \quad (2.12)$$

where  $m_\star$  is the typical stellar mass and  $\sigma$  the velocity dispersion.

The so-called "hard-soft" boundary is obtained by equating these two energies,

$$\frac{Gm_1m_2}{2a_{\text{hs}}} \sim \frac{1}{2}m_\star\sigma^2, \quad (2.13)$$

which gives the characteristic semi-major axis

$$a_{\text{hs}} \sim \frac{Gm_1m_2}{m_\star\sigma^2}. \quad (2.14)$$

Binaries with  $a < a_{\text{hs}}$  are called *hard*: their binding energy exceeds the typical kinetic energy of field stars. Binaries with  $a > a_{\text{hs}}$  are *soft*. This distinction is dynamically crucial. Encounters tend to drive systems away from the boundary: soft binaries are easily disrupted, while hard binaries tend to become more tightly bound through interactions [61]. This empirical rule is often summarized as Heggie's law: "hard binaries get harder; soft binaries get softer."

In interactions involving hard binaries, part of the binary's binding energy is transferred to passing stars as kinetic energy. The intruding star leaves the encounter with increased velocity, while the binary becomes more tightly bound (smaller  $a$ ). This process injects energy into the surrounding stellar population and is therefore referred to as binary heating. Binary heating can partially counteract relaxation-driven core contraction and plays a central role in regulating long-term cluster evolution.

## 2.3 Neutron Stars and the Dynamical Formation of Recycled Pulsars

Section 2.2 established that globular cluster cores are dynamically evolved and encounter-rich: their relaxation times are typically far shorter than the cluster age and gravitational focusing enhances close-interaction cross sections at the low veloc-

ity dispersions characteristic of globular clusters. In this section we connect these dynamical ingredients to neutron stars and pulsars. Two questions must be answered before discussing millisecond pulsars (MSPs) in clusters: (i) how neutron stars are retained in globular clusters despite natal kicks (velocity impulses imparted to neutron stars at birth due to asymmetries in supernova explosions.), and (ii) how dense cluster environments assemble and process the binaries required for recycling NSs into MSPs.

### 2.3.1 Natal kicks and the neutron-star retention problem

Neutron stars are born in supernova explosions. Observations of pulsar proper motions imply that many neutron stars receive substantial *natal kicks* at birth, with characteristic speeds of hundreds of  $\text{km s}^{-1}$  [4]. This immediately creates a tension with globular cluster potentials: the escape velocities of globular clusters are typically only tens of  $\text{km s}^{-1}$ , so a naive expectation would be that most newborn neutron stars are ejected.

A simple way to see the scale of the problem is to compare a characteristic escape speed to a characteristic kick speed. For a cluster of mass  $M$  and characteristic radius  $r$ , the escape speed scales as

$$v_{\text{esc}} \sim \left( \frac{2GM}{r} \right)^{1/2}. \quad (2.15)$$

For  $M \sim 10^5\text{--}10^6 M_{\odot}$  and  $r \sim$  a few pc, this yields  $v_{\text{esc}}$  of order  $\sim 30\text{--}80 \text{ km s}^{-1}$ , depending on concentration and where in the cluster the neutron star is born. This is far below the high-velocity tail of the natal-kick distribution.

The *retention fraction* is the fraction of neutron stars formed with speeds below the relevant escape speed. If the birth-speed distribution  $f(v)$  is known, one may write schematically

$$f_{\text{ret}} \simeq \int_0^{v_{\text{esc}}} f(v) dv, \quad (2.16)$$

emphasizing that retention is highly sensitive to the low-velocity end of the kick distribution and to the cluster escape speed. In practice, additional factors matter, including whether the supernova occurs in a binary, because the post-supernova centre-of-mass velocity depends on both the natal kick and the binary orbital motion.

To make Eq. 2.16 intuitive, it is useful to think of natal kicks as a *probability distribution of birth speeds*. When a neutron star is formed, it may receive a recoil velocity (a “kick”) due to asymmetries in the supernova explosion. We generally do not know the kick of an individual neutron star a priori, but we can describe

the population statistically by a speed distribution  $f(v)$ . A simple and widely used choice is a Maxwellian distribution for the three-dimensional speed, which arises if the kick velocity components are independent Gaussian random variables. This form is also supported empirically by pulsar proper motion measurements [4]. In this case, a single parameter  $\sigma_k$  (the one-dimensional dispersion) sets the overall scale of the kicks.

The retained fraction in a globular cluster is then simply the probability that the birth speed falls below the relevant escape speed  $v_{\text{esc}}$ , i.e. the cumulative probability  $f_{\text{ret}}(v_{\text{esc}}) = \int_0^{v_{\text{esc}}} f(v) dv$ . Figure 2.4 plots this cumulative retained fraction as a function of  $v_{\text{esc}}$  for two illustrative kick distributions. The “high-kick” curve adopts the Maxwellian model inferred for ordinary field pulsars by Hobbs et al. [4]. The “low-kick” curve is not presented as a unique physical model, but as a schematic example showing how retention changes if a subset of neutron stars form with substantially smaller kicks, as expected in some alternative supernova channels (e.g. electron-capture or ultra-stripped supernovae) [5, 6]. The key message is that neutron-star retention is controlled by the low-velocity tail of the kick distribution: modest changes to the kick physics (or the presence of a low-kick component) can increase the retained population by orders of magnitude for typical globular-cluster escape speeds.

Several mechanisms can enhance neutron-star retention in globular clusters. A key idea is that not all neutron stars are formed in the same way. Alternative formation channels such as electron-capture supernovae (ECS) and ultra-stripped supernovae can produce significantly lower kick velocities than canonical core-collapse explosions. In ECS, a degenerate O–Ne–Mg core collapses once electron captures onto nuclei remove pressure support, leading to a relatively symmetric and low-energy explosion. Ultra-stripped supernovae, on the other hand, occur in close binaries where prior mass transfer removes most of the stellar envelope, leaving only a small amount of ejecta at collapse; the reduced ejecta mass leads to weaker recoil kicks.[5, 6]. Even a modest low-kick component can greatly increase the number of neutron stars retained by a cluster, because  $f_{\text{ret}}$  depends on the *integrated* probability below  $v_{\text{esc}}$  (Eq. 2.16).

The observational existence of large pulsar populations in some clusters therefore suggests that (i) retention is non-negligible in at least a subset of clusters, and (ii) the neutron stars that are retained are subsequently delivered to the cluster core and efficiently recycled through binary interactions, as discussed below.

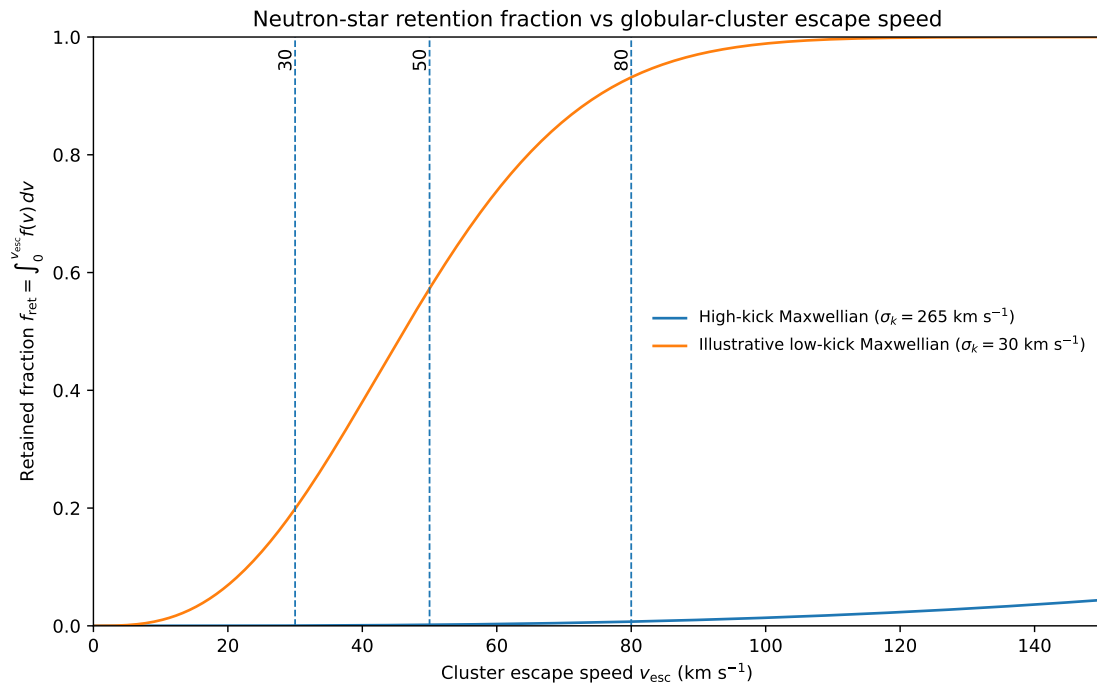


Figure 2.4: Retained neutron-star fraction as a function of cluster escape speed,  $f_{\text{ret}}(v_{\text{esc}}) = \int_0^{v_{\text{esc}}} f(v) dv$ , for two illustrative natal-kick distributions. The high-kick curve adopts a Maxwellian with one-dimensional dispersion  $\sigma_k = 265 \text{ km s}^{-1}$  inferred for ordinary field pulsars [4]. The low-kick curve is included for comparison to illustrate how the presence of low-kick formation channels (e.g. electron-capture or ultra-stripped supernovae) can substantially increase retention in globular clusters [5, 6]. Vertical lines indicate representative globular-cluster escape speeds.

### 2.3.2 Mass segregation and the central concentration of compact remnants

Once neutron stars are retained, they do not remain uniformly distributed. As described in Section 2.2, two-body relaxation drives *mass segregation*: heavier objects tend to lose kinetic energy to lighter stars and drift inward on a timescale comparable to the local relaxation time [60, 67]. Since neutron stars are more massive than the typical cluster star, they are preferentially concentrated in the core.

This central concentration has two important consequences. First, it increases the interaction probability of neutron stars by placing them in the densest region where the encounter rate parameter  $\Gamma$  is largest. Second, it increases the likelihood that neutron stars participate in binary-mediated interactions, because binaries in cluster cores act as large gravitational targets and are continuously processed by encounters.

Thus, even if the global neutron-star retention fraction is not large, the neutron stars that *are* retained can form a dynamically active central reservoir. This reservoir is the starting point for producing compact binaries and recycled pulsars.

### 2.3.3 Binary processing and the pathways to recycling in clusters

Recycling requires sustained mass transfer in a binary system. In the Galactic field, the canonical pathway to form a fully recycled millisecond pulsar proceeds through the evolution of a primordial binary, ultimately producing a neutron star with a low-mass companion that becomes a white dwarf after the accretion phase [27, 28]. In globular clusters, the crucial difference is that binaries are not only evolved, but also *assembled and modified dynamically*.

Hard binaries tend to survive and become harder in encounters, while soft binaries are preferentially disrupted [61, 69]. This explains why cluster cores can sustain a population of compact, tightly bound binaries that are particularly relevant for recycling neutron stars.

#### Exchange interactions and dynamical assembly

In a dense core, encounters between a binary and a single star are common. One of the most important outcomes is an *exchange interaction*: the incoming star replaces one of the binary components. Because gravitational interactions favour more massive objects binding more strongly, exchange interactions tend to place the most

massive participant into the final binary. As a result, neutron stars (and other remnants) can efficiently acquire companions even if they were born single, or can swap into binaries in place of lighter stars [69, 70].

This mechanism provides a natural route to producing binaries in which a neutron star is paired with a companion capable of driving an accretion episode. It also produces systems whose companions and orbital configurations need not resemble the outcomes of isolated binary evolution. The same dynamical processing can yield repeated episodes of binary formation and disruption, making the cluster environment capable of “recycling” neutron stars through multiple dynamical stages.

### **Hardening and the onset of mass transfer**

Even when a neutron star is already in a binary, repeated encounters can change the orbital separation. Hardening interactions shrink  $a$ , increasing the likelihood that the companion fills its Roche lobe and initiates mass transfer. In this way, stellar dynamics can accelerate the formation of interacting binaries compared to the field. This is one reason globular clusters are efficient producers of low-mass X-ray binaries, widely interpreted as direct progenitors of MSPs [63, 71].

### **Implications for pulsar populations**

The combined effect of (i) non-negligible retention, (ii) mass segregation into dense cores, and (iii) dynamical assembly and processing of binaries, produces pulsar populations that differ from the Galactic field. Clusters preferentially host recycled pulsars and a rich variety of binary configurations, including systems that are difficult to produce through isolated evolution alone.

## **2.3.4 Globular Clusters as Laboratories for Exotic Pulsar Systems**

The dynamical processes described in the previous sections do not merely modify the pulsar population quantitatively; they qualitatively reshape it. Globular clusters host a collection of pulsar systems that are rare or dynamically difficult to produce in the Galactic field. This diversity reflects the interplay between neutron-star retention, mass segregation, and repeated dynamical interactions in dense cluster cores [72–74].

One of the most striking observational facts is that globular clusters are extraordinarily efficient factories of millisecond pulsars (MSPs). Although globular clusters contain less than one percent of the stellar mass of the Milky Way, they

host a disproportionately large fraction of the known MSP population [72, 73]. This overabundance is widely interpreted as evidence that dynamical formation channels—particularly exchange interactions and binary hardening—are highly effective at producing recycled systems in dense stellar environments.

Beyond their sheer numbers, the *types* of pulsars found in clusters are particularly revealing. Many cluster MSPs reside in compact binaries with very low-mass companions undergoing active ablation by the pulsar wind. These systems are commonly referred to as “black widow” and “redback” pulsars and are thought to represent transitional evolutionary phases between accretion-powered and rotation-powered states [75]. The high stellar interaction rates in cluster cores likely facilitate the formation, destruction, and reshuffling of such tight binaries.

Clusters also host systems with unusually high orbital eccentricities for recycled pulsars. In the Galactic field, prolonged mass transfer during the recycling process typically circularises binary orbits. The presence of eccentric MSP–white dwarf systems in globular clusters is therefore interpreted as a signature of post-recycling dynamical encounters, such as close fly-bys or exchange interactions with other cluster members [36, 76]. These systems provide direct observational evidence that binaries in clusters can continue to evolve dynamically long after the initial recycling phase.

Even more extreme systems involve unusually massive or compact companions. A particularly intriguing example is PSR J0514–4002E in the globular cluster NGC 1851, whose timing measurements suggest the presence of a companion in the stellar-mass black-hole regime [77]. If confirmed, such systems would provide unique probes of black-hole populations in globular clusters and offer laboratories for studying relativistic dynamics in strongly interacting stellar environments.

Globular clusters have also produced some of the most relativistic binary pulsars known. The discovery of the double neutron star system PSR J0737–3039 in the Galactic field demonstrated the extraordinary precision with which pulsars can be used to test general relativity [78, 79]. In dense stellar systems, exchange interactions can in principle assemble similarly compact binaries containing two neutron stars or other massive companions. Such systems would provide exceptionally clean laboratories for relativistic gravity and binary evolution under dynamical conditions.

In addition to compact binaries, globular clusters host a population of isolated MSPs whose formation likely involved past binary disruption. These objects may represent the remnants of exchange interactions or binary destruction events, again highlighting the importance of cluster dynamics in shaping the observable pulsar population.

Taken together, the pulsar systems observed in globular clusters illustrate a fundamental point: dense stellar environments do not simply accelerate canonical binary evolution; they open entirely new evolutionary pathways. Exchange interactions can replace companions, perturb orbital parameters, and assemble configurations that are statistically improbable in the Galactic field. As a result, globular clusters function as natural laboratories for studying stellar dynamics, compact-object retention, relativistic binaries, and the late stages of stellar evolution.

This rich phenomenology provides a strong observational motivation for targeted pulsar searches in globular clusters. By probing deeper into these systems—particularly at higher radio frequencies where dispersion and scattering are mitigated—we gain access to a pulsar population shaped not only by stellar evolution, but also by the full complexity of long-term gravitational dynamics.

## 2.4 Pulsars as Probes of Globular Cluster Gravitational Potentials

One of the most powerful reasons that pulsars are useful in astrophysics is that they behave as exceptionally stable clocks. In practice, they are not clocks in the sense of emitting identical single pulses one after another—individual pulses often vary significantly in amplitude and morphology—but the *average* pulse profile formed by integrating many successive rotations is remarkably stable. This stability allows the pulse train of a pulsar to be used as a precise timing signal. Once a pulsar is detected repeatedly, one can build a phase-coherent model for its rotation and, where relevant, its binary motion. Small discrepancies between the predicted and measured arrival times of the pulses then reveal dynamical effects of extraordinary subtlety [12, 38, 40].

For pulsars located in globular clusters, this timing precision opens an especially interesting possibility. A pulsar in a globular cluster does not move freely in empty space: it resides within the gravitational potential of the cluster, and therefore experiences a non-negligible acceleration. This acceleration alters the apparent spin evolution of the pulsar and can also affect the apparent orbital evolution in binary systems. At first sight this is a nuisance, because it complicates the interpretation of the measured spin period derivative. However, the same effect makes pulsars valuable probes of the gravitational field of the cluster. In other words, pulsars in globular clusters can be used as precision accelerometers embedded within the cluster potential [36, 65, 80].

### 2.4.1 Minimal timing concepts: pulse times of arrival and phase-coherent models

Because pulsar timing is central to this discussion, it is useful to introduce the minimum set of concepts needed here, without attempting a full review of the subject. The fundamental observable in pulsar timing is the *time of arrival* (TOA). A TOA is the measured arrival time of a fiducial phase point in the averaged pulse profile recorded during a given observation. In practice, the observed profile is cross-correlated with a high signal-to-noise template, and the phase offset between the two is converted into an arrival time. The TOA is therefore not the arrival time of a single radio burst, but rather the arrival time of a well-defined reference phase in the stable integrated pulse profile [38, 40].

A timing model is then constructed to predict the rotational phase of the pulsar at any given time. If the pulsar is isolated and its spin evolution is smooth, the pulse phase can be expanded as a Taylor series around a reference epoch  $t_0$ ,

$$\phi(t) = \phi_0 + f_0(t - t_0) + \frac{1}{2}\dot{f}_0(t - t_0)^2 + \frac{1}{6}\ddot{f}_0(t - t_0)^3 + \dots, \quad (2.17)$$

where  $\phi(t)$  is the rotational phase,  $f$  is the spin frequency, and  $\dot{f}$ ,  $\ddot{f}$  are its time derivatives. A *phase-coherent* timing solution is one that accounts correctly for every single rotation of the neutron star across the entire data span. When such a solution is achieved, timing residuals—the differences between the measured TOAs and those predicted by the model—can be driven down to the microsecond level or better for many millisecond pulsars. This makes pulsars exquisitely sensitive probes of small accelerations, orbital perturbations, and propagation effects [40, 81, 82].

The parameter of immediate relevance for the present discussion is the spin period derivative  $\dot{P}$ , or equivalently the spin frequency derivative  $\dot{f}$ . In the simplest case of an isolated pulsar in the Galactic field (where the galactic potential affect the measured  $\dot{P}$  to a much lesser extent than GCs)  $\dot{P}$  is interpreted as the secular spin-down of the neutron star due to the loss of rotational energy. In a globular cluster, however, the observed  $\dot{P}$  need not be purely intrinsic. It can instead contain a substantial dynamical contribution due to the acceleration of the pulsar in the cluster potential.

### 2.4.2 From Doppler shifts to apparent spin derivatives

The origin of this dynamical contribution can be understood from a simple Doppler argument. Let the intrinsic spin frequency of the pulsar be  $f_{\text{int}}$ , and let the pulsar

have a line-of-sight velocity  $v_{\parallel}$  relative to the observer. In the non-relativistic limit  $|v_{\parallel}| \ll c$ , the observed spin frequency is

$$f_{\text{obs}} \simeq f_{\text{int}} \left(1 - \frac{v_{\parallel}}{c}\right), \quad (2.18)$$

where the sign convention is such that positive  $v_{\parallel}$  corresponds to motion away from the observer.

Differentiating Eq. (2.18) with respect to time gives

$$\dot{f}_{\text{obs}} = \dot{f}_{\text{int}} \left(1 - \frac{v_{\parallel}}{c}\right) - f_{\text{int}} \frac{a_{\parallel}}{c}, \quad (2.19)$$

where  $a_{\parallel} = \dot{v}_{\parallel}$  is the line-of-sight acceleration. Since  $|v_{\parallel}|/c \ll 1$ , we may write

$$\dot{f}_{\text{obs}} \simeq \dot{f}_{\text{int}} - f \frac{a_{\parallel}}{c}. \quad (2.20)$$

Dividing by  $f$  yields

$$\frac{\dot{f}_{\text{obs}}}{f} \simeq \frac{\dot{f}_{\text{int}}}{f} - \frac{a_{\parallel}}{c}. \quad (2.21)$$

Because  $P = 1/f$ , the relation between period and frequency derivatives is

$$\dot{P} = -\frac{\dot{f}}{f^2}, \quad \Rightarrow \quad \frac{\dot{P}}{P} = -\frac{\dot{f}}{f}. \quad (2.22)$$

Using this identity gives

$$\left(\frac{\dot{P}}{P}\right)_{\text{obs}} \simeq \left(\frac{\dot{P}}{P}\right)_{\text{int}} + \frac{a_{\parallel}}{c}. \quad (2.23)$$

This same reasoning applies to binary orbital periods,

$$\left(\frac{\dot{P}_b}{P_b}\right)_{\text{obs}} \simeq \left(\frac{\dot{P}_b}{P_b}\right)_{\text{int}} + \frac{a_{\parallel}}{c}. \quad (2.24)$$

### 2.4.3 Observed spin derivatives: intrinsic and kinematic contributions

For a pulsar in a globular cluster, the observed spin period derivative generally contains several contributions,

$$\frac{\dot{P}_{\text{obs}}}{P} = \frac{\dot{P}_{\text{int}}}{P} + \frac{\mu^2 D}{c} + \frac{a_{\parallel, \text{GC}}}{c} + \frac{a_{\text{Gal}}}{c}. \quad (2.25)$$

where  $\dot{P}_{\text{int}}$  is the intrinsic spin-down of the pulsar,  $\mu$  is the total proper motion,  $D$  is the cluster distance,  $a_{\parallel, \text{GC}}$  is the line-of-sight acceleration produced by the globular cluster potential, and  $a_{\text{Gal}}$  is the differential Galactic acceleration between the Solar System and the cluster [36, 83, 84].

#### 2.4.4 Pulsars as accelerometers in the cluster potential

Rearranging Eq. (2.23) immediately gives

$$a_{\parallel} = c \left[ \left( \frac{\dot{P}}{P} \right)_{\text{obs}} - \left( \frac{\dot{P}}{P} \right)_{\text{int}} \right]. \quad (2.26)$$

Thus, once the intrinsic and kinematic contributions are estimated or bounded, the line-of-sight gravitational acceleration can be inferred directly from pulsar timing. In this sense, a pulsar functions as a precision accelerometer.

To connect this measured acceleration to the cluster mass distribution, consider a spherically symmetric gravitational potential. Let the pulsar lie at a true distance  $r$  from the cluster centre. The magnitude of the gravitational acceleration is then

$$a(r) = \frac{GM(< r)}{r^2}, \quad (2.27)$$

where  $M(< r)$  is the mass enclosed within radius  $r$ . However, timing does not measure the full three-dimensional acceleration vector; it measures only the component along the line of sight. If the pulsar is displaced from the plane of the sky containing the cluster centre by a line-of-sight coordinate  $z$ , and has projected separation  $R_{\perp}$ , then

$$r^2 = R_{\perp}^2 + z^2. \quad (2.28)$$

For a spherically symmetric field directed toward the cluster centre, the line-of-sight component is

$$a_l(R_{\perp}, z) = -\frac{GM(< r)}{r^2} \frac{z}{r} = -\frac{GM(< r)z}{r^3}. \quad (2.29)$$

This expression is useful because it immediately clarifies two points.

First, the *sign* of the acceleration depends on the sign of  $z$ . A pulsar lying behind the cluster centre ( $z > 0$ , with the positive direction defined away from the observer) will be accelerated toward the centre and hence toward the observer, leading to a negative  $a_l$ . A pulsar lying in front of the cluster centre ( $z < 0$ ) will be accelerated away from the observer, leading to a positive  $a_l$ . Thus, the sign of the observed  $\dot{P}$

contribution reflects the pulsar’s unknown position along the line of sight.

Second, for a given projected separation  $R_{\perp}$ , the exact value of  $a_l$  cannot be known unless  $z$  is known. This is generally not the case. Therefore, for an individual pulsar one usually infers either the acceleration itself or an upper limit on it, but not its full three-dimensional location inside the cluster.

Nevertheless, Eq. (2.29) shows that the acceleration depends directly on the enclosed mass. In the idealised limit where  $r$  is known and the intrinsic spin-down is negligible, one may formally write

$$M(< r) = \frac{a(r) r^2}{G} = \frac{c r^2}{G} \left( \frac{\dot{P}}{P} \right)_{\text{cluster}}. \quad (2.30)$$

Although this expression is usually too idealised to apply literally to a real globular cluster pulsar, it captures the essential physical point: pulsar timing provides direct dynamical information about the cluster mass distribution.

### 2.4.5 Projected separations, angular offsets, and acceleration envelopes

Observationally, the pulsar position relative to the cluster centre is measured as an angular offset on the sky. If the angular separation is  $\theta_{\perp}$  and the cluster lies at distance  $D$ , then the projected physical separation is

$$R_{\perp} = D \theta_{\perp}. \quad (2.31)$$

Because the line-of-sight depth  $z$  is not known, the true distance from the cluster centre remains uncertain even when  $R_{\perp}$  is measured precisely. This uncertainty motivates the use of acceleration envelopes. For a given cluster mass model, one can calculate the maximum possible line-of-sight acceleration that a pulsar could experience at a given projected separation.

This idea is particularly intuitive if one imagines the cluster as generating a “gravitational bowl”. Near the centre the potential is deepest, and pulsars can experience larger accelerations. Farther out, the potential becomes shallower and the maximum allowed line-of-sight acceleration decreases. A plot of  $|a_l|$ —or equivalently  $|\dot{P}/P|$  after the necessary corrections—against projected radius therefore provides a compact visual summary of how the pulsar ensemble samples the cluster potential [65, 80, 84].

The use of the absolute value deserves brief comment. Since the sign of  $a_l$  depends on whether the pulsar lies in front of or behind the cluster centre, and this

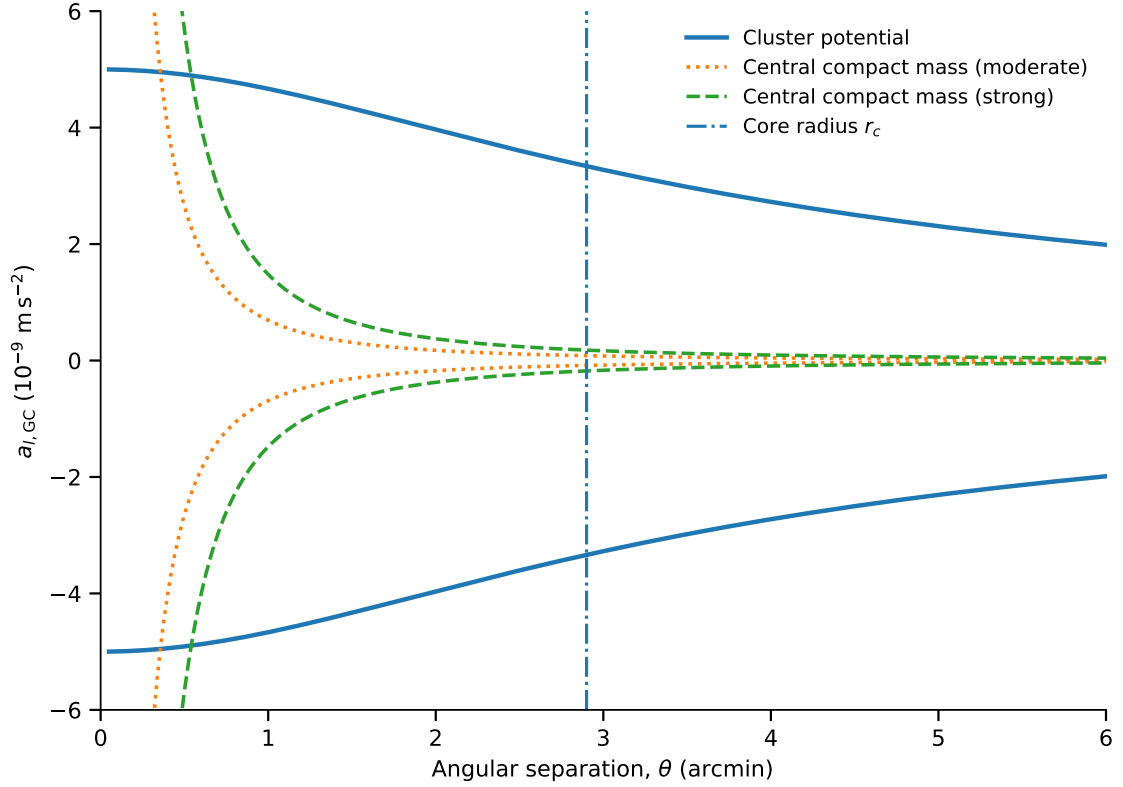


Figure 2.5: Conceptual illustration of the maximum and minimum line-of-sight accelerations expected for pulsars as a function of angular separation from the cluster centre. The solid curves represent the LOS acceleration envelope produced by a standard cluster potential, while the dotted and dashed curves illustrate the additional inner contribution that could arise from increasingly compact central mass components. The vertical line marks the cluster core radius  $r_c$ . Comparisons between timing-derived pulsar accelerations and such envelopes provide a direct dynamical probe of the cluster mass distribution.

depth is generally unknown, the most useful comparison with a static mass model is often made using  $|a_l|$  or  $|\dot{P}/P|$ . The absolute value therefore captures the *strength* of the gravitational field while avoiding degeneracy associated with the unknown line-of-sight position.

In practice, pulsars do not generally sit exactly on the maximum envelope. Most lie below it, because they sample a range of depths  $z$  and therefore only rarely attain the maximum possible line-of-sight acceleration at their projected separation. However, the outer boundary traced by the population can place powerful constraints on the central mass density and on the overall cluster potential.

## 2.4.6 Central mass concentrations and possible black holes

The inner regions of a globular cluster are of particular interest because they can host concentrated populations of dark remnants, including white dwarfs, neutron stars, and stellar-mass black holes. Any such unseen component steepens the central potential and can increase the maximum allowed accelerations near the cluster centre. Pulsars projected close to the core are therefore especially valuable probes of the central mass distribution.

If the observed accelerations of the innermost pulsars can be explained by a standard cluster model, then no additional mass component is required. If, however, the measured values exceed the expectations from a conventional model, one possible interpretation is that the cluster contains an additional compact mass concentration. This has motivated long-standing interest in using pulsars to test for intermediate-mass black holes in cluster cores [65, 84].

The logic is straightforward. A central black hole would steepen the potential within its radius of influence and raise the acceleration envelope for pulsars located near the cluster centre. By comparing the measured pulsar accelerations with those predicted by models with and without such a component, one can test whether a black hole is dynamically required. In this way, pulsar timing provides a complementary probe of cluster cores alongside optical proper motions, radial velocities, and surface-brightness modelling.

This method has already been applied in detail to several globular clusters that host multiple precisely timed pulsars. In 47 Tucanae, for example, measurements of spin and orbital period derivatives have been used to constrain the cluster potential and to test whether an additional central black-hole component is required [80, 84]. Similar ideas have been applied more broadly to other pulsar-rich systems. Whether or not an intermediate-mass black hole is ultimately present, the important point for the present discussion is that pulsars provide a uniquely direct dynamical probe

of the cluster gravitational field.

Taken together, these arguments complete a useful conceptual loop. Earlier in this chapter, globular clusters were introduced as environments that efficiently *produce* exotic pulsars through dynamical interactions and repeated binary processing. Here the perspective is reversed: once formed, those same pulsars become precision tools for studying the cluster itself. They are therefore not only products of cluster dynamics, but also tracers of the underlying gravitational structure that drives those dynamics in the first place.

# Chapter 3

## Observational Techniques for Radio Pulsar Searches

### 3.1 Introduction

Pulsars are among the most remarkable objects discovered in modern astronomy, yet they are also among the most difficult to detect. In the radio band they appear as extremely faint sources whose signals must be distinguished from a background dominated by instrumental noise, human made interferences and natural radio emission from the sky. The detection of pulsars therefore relies not only on their unique periodic emission but also on the sensitivity of the observing instruments and the techniques used to analyse the recorded data.

The discovery of the first pulsar in 1967 by Jocelyn Bell Burnell during her doctoral research at the University of Cambridge marked the beginning of a new field of astrophysics [11]. The instrument that revealed these signals was not originally designed to search for neutron stars (only a theorised, at the time). Instead, it was a large radio array built to study the scintillation of distant radio sources caused by the solar wind. During routine observations, the telescope recorded a sequence of highly regular radio pulses with a period of roughly 1.3 s. The remarkable stability of these signals soon led to the interpretation that they originated from rapidly rotating neutron stars [12].

In the years that followed, pulsar astronomy developed in close connection with advances in radio telescope technology. Larger antennas and more sensitive receivers allowed astronomers to detect weaker sources and survey larger regions of the sky. Early pulsar discoveries were made primarily with large single-dish radio telescopes such as the Jodrell Bank telescope in the United Kingdom and the Arecibo Observatory in Puerto Rico, which for decades remained one of the most sensitive radio

telescopes in the world [3]. These facilities played a central role in establishing the basic observational properties of pulsars and in building the first substantial pulsar catalogues.

The rapid growth of the known pulsar population was driven not only by improvements in telescope sensitivity but also by the development of new methods for analysing radio observations. As computing technology advanced, it became possible to examine large volumes of observational data in search of the characteristic periodic signals produced by pulsars. These developments transformed pulsar searches from small targeted observations into large systematic surveys of the radio sky.

A major milestone occurred in the early 1980s with the discovery of the first millisecond pulsar, PSR B1937+21, whose rotational period is almost one thousand times shorter than that of the first pulsar discovered [85]. The existence of such rapidly rotating neutron stars revealed an entirely new population of pulsars and required further advancements in observing systems capable of recording radio signals with very high time resolution.

Over the past several decades, pulsar searches have continued to evolve as new radio facilities have come online. Modern observatories such as the Green Bank Telescope, the Five-hundred-meter Aperture Spherical Telescope (FAST), and the MeerKAT radio array combine large collecting areas with wide observing bandwidths and advanced digital instrumentation. These capabilities allow astronomers to conduct pulsar surveys with unprecedented sensitivity and to explore regions of parameter space that were previously inaccessible.

Today, the search for pulsars represents a global effort involving radio telescopes, high-performance computing systems, and sophisticated data analysis techniques. The methods used to identify these objects have become increasingly refined as astronomers attempt to detect ever fainter and more complex signals.

In the following sections, the observational techniques used in radio pulsar searches are described in detail. We begin by discussing the radio telescopes and observing systems that collect pulsar signals. Subsequent sections then introduce the methods used to record, process, and analyse these data in order to identify candidate pulsars.

## 3.2 Radio Telescopes and Pulsar Observing Systems

Radio telescopes are the primary instruments used to detect pulsars. Because pulsars emit broadband radio emission with typical flux densities ranging from a few  $\mu\text{Jy}$  to a few  $\text{mJy}$ , high sensitivity is required to detect their signals. The sensitivity of a radio telescope depends primarily on its effective collecting area and system noise temperature, which together determine the minimum detectable signal.

In its simplest form, a radio telescope consists of a large reflecting dish that collects incoming electromagnetic radiation and focuses it onto a receiver located at the focal point. The receiver converts the incident radio waves into electrical voltages, which are subsequently amplified, digitized, and processed by digital backends. The fundamental observable recorded by a radio telescope is therefore a time-varying voltage proportional to the electric field of the incoming radiation.

The sensitivity of a pulsar observation is commonly described by the radiometer equation [2]:

$$S_{\min} = \frac{(S/N)_{\min} T_{\text{sys}}}{G \sqrt{n_p} B t_{\text{obs}}} \sqrt{\frac{W}{P - W}}, \quad (3.1)$$

where  $S_{\min}$  is the minimum detectable flux density,  $(S/N)_{\min}$  is the required signal-to-noise threshold,  $T_{\text{sys}}$  is the system temperature,  $G$  is the telescope gain,  $n_p$  is the number of summed polarizations,  $B$  is the observing bandwidth,  $t_{\text{obs}}$  is the integration time,  $P$  is the pulsar spin period, and  $W$  is the observed pulse width.

Each term in this equation has a clear physical interpretation. The system temperature  $T_{\text{sys}}$  represents the total noise contribution from the receiver electronics, the atmosphere, and the sky background. The telescope gain  $G$  converts flux density into system temperature and is proportional to the effective collecting area of the antenna. Increasing the observing bandwidth  $B$  or the integration time  $t_{\text{obs}}$  improves sensitivity by reducing the statistical noise in the measured signal. The factor  $\sqrt{W/(P - W)}$  reflects the fact that pulsar signals are typically present only during a small fraction of the rotational cycle, known as the duty cycle.

In addition to sensitivity, the angular resolution and field of view of a radio telescope play important roles in pulsar surveys. For a circular dish of diameter  $D$  observing at wavelength  $\lambda$ , the full width at half maximum (FWHM) of the telescope beam is approximately

$$\theta_{\text{FWHM}} \approx 1.2 \frac{\lambda}{D}. \quad (3.2)$$

This relation illustrates the trade-off between angular resolution and field of view: larger telescopes provide higher sensitivity and finer angular resolution but observe a smaller region of the sky at a time.

Modern pulsar surveys often employ large single-dish telescopes or interferometric arrays operating in beamforming mode. In interferometric arrays, signals from multiple antennas are combined coherently to form one or more “tied-array” beams that act as highly sensitive virtual telescopes [2]. This approach allows arrays to achieve both high sensitivity and flexible beam placement across the sky.

The receiver systems used for pulsar observations are designed to record signals with high time resolution and high frequency resolution over large bandwidths. After amplification, the analog voltages produced by the receiver are digitized and passed through digital spectrometers that divide the observing band into many narrow frequency channels. The resulting data are typically stored as *filterbank* or search-mode data products, which contain the detected signal power as a function of time and radio frequency.

High time resolution is essential for detecting pulsars, particularly millisecond pulsars (MSPs), whose spin periods can be as short as a few milliseconds. To resolve the pulse structure of such rapidly rotating objects, sampling times of tens of microseconds or better are often required. At the same time, the observing bandwidth is divided into many frequency channels in order to mitigate the effects of dispersion in the ISM, which introduces frequency-dependent delays across the observing band.

The combination of large collecting areas, wide bandwidth receivers, and high-time-resolution digital backends has dramatically increased the sensitivity of modern pulsar surveys. These technological advances have enabled the discovery of thousands of pulsars, including large populations of MSPs and binary systems that play a crucial role in increasing the pulsar population.

### 3.3 Propagation Effects in the Interstellar Medium

Before reaching a radio telescope on Earth, the radio emission produced by a pulsar must propagate through the ISM. Although interstellar space is often described loosely as empty, in reality it contains a dilute mixture of gas, dust, magnetic fields, and cosmic rays. A significant fraction of the gaseous component is ionised, forming a tenuous plasma composed primarily of free electrons and ions. Typical electron densities in the warm ionised phase of the ISM are very small by terrestrial standards, but pulsar signals travel over path lengths of hundreds to thousands of parsecs. Over

such enormous distances, even a very dilute plasma produces measurable and often substantial effects on the propagation of radio waves.

For pulsar astronomy, these propagation effects are not merely secondary complications. They are central to both detection and interpretation. On the one hand, they distort the signal and can reduce detectability if not properly accounted for. On the other hand, they encode information about the electron density, magnetic field, and turbulent structure of the ISM itself. Pulsars therefore act not only as sources to be detected, but also as beacons that probe the ionised Galactic medium along the line of sight [2, 3, 12].

The principal propagation effects relevant for radio pulsars are *dispersion*, *Faraday rotation*, *scattering*, and *scintillation*. These arise from different physical properties of the plasma. A homogeneous ionised plasma produces dispersion. If the plasma is also magnetised, it gives rise to Faraday rotation. If the electron density is inhomogeneous and turbulent, it produces scattering and scintillation. This classification is particularly useful because it makes clear that different features of pulsar observations trace different aspects of the ISM [2]. In what follows, each of these effects is described in turn.

### 3.3.1 Dispersion

Dispersion is the most important propagation effect for pulsar searches. In a vacuum, electromagnetic waves of all frequencies propagate at the same speed,  $c$ . In a plasma, however, the electric field of the wave drives oscillations of the free electrons, and this modifies the propagation speed of the wave in a frequency-dependent way. The relevant quantity is the group velocity, since it describes the propagation of the pulse envelope.

The characteristic oscillation frequency of the electrons in a plasma is the *plasma frequency*, which for an electron number density  $n_e$  is

$$f_p = \left( \frac{n_e e^2}{\pi m_e} \right)^{1/2}, \quad (3.3)$$

where  $e$  is the electron charge and  $m_e$  is the electron mass. For radio waves of frequency  $f$  propagating through a cold plasma, the group velocity is

$$v_g = c \sqrt{1 - \left( \frac{f_p}{f} \right)^2}. \quad (3.4)$$

For typical pulsar observations, the observing frequency is much larger than the plasma frequency of the ISM, so that  $f \gg f_p$ . In this limit the group velocity may

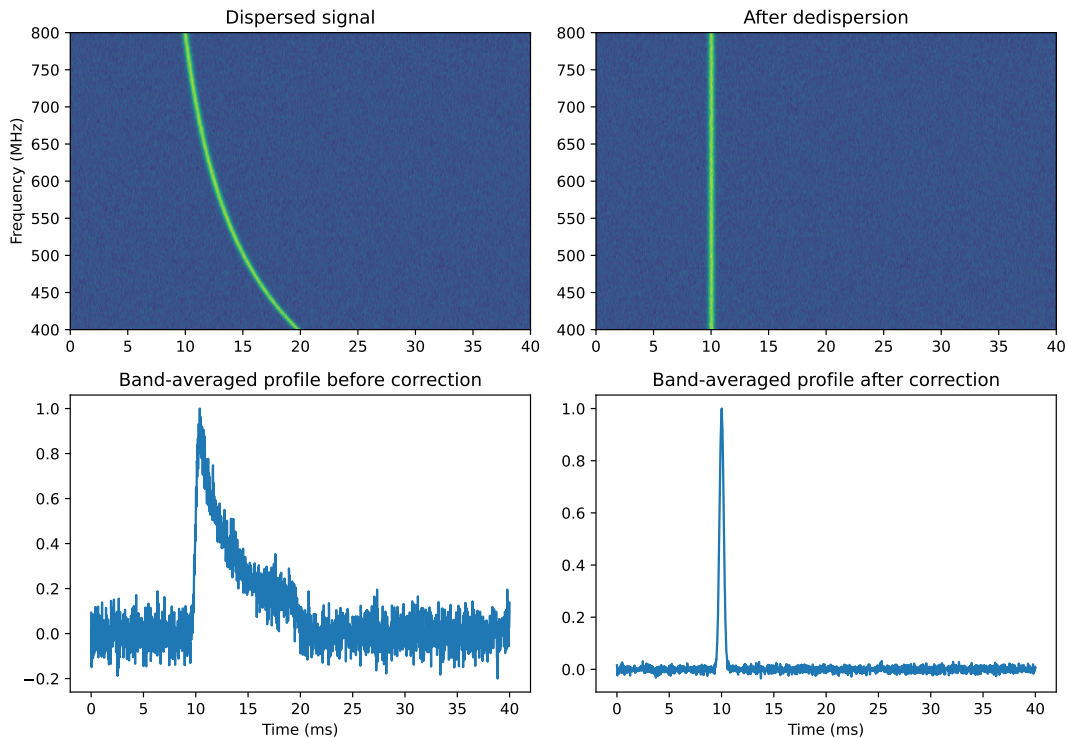


Figure 3.1: Illustration of dispersive delay in a pulsar signal across a radio observing band. *Top left*: Simulated dynamic spectrum showing the characteristic frequency-dependent arrival time of a dispersed pulse in the ionised interstellar medium (ISM), where lower-frequency radiation arrives later than higher-frequency radiation following the cold-plasma dispersion law ( $t \propto f^{-2}$ ). *Top right*: The same signal after correcting for the dispersive delay using the correct dispersion measure (DM), resulting in the pulse aligning across all frequency channels. *Bottom panels*: Band-averaged time series before and after dedispersion. Before correction the pulse is significantly smeared due to the frequency dependent delays, while after dedispersion the intrinsic narrow pulse is recovered with substantially higher peak signal-to-noise ratio.

be expanded to first order as

$$v_g \approx c \left[ 1 - \frac{1}{2} \left( \frac{f_p}{f} \right)^2 \right]. \quad (3.5)$$

This immediately shows the essential result: lower-frequency radio waves propagate slightly more slowly through the ISM than higher-frequency waves. Consequently, if a pulsar emits a broadband pulse, the high-frequency part of the pulse arrives first and the low-frequency part arrives later see 3.1.

The propagation delay may be obtained by integrating along the line of sight. The resulting dispersive delay relative to an infinitely high frequency is

$$\Delta t(f, \infty) = K \text{DM} f^{-2}, \quad (3.6)$$

where

$$K = 4.15 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ ms} \quad (3.7)$$

and the quantity

$$\text{DM} = \int_0^d n_e(l) dl \quad (3.8)$$

is the *dispersion measure*. DM is one of the most fundamental observables in pulsar astronomy. Physically, it represents the integrated column density of free electrons along the line of sight between the pulsar and the observer. Operationally, it quantifies the amount of frequency-dependent dispersive delay introduced by the ISM [2, 12].

The delay between two observing frequencies,  $f_1$  and  $f_2$ , is therefore

$$\Delta t(f_1, f_2) = K \text{DM} (f_1^{-2} - f_2^{-2}). \quad (3.9)$$

This relation is of direct practical importance. In raw pulsar search data, a real astrophysical signal appears as a characteristic sweep across the band, with arrival time varying as  $f^{-2}$ . In modern pulsar search pipelines the data are therefore corrected for dispersion over a large set of trial DM values in order to align the pulse across frequency and recover its full signal-to-noise ratio. This process, known as dedispersion, is discussed in the following section.

Beyond its role in signal recovery, DM is also astrophysically informative. Given a model for the Galactic distribution of free electrons, the measured DM can be used to estimate a pulsar distance. Conversely, if the pulsar distance is known independently, the DM constrains the integrated electron content along the line of sight. Large pulsar samples with measured DMs have therefore played a central role in constructing Galactic free-electron models such as NE2001 and YMW16 [86? ].

### 3.3.2 Faraday Rotation

If the ionised plasma is permeated by a magnetic field with a non-zero line-of-sight component, the two circularly polarised modes of the radiation propagate with slightly different phase velocities. A linearly polarised wave can be decomposed into the sum of left- and right-hand circularly polarised components, so this differential propagation causes the plane of linear polarisation to rotate during propagation. This effect is known as *Faraday rotation* [2, 87].

The change in polarisation position angle,  $\Delta\psi$ , is given by

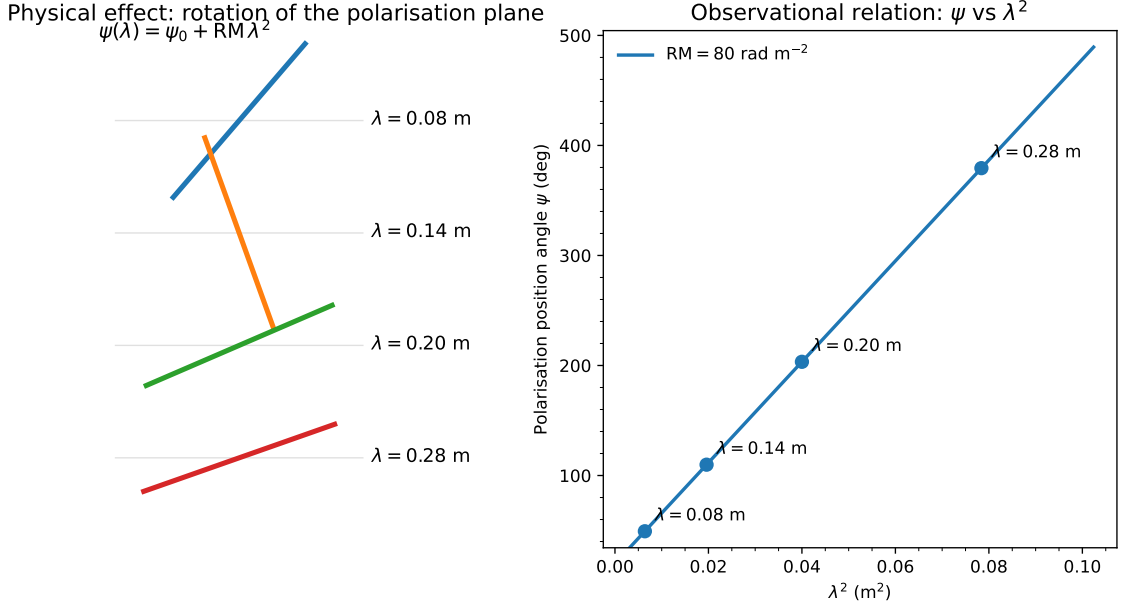


Figure 3.2: Illustration of Faraday rotation in a magnetised ionised medium. *Left panel:* the physical effect of Faraday rotation. As linearly polarised radiation propagates through the interstellar medium (ISM), the plane of polarisation rotates by an angle that increases with the square of the observing wavelength. *Right panel:* the corresponding observational relation between the polarisation position angle  $\psi$  and  $\lambda^2$ . According to  $\psi = \psi_0 + \text{RM} \lambda^2$ , the slope of this relation yields the rotation measure (RM). The labelled points correspond to the wavelengths illustrated in the left panel, demonstrating how the physical rotation of the polarisation vector translates into the linear  $\psi$ - $\lambda^2$  relation measured in pulsar polarimetric observations.

$$\Delta\psi = \text{RM} \lambda^2, \quad (3.10)$$

where  $\lambda$  is the observing wavelength and RM is the *rotation measure* (Fig. 3.2). In practical astrophysical units,

$$\text{RM} = 0.81 \int_0^d n_e B_{\parallel} dl \quad \text{rad m}^{-2}, \quad (3.11)$$

where  $n_e$  is in  $\text{cm}^{-3}$ ,  $B_{\parallel}$  is the line-of-sight magnetic field in  $\mu\text{G}$ , and  $dl$  is in pc. RM is therefore the line-of-sight integral of the product of electron density and magnetic field.

RM is another core observable in pulsar astronomy. While DM measures the integrated free-electron column density, RM measures the magnetised electron content along the line of sight. Because DM and RM depend on related but different path integrals, they can be combined to estimate the average line-of-sight magnetic field weighted by the electron density. One obtains

$$\langle B_{\parallel} \rangle = \frac{\int_0^d n_e B_{\parallel} dl}{\int_0^d n_e dl} \simeq 1.23 \mu\text{G} \left( \frac{\text{RM}}{\text{rad m}^{-2}} \right) \left( \frac{\text{DM}}{\text{pc cm}^{-3}} \right)^{-1}. \quad (3.12)$$

This simple relation is extremely important conceptually: pulsars provide a way to probe not only the free-electron content of the Galaxy through DM, but also its large-scale magnetic field structure through the combination of RM and DM [2, 12]. For this reason pulsars are among the most valuable probes of the magneto-ionic ISM.

Faraday rotation is not generally a limiting factor for the detectability of a pulsar in total-intensity searches, but it is crucial in polarimetric analyses and in any attempt to interpret the intrinsic polarisation properties of pulsar emission.

### 3.3.3 Scattering

The ISM is not perfectly homogeneous. Instead, it contains irregularities in the electron density over a wide range of spatial scales, commonly interpreted as the result of turbulence. These density fluctuations produce local variations in the refractive index, causing radio waves to propagate along multiple paths from the pulsar to the observer rather than along a single geometric path [2, 43].

The consequence is *multipath propagation*. Different parts of the wavefront traverse slightly different path lengths and therefore arrive at different times. In the time domain, this causes a sharp intrinsic pulse to be broadened. Rather than appearing as a narrow symmetric pulse, the observed signal develops an asymmetric shape, usually characterised by a sharp leading edge and an extended trailing tail. In the simplest thin-screen picture, the observed pulse profile is the intrinsic pulse convolved with a one-sided exponential response function,

$$I(t) \propto \exp\left(-\frac{\Delta t}{\tau_{\text{sc}}}\right), \quad (3.13)$$

where  $\tau_{\text{sc}}$  is the *scattering timescale* [40, 88]. In this sense, scattering acts as a temporal smearing process.

The scattering timescale depends strongly on frequency (see Fig. 3.3). For a Kolmogorov-like spectrum of density fluctuations, a commonly used approximation is

$$\tau_{\text{sc}} \propto f^{-4}, \quad (3.14)$$

although the exact exponent may differ somewhat depending on the distribution and spectrum of the irregularities [43]. The key practical point is that scattering

Simulated frequency evolution of pulse broadening by scattering  
 $\tau_{sc} \propto f^{-4}$  (individual normalization)




Figure 3.3: Simulated frequency evolution of pulse broadening caused by multipath scattering in the ISM. The same intrinsic pulse profile is shown at three observing frequencies assuming a scattering timescale scaling as  $\tau_{sc} \propto f^{-4}$ . At higher observing frequencies the pulse remains narrow, while at lower frequencies the pulse develops an increasingly extended exponential tail. This broadening reduces the peak signal-to-noise ratio and can strongly limit the detectability of pulsars in low-frequency surveys.

grows rapidly toward lower frequencies. This is one of the main reasons why low-frequency pulsar searches can become strongly sensitivity limited, especially for fast-spinning MSPs: a pulse that is intrinsically narrow may be broadened until it occupies a substantial fraction of the pulse period, greatly reducing the peak signal-to-noise ratio.

Scattering is particularly severe along lines of sight through the Galactic plane or toward regions of enhanced turbulence. It is therefore not determined by DM alone. Two sightlines with similar DM may exhibit very different levels of scattering depending on the distribution of turbulent structures between the source and the observer. This is an important practical point when interpreting survey detectability: DM measures integrated electron content, while scattering traces fluctuations in that electron distribution, not simply its total amount.

### 3.3.4 Interstellar Scintillation

Closely related to scattering is *interstellar scintillation*. The same density irregularities that produce multipath propagation also cause interference between the scattered wavefronts at the observer. Depending on their relative phase, the interference may be constructive or destructive, producing variations in the observed intensity of the pulsar as a function of both time and frequency [2, 43].

In dynamic spectra, scintillation appears as a patchy pattern of bright and dim regions, commonly called *scintles*. It is in some sense the radio analogue of the twinkling of optical stars, although the physical medium involved is the ionised ISM rather than the Earth’s atmosphere.

A useful quantity here is the *scintillation bandwidth*,  $\Delta f_{\text{sc}}$ , which characterises the frequency scale over which the intensity remains correlated. The condition for interference to be important is approximately

$$2\pi \Delta f_{\text{sc}} \tau_{\text{sc}} \sim 1. \quad (3.15)$$

This relation expresses the fact that coherence is lost once the phase difference between scattered paths changes substantially across the bandwidth considered. It immediately implies that

$$\Delta f_{\text{sc}} \propto \tau_{\text{sc}}^{-1}. \quad (3.16)$$

Since the scattering timescale decreases rapidly with increasing frequency, the scintillation bandwidth correspondingly increases with frequency. Using the approximate scaling  $\tau_{\text{sc}} \propto f^{-4}$ , one obtains

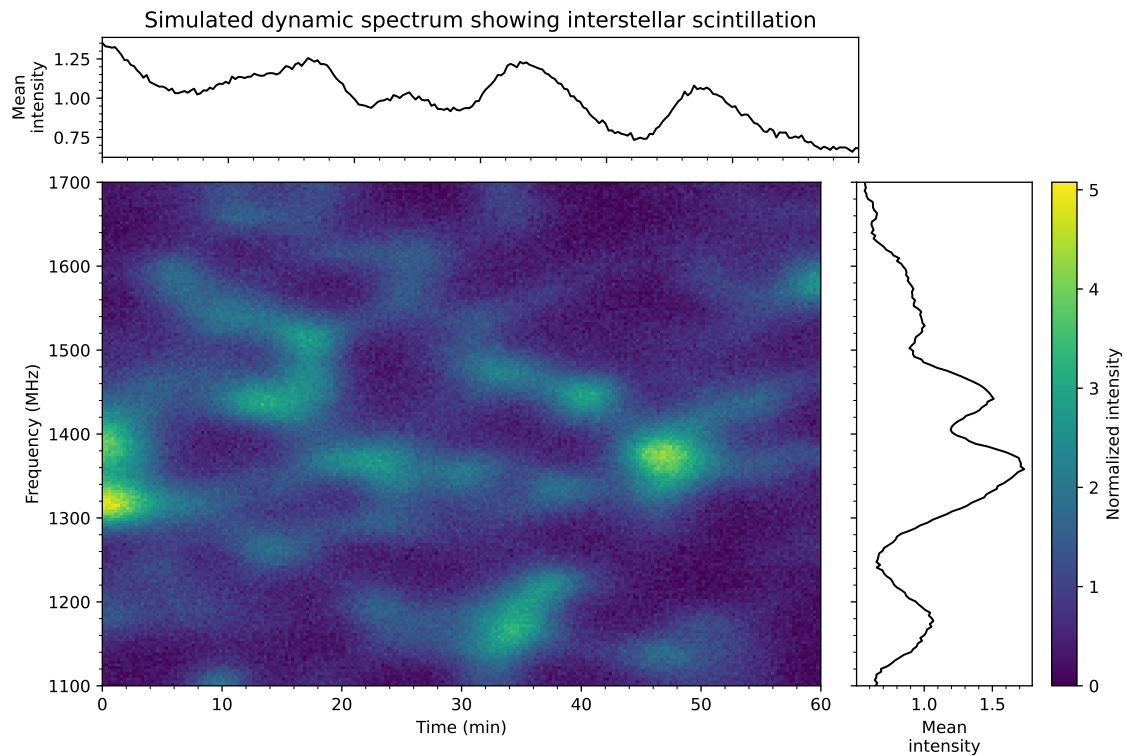


Figure 3.4: Simulated dynamic spectrum illustrating interstellar scintillation of a pulsar signal. The main panel shows the observed intensity as a function of time and observing frequency. The patchy bright and dim regions (“scintles”) arise from interference between multiple propagation paths through the turbulent ionised interstellar medium. The top and right panels show the mean intensity variations with time and frequency, respectively. Scintillation can strongly modulate the apparent brightness of pulsars and may therefore affect their detectability in individual observations.

$$\Delta f_{\text{sc}} \propto f^4. \quad (3.17)$$

Thus, at low frequencies the scintillation bandwidth is small and the intensity pattern varies rapidly across frequency, whereas at higher frequencies the correlation bandwidth becomes larger (see Fig. 3.4).

Scintillation also occurs over a characteristic timescale, set by the relative motion between the pulsar, the scattering medium, and the observer. As this pattern drifts across the line of sight, the apparent brightness of the pulsar changes with time. Observationally, this means that a pulsar can sometimes appear strongly enhanced and at other times substantially dimmed. In practical survey work, scintillation can therefore make a faint pulsar detectable in one epoch but not in another. It does not alter the intrinsic periodicity nor the duration (unlike scattering) of the signal, but it can modulate its observability in a significant way.

Taken together, dispersion, Faraday rotation, scattering, and scintillation show that pulsar signals do not arrive at the telescope in their pristine form. Instead, they carry the imprint of the entire magneto-ionic medium through which they have propagated. For pulsar searching, the most immediate and essential consequence is the need to correct for dispersion in order to recover narrow periodic signals efficiently. Scattering, however, is not routinely removed in standard pulsar search-mode processing. In principle, the effects of scattering can be deconvolved if the scattering response is sufficiently well constrained, for example through cyclic spectroscopy using high-S/N voltage or baseband data [89]. In ordinary survey pipelines, however, scattering is usually treated as an irreversible broadening effect that reduces the peak signal-to-noise ratio. The methods used to remove dispersive delay from observational data are described in section 3.5.7.

The interplay between these propagation effects also has important implications for the choice of observing frequency in pulsar surveys. Since dispersive smearing and scattering both increase strongly toward lower radio frequencies, searches targeting fast-spinning pulsars, particularly millisecond pulsars in high-density or high-DM environments such as globular clusters or the Galactic centre, are often carried out at higher frequencies to preserve pulse detectability. Conversely, lower-frequency observations benefit from the typically steep pulsar spectra and can be more sensitive to slow, nearby pulsars with relatively low dispersion and scattering. The optimal observing frequency therefore reflects a trade-off between intrinsic pulsar brightness and propagation-induced signal degradation, and depends on both the target population and the line of sight through the interstellar medium.

### 3.4 Radio Telescopes and Pulsar Search Data Acquisition

The detection of pulsars begins with the collection of radio emission by large radio telescopes and the subsequent conversion of this signal into digital data products that can be analysed using pulsar search techniques. Between the arrival of the electromagnetic radiation at the telescope and the moment when an astronomer receives a data file, the signal undergoes several stages of amplification, digitisation, and digital signal processing. Understanding this chain of processes is important for interpreting the structure and limitations of pulsar search data.

### 3.4.1 Radio telescopes for pulsar observations

Radio telescopes detect the electric field associated with incoming electromagnetic radiation. A reflecting surface collects the radiation and focuses it onto a receiver system where it is converted into electrical voltages. The sensitivity of a radio telescope is primarily determined by its effective collecting area and system noise temperature. Larger collecting areas allow weaker signals to be detected, while lower receiver noise improves the signal-to-noise ratio of the measured signal.

Two broad classes of radio telescope systems are widely used in pulsar astronomy: large single-aperture telescopes and interferometric arrays. Both approaches play important and complementary roles in pulsar searches.

Single-dish radio telescopes consist of a large reflecting surface that focuses incoming radiation onto a feed located near the focal point. The sensitivity of such instruments is determined largely by the diameter of the dish. Several of the most productive pulsar telescopes in operation today belong to this category. Examples include the Five-hundred-meter Aperture Spherical Telescope (FAST) in China, the Green Bank Telescope (GBT) in the United States, and the Effelsberg 100-m telescope in Germany. The enormous collecting area of FAST, in particular, makes it one of the most sensitive instruments ever built for pulsar searches.

Interferometric arrays, on the other hand, consist of multiple antennas distributed over a region ranging from hundreds of metres to several kilometres. Instead of relying on a single large reflector, the signals received by the individual antennas are combined digitally. By correlating the signals from pairs of antennas, interferometers can synthesize an effective aperture comparable to the maximum separation between antennas, thereby achieving extremely high angular resolution [? ].

Although interferometers are often used for imaging observations, pulsar searches typically exploit a technique known as *beamforming*. In beamforming mode the voltage signals from multiple antennas are combined after applying appropriate geometric delay corrections. The resulting signal forms a highly sensitive narrow beam on the sky known as a *tied-array beam*. Interferometric arrays can form many such beams simultaneously, allowing large regions of the sky to be searched with high sensitivity. The MeerKAT telescope is a prominent example of an array that has been extensively used for pulsar searches using beamformed observing modes.

### 3.4.2 Front-end receivers and signal amplification

Once the radio signal has been collected by the telescope, it is directed to the front-end receiver system. The receiver is responsible for converting the incoming

electromagnetic radiation into electrical signals that can be amplified and processed electronically.

The front-end typically includes a cryogenically-cooled feed, low-noise amplifiers (LNAs), and receiver electronics that define the observing band. The feed converts the incoming electric field into voltages in one or more polarisation channels. In most modern receivers two orthogonal polarisations are recorded, allowing the full polarisation properties of the signal to be preserved.

The signals from the feed are extremely weak and must be amplified before further processing. Low-noise amplifiers are therefore placed as close as possible to the feed in order to minimise the introduction of additional noise. The receiver electronics then filter and down-convert the signal into an intermediate frequency range suitable for digitisation.

### 3.4.3 Digital back-ends and channelisation

After amplification and frequency conversion, the analogue voltages are sent to the digital back-end systems of the telescope. At this stage the signals are digitised using high-speed analogue-to-digital converters. Digitisation converts the continuous voltage signal into a discrete time series that can be processed by digital electronics.

One of the most important steps in the digital processing chain is *channelisation*. In this process the observing bandwidth is divided into many narrow frequency channels. Channelisation is typically performed using digital filterbanks or polyphase filterbank algorithms, which allow the wide observing band to be separated into hundreds or thousands of channels with well-defined spectral responses.

The output of this stage is a two-dimensional data stream representing the signal intensity as a function of both time and frequency. Each frequency channel contains a time series sampled at regular intervals. The sampling time must be sufficiently short to preserve the rapid variations expected from pulsar signals, which can have spin periods ranging from seconds down to a few milliseconds.

### 3.4.4 Search-mode observations

Pulsar observations can be recorded in several different modes depending on the scientific objective. For pulsar surveys and discovery observations the telescope typically operates in *search mode*. In this mode the signal intensity is recorded as a function of time and frequency without assuming any prior knowledge of the pulsar period. The resulting data retain the full time resolution necessary for searching for periodic signals over a wide range of trial periods.

This approach contrasts with *fold mode* observations, in which the data are folded i.e summed in pulse phase in real time at a known pulsar period in order to produce an integrated pulse profile. Fold-mode observations are commonly used in precision pulsar timing experiments once the spin period of a pulsar is already known [2].

Because search-mode observations retain high time and frequency resolution across large bandwidths, the resulting data rates can be extremely large. For modern wideband receivers the raw data rates can reach many gigabytes per second. Consequently, some degree of preliminary processing is often applied before the data are written to disk. Such operations may include initial mitigation of radio-frequency interference (RFI), integration over short time intervals, or optional subbanding of neighbouring frequency channels in order to reduce the total data volume.

### 3.4.5 Pulsar search data formats

The processed data streams produced by the digital back-end are stored in specialised file formats designed for time–frequency data. Two widely used formats in pulsar astronomy are filterbank files and PSRFITS files.

Filterbank files, commonly identified by the extension `.fil`, contain the signal intensity recorded as a function of time and frequency. This format originated in pulsar signal processing software developed for pulsar search experiments and has become widely used for storing search-mode observations. Each file begins with a header that describes the observing parameters, including the telescope name, observing frequency, bandwidth, number of frequency channels, and sampling time. The remainder of the file consists of a sequential array of intensity measurements corresponding to each frequency channel and time sample [2, 90].

Another commonly used format is PSRFITS, which is based on the Flexible Image Transport System (FITS) standard widely used throughout astronomy [91, 92]. PSRFITS provides a structured representation of pulsar data and includes extensive metadata describing the observing setup, telescope configuration, and instrumental parameters. The time–frequency data themselves are stored within the FITS structure in a form specifically designed for pulsar observations [93].

### 3.4.6 Pulsar search software: the PRESTO toolkit

Once pulsar search observations have been recorded and stored in time–frequency data formats such as filterbank or PSRFITS files, the next stage involves a sequence of processing steps designed to identify periodic or transient astrophysical signals within the data. These steps typically include radio-frequency interference (RFI)

mitigation, dedispersion across a range of trial dispersion measures (DMs), Fourier-domain searches for periodic signals, and candidate evaluation. In modern pulsar astronomy, these tasks are usually performed using specialised software packages developed specifically for pulsar signal processing.

One of the most widely used pulsar search software suites is the PRESTO package, originally developed by Scott M. Ransom for the analysis of high-time-resolution radio data [94? ]. PRESTO provides a comprehensive collection of programs and libraries designed to process pulsar search observations from raw time-frequency data through to candidate identification. Over the past two decades, it has become one of the standard toolkits used in many large pulsar surveys and targeted search projects conducted with major radio telescopes.

The PRESTO software suite implements many of the core algorithms required for pulsar discovery. These include routines for identifying and mitigating radio-frequency interference, dedispersing the data across a grid of trial dispersion measures, performing Fourier-domain searches for periodic signals, and carrying out acceleration searches to detect pulsars in compact binary systems [2? ]. The package also contains tools for candidate sifting, visualisation, and further analysis of promising pulsar detections.

Because of its flexibility and extensive functionality, PRESTO has become a foundational tool in modern pulsar search pipelines. Many contemporary surveys employ automated processing pipelines built around PRESTO routines to handle the large data volumes produced by modern radio telescopes. Detailed documentation and tutorials for the software are available through the official repository and user guides.<sup>1</sup>

In the following sections, the individual stages of the pulsar search process are described in detail, beginning with the identification and mitigation of radio-frequency interference in the observational data.

## 3.5 Pulsar pre-search Processing

Once pulsar observations have been recorded in search-mode formats, the task of the astronomer is to transform these time-frequency data into a set of candidate astrophysical signals. In practice this begins not with the search itself, but with the removal or mitigation of non-astrophysical contamination. Modern pulsar search data are almost invariably affected by RFI, instrumental systematics, and non-

---

<sup>1</sup>The PRESTO software package and documentation are available at <https://github.com/scottransom/presto>. Tutorials and example workflows are provided within the repository and associated user documentation.

stationary noise processes. If left untreated, these effects can dominate the statistics of the data, generate large numbers of false candidates, and reduce the sensitivity of the search to real pulsars. For this reason, RFI mitigation is typically the first major processing step in a pulsar search pipeline.

### 3.5.1 Radio-frequency interference and noise in pulsar search data

Pulsar search data contain several distinct forms of contamination, and it is useful to distinguish them conceptually before discussing mitigation strategies.

The ideal background noise of a radio observation is often approximated as *white noise*. In the time domain, white noise is a random process with uncorrelated samples and approximately Gaussian amplitude statistics. In the Fourier domain, it corresponds to a flat power spectrum. Thermal receiver noise and sky noise contribute approximately white components to the measured signal over sufficiently small time and frequency intervals [2].

Real observations, however, are seldom governed by white noise alone. Pulsar search data commonly exhibit *red noise*, meaning excess power at low fluctuation frequencies. Red noise can arise from slow gain variations, baseline drifts, atmospheric or instrumental instabilities, imperfect bandpass removal, and other non-stationary processes. In the Fourier domain it appears as enhanced power at low Fourier frequencies and can significantly reduce sensitivity to long-period pulsars [95, 96].

In addition to these stochastic noise processes, radio observations are frequently contaminated by RFI of human made origin. Such interference may be persistent or intermittent, narrow-band or broadband, impulsive or periodic. Common sources include telecommunication signals, radar, electrical switching, digital electronics, power-line harmonics, satellite transmissions, and impulsive discharges from nearby infrastructure [97, 98]. The observational manifestation depends on the source. Narrow-band RFI tends to contaminate specific frequency channels and may remain stable over long time spans. Broadband impulsive RFI can contaminate many or all frequency channels simultaneously for short intervals. Periodic terrestrial signals may appear as discrete features in Fourier space and are often referred to in pulsar-search practice as *birdies*. Figure 3.5 shows various forms of RFI commonly seen in radio data.

The central challenge of RFI mitigation is therefore twofold. First, one must identify which parts of the data are likely to be contaminated. Second, one must suppress or mask those data without unnecessarily removing astrophysical signal. This is not a trivial problem: overly aggressive excision can reduce sensitivity to

Illustration of common forms of radio-frequency interference in pulsar search data

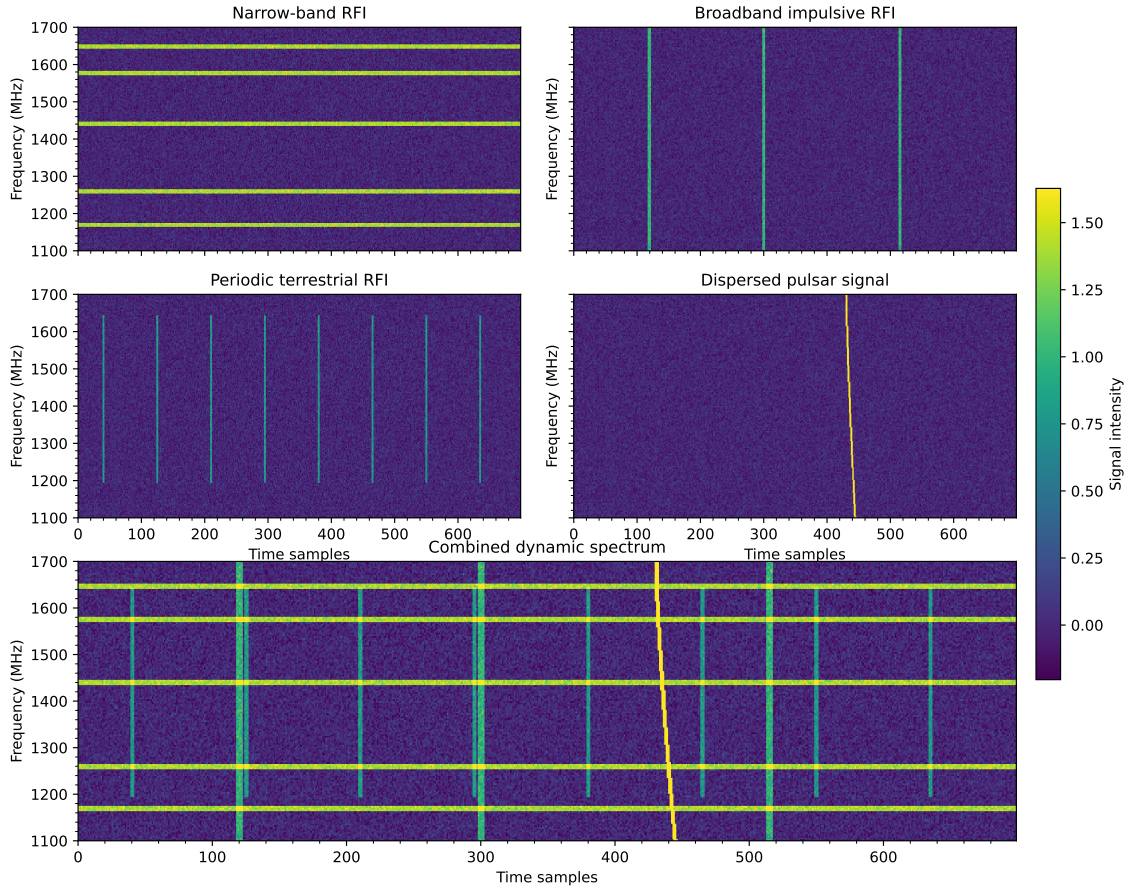


Figure 3.5: Illustration of common forms of radio-frequency interference (RFI) encountered in pulsar search data. The upper panels show simplified dynamic spectra illustrating individual types of contamination: *narrow-band RFI*, which appears as persistent horizontal features affecting specific frequency channels; *broadband impulsive RFI*, which produces vertical features affecting many channels simultaneously for short time intervals; and *periodic terrestrial RFI*, which appears as regularly repeating bursts in time. For comparison, the lower-right panel shows a simulated dispersed pulsar signal, which produces a frequency-dependent sweep across the observing band due to interstellar dispersion. The bottom panel shows a combined dynamic spectrum containing all of these components simultaneously, illustrating how different forms of interference and astrophysical signals can coexist in real pulsar search observations.

real pulsars, while insufficient excision leads to large numbers of false candidates and distorts the statistical properties of the search.

### 3.5.2 Why RFI can mimic pulsars

Pulsars are sought as weak periodic or transient signals embedded in noise. Unfortunately, many forms of terrestrial interference can also appear periodic or impulsive. Narrow-band transmitters can generate strong lines in Fourier spectra. Periodic electrical interference can produce harmonics that resemble the harmonic structure of real pulsars. Broadband impulsive signals can mimic bright single pulses. The problem becomes more severe in large blind surveys, where even a small false-positive rate per observation leads to an enormous number of spurious candidates.

The crucial physical discriminator is that most terrestrial RFI is *undispersed*. A genuine pulsar signal propagating through the ISM accumulates a dispersive delay described by its DM, whereas local interference generally arrives simultaneously across the observing band and therefore peaks at  $DM \approx 0$ .

### 3.5.3 RFI mitigation in a PRESTO-based workflow

In PRESTO, the first step after verifying the metadata of the input search-mode file is usually RFI identification with the routine `rfifind`. In the official PRESTO workflow, `rfifind` is used before dedispersion and periodicity searching, and the masks it generates are then applied in the subsequent processing steps [99, 100].

The basic philosophy of `rfifind` is statistical. The data are divided into short time blocks, and statistics are computed in both the time and frequency domains. Chunks of data whose behaviour deviates significantly from the expected noise statistics are flagged as interference candidates. The documentation exposes several controls for this procedure, including a time-domain sigma threshold, a frequency-domain sigma threshold, and fractional thresholds specifying how many bad channels are required before a full time interval is masked and how many bad intervals are required before a full channel is masked. In other words, `rfifind` does not merely look for obviously bright events; it builds a statistical description of the observation and identifies portions of the data that are inconsistent with the expected baseline behaviour.

In the PRESTO tutorial, the typical block size used for RFI statistics is of order a few seconds, and the user is encouraged to inspect the resulting mask and adjust the integration time if too much data are being flagged [99]. This is an important practical point: RFI mitigation is partly automated, but it still requires judgement.

### 3.5.4 Clipping, masking, and DM= 0 filtering

Several complementary RFI mitigation strategies are commonly used in pulsar search pipelines, either within PRESTO or in related survey processing systems.

A first line of defence is *clipping* in the time domain. Strong impulsive events can dominate the time series and contaminate large portions of the Fourier spectrum. Time-domain clipping identifies samples or short intervals with amplitudes far above the expected noise level and either replaces them or masks them. In PRESTO, clipping of strong transient DM= 0 signals is enabled by default in the early preparation steps unless explicitly disabled [99? ].

A second approach is *masking*. Once channels or intervals have been identified as contaminated, they are excluded from further processing. Masking is one of the most widely used and least assumption-dependent RFI mitigation techniques. It is especially effective for persistent narrow-band interference and for time ranges dominated by strong broadband bursts [97? ]. The cost, of course, is a reduction in effective bandwidth or integration time.

A third and particularly important strategy is *zero-DM filtering*. The principle is simple: at each time sample, the mean across all frequency channels is subtracted from every channel. Because undispersed broadband terrestrial signals arrive simultaneously across the band, they contribute strongly to this mean and are therefore suppressed (see Fig. 3.6). In contrast, a dispersed celestial signal is distributed across different time samples in different channels and is therefore affected much less strongly [98]. Zero-DM filtering has proved especially useful in mitigating impulsive broadband RFI in pulsar surveys, and it is supported directly in PRESTO through the `-zerodm` option. However, it is not without cost: because it subtracts a common mode across the band, it can also remove some power from low-DM astrophysical signals and must therefore be applied with care [98? ].

### 3.5.5 Birdies and periodic terrestrial signals

Not all problematic RFI is impulsive. Some interference sources are periodic or quasi-periodic and therefore appear as narrow spectral lines in the Fourier domain. These spurious periodicities are widely referred to as birdies in pulsar-search practice. They can arise from electrical mains harmonics, clocking electronics, digitiser artefacts, and other periodic terrestrial systems.

Birdies are especially troublesome because pulsars themselves are sought as periodic signals. A strong birdie can generate a family of harmonics that closely resembles the harmonic structure of a real pulsar, especially for narrow pulse duty

Effect of zero-DM subtraction on undispersed terrestrial interference

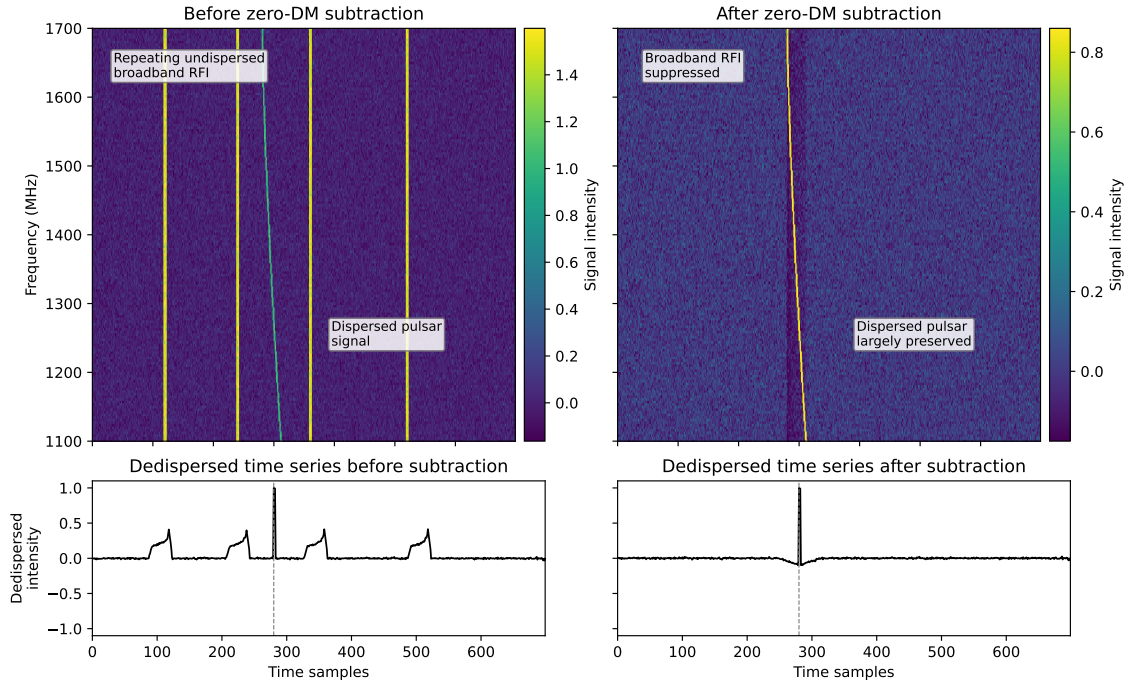


Figure 3.6: Illustration of the effect of zero-DM subtraction on undispersed terrestrial interference in pulsar search data. The upper panels show the dynamic spectrum before and after zero-DM subtraction. Broadband terrestrial RFI, which arrives simultaneously across all frequency channels and therefore peaks at  $DM \approx 0$ , appears as vertical stripes in the dynamic spectrum. Subtracting the mean across frequency channels at each time sample strongly suppresses this undispersed component while largely preserving the dispersed pulsar signal. The lower panels show the time series after dedispersion at the pulsar dispersion measure. Prior to mitigation, terrestrial bursts contaminate the dedispersed time series and obscure the astrophysical signal. After zero-DM subtraction, the RFI contribution is significantly reduced and the pulsar peak becomes clearly identifiable.

cycles. In the standard PRESTO workflow, one therefore inspects a  $DM=0$  time series and its Fourier transform before performing the full dedispersion search. The logic is straightforward: if a periodic signal is strongest at  $DM=0$ , then it is almost certainly local interference rather than a dispersed pulsar signal. The tutorial workflow then uses Fourier-domain inspection tools to identify these periodic features and construct a zap list for subsequent searches [99]. In PRESTO’s own description of the data-preparation stage, this role is served by routines such as `zapbirds` [100].

In practice, birdie removal is usually performed by excluding narrow frequency ranges around known or measured interference frequencies in the Fourier domain. This procedure is highly effective when the offending frequencies are stable, but less so when the RFI drifts or when the interference is only intermittently present. For this reason, birdie identification is commonly combined with the broader statistical masking performed by `rfifind`.

### 3.5.6 White noise, red noise, and the limits of mitigation

It is tempting to think of RFI mitigation as a problem of simply removing obviously artificial signals, but the practical reality is more subtle. Even after obvious interference has been excised, the residual statistics of pulsar search data may still depart strongly from the ideal white-noise case. Slow baseline variations and red-noise processes can leave significant low-frequency power in the Fourier spectrum, reducing sensitivity to long-period pulsars. This issue has been demonstrated in survey analyses showing that red noise can substantially degrade pulsar detectability at long periods and low DMs [95, 96]. In other words, the limiting background for pulsar searching is often not simply thermal noise, but a combination of instrumental, terrestrial, and statistical non-stationarity.

This distinction matters because different mitigation techniques target different failure modes. Clipping helps with strong impulsive interference. Masking helps with contaminated channels or intervals. Zero-DM filtering helps with broadband undispersed signals. Fourier-domain zapping helps with stable periodic birdies. Red-noise suppression, which becomes important in the search stage itself, targets slow low-frequency spectral contamination. No single method solves all problems, and modern pipelines therefore rely on multiple complementary layers of mitigation.

Once this stage is complete, the next task is to reverse the most important astrophysical distortion introduced by the ISM itself, namely dispersion. The RFI-mitigated search-mode data are therefore ready for the dedispersion stage, which is usually the next major operation in a pulsar search pipeline.

### 3.5.7 De-dispersion

As discussed in Section 3.3.1, dispersive propagation through the ionised interstellar medium causes lower-frequency radiation to arrive later than higher-frequency radiation. In pulsar search data this produces the characteristic  $f^{-2}$  sweep across the observing band shown in Fig. 3.1. If the frequency channels are summed without correcting for this delay, the pulse energy becomes spread over many time samples and the peak signal-to-noise ratio of the pulse is significantly reduced (or the pulse even completely washed out, if the delay is larger than the pulse period).

The purpose of de-dispersion is therefore to compensate for these delays and reconstruct the original broadband pulse by aligning the signal across the observing band. In a pulsar search the true DM of a source is not known *a priori*, so the correction must be applied for a large set of trial DM values. For each trial DM, the frequency channels are shifted in time according to the expected dispersive delay and then summed to produce a single time series. When the assumed DM matches the true value, the pulses align across the band and the signal is recovered with maximum signal-to-noise ratio.

If the assumed DM differs from the true DM, a residual dispersive smearing remains in the dedispersed time series. The characteristic broadening due to a DM error  $\Delta\text{DM}$  across an observing bandwidth  $\Delta f$  centred at frequency  $f$  can be approximated as

$$t_{\text{DM}} \approx 8.3 \mu\text{s} \left( \frac{\Delta\text{DM}}{\text{pc cm}^{-3}} \right) \left( \frac{\Delta f}{\text{MHz}} \right) \left( \frac{f}{\text{GHz}} \right)^{-3}. \quad (3.18)$$

This strong  $f^{-3}$  dependence makes dispersion particularly problematic at low radio frequencies. For metre-wavelength facilities such as LOFAR or NenuFAR, the dispersive delay across even modest bandwidths can reach many milliseconds. Such smearing is sufficient to wash out the signal from fast millisecond pulsars unless extremely fine DM spacing or coherent de-dispersion techniques are employed (see below).

An additional limitation arises from the finite width of the individual frequency channels in channelised pulsar search data. In standard search pipelines the data are recorded as detected power in many narrow channels. In this case the dispersive delay can only be corrected between channels, not within them. The resulting residual broadening, known as *intra-channel dispersive smearing*, is given approximately by

$$t_{\text{chan}} \approx 8.3 \mu\text{s} \left( \frac{\text{DM}}{\text{pc cm}^{-3}} \right) \left( \frac{\Delta f_{\text{chan}}}{\text{MHz}} \right) \left( \frac{f}{\text{GHz}} \right)^{-3}, \quad (3.19)$$

where  $\Delta f_{\text{chan}}$  is the channel bandwidth. For large DMs or wide channels this term alone can exceed the intrinsic pulse width of a millisecond pulsar, significantly reducing detectability.

Two principal approaches exist for correcting dispersion in pulsar data. In *incoherent de-dispersion*, the detected power in each frequency channel is shifted in time according to the expected dispersive delay and the channels are then summed. This approach is computationally efficient and is therefore widely used in large pulsar surveys. However, as noted above, it cannot remove dispersive smearing within individual channels.

In contrast, *coherent de-dispersion* operates directly on the complex voltage data, before channel formation, and removes the dispersive transfer function of the interstellar plasma in the Fourier domain. This method eliminates dispersive smearing across the entire observing band, but it requires access to the raw voltage data and significantly greater computational resources. Coherent de-dispersion is therefore typically used in pulsar timing observations of known sources rather than in wide-field pulsar searches.

In this work the de-dispersion step was performed using the `prepdata` routine from the PRESTO pulsar-search software suite<sup>2</sup>. `prepdata` reads the cleaned filter-bank data and generates a dedispersed time series for a specified trial DM. The routine computes the appropriate delay for each frequency channel and applies the corresponding time shifts before summing the channels to produce a single dedispersed time series. Repeating this process for a grid of trial DM values produces a collection of dedispersed time series that are subsequently searched for periodic signals.

The dedispersed time series generated by `prepdata` form the input for the periodicity and acceleration searches described in the following sections.

### 3.5.8 Red-noise mitigation

Following dedispersion, the data consist of a set of time series, each corresponding to a different trial dispersion measure. In an ideal observation these time series would contain only white noise together with any periodic pulsar signal present in the data. In practice, however, as mentioned above, they often exhibit slow variations in the baseline arising from instrumental gain fluctuations, atmospheric effects, residual radio-frequency interference, or imperfections in the receiver system giving rise to the so-called *red noise* in the power spectra [2].

---

<sup>2</sup><https://www.cv.nrao.edu/~sransom/presto/>

Red noise can significantly degrade the sensitivity of pulsar searches, particularly for long-period pulsars. In the Fourier domain, slow baseline fluctuations generate strong power in the lowest frequency bins of the spectrum. Since pulsars are detected as narrow peaks in the Fourier power spectrum, this excess low-frequency power can mask real signals or distort the statistical distribution of the noise background, making it more difficult to identify genuine pulsar candidates during searches. [101, 102].

To mitigate these effects, pulsar search pipelines apply procedures that suppress slow variations in the time series before or after performing the Fourier transform. One common approach operates in the time domain, where a running median or low-order polynomial is subtracted from the data in order to remove long-timescale trends in the baseline. Alternatively, red-noise mitigation may be performed in the Fourier domain by locally normalising or “whitening” the power spectrum so that the background noise becomes approximately flat across frequency.

In the PRESTO pulsar-search software suite, red-noise suppression is performed in the frequency domain. The Fourier spectrum of each dedispersed time series is computed and then locally normalised using algorithms implemented in routines such as `realfft` and `rednoise`. These procedures reduce the excess power at low frequencies while preserving the narrow spectral features associated with genuine periodic signals [103, 104].

The resulting whitened spectra provide a statistically well-behaved background against which periodic signals can be identified. With the effects of dispersion and red noise mitigated, the dedispersed time series are then ready for the Fourier-domain periodicity searches.

## 3.6 Periodic Pulsar Search Techniques

### 3.6.1 Fourier-domain periodicity searches

After dedispersion and red-noise mitigation, each trial-DM data product is a one-dimensional time series that may contain a weak periodic pulsar signal buried in stochastic noise. The central idea of Fourier-domain searching is that a strictly repeating signal, even when not visually apparent in the time domain, can be transformed into a small number of localized peaks in frequency space. In this way, the Fourier transform acts as a highly efficient matched filter for periodic phenomena [2, 103].

Let the dedispersed time series be sampled uniformly with sampling interval  $t_{\text{samp}}$ , so that the total observing duration is

$$T_{\text{obs}} = N t_{\text{samp}}, \quad (3.20)$$

where  $N$  is the total number of samples. If the discrete time series is denoted  $x_n \equiv x(t_n)$  with  $t_n = n t_{\text{samp}}$  and  $n = 0, \dots, N-1$ , then its discrete Fourier transform (DFT) is

$$\tilde{x}_k = \sum_{n=0}^{N-1} x_n \exp\left(-2\pi i \frac{kn}{N}\right), \quad k = 0, \dots, N-1. \quad (3.21)$$

Each index  $k$  corresponds to a Fourier frequency

$$\nu_k = \frac{k}{T_{\text{obs}}}, \quad (3.22)$$

so that the spacing between adjacent Fourier bins is

$$\delta\nu = \frac{1}{T_{\text{obs}}}. \quad (3.23)$$

Equation 3.23 is of fundamental practical importance: longer observations produce narrower Fourier bins, allowing periodic signals to be measured with finer frequency resolution. However, if the signal frequency is not perfectly constant over the observation (e.g. due to acceleration), the power can spread across multiple bins, effectively reducing the signal-to-noise ratio.

To see why the Fourier transform is so effective for pulsar searches, consider first a noiseless sinusoidal signal of amplitude  $A$  and frequency  $\nu_0$ ,

$$x_n = A \cos(2\pi\nu_0 t_n + \phi), \quad (3.24)$$

where  $\phi$  is an arbitrary phase. If  $\nu_0$  lies exactly at one of the Fourier frequencies  $\nu_k$ , then the sum in Eq. 3.21 adds coherently and the signal power is concentrated into the corresponding Fourier bin. In contrast, random noise does not add coherently. Its contribution is distributed over many Fourier bins, producing a comparatively broad background rather than a localized spectral peak.

This is the essential reason why periodic signals that are almost invisible in the time domain can become obvious in the Fourier domain. In a sufficiently long observation, the contribution of a periodic signal to the correct Fourier bin grows coherently with the number of samples, whereas the stochastic noise grows only as a random walk. Qualitatively, if the coherent signal amplitude scales as  $\propto N$  in the correct bin, the corresponding noise amplitude scales only as  $\propto \sqrt{N}$ . The transform therefore acts as an integrator of periodic power, greatly enhancing sensitivity to

repetition.

In practical searches, the quantity inspected is usually the Fourier power spectrum,

$$P_k = |\tilde{x}_k|^2. \quad (3.25)$$

For a time series containing only white Gaussian noise, the real and imaginary parts of  $\tilde{x}_k$  are themselves Gaussian-distributed random variables, and the Fourier powers follow an exponential distribution (equivalently a  $\chi^2$  distribution with two degrees of freedom) after appropriate normalisation [2]. A genuine periodic signal is therefore identified as a statistically significant excess above this noise background.

The role of the observing duration can be visualised directly. A short observation contains fewer cycles of a repeating signal, and the corresponding Fourier peak is less prominent relative to the noise background. A longer observation contains more repetitions of the same signal, and the coherent power grows accordingly. This is why, all else being equal, longer integrations substantially improve sensitivity to faint periodic sources. Figure 3.7 illustrates this principle by comparing noisy time series and their corresponding Fourier spectra for short and long integrations.

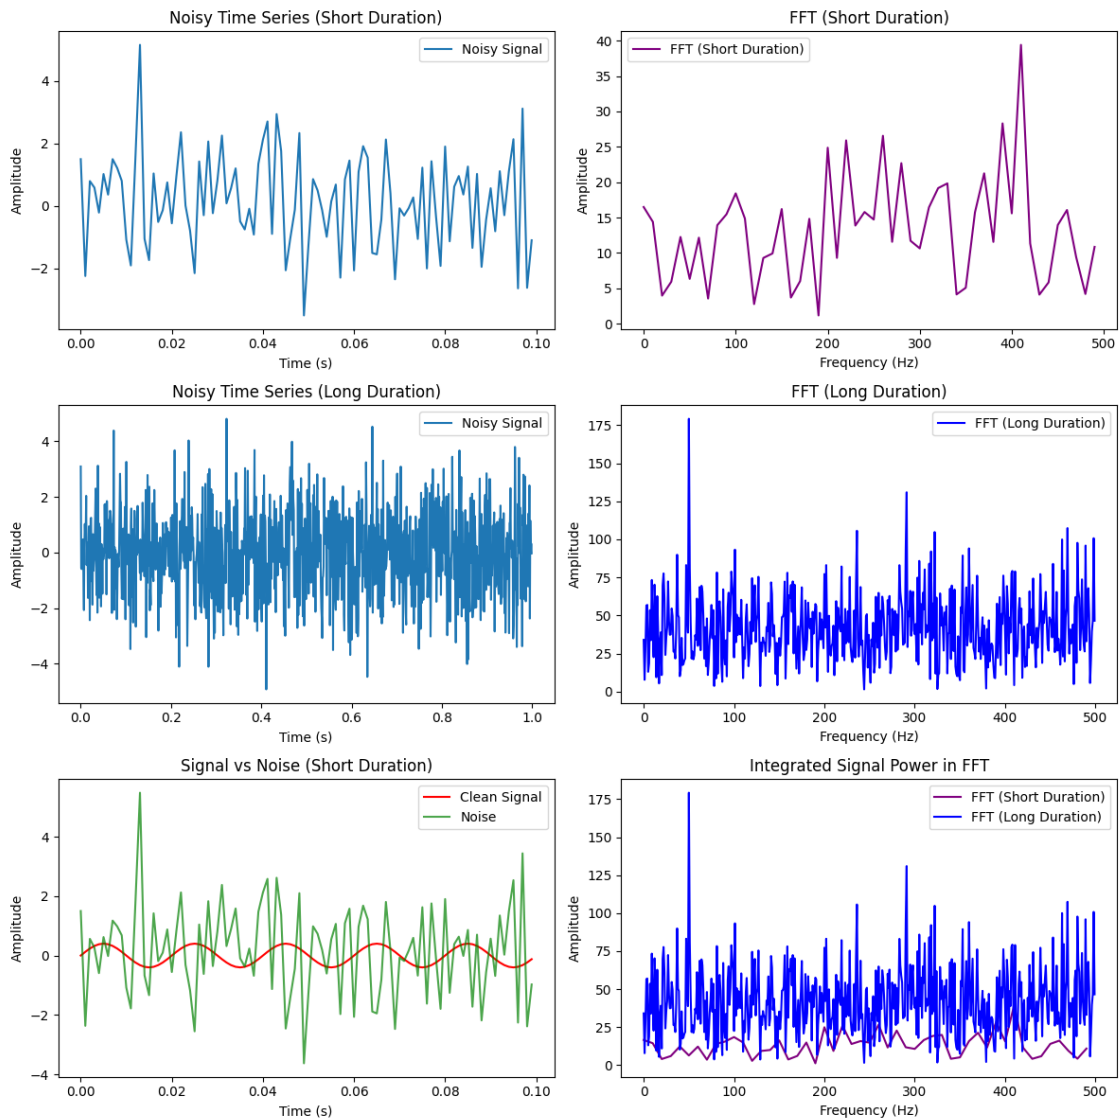


Figure 3.7: Illustration of the power of Fourier-domain periodicity searches. *Top panels:* a short-duration noisy time series and its Fourier spectrum. Although the periodic signal is difficult to identify by eye in the time domain, the FFT already reveals a localized excess at the signal frequency. *Middle panels:* the same experiment for a longer observation. The time-domain signal still appears noise dominated, but the Fourier peak becomes much more significant because the periodic component adds coherently over a larger number of cycles. *Bottom panels:* comparison of the clean periodic component with the noise realization in the time domain, and the corresponding growth of coherent signal power in the Fourier spectrum. This demonstrates why the FFT is such a powerful tool for uncovering repeating signals buried in noise.

For pulsar searches, the relevant repetition frequency is the spin frequency of the neutron star, which for an isolated pulsar is

$$f_0 = \frac{1}{P}, \quad (3.26)$$

where  $P$  is the pulsar spin period. In the simplest case, an isolated pulsar will therefore produce a peak in the Fourier spectrum at  $\nu = f_0$  together with additional power at integer multiples of this frequency, (see discussion in Section 3.6.3).

### 3.6.2 Implementation in PRESTO: the realfft stage

Although the DFT in Eq. 3.21 provides the formal basis of Fourier-domain searching, a direct evaluation of that expression is computationally prohibitive for modern pulsar datasets. The naive calculation requires of order  $N^2$  operations, which is impractical when  $N$  may be of order  $10^7$ – $10^9$  samples in a single dedispersed time series. The practical solution is the fast Fourier transform (FFT), which exploits symmetries in the DFT to reduce the computational cost to order  $N \log N$  [2, 105]. This reduction is precisely what makes large modern pulsar surveys feasible.

In a PRESTO-based pulsar search workflow, the dedispersed time series generated by `prepdata` are transformed into the Fourier domain using the routine `realfft`<sup>3</sup>. The name reflects an important property of pulsar search data: the dedispersed time series are real-valued. For a real input sequence, the Fourier transform obeys the Hermitian symmetry

$$\tilde{x}_{N-k} = \tilde{x}_k^*, \quad (3.27)$$

where the asterisk denotes complex conjugation. This means that the negative- and positive-frequency halves of the spectrum contain redundant information. In practice, only the non-redundant positive-frequency half of the Fourier spectrum must be retained, making the calculation more efficient in both time and storage.

Operationally, `realfft` takes each dedispersed time series and computes the complex Fourier amplitudes  $\tilde{x}_k$ . These are then converted into a power spectrum through Eq. 3.25. The resulting spectrum is the fundamental search product: any strictly periodic or quasi-periodic signal present in the time series will appear as one localized peak ( and potentially, a number of harmonics) above the noise floor.

A useful way to interpret this computational step is to view the FFT as a bank of simultaneous sinusoidal filters. Rather than testing one trial frequency at a time, the transform evaluates the response of the data to all Fourier frequencies  $\nu_k = k/T_{\text{obs}}$  at once. This parallelism is one of the main reasons why FFT-based pulsar searches are

---

<sup>3</sup><https://www.cv.nrao.edu/~sransom/presto/>

so powerful. A single transform of a dedispersed time series searches an enormous range of candidate spin periods in one operation.

In practice, the raw Fourier power spectrum is not yet ready for candidate selection. Instrumental effects, residual red noise, and baseline variations can distort the low-frequency end of the spectrum and modify the underlying noise statistics. For this reason, PRESTO applies additional processing steps, including local normalisation and spectral whitening, after the FFT stage and before the final peak selection [103, 104]. These operations were introduced conceptually in Section 3.5.9 and are part of the reason that genuine pulsar signals can be distinguished from broad low-frequency noise features.

Nevertheless, the `realfft` step is the moment at which the search problem changes character. Before the FFT, the data are a noisy time series in which a weak repeating pulsar signal may be impossible to see by eye. After the FFT, the same repetition becomes a compact spectral feature. In this sense, `realfft` provides the central bridge between time-domain radio data and Fourier-domain periodicity searches.

### 3.6.3 Harmonic structure and harmonic summing

A real pulsar does not emit a sinusoidal waveform. Instead, it produces a train of narrow pulses separated by one spin period  $P$ . This distinction is crucial, because the Fourier transform of a non-sinusoidal periodic signal is not confined to a single frequency. Instead, its power is distributed among the fundamental spin frequency and a series of harmonics at integer multiples of that frequency [2, 103].

This can be demonstrated explicitly with a simple model. Consider a periodic train of top-hat pulses of amplitude  $A$ , width  $W$ , and period  $P$ . Over one period, the signal may be written as

$$s(t) = \begin{cases} A, & 0 \leq t < W, \\ 0, & W \leq t < P, \end{cases} \quad s(t + P) = s(t). \quad (3.28)$$

The complex Fourier coefficient of the  $m$ th harmonic is

$$c_m = \frac{1}{P} \int_0^P s(t) \exp\left(-2\pi im \frac{t}{P}\right) dt. \quad (3.29)$$

Substituting Eq. 3.28 gives

$$c_m = \frac{A}{P} \int_0^W \exp\left(-2\pi im \frac{t}{P}\right) dt, \quad (3.30)$$

which evaluates to

$$c_m = A \frac{W}{P} \operatorname{sinc}\left(m\pi \frac{W}{P}\right) \exp\left(-i\pi m \frac{W}{P}\right), \quad (3.31)$$

where  $\operatorname{sinc}(x) \equiv \sin x/x$ . The corresponding harmonic power is therefore

$$|c_m|^2 \propto \left(\frac{W}{P}\right)^2 \operatorname{sinc}^2\left(m\pi \frac{W}{P}\right). \quad (3.32)$$

Equation 3.32 captures an important practical fact about pulsar searches. If the duty cycle  $W/P$  is large, the signal is relatively sinusoidal and most of its power is concentrated in the fundamental frequency. If the duty cycle is small, corresponding to a narrow pulse, the signal contains substantial power in many harmonics. In other words, narrow pulses produce broader harmonic content in the Fourier domain.

This immediately explains why pulsar searches do not simply look for one strong Fourier bin. A narrow-pulsed pulsar may have only modest power at the fundamental, while a substantial fraction of the detectable signal is distributed among higher harmonics. To recover this power, pulsar search algorithms perform *harmonic summing*. The basic idea is to stretch the Fourier spectrum by integer factors and add it back to the original spectrum so that harmonically related peaks align at the location of the fundamental.

In idealized notation, if  $P_k$  denotes the original power spectrum, then an  $n_{\text{harm}}$  harmonic sum may be written schematically as

$$H_{n_{\text{harm}}}(k) = \sum_{m=1}^{n_{\text{harm}}} P(mk), \quad (3.33)$$

where the rebinned or stretched spectra are added after being mapped onto the bin index corresponding to the fundamental. In practical software, this procedure is implemented through resampling and interpolation in Fourier space rather than by the simple symbolic form of Eq. 3.33, but the physical meaning is the same.

The gain in sensitivity can be substantial for narrow-pulse signals. If a pulsar has appreciable power in, for example, the first 2, 4, 8, or 16 harmonics, then summing those harmonics can increase the significance of the candidate relative to the noise background. This is why PRESTO and similar pulsar-search packages routinely search multiple harmonic sums rather than only the unsummed spectrum [103, 104].

There is, however, a trade-off. Summing too many harmonics can begin to add more noise than signal if the true harmonic content falls off rapidly. For this reason search pipelines typically examine several harmonic sums separately and retain the maximum significance obtained across them. A nearly sinusoidal signal is usually

best detected in the unsummed or low-order summed spectrum, whereas a narrow duty cycle pulsar is often best detected after summing many harmonics.

Harmonic summing is therefore a natural extension of the FFT search itself. The FFT reveals periodicity by concentrating coherent signal power into Fourier space, and harmonic summing recovers the additional coherent power distributed among integer multiples of the spin frequency. Together, these two steps make Fourier searches highly sensitive to the non-sinusoidal pulse trains produced by real pulsars.

### 3.6.4 Binary motion and Doppler modulation

The Fourier-domain search methods described in the previous sections assume that the pulsar spin frequency remains constant throughout the observation. This is a good approximation for isolated pulsars, but it breaks down for pulsars in binary systems. As discussed earlier in Section 1.6.2, orbital motion introduces a time-dependent Doppler shift in the observed pulsar spin frequency. Even if the intrinsic spin of the neutron star is perfectly stable, the changing line-of-sight velocity causes the observed frequency to vary during the observation [2, 34, 103].

In the non-relativistic limit, appropriate for binary pulsars with  $|v_{\parallel}| \ll c$ , the observed spin frequency may be written as

$$f_{\text{obs}}(t) = f_0 \left( 1 + \frac{v_{\parallel}(t)}{c} \right), \quad (3.34)$$

where  $f_0$  is the intrinsic pulsar spin frequency,  $v_{\parallel}(t)$  is the line-of-sight component of the orbital velocity, and  $c$  is the speed of light. Equation 3.34 immediately shows that the observed spin frequency is no longer constant, but instead tracks the orbital motion of the pulsar around the binary centre of mass.

For a circular orbit, the line-of-sight velocity varies sinusoidally with orbital phase. It is therefore convenient to write

$$v_{\parallel}(t) = K \sin \left( \frac{2\pi t}{P_b} + \phi_0 \right), \quad (3.35)$$

where  $K$  is the projected orbital velocity amplitude,  $P_b$  is the binary orbital period, and  $\phi_0$  is a constant orbital phase. Substituting Eq. 3.35 into Eq. 3.34 gives

$$f_{\text{obs}}(t) = f_0 \left[ 1 + \frac{K}{c} \sin \left( \frac{2\pi t}{P_b} + \phi_0 \right) \right]. \quad (3.36)$$

This relation shows that orbital motion modulates the apparent spin frequency of the pulsar even when the intrinsic spin frequency  $f_0$  is constant. The observed signal is therefore not strictly monochromatic over long integrations. Instead, the

instantaneous frequency drifts during the observation, and this drift reduces the effectiveness of a standard FFT search, which assumes that the signal remains confined to a single Fourier bin.

The physical origin of this effect is illustrated by the line-of-sight acceleration and jerk curves shown in Fig. 3.8. These quantities are the first and second time derivatives of the orbital velocity and provide the most direct bridge between binary orbital dynamics and pulsar search algorithms.

### 3.6.5 Line-of-sight acceleration and jerk

To understand how binary motion affects pulsar searches, it is useful to expand the Doppler modulation in terms of successive time derivatives. Differentiating Eq. 3.34 with respect to time gives

$$\dot{f}_{\text{obs}}(t) = f_0 \frac{1}{c} \frac{dv_{\parallel}}{dt}. \quad (3.37)$$

Since the time derivative of the line-of-sight velocity is the line-of-sight acceleration,

$$a_{\parallel}(t) \equiv \frac{dv_{\parallel}}{dt}, \quad (3.38)$$

we obtain the important relation

$$\dot{f}_{\text{obs}}(t) = \frac{f_0}{c} a_{\parallel}(t). \quad (3.39)$$

Equation 3.39 is central to acceleration searches. It shows that a pulsar undergoing a non-zero line-of-sight acceleration has a non-zero observed frequency derivative, even when its intrinsic spin is constant over the observation.

For the circular-orbit velocity model in Eq. 3.35, the line-of-sight acceleration is obtained directly by differentiation:

$$a_{\parallel}(t) = \frac{dv_{\parallel}}{dt} = \frac{2\pi K}{P_b} \cos\left(\frac{2\pi t}{P_b} + \phi_0\right). \quad (3.40)$$

The maximum magnitude of the line-of-sight acceleration is therefore

$$|a_{\parallel}|_{\text{max}} = \frac{2\pi K}{P_b}. \quad (3.41)$$

This expression immediately shows that compact binaries with short orbital periods produce the largest accelerations and are therefore the most strongly affected by Doppler smearing during a pulsar search.

The next derivative introduces the line-of-sight jerk, defined as

$$j_{\parallel}(t) \equiv \frac{da_{\parallel}}{dt} = \frac{d^2v_{\parallel}}{dt^2}. \quad (3.42)$$

Differentiating Eq. 3.39 with respect to time gives

$$\ddot{f}_{\text{obs}}(t) = \frac{f_0}{c} j_{\parallel}(t), \quad (3.43)$$

so the second derivative of the observed spin frequency is directly proportional to the line-of-sight jerk.

Using Eq. 3.40, the jerk for a circular orbit becomes

$$j_{\parallel}(t) = \frac{da_{\parallel}}{dt} = - \left( \frac{2\pi}{P_b} \right)^2 K \sin \left( \frac{2\pi t}{P_b} + \phi_0 \right). \quad (3.44)$$

Its maximum magnitude is therefore

$$|j_{\parallel}|_{\text{max}} = \left( \frac{2\pi}{P_b} \right)^2 K. \quad (3.45)$$

Equations 3.40 and 3.44 are illustrated in Fig. 3.8, which shows the variation of  $a_{\parallel}(t)$  and  $j_{\parallel}(t)$  over one full binary orbit. The upper panel shows that the line-of-sight acceleration varies smoothly and periodically, changing sign as the pulsar moves from one side of the orbit to the other. The lower panel shows the corresponding jerk, which describes how rapidly the acceleration itself changes with orbital phase. In other words, acceleration controls the instantaneous frequency drift of the signal, while jerk controls how that drift rate evolves during the observation.

This distinction is important for pulsar searches. If the observation is short compared with the orbital timescale, the acceleration may be treated as approximately constant over the duration of the observation. In that case,  $\ddot{f}_{\text{obs}}$  may be neglected and the signal can be modelled with a constant frequency derivative. If, however, the observation is long enough (or the orbital phase unfavourable enough) that  $a_{\parallel}(t)$  changes significantly during the integration, then the jerk term becomes non-negligible and the constant-acceleration approximation begins to fail [103, 106, 107].

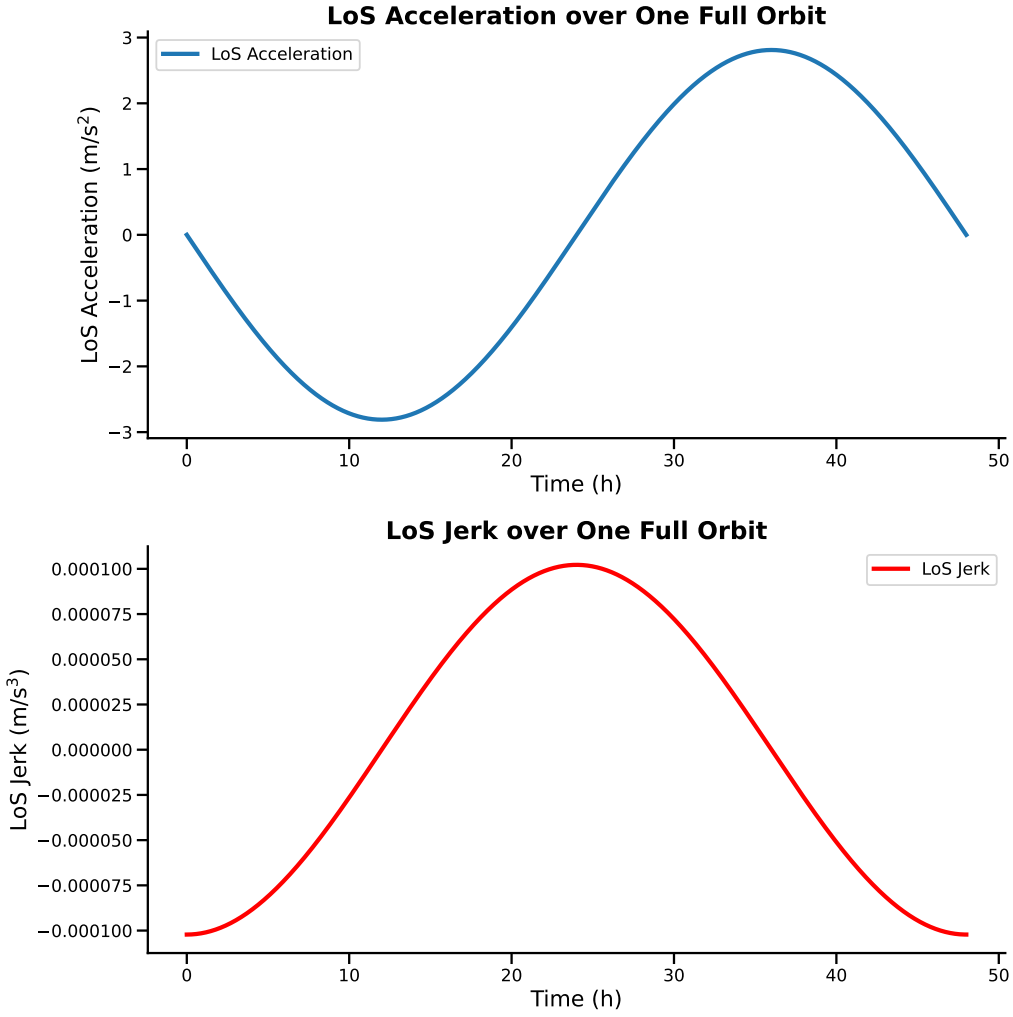


Figure 3.8: Illustration of the line-of-sight orbital acceleration  $a_{\parallel}(t)$  and line-of-sight jerk  $j_{\parallel}(t)$  over one full binary orbit for a circular system. The upper panel shows the periodic variation of the line-of-sight acceleration, which determines the instantaneous drift rate of the observed pulsar spin frequency through  $\dot{f}_{\text{obs}} = (f_0/c) a_{\parallel}$ . The lower panel shows the corresponding line-of-sight jerk, which determines the rate of change of that drift through  $\ddot{f}_{\text{obs}} = (f_0/c) j_{\parallel}$ . Together, these quantities provide the physical origin of Doppler smearing in Fourier-domain pulsar searches.

### 3.6.6 Signal drift in the Fourier domain

The relations derived above connect orbital dynamics directly to pulsar search behaviour. A standard FFT search assumes that the signal remains at a constant frequency throughout the observing time  $T_{\text{obs}}$ . In that case, the signal power accumulates coherently in a single Fourier bin of width

$$\delta\nu = \frac{1}{T_{\text{obs}}}. \quad (3.46)$$

However, for a pulsar in a binary system the observed frequency is time-dependent. Expanding  $f_{\text{obs}}(t)$  about the start or midpoint of the observation gives

$$f_{\text{obs}}(t) \approx f_0 + \dot{f} t + \frac{1}{2} \ddot{f} t^2 + \dots, \quad (3.47)$$

where, over a sufficiently short observation,  $\dot{f}$  and  $\ddot{f}$  may be treated as approximately constant.

If the jerk is negligible over the observation, then Eq. 3.47 reduces to

$$f_{\text{obs}}(t) \approx f_0 + \dot{f} t. \quad (3.48)$$

This is the constant-acceleration approximation. The total frequency excursion during the observation is then approximately

$$\Delta f \approx \dot{f} T_{\text{obs}}. \quad (3.49)$$

Dividing this by the Fourier-bin width from Eq. 3.46 gives the number of Fourier bins drifted by the signal:

$$z \approx \frac{\Delta f}{\delta \nu} = \dot{f} T_{\text{obs}}^2. \quad (3.50)$$

Using Eq. 3.39, this may be written as

$$z \approx \frac{f_0 a_{\parallel}}{c} T_{\text{obs}}^2. \quad (3.51)$$

Equation 3.51 is one of the most useful practical expressions in binary pulsar searching. It shows that the degree of Fourier smearing increases with spin frequency, line-of-sight acceleration, and the square of the observing time. When  $|z| \ll 1$ , the signal remains effectively confined to one Fourier bin and a standard FFT search retains most of its sensitivity. When  $|z| \gtrsim 1$ , the signal power is spread across several bins and the peak significance is reduced. This effect is illustrated in Fig. 3.9, where the accelerated pulsar produces a broadened and weakened Fourier response compared to the unaccelerated case.

If the jerk is not negligible, then the quadratic term in Eq. 3.47 must also be retained:

$$f_{\text{obs}}(t) \approx f_0 + \dot{f} t + \frac{1}{2} \ddot{f} t^2. \quad (3.52)$$

In this regime the drift is no longer linear in time. Instead, the signal moves through Fourier space with a drift rate that itself changes during the observation.

Using Eq. 3.43, the curvature term is set directly by the line-of-sight jerk:

$$\ddot{f} = \frac{f_0}{c} j_{\parallel}. \quad (3.53)$$

Acceleration therefore produces a *linear* drift of the signal across Fourier bins, whereas jerk introduces *curvature* in that drift [106]. In physical terms, acceleration determines how fast the apparent spin frequency moves through the FFT, while jerk determines how rapidly that drift rate changes during the observation. This is why acceleration searches are effective when the constant-acceleration approximation is valid, but become suboptimal for very compact binaries observed over a significant fraction of their orbit.

A useful estimate of the additional frequency excursion introduced by jerk is

$$\Delta f_{\text{jerk}} \sim \frac{1}{2} \ddot{f} T_{\text{obs}}^2, \quad (3.54)$$

which corresponds to an additional non-linear distortion of the Fourier response. Although this effect is harder to visualise directly than the linear bin drift caused by acceleration, its practical consequence is the same: the signal power is distributed over a larger region of Fourier space, reducing the peak significance in a standard periodicity search.

This behaviour motivates the two broad classes of correction methods used in binary pulsar searches. In the time-domain approach, the time series is resampled to remove the effects of orbital motion before the FFT is computed. In the frequency-domain approach, one first performs the FFT and then searches for the characteristic drift patterns produced by non-zero  $\dot{f}$  and, where relevant,  $\ddot{f}$ . These methods are described in the following sections [34, 103, 104].

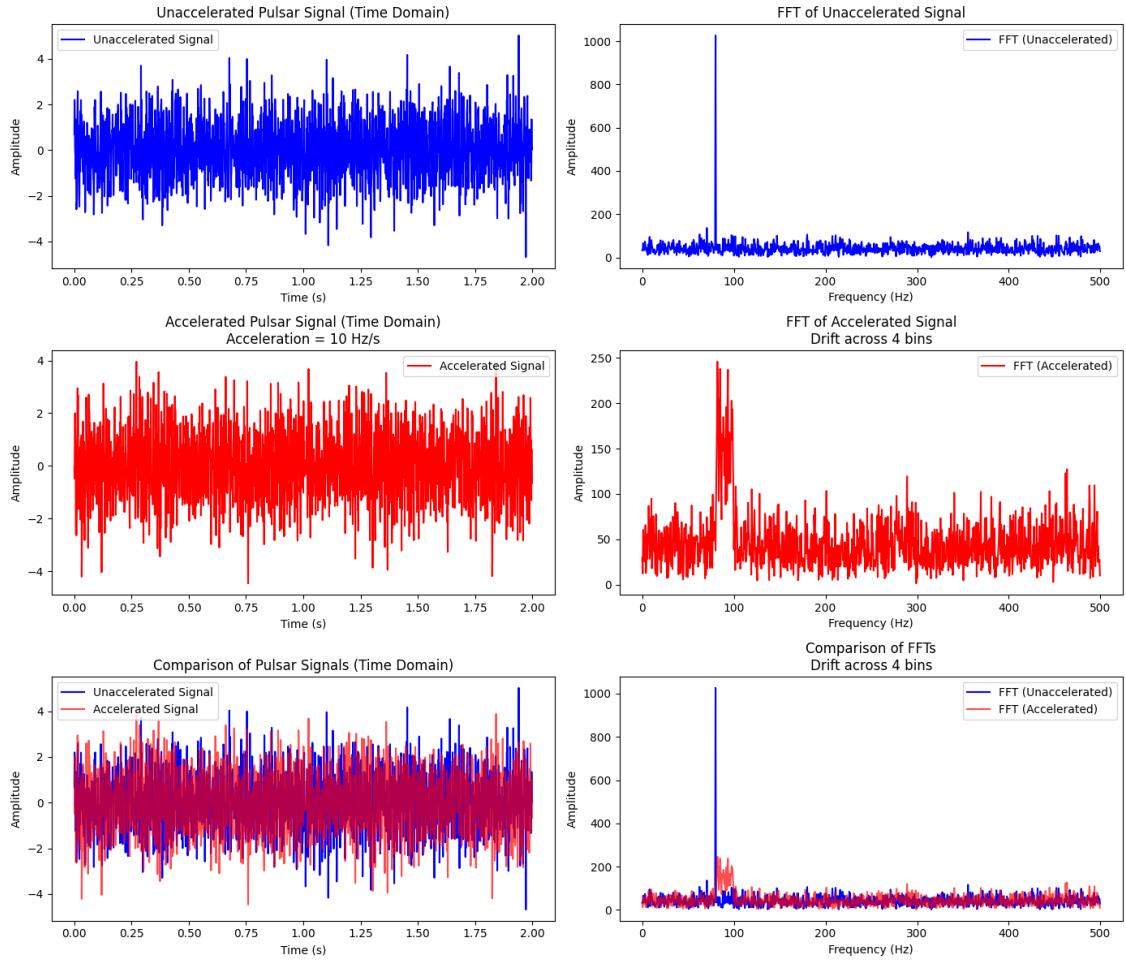


Figure 3.9: Illustration of the effect of binary acceleration on a pulsar signal in the Fourier domain. In the absence of binary motion (top panels), the periodic signal remains at a fixed frequency throughout the observation and appears as a sharp peak in the FFT. Under non-zero line-of-sight acceleration (central panels), the observed spin frequency drifts during the observation and the signal power is spread across multiple Fourier bins. This Doppler smearing reduces the peak significance in a standard FFT search and motivates the use of acceleration-search techniques (see comparison in bottom panels).

### Phase evolution of an accelerated signal

It is useful to begin from the signal phase rather than the instantaneous frequency. Since frequency is the time derivative of phase divided by  $2\pi$ , the phase of the observed signal is

$$\Phi(t) = 2\pi \int_0^t f_{\text{obs}}(t') dt'. \quad (3.55)$$

Substituting Eq. 3.52 into Eq. 3.55 gives

$$\Phi(t) \approx 2\pi \left( f_0 t + \frac{1}{2} \dot{f} t^2 + \frac{1}{6} \ddot{f} t^3 \right). \quad (3.56)$$

Equation 3.56 shows explicitly how binary motion destroys the simple phase model of an isolated pulsar. In the absence of orbital effects,  $\Phi(t) = 2\pi f_0 t$ , so successive pulses remain perfectly phase coherent and all signal power accumulates into a narrow set of Fourier bins. The additional quadratic and cubic terms represent phase errors relative to this simple model. The quadratic term arises from a non-zero acceleration and the cubic term from a non-zero jerk. If these terms become appreciable over the interval  $T_{\text{obs}}$ , the signal power is smeared and the sensitivity of a standard FFT search degrades.

Using the results from the previous section, one may write

$$\dot{f} = \frac{f_0}{c} a_{\parallel}, \quad \ddot{f} = \frac{f_0}{c} j_{\parallel}, \quad (3.57)$$

so that Eq. 3.56 becomes

$$\Phi(t) \approx 2\pi f_0 \left[ t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 \right]. \quad (3.58)$$

This form is especially useful because it shows that the search problem may be viewed either as one of correcting the time coordinate or as one of correcting the phase evolution. These two viewpoints lead naturally to time-domain and frequency-domain search methods, respectively.

### Time-domain methods: resampling before the FFT

The guiding idea of the time-domain approach is simple: if orbital motion causes the pulse train to arrive too early or too late relative to a uniformly spaced clock, then one should undo that distortion before carrying out the Fourier analysis. In other words, one seeks a new time coordinate in which the pulsar once again appears to spin at a constant frequency.

Suppose that a corrected time coordinate  $\tau$  is defined such that the phase is written in the isolated-pulsar form

$$\Phi(t) = 2\pi f_0 \tau. \quad (3.59)$$

Comparing Eq. 3.59 with Eq. 3.58 yields

$$\tau = t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 + \dots. \quad (3.60)$$

Equation 3.60 is the essence of the time-domain correction: binary motion is removed by redefining the time axis so that the pulsar signal is demodulated to a nearly constant frequency. If only the acceleration term is retained, one obtains the constant-acceleration approximation; if the jerk term is also included, one obtains the next-order correction. In the most general case, one may instead use a full orbital timing model and remove the Roemer delay due to the orbit directly.

In actual search pipelines the data are not continuous but discretely sampled. Hence, the corrected time series must be obtained by interpolation or nearest-neighbour assignment. Let the observed time series be sampled at uniform intervals  $t_{\text{samp}}$ , and denote its samples by  $s_i$ . A resampled time series  $s_j^{\text{rs}}$  is then constructed by evaluating which observed sample corresponds to each corrected time  $\tau_j = j t_{\text{samp}}$ . In the simplest implementation one identifies the nearest sample in the original series and copies its value into the resampled sequence. More sophisticated approaches use interpolation between neighbouring samples. The underlying idea is the same in either case: the signal is demodulated in the time domain before any Fourier transform is computed [34, 103].

The important computational consequence is that this procedure must be repeated for every trial correction. If one searches over a set of trial accelerations, then for each trial one must construct a new resampled time series and compute its FFT. If one extends the search to include jerk, or even a full orbital template bank, the number of resampling trials increases rapidly. Time-domain methods are therefore very natural and physically intuitive, and they can remain fully coherent over a large class of signal models, but they can also become computationally expensive because the FFT must be repeated for each template.

Despite this cost, time-domain methods have an important advantage: they are conceptually very general. If the orbital model is known or can be parameterised accurately, the signal can in principle be corrected all the way back to a simple periodic train before the FFT is applied. This is the logic behind coherent demodulation schemes and this type of time-domain resampling.

### **Frequency-domain methods: FFT first, correction afterwards**

The alternative strategy is to accept that the observed time series is already modulated by orbital motion, compute its FFT once, and then search the Fourier spectrum for the signature of that modulation. In this approach the orbital correction is not applied to the data samples themselves, but to the spectral response of the signal.

The idea of a frequency-domain acceleration search is to compute the FFT once and then correlate the Fourier amplitudes with a family of templates corresponding

to different trial values of  $z$ , the number of drifted bins defined in equation 3.4.8. Each template represents the expected Fourier response of a signal whose apparent frequency drifts linearly during the observation. By scanning over a range of  $z$ , the algorithm attempts to reconstruct the coherent power that has been spread across the spectrum by acceleration [103, 104].

This is the basis of the standard Fourier-domain acceleration search implemented in PRESTO through `accelsearch`. This is the basis of the standard Fourier-domain acceleration search implemented in PRESTO through `accelsearch`. In practical terms, one first computes the FFT of the dedispersed time series. For each harmonic-summing stage, the Fourier amplitudes are then correlated with matched-filter templates over a range of trial  $z$  values, and the resulting responses from harmonically related signals are combined. Thus, the acceleration matched filtering is applied at each harmonic stage, rather than being performed only after a single final harmonic sum. Because the FFT is computed only once per dedispersed time series, the method is much faster than a brute-force time-domain resampling search over many trial accelerations. This computational efficiency is the main reason why frequency-domain acceleration searches have become standard in large pulsar surveys.

The price paid for this efficiency is that the method is usually derived under the assumption that the acceleration remains approximately constant over the observation. This requirement is the reason why standard acceleration searches work best when the observation spans only a modest fraction of the binary orbit. If the observation becomes too long, the jerk term can no longer be neglected, the drift through Fourier space is no longer linear, and the standard constant-acceleration templates become progressively mismatched to the signal. A commonly used rule of thumb is that the constant-acceleration approximation remains adequate when  $T_{\text{obs}} \lesssim P_b/10$ , although the exact limit depends on the orbital phase and system parameters [34, 35].

When the acceleration is not constant over the duration of the observation, higher-order derivatives of the line-of-sight motion become important, in particular the *jerk*. In this regime, the apparent spin frequency no longer drifts linearly in time, but follows a curved trajectory in Fourier space. As a result, standard acceleration searches, which correct only for a constant frequency derivative, become increasingly ineffective. To recover sensitivity, one must extend the search to include the jerk term, introducing an additional parameter (commonly denoted  $w$ ) that accounts for the second-order frequency derivative. This leads to a two-dimensional search over  $(z, w)$ , allowing the templates to better match signals from compact binaries where orbital motion varies significantly over the observation timescale [106].

### 3.6.7 Candidate generation, sifting, and folding

The Fourier-domain search techniques described in the previous sections produce a very large number of potential periodic signals. In practice, the vast majority of these detections do not correspond to astrophysical pulsars but instead arise from random noise fluctuations, residual radio-frequency interference (RFI), or instrumental artefacts. Modern pulsar search pipelines therefore include several additional stages that progressively reduce the number of candidates before detailed inspection is performed. These stages typically include *candidate generation*, *candidate sifting*, and finally *candidate folding and diagnostic analysis* [2, 103, 104].

#### Candidate generation in Fourier searches

Following the FFT and harmonic summing steps, the search algorithm scans the Fourier spectrum for peaks exceeding a predefined statistical threshold. Each peak corresponds to a potential periodic signal characterised by a trial spin frequency  $f$ , dispersion measure (DM), and, in acceleration searches (and also a trial second derivative in case of Jerk), a trial frequency derivative  $\dot{f}$ .

In practice, the search is performed over many trial dispersion measures and, where applicable, trial accelerations and jerks. A typical observation may therefore produce tens of thousands or even millions of raw candidates. For each detection above the chosen spectral signal-to-noise threshold, the pipeline records several parameters including

- the candidate spin period  $P$ ,
- the trial dispersion measure,
- the signal-to-noise ratio of the detection,
- the harmonic number of the detection,
- the trial acceleration (or frequency derivative).
- the trial jerk (the second derivative of frequency).

These raw detections form the initial candidate list that must be further processed before meaningful astrophysical interpretation is possible.

## Candidate sifting

The next stage is *candidate sifting*, whose purpose is to reduce the large list of raw detections to a manageable number of promising pulsar candidates. Candidate sifting typically involves several filtering operations designed to eliminate obvious duplicates and non-astrophysical signals [103, 108].

A common feature in Fourier searches is that a single pulsar can appear multiple times in the candidate list. For example, detections may occur at the fundamental frequency as well as at several harmonics. Similarly, the same signal may appear at nearby trial dispersion measures or accelerations. Candidate sifting algorithms therefore identify groups of detections that correspond to the same underlying periodic signal and retain only the most significant representative of each group.

Additional filtering criteria are often applied during this stage. Candidates with extremely low dispersion measures are frequently rejected because terrestrial interference is not dispersed by the interstellar medium. Similarly, candidates with periods corresponding to known instrumental frequencies or power-line harmonics may be discarded. Other heuristics may remove candidates whose signal-to-noise ratios are below a specified threshold or whose detection significance decreases rapidly when harmonics are combined.

In software packages such as PRESTO, the sifting procedure is implemented by dedicated scripts that analyse the candidate lists produced by the periodicity search. The output of this stage is a much smaller set of unique candidate signals that warrant further investigation [103, 104].

## Folding of pulsar candidates

Once the candidate list has been reduced through sifting, the remaining signals are examined in greater detail by *folding* the time series at the candidate period and dispersion measure. Folding is the process of aligning successive rotations of the pulsar signal and summing them coherently to produce an integrated pulse profile.

For a candidate with period  $P$ , each sample in the dedispersed time series is assigned a pulse phase

$$\phi(t) = \left( \frac{t}{P} \right) \bmod 1, \quad (3.61)$$

where  $\phi \in [0, 1)$  represents the rotational phase within a single pulse period. The data are then accumulated into a set of phase bins to construct the average pulse profile. If a genuine pulsar is present, the periodic pulses add coherently and produce a stable pulse shape, while random noise averages out.

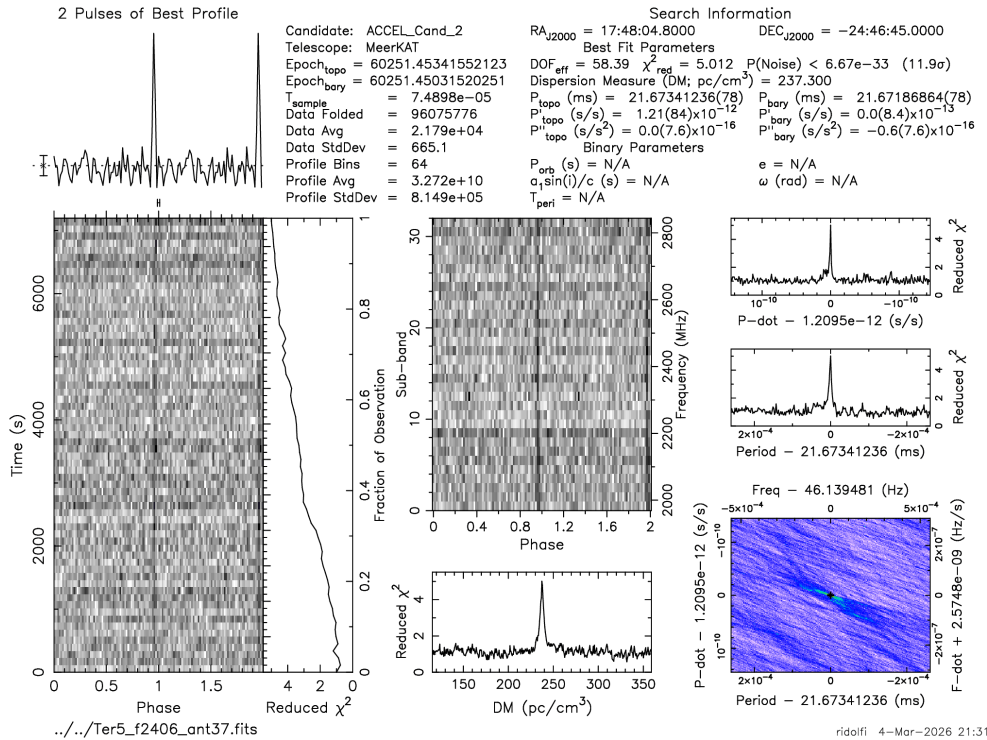


Figure 3.10: A typical pulsar diagnostic plot prepared by PRESTO’s Prepfold routine.

In practice the candidate period obtained from the FFT search is not perfectly precise. Folding routines therefore perform a local optimisation of several parameters, including the period  $P$ , its derivative  $\dot{P}$ , and the dispersion measure. The optimisation is typically carried out by searching over small ranges around the candidate parameters and selecting the combination that maximises a detection statistic such as the reduced  $\chi^2$  of the folded profile.

### Interpretation of pulsar diagnostic plots

After candidate sifting, the remaining signals are folded and summarised in diagnostic plots that allow rapid visual assessment of their astrophysical credibility. Figure 3.10 shows a typical example of such a diagnostic plot produced by PRESTO. These plots combine several independent views of the data that together provide strong evidence for or against the presence of a genuine pulsar signal.

Each panel of the diagnostic plot probes a different aspect of the signal, and experienced pulsar searchers use the combined behaviour of these panels to evaluate candidates.

**Integrated pulse profile** The integrated pulse profile represents the average signal intensity as a function of pulse phase. This profile is obtained by summing many thousands or millions of individual rotations of the candidate signal.

For a genuine pulsar, the integrated profile typically shows a stable and well-defined pulse shape, often consisting of one or more components. The signal-to-noise ratio of this profile increases as the square root of the number of pulses integrated [2].

Noise fluctuations usually produce irregular or poorly defined profiles, while narrow spikes confined to a single phase bin often indicate radio-frequency interference (RFI). Real pulsar profiles tend to have finite width and remain consistent when the data are divided into smaller time segments.

**Pulse phase as a function of time** The time–phase diagram shows how the pulse phase evolves throughout the observation. Each horizontal slice corresponds to a short segment of the data folded at the candidate period.

For a real pulsar, the pulse appears as a vertical stripe that remains at a constant phase throughout the observation. This indicates that the period, the first as well as the second period derivative used for folding accurately tracks the rotational phase of the pulsar.

If the candidate period is slightly incorrect, the pulse will drift diagonally across the plot as the observation progresses. If the first period derivative ( $\dot{P}$ ) is incorrect, the pulse track follows a quadratic curvature, while an unaccounted or incorrect second derivative ( $\ddot{P}$ ) produces a characteristic S-shaped structure. Random noise produces no coherent pattern, while intermittent interference may appear only in isolated time segments.

**Pulse phase as a function of observing frequency** The frequency–phase panel displays the folded pulse signal across the observing bandwidth, typically divided into sub-bands.

A genuine pulsar signal should appear aligned across frequency channels when the correct dispersion measure is applied. Because the interstellar medium disperses radio signals, an incorrect DM will cause the pulse to appear slanted across the frequency axis.

Terrestrial interference, which is not dispersed, often produces signals that align at  $DM \approx 0$ , or that appear only in restricted frequency ranges (and have therefore similarly high S/N over large ranges of DM).

**Dispersion measure curve** The dispersion-measure curve shows the detection significance as a function of trial DM. A real pulsar should produce a clear peak at a well-defined non-zero DM corresponding to the integrated electron column density along the line of sight.

Noise signals usually produce curves with no clear maximum. RFI signals often peak near  $DM \approx 0$  because they originate within the Earth’s atmosphere or instrumentation and therefore do not experience interstellar dispersion.

**Period and acceleration optimisation** Additional panels display the detection statistic as a function of small adjustments to the candidate period and its derivative (or acceleration). These curves indicate whether the signal strength increases when the folding parameters are refined.

Real pulsar candidates typically produce sharp peaks centred on the optimal parameter values. In contrast, noise candidates often produce broad or irregular structures without a well-defined maximum.

**Period–acceleration search plane** The two-dimensional panel in the lower-right portion of the diagnostic plot shows the detection statistic across a grid of trial periods and period derivatives. The candidate location is marked within this parameter space.

For accelerated pulsars in binary systems, this panel often shows a distinct peak or ridge corresponding to the best-fit combination of period and acceleration. The surrounding pattern reflects correlations between these parameters during the optimisation process.

**Common false-positive signatures** In practice, many candidates that initially appear significant are later identified as non-astrophysical signals. Several common patterns can often be recognised directly in the diagnostic plots:

- **Broadband RFI:** Narrow spikes in the integrated profile, often accompanied by signals that appear simultaneously across all frequencies and peak near  $DM \approx 0$ .
- **Narrow-band RFI:** Signals confined to a small number of adjacent frequency channels, often appearing with high S/N and at DMs different from zero, as they are largely unaffected by the de-dispersion process.
- **Noise fluctuations:** Weak peaks that disappear when the data are subdivided in time or frequency.

- **Incorrect DM solutions:** Signals that appear strong only over a limited frequency range or whose pulse phase varies systematically across the observing band.

Identifying these failure modes is an essential part of pulsar candidate evaluation. Although automated classification algorithms are increasingly used in modern surveys, visual inspection of diagnostic plots remains an important step in validating new pulsar discoveries [108, 109].

## 3.7 Additional pulsar search techniques

The Fourier-based search methods described in the previous sections form the backbone of most modern pulsar surveys. Fast Fourier Transform (FFT) searches combined with harmonic summing and acceleration corrections are highly effective for detecting periodic signals whose spin frequency remains approximately stable over the duration of the observation. However, not all pulsars satisfy these assumptions. In particular, signals from very long-period pulsars, highly accelerated binary systems, or sources that emit sporadic pulses can be difficult to detect using standard Fourier techniques.

To address these limitations, a number of complementary search methods have been developed. These approaches target different signal regimes and can significantly improve sensitivity in situations where traditional FFT-based searches become inefficient.

### 3.7.1 Fast Folding Algorithm (FFA) searches

One important alternative to Fourier-based periodicity searches is the Fast Folding Algorithm (FFA). The FFA operates directly in the time domain by folding the time series at a large number of trial periods and evaluating the statistical significance of the resulting profiles [2, 110].

Unlike FFT searches, which rely on harmonic summing to recover narrow pulse profiles, the FFA preserves the full coherence of the signal throughout the folding process. As a result, the FFA can be significantly more sensitive to long-period pulsars and pulsars with broad pulse profiles. This makes it particularly useful for detecting sources whose periods are longer than a few seconds, where the number of observed rotations during the observation is relatively small.

Historically, the computational cost of FFA searches limited their use in large surveys. However, modern implementations and increased computing power have

renewed interest in this technique, and several recent pulsar discoveries have been made using FFA-based pipelines [111, 112].

### 3.7.2 Single-pulse searches

Some neutron stars emit detectable radio bursts that do not form a stable periodic signal during a typical observation. Such sources include rotating radio transients (RRATs), intermittent pulsars, and some magnetars. These objects may produce individual bright pulses separated by long intervals of weak or undetectable emission.

Single-pulse search techniques therefore look for isolated dispersed bursts rather than periodic signals. Some of the very first pulsars were discovered through this technique. After dedispersion, the time series is searched for peaks exceeding a specified signal-to-noise threshold. The search is typically performed over a range of boxcar widths in order to detect pulses of different durations [113, 114].

A key diagnostic for identifying astrophysical bursts is the dispersion-measure dependence of the signal. Genuine pulses from distant sources appear at multiple frequencies with a delay that follows the cold-plasma dispersion relation. In contrast, terrestrial interference usually lacks this characteristic behaviour.

Single-pulse searches have become an essential component of modern pulsar pipelines and have led to the discovery of fast radio bursts (FRBs), which are bright millisecond-duration radio transients of extragalactic origin [115].

### 3.7.3 Stacking and phase-modulation techniques

In addition to standard periodicity searches, several techniques have been developed to improve sensitivity to faint or highly accelerated pulsars by combining information across time or frequency in different ways.

A simple approach is *incoherent stacking* of multiple independent observations. In this method, power spectra from separate observing epochs are summed without preserving phase information. This is particularly effective for faint isolated pulsars, where each individual observation may be below the detection threshold, but the combined signal becomes significant when multiple observations are added.

For binary systems, where the apparent spin frequency evolves during the observation, more specialised methods are required. One such approach is the *stack-and-slide* technique. A long observation is divided into shorter segments, within which the pulsar period can be approximated as nearly constant. A Fourier transform is computed for each segment, producing a set of power spectra in which the signal is less smeared across frequency bins. These spectra are then combined by applying

trial frequency drifts corresponding to different accelerations, effectively “sliding” the spectra before summing them. This allows partial recovery of signals from binaries that would otherwise be undetectable in a single long Fourier transform.

In the regime of extremely compact binaries ( $P_b \ll T_{\text{obs}}$ ), the *phase-modulation search* [35] provides an alternative strategy. In this case, the orbital motion imprints a characteristic phase modulation on the pulsar signal, producing a series of evenly spaced sidebands around the intrinsic spin frequency in the Fourier domain. By searching for this distinctive pattern, one can detect pulsars in very tight orbits where both standard acceleration and stack-and-slide techniques lose sensitivity.

### 3.7.4 Template-bank searches for compact binaries

For very short orbital period systems, the constant-acceleration approximation used in standard acceleration searches may no longer be valid. In such cases, the observed pulsar frequency can evolve rapidly during the observation due to higher-order orbital effects.

Template-bank searches attempt to address this problem by constructing a grid of trial orbital parameters and performing coherent demodulation of the time series for each template [116, 117]. This approach is similar in spirit to matched-filter searches used in gravitational-wave data analysis [118, 119].

Although computationally expensive, template-bank methods provide a powerful way to search for pulsars in extremely compact binary systems [120], including systems that may host neutron star–black hole pairs or ultra-short-period double neutron star binaries.

Together, these complementary search techniques extend the discovery space of pulsar surveys beyond the regime where standard FFT-based methods are most effective. This technique is discussed in detail in 5.4.5

# Chapter 4

## Results from the first MeerKAT S-band pulsar survey

### 4.1 Introduction

Globular clusters (GCs) are among the most productive environments for pulsar searches. These dense, gravitationally bound systems contain  $\sim 10^4$ – $10^6$  stars packed into compact volumes, creating ideal conditions for frequent dynamical encounters. Such interactions can exchange companions, harden binaries, and recycle neutron stars through the transfer of mass and angular momentum from a companion star, efficiently producing binary millisecond pulsars (MSPs) that are otherwise comparatively rare in the Galactic field [27, 31]. For this reason, GCs are attractive targets for deep observations, as a single long pointing is likely to host multiple pulsars and can build signal-to-noise for intrinsically faint systems or those affected by scattering and eclipses. At the same time, long dwell times can be detrimental to the detection of pulsars in compact binaries, since orbital acceleration and higher-order motion can smear the signal over extended integrations if not properly accounted for.

The scientific motivation for globular-cluster pulsar surveys is correspondingly broad. Recycled pulsars in compact binaries provide precision laboratories for tests of relativistic gravity [e.g. 121, 122], constraints on the neutron-star equation of state [123], and the study of exotic binaries formed through dynamical interactions in dense stellar systems [e.g. 124]. Ensembles of stable millisecond pulsars contribute to pulsar timing array experiments targeting nanohertz gravitational waves [e.g. 125], while globular-cluster pulsars, by virtue of their common distance and shared environment, serve as sensitive probes of cluster gravitational potentials and the intervening interstellar medium [126]. These drivers have motivated sustained

pulsar-search campaigns over several decades with facilities including Murriyang, the radio telescope at Parkes (New South Wales, Australia), the Green Bank Telescope (GBT; West Virginia, USA), and more recently with upgraded instrumentation and new high-sensitivity facilities such as the Giant Metrewave Radio Telescope (GMRT; India) and the Five-hundred-metre Aperture Spherical Telescope (FAST; China) [127–130].

MeerKAT has become a cornerstone facility for pulsar science in the southern hemisphere. Its combination of collecting area, modern receivers, and wide instantaneous bandwidth has made it a highly effective instrument for deep pulsar searches and follow-up timing observations, including in globular clusters (see Bailes et al. 131). The impact of MeerKAT on globular-cluster pulsar studies is already evident through the Transients and Pulsars with MeerKAT (TRAPUM) and MeerTime programmes, the two pulsar-focused Large Survey Projects carried out with MeerKAT over the last five years. Together, these projects have delivered a substantial fraction of recent globular-cluster pulsar discoveries and continue to expand the accessible parameter space for faint systems and compact binaries [132–134].

Most globular-cluster pulsar surveys to date have been conducted at frequencies near 1–1.4 GHz, where pulsar flux densities are typically higher compared to higher frequency observations and receiver performance has historically been optimized. However, for distant clusters with large foreground dispersion measures (DMs), pulse detectability at such frequencies can be significantly limited by dispersive smearing and, more critically, multi-path scattering in the interstellar medium. The scattering timescale decreases steeply with observing frequency, approximately as  $\tau_{\text{sc}} \propto \nu^{-4}$  [135], such that millisecond pulsars in high-DM systems may suffer substantial pulse broadening at L-band. Higher-frequency searches have therefore been explored with single-dish facilities including Effelsberg, the GBT, and Parkes in targeted studies of heavily scattered or high-DM pulsars, leading to successful detections in otherwise challenging systems [136, 137]. Over the past two decades, several globular clusters have also been observed at S-band with the GBT in dedicated campaigns [e.g. 138? , 139], demonstrating the utility of higher observing frequencies for mitigating dispersive smearing and scattering in dense Galactic-plane sightlines.

Nevertheless, most large globular-cluster pulsar surveys have historically been conducted near 1–1.4 GHz, where pulsar flux densities are typically higher and survey instrumentation has been optimised. Systematic, uniformly processed multi-cluster surveys above 2 GHz remain comparatively uncommon. MeerKAT’s large collecting area, coherent tied-array beamforming capability, and wide instantaneous bandwidth provide a complementary high-sensitivity platform for conducting such

systematic S-band searches across a broader cluster sample.

The work presented here reports the preliminary results of a MeerKAT S-band ( $\nu \simeq 2.4$  GHz) globular cluster pulsar survey of 14 clusters observed during 2023–2024 as part of TRAPUM. We present the outcomes of the search and folding campaign, including the discovery of four millisecond pulsars in the globular cluster Glimpse-C01 (see Fig. 4.2) and the re-detection of 39 previously known pulsars across the full sample (see Table 4.1 and Section 4.4.1). Beyond reporting detections, this chapter also aims to quantify practical S-band performance for future studies. To that end, we (i) estimate S-band flux densities for the detected pulsars, (ii) place these measurements in context by comparing a subset to representative MeerKAT L-band flux-density values for the same objects (see Table 4.4), and (iii) derive illustrative spectral-index constraints for selected non-detections using S-band upper limits (see Section 4.4.5).

## 4.2 Target selection and observations

For this survey, 14 globular clusters were observed with the S1-band of the MeerKAT S-band receiver system, covering a frequency range of 1750–3500 MHz, with an effective usable bandwidth of  $\Delta\nu \simeq 875$  MHz centred near 2406 MHz over a one-year period starting in March 2023. These GCs are listed in Table 4.1. All targets host at least one previously known pulsar and had previously been observed with MeerKAT at L-band as part of the TRAnsients and PULsars with MeerKAT (TRAPUM) survey [132, 133, 140–143], except NGC 6316, Glimpse-C01, NGC 5986, NGC 6539, and NGC 6712.

Several clusters in our sample have previously been searched at S-band or comparable frequencies with other facilities, including Terzan 5, NGC 6440, NGC 6441,  $\omega$  Centauri, and M28 [e.g. 136, 138, 144–146]. NGC 6316 has also been targeted in a recent GBT S-band study [147]. For the remaining clusters in our sample, no dedicated S-band pulsar-search results have yet been reported.

Table 4.1: Summary of globular-cluster pointings in the first MeerKAT S-band pulsar survey. For each target we list the cluster identifier, J2000 right ascension and declination of the telescope pointing, the observation date in MJD, the nominal cluster dispersion measure (DM) used for dedispersing the data, the total integration length, and the sampling interval. The columns “Re-det” and “New-det” indicate whether any known pulsar was re-detected or a new pulsar discovered in that pointing. The DM range searched for each globular cluster was  $\pm 20\%$  of the nominal cluster DM (taken as the median for clusters with multiple pulsars and as the DM of pulsar A in clusters with only one known pulsar), with the DM step sizes calculated using PRESTO’s `DDplan.py`.

Source	RA (J2000)	Dec (J2000)	MJD	DM ( $\text{pc cm}^{-3}$ )	Length (s)	$T_{\text{samp}}$ ( $\mu\text{s}$ )	Re-det	New-det
Glimpse-C01 1km	18:48:49.70	-01:29:50.00	60158.752	491.1	7202	150	A,B	C,D,E
Glimpse-C01 FA	18:48:49.70	-01:29:50.00	60158.752	491.1	7201	150	A	C,F
Glimpse-C01A 1km	18:48:48.21	-01:29:58.27	60158.752	491.17	7200	150		
Glimpse-C01A FA	18:48:48.21	-01:29:58.27	60158.752	491.17	7199	150		
M28 1km	18:24:32.81	-24:52:11.20	60286.630	119.9	7201	75	A,B,C, E,G,H	–
M28 FA	18:24:32.81	-24:52:11.20	60286.630	119.9	7200	75	A,B,C, E,G	–
M22 1km	18:36:23.94	-23:54:17.10	60138.824	91.2	7202	150		–
M22 FA	18:36:23.94	-23:54:17.10	60138.824	91.2	7201	150		–
M22A 1km	18:36:25.4452	-23:54:52.39	60138.824	89.107	7199	150	A	–
M22B 1km	18:36:23.3760	-23:55:20.9129	60138.824	93.772	7200	150	B	–
NGC5986 1km	15:46:03.00	-37:47:11.10	60142.901	92	7201	75		–
NGC5986 FA	15:46:03.00	-37:47:11.10	60142.901	92	7198	75		–
NGC5986A 1km	15:46:03.44	-37:47:10.1	60142.901	92.17	7201	75	A	–
NGC5986A FA	15:46:03.44	-37:47:10.1	60142.901	92.17	7201	75	A	–
NGC6316 1km	17:16:37.30	-28:08:24.40	60398.057	172.1	7202	75	A	–
NGC6316 FA	17:16:37.30	-28:08:24.40	60398.057	172.1	7199	75		–
NGC6440 1km	17:48:52.70	-20:21:36.90	60314.496	222.51	7200	75	A	–
NGC6440 FA	17:48:52.70	-20:21:36.90	60314.496	222.51	7199	75	A,B,E	–
NGC6441 FA	17:50:13.80	-37:03:11.00	60296.589	236.6	7202	75		–
NGC6522 1km	18:03:34.02	-30:02:02.30	60259.665	192.6	7200	75	A	–
NGC6522 FA	18:03:34.02	-30:02:02.30	60259.665	192.6	7199	75	A	–
NGC6539 1km	18:04:49.68	-07:35:09.1	60138.693	186	7202	75	A	–
NGC6539 FA	18:04:49.68	-07:35:09.1	60138.693	186	7201	75	A	–
NGC6539A FA	18:04:49.8954	-07:35:24.69	60138.693	186.316	7200	75	A	–
NGC6544 1km	18:07:20.58	-24:59:50.40	60126.971	135.5	7202	75	A,B	–
NGC6544 FA	18:07:20.58	-24:59:50.40	60126.971	135.5	7201	75	A,B	–
NGC6544A FA	18:07:20.3556	-24:59:52.9015	60126.971	134.004	7199	75	A,B	–
NGC6544B FA	18:07:20.8712	-25:00:01.915	60126.971	137.15	7200	75	A,B	–
NGC6624 1km	18:23:40.51	-30:21:39.70	60329.502	86.2	7200	75	A,B,D	–
NGC6624 FA	18:23:40.51	-30:21:39.70	60329.502	86.2	7199	75	A,B,D	–
NGC6712 1km	18:53:04.30	-08:42:22.00	60144.755	155	7201	75		–
NGC6712 FA	18:53:04.30	-08:42:22.00	60144.755	155	7200	75		–
NGC6712A FA	18:53:04.07409	-08:42:28.254	60144.755	155.125	7199	75	A	–
OmegaCen3 1km	13:26:47.24	-47:28:46.50	60330.964	100.4	7201	75		–
OmegaCen3 1km	13:26:47.24	-47:28:46.50	60331.055	100.4	7201	75		–
Ter5 1km	17:48:04.80	-24:46:45.00	60251.453	237.75	7199	75	A,C,E,G,I,L,M, N,O,V,W,Y,Z, ae,ai,am	
Ter5 FA	17:48:04.80	-24:46:45.00	60251.453	237.75	7199	75	A,I,L,M, W,Z	

Targets were selected using complementary criteria intended to probe the strengths and limitations of MeerKAT S-band pulsar searching in globular clusters. Clusters with high DMs were prioritised because higher observing frequencies reduce the impact of dispersion and multi-path scattering, improving sensitivity to fast-spinning and strongly accelerated systems that can be difficult to detect at L-band (e.g. Glimpse-C01; see Section 4.4.2). In parallel, several clusters with already rich pulsar populations were re-observed to maximise the likelihood of re-detections, to confirm candidates, and to provide a consistent dataset for assessing S-band sensitivity for future surveys.

Each observation used a four-beam Pulsar Timing User Supplied Equipment (PTUSE) tied-array configuration [148]. For every pointing, one full-array beam (typically 56–64 antennas, subject to availability) and one “1-km” core beam (typically 38–44 inner-core antennas) were centred on the nominal cluster position (see Fig. 4.1). The two beams sample different regions of parameter space: the full-array beam provides higher sensitivity over a narrower field of view, while the 1-km core beam offers lower sensitivity but a larger on-sky coverage, improving robustness to positional uncertainties within the cluster.

Where sufficiently precise positions were available for previously known pulsars (e.g. NGC 6624, NGC 5986, NGC 6539, and NGC 6544), the remaining two beams (again, one full-array and one 1-km beam) were pointed at the known pulsar location(s). In these cases, re-detections were obtained through blind searches as well as coherent folding using published timing ephemerides, with the aim of confirming detectability, assessing signal stability, and, in selected cases, verifying the consistency of existing timing solutions (e.g. NGC 6712A; Yan et al. 149) or candidates reported in earlier surveys (e.g. Glimpse-C01A; McCarver et al. 150).

## 4.3 Data Analysis

### 4.3.1 Data Format for the S-band survey

The pulsar search was performed on search-mode PSRFITS data products (see Hotan et al. 93) generated as part of the standard MeerKAT pulsar-processing chain. The data consist of total-intensity filterbank files formed by summing the two recorded linear polarisations, and are channelised into 256 frequency channels spanning the full S-band bandwidth. The data are dedispersed at the nominal cluster dispersion measure (DM) adopted for the search and have a final time resolution appropriate for millisecond-pulsar detection for the corresponding cluster DMs (Table 4.1). Radio-frequency interference (RFI) mitigation and standard quality checks were applied

prior to the search analysis.

### 4.3.2 Search operation

All 14 globular clusters in this survey were searched using `PULSAR_MINER` v2.0-beta<sup>1</sup>, an automated pulsar-search pipeline built on the `PRESTO` software suite [?]. The overall workflow follows the standard search logic described in Chapter 3, in particular the steps outlined in Sections 3.5.7, 3.6, and 3.6.7. All `PSRFITS` files were first processed to further mitigate RFI. An RFI mask was generated for each observation using the `rfifind` routine, with parameters optimised on a per-observation basis to identify contaminated time intervals and frequency channels. These data segments were subsequently excluded from further analysis.

Following RFI masking, a zero-dispersion-measure (0-DM) time series was generated using `prepdata` and examined for periodic terrestrial signals. Repeating features identified in the 0-DM spectrum were flagged as instrumental or anthropogenic interference (“birdies”), and the corresponding frequencies were removed from the Fourier spectra prior to the pulsar search. As in standard pulsar-search pipelines, a de-dispersion plan tailored to each observation was then constructed using the `DDplan.py` utility in `PRESTO`. This procedure defines a set of discrete trial DMs such that the residual dispersion smearing arising from finite channel bandwidth, finite time resolution, and from the true DM lying between adjacent trials remains within acceptable limits across the observing band.

For this survey, the overall DM search range was chosen to extend to within  $\pm 20\%$  of the nominal cluster DM. This was a deliberately conservative choice intended to retain sensitivity to possible outlier pulsars lying along sightlines with enhanced free-electron content, especially in high-DM systems such as Glimpse-C01. Adopting a broad fractional DM range therefore ensured sensitivity to pulsars with substantial DM offsets from the nominal cluster value.

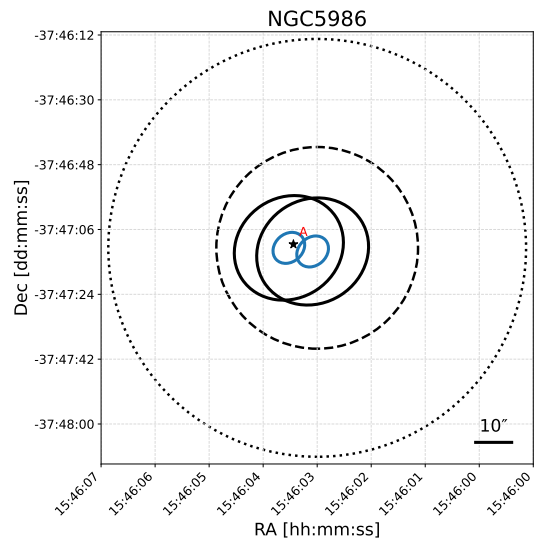
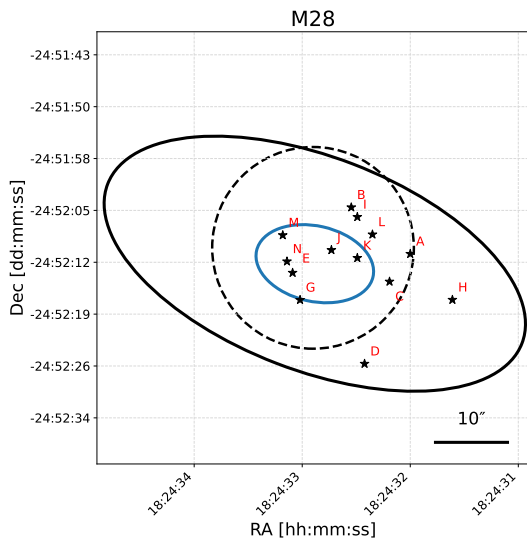
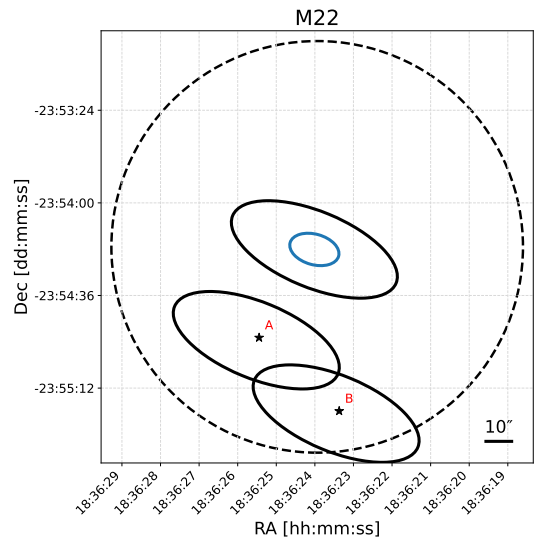
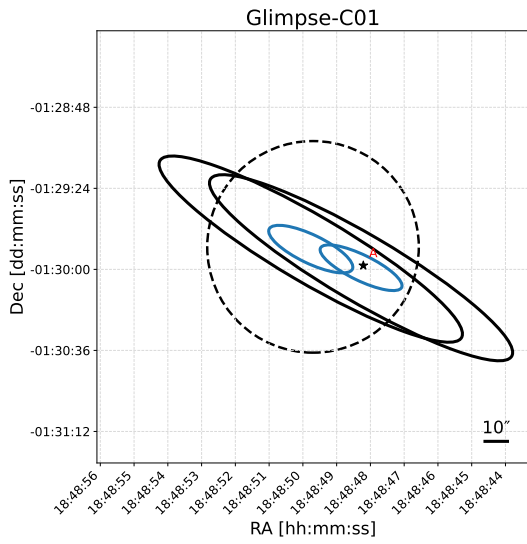
To recover pulsars in compact binary systems, the data were searched for periodic signals in the Fourier domain using both acceleration and jerk search techniques. In a binary system, orbital motion induces a line-of-sight acceleration,  $a_{\parallel}$ , that Doppler-shifts the apparent spin frequency of the pulsar. As discussed in Section 3.6.5, this causes the signal power to drift across multiple Fourier bins over the duration of an observation. For approximately constant acceleration, the resulting frequency derivative may be written as


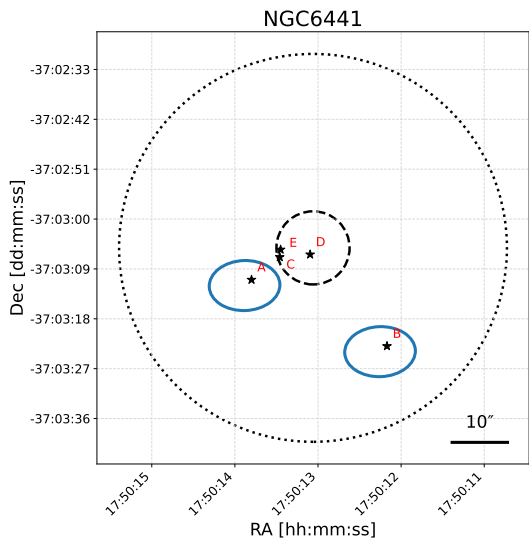
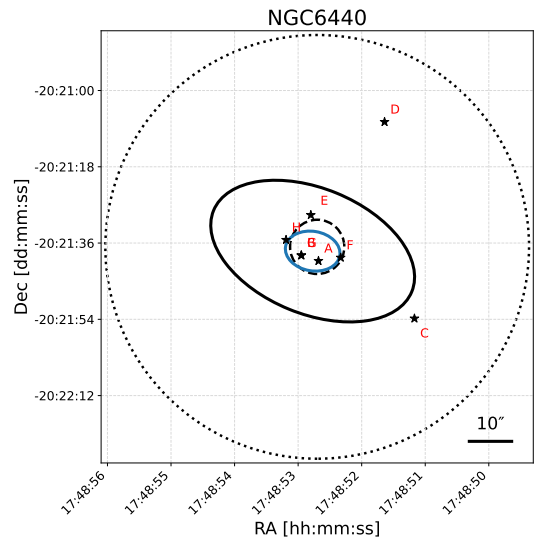

$$\dot{f}_{\text{obs}} \simeq \left(\frac{a_{\parallel}}{c}\right) f_0,$$


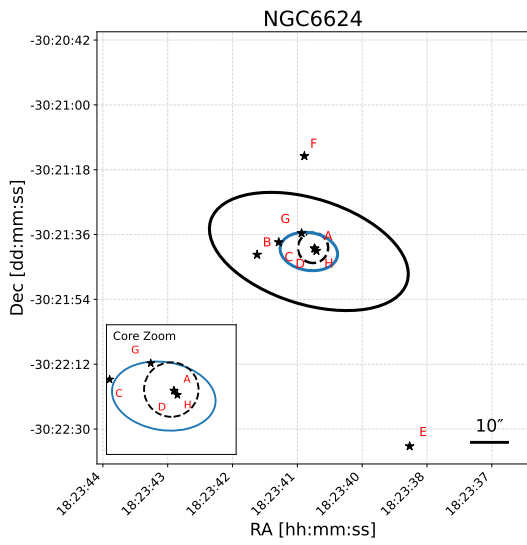
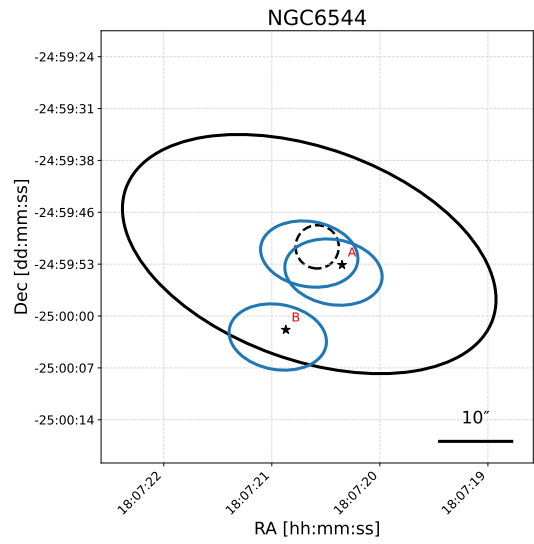

---

<sup>1</sup>[https://github.com/alex88ridolfi/PULSAR\\_MINER.git](https://github.com/alex88ridolfi/PULSAR_MINER.git)

<sup>2</sup>[https://wchenastro.github.io/mosaic\\_web/](https://wchenastro.github.io/mosaic_web/)









where  $f_0$  is the intrinsic pulsar spin frequency and  $c$  is the speed of light [104]. If uncorrected, this frequency drift reduces sensitivity in standard Fourier searches.

The acceleration search, implemented using `accelsearch`, assumes that the pulsar experiences a constant line-of-sight acceleration over the integration time and attempts to recover signal power by summing across adjacent Fourier bins corresponding to trial acceleration values [104]. The jerk search extends this approach by allowing for a linear change in acceleration, that is, a constant line-of-sight jerk  $j_{\parallel}$ , during the observation. This increases sensitivity to systems in which the constant-acceleration approximation breaks down, such as very compact binaries or observations taken near orbital phases where the acceleration changes rapidly [? ].

While acceleration and jerk searches substantially improve sensitivity to binary pulsars, their underlying assumptions remain approximations and can still lead to residual signal smearing over long integrations. To mitigate this effect and to control computational cost, the data were also searched at multiple effective integration lengths [? ]. For each observation, time series were prepared for 15, 30, and 60 min sub-integrations, in addition to the full  $\sim 2$  h observation. Shorter integrations reduce the accumulated frequency drift and improve detectability for the most compact or highly accelerated systems, at the expense of raw signal-to-noise ratio.

### 4.3.3 Sifting and folding

At the end of the periodicity search, candidate lists were generated for each time-series chunk and trial DM. Here, a candidate corresponds to a statistically significant peak in the Fourier power spectrum, or in a harmonically summed spectrum, exceeding a predefined detection threshold. Candidates were grouped according to harmonic relationships and sifted using standard criteria, requiring that a signal be significant in at least three adjacent DM trials. Only candidates with a detection significance of  $\geq 4\sigma$  were retained [104]. Depending on the level of residual RFI, the presence of bright or long-period pulsars, and the integration length, this procedure typically produced  $\sim 200$ – $500$  candidates per beam.

All surviving candidates were then folded using the `prepfold` routine, taking the detected spin period, DM, and trial acceleration as initial parameters. During folding, these parameters were refined to maximise the  $\chi^2$  of the folded pulse profile relative to an average unpulsed background. The resulting diagnostic outputs include integrated pulse profiles, time–phase and frequency–phase plots, DM response curves, and frequency–frequency-derivative planes.

The folded candidates were subsequently inspected visually to identify strong, broadband signals that are persistent over a significant fraction of the observation

and display characteristics consistent with astrophysical pulsars. Promising candidates were cross-checked against the catalogue of known globular-cluster pulsars maintained by the Max Planck Institute for Radio Astronomy (MPIfR)<sup>3</sup>. For each cluster, known pulsars were searched for explicitly, including up to 16 harmonics of the fundamental spin frequency, particularly in systems hosting bright or long-period pulsars.

In addition to the blind search described above, targeted folding was performed for all known pulsars covered by the S-band beams using the best existing timing ephemerides available to us. For this purpose, available MeerKAT parameter files were used to fold both the S-band census observations presented in this work and multiple epochs of archival MeerKAT L-band data using `dspsr`<sup>4</sup>. These coherently folded profiles form the basis of the flux-density measurements and spectral-index estimates discussed later in Sections 4.4.4 and 4.4.5.

#### 4.3.4 Flux density measurements

An important component of this work is the measurement of mean flux densities for a subset of re-detected pulsars, with the aim of providing reference values and empirical sensitivity estimates for future observations with MeerKAT or other telescopes at S-band. Flux-density measurements are restricted to pulsars for which reliable folding ephemerides are available and for which the detections yield sufficient signal-to-noise ratio to permit robust baseline and on-pulse flux estimation. Systems exhibiting strong eclipses or highly variable detectability, such as spider binaries, are excluded where the observed signal is not representative of the intrinsic pulsar emission.

Rather than relying on direct calibration with the MeerKAT noise-diode system, which may be affected by backend non-linearities, we estimate mean flux densities using the standard pulsar radiometer equation applied to folded pulse profiles, following the methodology described by [? ]. Throughout this work, we express the radiometer equation in terms of the effective system-equivalent flux density (SEFD), defined as  $\text{SEFD} = T_{\text{sys}}/G$ , where  $T_{\text{sys}}$  is the total system temperature and  $G$  is the telescope gain. The determination of the effective tied-array SEFD for our S-band observations is described in Appendix B.

For a folded pulsar observation with integration time  $t_{\text{obs}}$  and usable bandwidth

---

<sup>3</sup><https://www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html>

<sup>4</sup><https://dspsr.sourceforge.net/>

$\Delta\nu$ , the phase-averaged mean flux density  $S$  is given by

$$S = \frac{\text{SNR} \times \text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu t_{\text{obs}}}} \sqrt{\frac{N_{\text{on}}}{N_{\text{off}}}}, \quad (4.1)$$

where SNR is the signal-to-noise ratio of the integrated pulse profile,  $n_{\text{pol}} = 2$  is the number of summed orthogonal polarisations,  $N_{\text{on}}$  is the number of phase bins containing pulsed emission, and  $N_{\text{off}} = N_{\text{bin}} - N_{\text{on}}$  is the number of off-pulse bins used to estimate the baseline noise level.

The signal-to-noise ratio is measured directly from the folded profile as

$$\text{SNR} = \frac{A}{\sigma_A}, \quad (4.2)$$

where  $A$  is the integrated pulse amplitude above the baseline and  $\sigma_A$  is its associated uncertainty derived from the off-pulse root-mean-square level (see Appendix ??).

The on-pulse region was identified following the procedure described by ? ]. An initial selection of contiguous phase bins exceeding the off-pulse baseline root-mean-square by a fixed threshold was performed, after which the selection was visually inspected to ensure that only statistically significant pulsed emission was included. The remaining bins were designated as the off-pulse region and used to determine the baseline level and noise properties. This approach avoids circular dependence on the total profile SNR and ensures consistent estimation of  $N_{\text{on}}$  and  $N_{\text{off}}$  across pulsars.

Uncertainties on the derived flux densities are obtained by propagating the uncertainty in the measured SNR and the off-pulse root-mean-square through Eq. (1). Since the SEFD, bandwidth, number of polarisations, and integration time are fixed for a given observation, the dominant source of statistical uncertainty arises from the profile-noise estimate.

## 4.4 Results

### 4.4.1 Re-detections

This survey resulted in the re-detection of 39 pulsars across 14 globular clusters. Details of all re-detections, including the specific tied-array beams in which each pulsar was detected, are summarized in Table 4.1. The beam information, together with pulsar parameters such as spin period  $P$  and dispersion measure (DM) obtained from the MPIfR pulsar catalogue, provides a practical basis for subsequent timing analysis using the S-band observation data archived in the long-term storage facilities

of the Astronomical Observatory of Cagliari.

#### 4.4.2 New Discoveries and other search results from Glimpse-C01

Four new pulsars were discovered in the globular cluster Glimpse-C01, which previously hosted two known pulsars: J1848–0129A [150] and J1848–0129B discovered by the Five-hundred-meter Aperture Spherical Telescope (FAST) [151]. The full set of pulsars currently known in this cluster is listed in Table 4.2, together with their basic measured parameters.

All four newly discovered pulsars are members of binary systems and have been added to the pulsar catalogue maintained by MPIfR as well as to the TRAPUM discovery page<sup>5</sup>. Time–phase plots for these pulsars are shown in Fig. 4.2.

This globular cluster had previously been searched at radio S-band (1850 MHz) using the Green Bank Telescope (GBT) and at L-band using FAST. Following the discovery of the four pulsars in the MeerKAT data, the archival GBT observations were re-examined and J1848–0129D was successfully recovered through acceleration searches performed on a 7.36 h observation. J1848–0129A was re-detected in the MeerKAT data, while J1848–0129B, originally discovered by FAST, could not be detected through blind searches in the MeerKAT dataset. It was later recovered by folding the data using ephemerides provided by the FAST team.

After the initial discoveries in the MeerKAT data, Glimpse-C01 was observed on multiple occasions with MeerKAT in two dedicated orbital monitoring campaigns separated by approximately one year. These data, together with the GBT and FAST observations, enabled preliminary characterisation of the newly discovered systems, which we now discuss.

##### J1848–0129C

J1848–0129C is a 6.44 ms binary millisecond pulsar detected at a dispersion measure of  $491.8 \text{ pc cm}^{-3}$ . The pulsar was independently identified in both the 1 km core and full-array PTUSE beams and was subsequently recovered in archival FAST observations.

The combined dataset spans a wide range of orbital phases and enabled the construction of a Keplerian orbital solution for J1848–0129C. The derived orbital parameters are given in Table 4.3. The solution consistently folds all available

---

<sup>5</sup><https://www.trapum.org/discoveries>

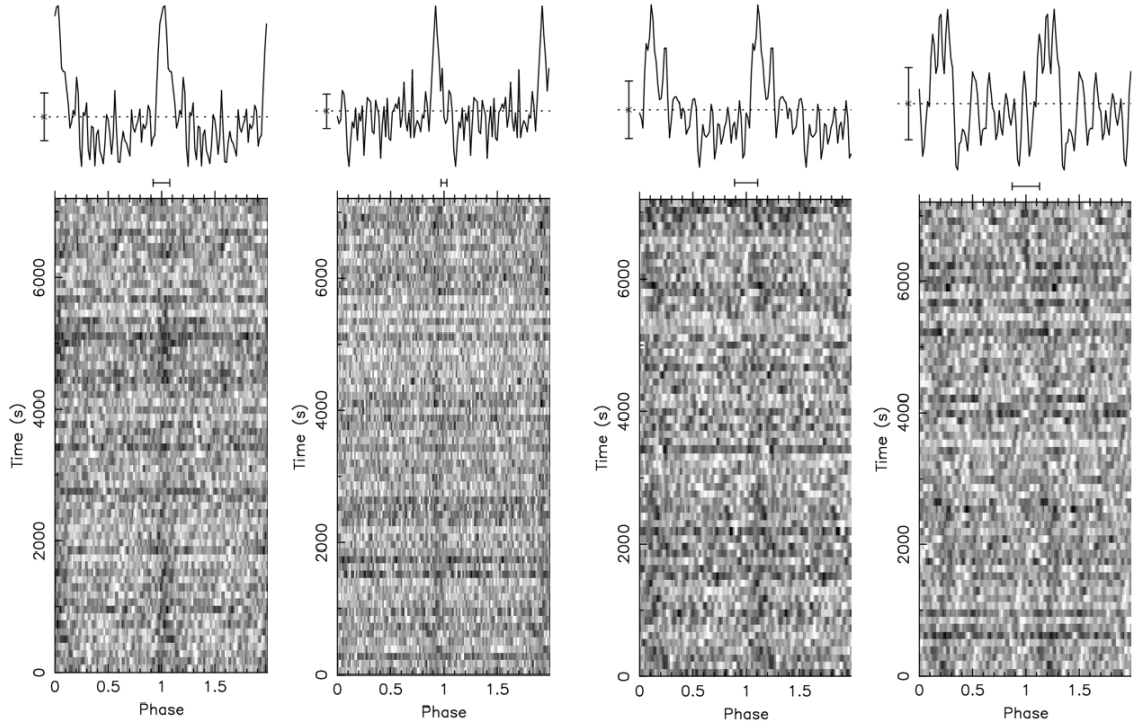


Figure 4.2: Time–phase plots and two pulse periods of the integrated pulse profiles for the four newly discovered pulsars in Glimpse-C01 (J1848–0129C, D, E and F), shown in order.

Table 4.2: Parameters for the pulsars in the globular cluster Glimpse-C01. Newly discovered pulsars are shown in boldface. Listed are the barycentric spin period  $P$ , dispersion measure (DM), and mean flux density  $S$  measured at MeerKAT S-band. Uncertainties reported in parentheses correspond to  $1\sigma$  errors on the final quoted digit(s). Flux-density uncertainties are derived from the off-pulse noise of the folded profiles, as described in Section A.1.

Pulsar	$P$ (ms)	DM ( $\text{pc cm}^{-3}$ )	$S$ (mJy)
J1848–0129A	19.78	491.1	0.032(0.002)
J1848–0129B	13.80	482.0	0.010(0.001)
<b>J1848–0129C</b>	<b>6.44</b>	491.8	0.035(0.002)
<b>J1848–0129D</b>	<b>17.10</b>	457.9	0.012(0.001)
<b>J1848–0129E</b>	<b>4.54</b>	479.9	0.008(0.001)
<b>J1848–0129F</b>	<b>4.17</b>	520.1	0.008(0.001)

detections and reproduces the expected orbital Doppler modulation, as illustrated in Fig. 4.3.

Assuming a neutron-star mass of  $M_p = 1.4 M_\odot$ , the mass function implies a minimum companion mass of  $M_{c,\min} = 0.3525 M_\odot$  for an edge-on orbit ( $i = 90^\circ$ ) and a median companion mass of  $M_c = 0.4169 M_\odot$  for a representative inclination of  $i = 60^\circ$ . The companion therefore lies in an intermediate mass regime, substantially heavier than the  $\sim 0.23 M_\odot$  helium white-dwarf companions predicted for this orbital period by [152].

A notable property of J1848–0129C is its orbital-phase-dependent detectability. Using the orbital solution in Table 4.3, we computed the orbital phase corresponding to each MeerKAT observation. Detections are confined to phases away from superior conjunction, while observations predicted to sample phases near superior conjunction (orbital phase  $\sim 0.25$ ) consistently resulted in non-detections.

In one case, two observations separated by less than 24 h sampled orbital phases immediately preceding superior conjunction. The earlier observation yielded a detection, whereas the subsequent observation, probing deeper into the predicted conjunction phase interval, resulted in a non-detection. No clear ingress or egress structure is resolved within individual  $\sim 2$  h integrations; however, the systematic confinement of non-detections to the phase range expected for superior conjunction is reproducible across campaigns separated by approximately one year.

The stability of this phase-locked non-detection interval disfavors refractive scintillation or acceleration-search limitations as the dominant cause of intermittency. Instead, the behaviour is consistent with eclipse or absorption in an intrabinary medium.

To assess whether the eclipse interpretation is geometrically plausible, we performed a simple consistency check using the derived orbital parameters. Adopting the phase interval  $\Delta\phi \simeq 0.24$  (phases  $\sim 0.13$ – $0.37$ ) as the putative eclipse window, the corresponding duration is  $\Delta t \simeq 1.2$  d for  $P_b = 5.03$  d. For a circular edge-on orbit, this phase width implies a minimum eclipsing scale  $R_{\text{ecl}} \gtrsim a \sin(\pi\Delta\phi)$ .

Using Kepler’s third law with the parameters in Table 4.3 and  $M_p = 1.4 M_\odot$ , the orbital separation is  $a \simeq 15 R_\odot$ . Substituting  $\Delta\phi \simeq 0.24$  yields  $R_{\text{ecl}} \gtrsim 10 R_\odot$  as a lower limit on the characteristic size of the absorbing region in the edge-on case. For comparison, the companion Roche-lobe radius for  $q \simeq 0.25$ – $0.30$  is  $R_{L,c} \simeq 4 R_\odot$ .

These order-of-magnitude estimates indicate that, if the non-detections are indeed caused by eclipse or absorption, the absorbing material must extend over several solar radii and cannot be explained by a compact geometric occultation alone. This calculation assumes a symmetric eclipse centred on superior conjunction and should

therefore be regarded as a preliminary geometric consistency check rather than a detailed physical model.

Under this interpretation, J1848–0129C is not compatible with a classical compact redback configuration and instead occupies the long-period spider regime associated with the proposed “Huntsman” subclass [see 153, 154]. Its orbital period places it directly alongside the Huntsman prototype PSR J1417–4402 ( $P_b \simeq 5.4$  d) and far beyond the orbital periods of known globular-cluster redbacks and candidates, whose longest currently reported value, held by J1740–5340B, is  $P_b \simeq 1.97$  d [155].

If confirmed as a Huntsman-like eclipsing system, J1848–0129C would constitute the first example of this long-period spider subclass in a Galactic globular cluster. Even under the more conservative classification of an ordinary redback, its  $P_b = 5.03$  d would still represent the longest orbital period measured for a globular-cluster redback system by a substantial margin.

### J1848–0129D

J1848–0129D is a binary pulsar discovered in an acceleration search of MeerKAT S-band observations targeting the core of the globular cluster Glimpse-C01 and subsequently re-detected in archival GBT data as well as in FAST and later MeerKAT follow-up observations. The pulsar has a spin period of  $P \simeq 17.1$  ms and a dispersion measure of  $DM \simeq 459$  pc cm<sup>-3</sup>.

We derived a Keplerian orbital solution by fitting a BT binary model to the observed spin-period modulation using detections from both MeerKAT and FAST observations. The resulting solution folds all available measurements, and the orbit is consistent with being circular ( $e \approx 0$ ). The corresponding spin-period modulation is also shown in Fig. 4.3.

Table 4.3: Best-fitting orbital parameters for J1848–0129C and J1848–0129D, together with their relative  $1\sigma$  uncertainties, obtained using a modified version of the PRESTO routine `fit_circular_orbit.py`.

Pulsar	J1848–0129C	J1848–0129D
Spin Period, $P$ (ms)	6.44835719(2)	17.110102567(7)
Orbital period, $P_b$ (d)	5.0293642(6)	3.39134255(7)
Projected semi-major axis, $x$ (lt-s)	6.95158(14)	12.1271(4)
Epoch of periastron, $T_0$ (MJD)	60591.731656(44)	59279.22166(3)
Eccentricity, $e$	0	0
Mass function, $f(M)$ ( $M_\odot$ )	0.01426	0.16649
Minimum companion mass, $M_c$ ( $M_\odot$ )	0.3525	0.9809

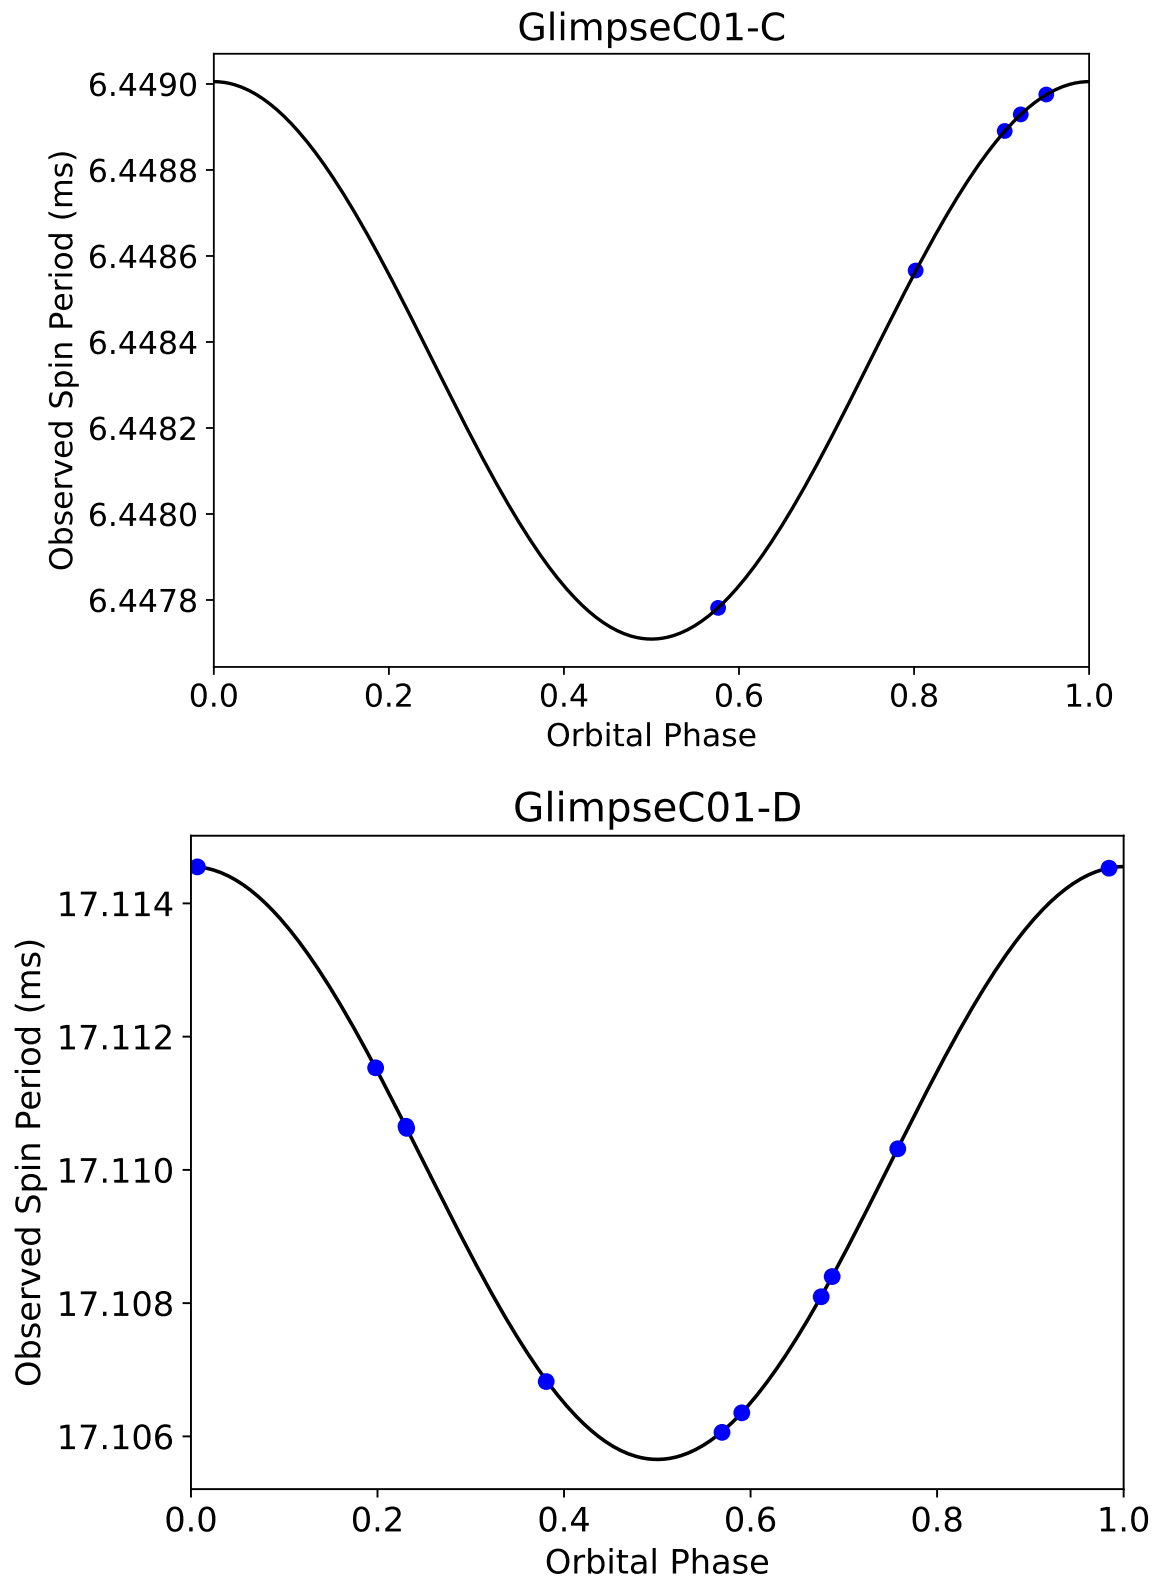


Figure 4.3: Observed spin period as a function of orbital phase for J1848–0129C (top) and J1848–0129D (bottom). Blue points represent measured spin periods from individual observations, while the solid black curves indicate the predictions obtained from the orbital solutions listed in Table 4.3.

Assuming again  $M_p = 1.4 M_\odot$ , the system’s mass function implies  $M_{c,\min} \simeq 0.98 M_\odot$  and  $M_c \simeq 1.20 M_\odot$  for a representative inclination of  $i = 60^\circ$ . While this mass range overlaps the lowest neutron-star masses, the absence of measurable eccentricity disfavors a double neutron-star interpretation [e.g. 156]. Instead, the system is more naturally explained as a mildly recycled pulsar with a massive CO or ONeMg white-dwarf companion formed through stable mass transfer [157].

### **J1848–0129E and J1848–0129F**

J1848–0129E and J1848–0129F are two faint binary millisecond pulsars. J1848–0129E was detected in the 1 km core PTUSE beam and has a spin period of 4.54 ms with a dispersion measure of  $479.9 \text{ pc cm}^{-3}$ , while J1848–0129F was detected in the full-array PTUSE beam and has a spin period of 4.17 ms with a dispersion measure of  $520.1 \text{ pc cm}^{-3}$ . Both pulsars exhibit non-zero line-of-sight accelerations, indicating that they are members of binary systems.

Owing to their intrinsic faintness and the limited number of detections, no further constraints on the orbital parameters of either system could be obtained. No additional re-detections were secured during subsequent follow-up observations, preventing the derivation of coherent orbital solutions. At the time of writing, J1848–0129F has the highest dispersion measure reported for any pulsar associated with a globular cluster. Further targeted observations will be required to confirm the binary properties of these systems and to better characterise their emission and propagation properties.

### **4.4.3 Dispersion measure discrepancy and cluster membership**

Two of the pulsars detected towards Glimpse-C01 exhibit dispersion measures that differ substantially from the nominal value measured towards the cluster core ( $\text{DM} \simeq 491.1 \text{ pc cm}^{-3}$ ). In particular, J1848–0129D is detected at  $\text{DM} \simeq 459.1 \text{ pc cm}^{-3}$ , while J1848–0129F is detected at  $\text{DM} \simeq 520.1 \text{ pc cm}^{-3}$ . Taken together, these pulsars span a total DM range of approximately  $61 \text{ pc cm}^{-3}$ . This unusually large range motivates a careful examination of whether the observed DM discrepancies are compatible with cluster membership, or whether one might instead be detecting MSPs unrelated to the cluster.

Several observational considerations support the interpretation that both pulsars are associated with Glimpse-C01 despite their displaced DMs. First, the detections were obtained in narrow, full-array tied-array beams centred on the cluster core,

which strongly reduces the probability of a chance alignment with unrelated field pulsars.

If the pulsars are indeed associated with Glimpse-C01, the most natural explanation for the DM differences is real structure in the ionized interstellar medium along the line of sight. Glimpse-C01 lies behind a heavily obscured Galactic-plane sightline, where significant small-scale variations in electron column density are expected [e.g. 86, 158]. In such environments, different sightlines across the angular extent of a globular cluster can sample distinct foreground ionized structures, leading to measurable DM differences among pulsars bound to the same cluster potential.

To place the Glimpse-C01 DM behaviour in a broader context, we compared its DM distribution with those of other Galactic globular clusters hosting multiple known pulsars. Published DM values were compiled from the MPIfR Globular Cluster Pulsar Catalogue, and only clusters with  $\geq 5$  known pulsars were included in order to ensure that intra-cluster statistics are meaningful. For each cluster, we computed the dispersion of the pulsar DMs,  $\sigma_{\text{DM}}$ , and adopted the median pulsar DM as the nominal cluster DM.

Figure 4.4 shows  $\sigma_{\text{DM}}$  as a function of nominal cluster DM. A power-law relation of the form  $\sigma_{\text{DM}} = A \text{DM}^\alpha$  was fitted in log–log space using ordinary least squares. To estimate uncertainties on the fitted parameters, we performed bootstrap resampling of the cluster sample (4000 realizations), refitting the relation for each resample and deriving the 16th–84th percentile intervals for the slope and normalization. The shaded region in the figure represents the root-mean-square scatter of the residuals in log space about the best-fitting relation.

A clear trend is observed: clusters with larger nominal DMs tend to exhibit larger intra-cluster DM dispersions. The best-fitting slope is statistically consistent with unity, indicating an approximately linear scaling,  $\sigma_{\text{DM}} \propto \text{DM}$ . This behaviour is qualitatively consistent with the trend noted by Freire [159], who investigated the spread of DMs in globular clusters using the full DM span,  $\Delta\text{DM} = \text{DM}_{\text{max}} - \text{DM}_{\text{min}}$ , as a proxy for internal variation. At that time, the number of clusters hosting multiple known pulsars was limited, and the root-mean-square dispersion could not be robustly determined for most systems. With the substantially expanded globular-cluster pulsar population now available, the use of  $\sigma_{\text{DM}}$  provides a more statistically stable measure of intra-cluster DM variation.

Glimpse-C01 lies at the high-DM end of the distribution and exhibits a large internal DM dispersion consistent with the overall population trend. Its relatively large DM offsets therefore appear compatible with the behaviour observed in other high-DM clusters, rather than indicating an anomalous or fundamentally distinct

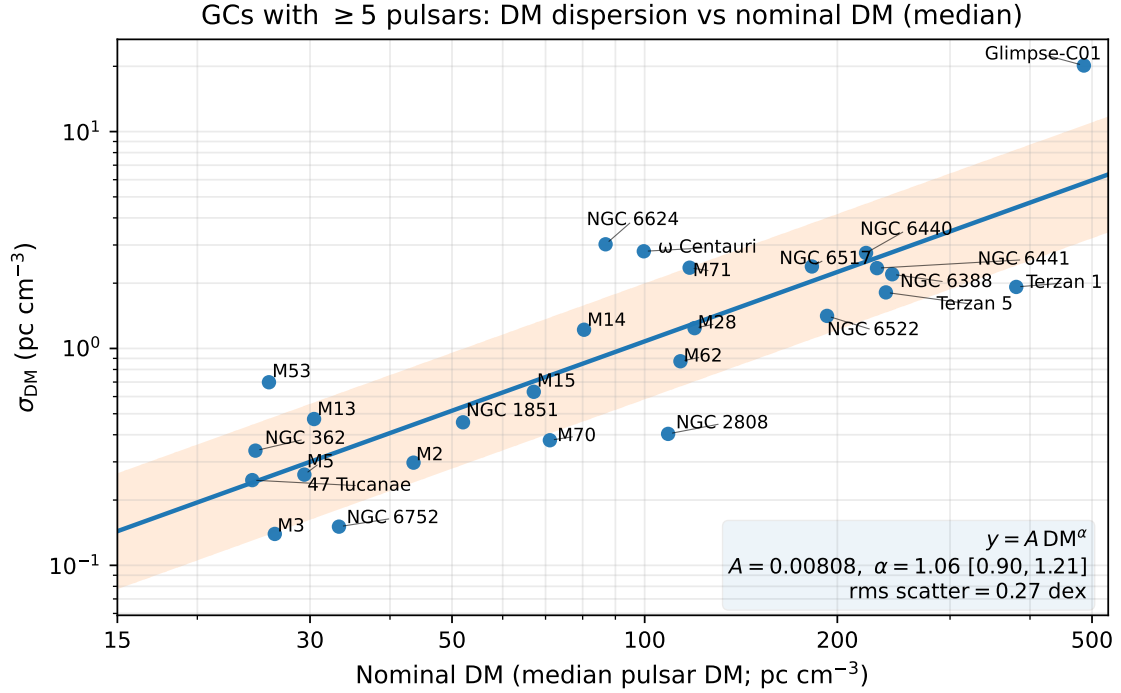


Figure 4.4: Intra-cluster dispersion-measure dispersion,  $\sigma_{\text{DM}}$ , as a function of the nominal cluster dispersion measure (defined as the median pulsar DM) for Galactic globular clusters hosting at least five known radio pulsars. The solid line shows the best-fitting power-law relation of the form  $\sigma_{\text{DM}} = A \text{DM}^\alpha$ , with  $A = 0.008$  and  $\alpha = 1.06$  (bootstrap 16–84% interval: 0.90–1.21). The shaded region represents the root-mean-square scatter in log space (0.27 dex) about the fit.

DM structure.

Taken together, the beam-centred detections and the placement of Glimpse-C01 within the established population trends of intra-cluster DM dispersion support the interpretation that the DM offsets observed for J1848–0129D and J1848–0129F arise from small-scale structure in the intervening interstellar medium along this high-DM Galactic-plane sightline. The observed offsets therefore do not, by themselves, require invoking a foreground or background origin.

### Position of J1848–0129A

The S-band survey observations of Glimpse-C01 employed two dedicated tied-array beams for J1848–0129A: one formed using the MeerKAT 1 km core and the other using the full array. These beam pointings were centred on the reported position of J1848–0129A as determined from the VLITE millisecond-pulsar survey, in which the MSP was discovered by targeting a VLA-detected point-like radio source [150]. However, the pulsar was not detected in either of these dedicated beams. Instead,

J1848–0129A was detected in both the 1 km core and full-array beams pointed at the centre of the globular cluster. This discrepancy raises the possibility that the position inferred from the VLITE survey does not correspond to J1848–0129A, and that the VLITE detection may be associated with a different radio source.

This result was confirmed by a follow-up observation of Glimpse-C01 conducted in April 2024 as part of an orbital-determination and timing campaign for the newly discovered cluster pulsars. These observations used the same beam configurations as the original S-band survey. Once again, J1848–0129A was not detected in the two beams centred on the VLITE position, but was clearly detected in the beams pointed at the cluster core, reinforcing the conclusion that the VLITE point source is unlikely to be associated with J1848–0129A.

To investigate further whether the VLITE point source might instead correspond to a previously unidentified pulsar, we performed full as well as segmented acceleration and jerk searches on both the original S-band survey data and the April 2024 follow-up observations at the VLITE beam pointings. These searches were carried out using both the 1 km core and full-array data and explored very high acceleration and jerk trial ranges. No additional pulsars were detected in any of these searches. As a result, the radio source detected in the VLITE survey<sup>6</sup>, originally associated with J1848–0129A, remains unidentified.

#### 4.4.4 S-band and L-band flux density measurements

Table 4.4 presents mean flux-density measurements for a subset of pulsars detected in the MeerKAT S-band globular-cluster survey that also have available MeerKAT L-band observations. The pulsars listed were selected based on the availability of reliable ephemerides, which allowed coherent folding of both the S-band survey data and archival or targeted L-band observations. The S-band flux densities reported here are derived from single-epoch observations, whereas the corresponding L-band flux densities were obtained by averaging measurements from multiple epochs where such data were available.

It is important to note that the S-band and L-band observations used to populate Table 4.4 were not obtained contemporaneously. In several cases, the observations are separated by months to years. Over such timescales, refractive interstellar scintillation is expected to introduce substantial variability in the observed flux densities [see, e.g. 7], potentially enhancing or suppressing the measured signal by factors of a few [2, 43]. As a result, the flux-density values reported in the two bands are not intended to be compared directly, and any apparent differences between the L-

---

<sup>6</sup><https://vlite.nrao.edu/>

band and S-band measurements should not necessarily be interpreted as indicative of intrinsic spectral behaviour.

The primary purpose of Table 4.4 is therefore to provide reference flux-density measurements for pulsars detected in the S-band survey. The use of multi-epoch averaging at L-band serves to mitigate the effects of scintillation and provides a more representative estimate of the typical L-band flux density for comparison at a population level. These measurements also act as a consistency check for the flux-density estimation procedure described in Appendix A.1, applied across different observing bands and datasets.

Terzan 5 is presented separately here because it contributes a particularly large number of pulsar detections within a single cluster, and because extensive literature flux density measurements are available for its known pulsars. Using the  $S_{1400}$  and  $S_{2000}$  values reported by Martsen et al. [7], we compare these with our single-epoch MeerKAT S-band flux density measurements at  $\sim 2406$  MHz. Table 4.5 lists the resulting  $S_{2400}$  values together with the literature measurements, and gives updated spectral indices derived using the literature  $S_{1400}$  values and our measured  $S_{2400}$  flux densities. This allows the Terzan 5 detections to be placed in the context of previous flux density measurements while keeping this unusually large single-cluster sample separate from the rest of the survey detections.

A more quantitative investigation of spectral properties would require either simultaneous multi-frequency observations or a sufficiently large number of epochs at each frequency to average over scintillation effects. Such an analysis is beyond the scope of the present work.

Beyond serving as reference measurements, the flux densities listed in Table 4.4 provide an empirical indication of the range of pulsar brightnesses accessible to MeerKAT S-band observations in globular clusters. The detection of several known pulsars with reliable orbital solutions at 2.4 GHz demonstrates the viability of S-band searches for both isolated and binary millisecond pulsars in dense cluster environments. These measurements therefore establish a practical sensitivity scale for the present survey, which is useful for interpreting subsequent non-detections and for informing the design of future high-frequency globular-cluster pulsar surveys.

#### 4.4.5 S-band non-detections and reference spectral-index constraints

A number of previously known globular-cluster pulsars fall within the MeerKAT S-band ( $\sim 2.4$  GHz) PTUSE tied-array beam coverage but are not detected in our S-band survey. Several of these pulsars are, however, reliably detected in MeerKAT

L-band ( $\sim 1.4$  GHz) observations reported in the literature. While our S-band survey consists of a single epoch per cluster, we examine whether the combination of published L-band detections and S-band non-detections can be used to derive *reference* upper limits on the radio spectral index.

This exercise is not intended to yield precise spectral measurements. Instead, the derived limits are meant to serve as indicative constraints that help contextualise S-band non-detections of otherwise established L-band pulsars. The resulting values are expected to carry substantial uncertainties due to the limited number of epochs, possible epoch-to-epoch variability driven by interstellar scintillation, and assumptions regarding the effective sensitivity of coherent tied-array beams. For the purpose of deriving upper limits, we adopt the nominal usable S1-band bandwidth of 875 MHz for the best-case scenario. This is very close to the typical effective bandwidth after RFI excision, which is  $\gtrsim 98\%$  of the processed band. The resulting limits therefore represent optimistic best-case sensitivity estimates.

From the full set of S-band non-detections, we restrict attention to pulsars that (i) are reported to be moderately to highly bright at L-band in published MeerKAT or comparable observations, (ii) are not known spider systems (black widows or redbacks) or strongly eclipsing binaries, and (iii) have sufficiently precise positional localisations to confirm that they lie within the S-band PTUSE beam response. Pulsars clearly outside the S-band beam footprint, as well as systems with poor L-band localisation for which beam coverage cannot be reliably assessed, are excluded from this analysis. This selection is intended to minimise ambiguity arising from eclipses, orbital modulation, and positional uncertainties.

Assuming a power-law radio spectrum,

$$S_\nu \propto \nu^\alpha, \quad (4.3)$$

an S-band non-detection can be combined with a published L-band flux density to place an upper limit on the spectral index  $\alpha$ . The minimum detectable flux density at S-band is estimated using the standard radiometer equation [2],

$$S_{\min} = \text{SNR}_{\min} \frac{\text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu t_{\text{obs}}}} \sqrt{\frac{\delta}{1 - \delta}}, \quad (4.4)$$

where  $\text{SNR}_{\min}$  is the adopted detection threshold,  $\text{SEFD}_{\text{array}}$  is the effective system-equivalent flux density of the tied-array beam<sup>7</sup>,  $n_{\text{pol}} = 2$  is the number of summed polarisations,  $\Delta\nu$  is the usable bandwidth,  $t_{\text{obs}}$  is the integration time, and  $\delta = W/P$

<sup>7</sup><https://skaafrika.atlassian.net/wiki/spaces/ESDKB/pages/1588854789/S-band+capability+and+status>

is the pulse duty cycle. For consistency, we adopt duty cycles comparable to those inferred at L-band, noting that this is a simplifying assumption.

Given an L-band flux density  $S_{\nu_L}$  at frequency  $\nu_L$  and an S-band non-detection at frequency  $\nu_S$ , the corresponding upper limit on the spectral index is

$$\alpha_{\max} = \frac{\ln(S_{\min}/S_{\nu_L})}{\ln(\nu_S/\nu_L)}. \quad (4.5)$$

In the following, we adopt  $\nu_S = 2.4$  GHz,  $t_{\text{obs}} = 2$  h,  $\Delta\nu = 875$  MHz, and  $\text{SNR}_{\min} = 8$  for our S-band search configuration. The effective  $\text{SEFD}_{\text{array}}$  is computed from published MeerKAT S-band receiver-performance curves, scaled by the number of antennas contributing to the tied-array beam.

For non-detections of pulsars with known positions, the effective sensitivity may be reduced relative to the ideal radiometer-equation prediction if the pulsar lies away from the beam boresight. Rather than modelling the detailed tied-array beam response and epoch-specific conditions for each case, we adopt a conservative sensitivity margin when deriving spectral-index limits. This approach is intended to bracket plausible detectability while avoiding over-interpretation of single-epoch non-detections.

Among the non-detections satisfying the selection criteria described above, only J1823–3021G and J1824–2452D have sufficiently reliable L-band flux-density measurements and positional information to permit meaningful spectral-index constraints. We therefore restrict the following analysis to these two systems.

For J1823–3021G, Ridolfi et al. [132] report a mean L-band flux density of  $S_{1.3\text{GHz}} = 47 \pm 10 \mu\text{Jy}$ . Combining this with the S-band non-detection yields an indicative constraint of  $\alpha \lesssim -1.6$  to  $-1.2$ , where the range reflects conservative sensitivity assumptions and the quoted uncertainty on the L-band flux density.

A similar constraint can be derived for J1824–2452D, which lies within our 1 km S-band PTUSE beam coverage but is not detected in our S-band data. A published L-band reference flux density for J1824–2452D is  $S_{1.4\text{GHz}} \simeq 0.05$  mJy [2011MNRAS.418.477B]. Under the same conservative assumptions, the S-band non-detection implies an upper limit on the spectral index in the range  $\alpha \lesssim -2.5$  to  $-2.1$ .

A similar consideration applies to several of the non-detections in Terzan 5. Using the literature flux densities and spectral indices [7], extrapolated to our observing frequency of  $\sim 2406$  MHz, we find that PSRs J1748–2446H, J1748–2446R, J1748–2446U, J1748–2446ac, J1748–2446aa, J1748–2446ag, and J1748–2446ah are all expected to be very faint at MeerKAT S band, and their non-detection in our single-epoch observations is therefore not surprising. In the same context,

Table 4.4: Flux densities of pulsars measured in MeerKAT S-band and L-band observations. The  $S_S$  values are from single-epoch detections, whereas  $S_L$  values are averages over multiple epochs, where available. Spectral indices are computed assuming  $S_\nu \propto \nu^\alpha$  between 1.2 and 2.4 GHz. Uncertainties are  $1\sigma$ ; values in parentheses denote the uncertainty on the final quoted digit(s).

Pulsar Name	$S_S$ (mJy at 2.4 GHz)	$S_L$ (mJy at 1.2 GHz)	$\alpha$
J1836–2354A	0.074(12)	0.201(23)	−1.44(29)
J1836–2354B	0.071(11)	0.128(45)	−0.85(55)
J1803–3002A	0.094(15)	0.318(51)	−1.76(33)
J1804–0735A	0.090(14)	–	–
J1824–2452A	0.192(31)	0.328(53)	−0.77(33)
J1824–2452B	0.012(2)	0.019(3)	−0.66(33)
J1824–2452C	0.078(12)	0.089(14)	−0.19(32)
J1824–2452E	0.014(2)	–	–
J1824–2452G	0.032(5)	–	–
J1853–0842A	0.024(4)	–	–
J1823–3021A	0.124(34)	0.218(340)	−0.81(228)
J1823–3021B	0.084(14)	0.092(21)	−0.13(41)
J1823–3021D	0.040(14)	0.051(19)	−0.35(74)

PSRs J1748–2446ao and J1748–2446aw are already very faint in MeerKAT L-band observations [160], with  $S_{1284} = 0.012$  and  $0.010$  mJy, respectively, and their absence in our S-band data is likewise expected, indicating no evidence for an unusual inverted spectrum in either case. By contrast, the extrapolated S-band flux densities of PSRs J1748–2446F and J1748–2446K are  $\sim 0.029$  and  $\sim 0.030$  mJy, respectively, which places them within the practical sensitivity range of our MeerKAT S-band observations. Since both are isolated pulsars, their non-detection is unlikely to be related to orbital-phase effects, and is more plausibly explained by refractive scintillation.

Taken at face value, these limits imply that the pulsars would need to exhibit relatively steep radio spectra for their published L-band flux densities to fall below the S-band detection threshold at 2.4 GHz. In practice, single-epoch non-detections can also be influenced by refractive scintillation and possible duty-cycle evolution with observing frequency. The derived values should therefore be interpreted as conservative upper bounds intended to characterise expected detectability in a single-epoch S-band survey.

Finally, adopting the same survey configuration and tied-array assumptions, the effective S-band detection threshold implied by the radiometer equation corresponds to a practical reference level of order  $\sim 10$ – $20$   $\mu$ Jy at 2.4 GHz for the 1 km PTUSE beams.

Table 4.5: Flux density measurements for the Terzan 5 pulsars detected in our single-epoch MeerKAT S-band observations. Owing to the large number of pulsars detected in this cluster, we present the Terzan 5 measurements separately from the rest of the sample. The  $S_{1400}$  and  $S_{2000}$  values are taken from Martsen et al. [7], while  $S_{2400}$  gives our measured flux densities at a central frequency of  $\sim 2406$  MHz. The spectral index  $\alpha$  is computed using the literature  $S_{1400}$  value together with our measured  $S_{2400}$ , assuming a power-law spectrum,  $S_\nu \propto \nu^\alpha$ . No literature flux density measurement was available for PSR J1748–2446am.

Pulsar name	$S_{1400}$ (mJy)	$S_{2000}$ (mJy)	$S_{2400}$ (mJy)	$\alpha$
J1748–2446A	2.700	1.700	0.702(9)	–2.50
J1748–2446C	1.100	0.670	0.613(12)	–1.08
J1748–2446E	0.170	0.110	0.090(10)	–1.18
J1748–2446G	0.024	0.022	0.027(2)	+0.22
J1748–2446I	0.095	0.055	0.034(1)	–1.91
J1748–2446L	0.096	0.043	0.026(9)	–2.42
J1748–2446M	0.140	0.091	0.068(16)	–1.34
J1748–2446N	0.150	0.100	0.097(22)	–0.81
J1748–2446O	0.310	0.160	0.101(30)	–2.08
J1748–2446V	0.100	0.077	0.034(11)	–2.00
J1748–2446W	0.054	0.031	0.022(8)	–1.67
J1748–2446Y	0.037	0.029	0.021(7)	–1.05
J1748–2446Z	0.030	0.023	0.019(7)	–0.85
J1748–2446ae	0.056	0.050	0.043(11)	–0.49
J1748–2446ai	0.033	0.028	0.021(11)	–0.84
J1748–2446am	...	...	0.008(2)	...

## 4.5 Discussion and Conclusions from the survey

This work presents the first dedicated MeerKAT S-band ( $\nu \simeq 2.4$  GHz) pulsar survey of globular clusters, targeting 14 systems observed during 2023–2024. The survey re-detected 39 previously known pulsars and discovered four new millisecond pulsars in Glimpse-C01 (J1848–0129C, D, E, and F). These results demonstrate that coherent tied-array beams at S-band provide a powerful complement to lower-frequency searches, particularly for high-DM clusters where dispersion and scattering increasingly limit detectability at  $\sim 1$ –1.4 GHz.

A key practical outcome of the survey is an empirical single-epoch detectability scale for MeerKAT S-band tied-array observations. Across detections and non-detections, the results are consistent with a characteristic sensitivity of order  $\sim 10$ – $20$   $\mu\text{Jy}$  for  $\sim 2$  hr integrations under the survey configuration adopted here. This level should be interpreted as an operational planning benchmark for comparable S-band globular-cluster searches, rather than as a strict completeness threshold, since detectability will vary with scintillation, pulse duty cycle, and observing conditions.

J1848–0129C exhibits a combination of a multi-day orbit and reproducible orbital-phase-dependent detectability. The coherent BT solution yields  $P_b = 5.03$  d and a companion mass of  $M_c \simeq 0.35$ – $0.42 M_\odot$  for typical inclinations. Non-detections are confined to orbital phases near superior conjunction and recur across campaigns separated by approximately one year. This phase-locked behaviour disfavors stochastic propagation effects and is suggestive of eclipse or absorption within the binary system.

Interpreting the phase interval  $\Delta\phi \simeq 0.24$  as an eclipse window corresponds to a duration of  $\sim 1.2$  d. Under the assumption of a circular edge-on orbit, this phase width implies a characteristic absorbing scale of order  $\sim 10 R_\odot$ , based on an orbital separation of  $a \simeq 15 R_\odot$ . For comparison, the expected Roche-lobe radius of a  $\sim 0.4 M_\odot$  companion in such an orbit is  $\sim 4 R_\odot$ . While these estimates are intended as approximate geometric constraints, they indicate that a purely compact stellar occultation would be insufficient to account for an eclipse of this duration, and that extended intrabinary material would likely be required if the eclipse interpretation is correct.

With  $P_b = 5.03$  d, J1848–0129C lies at the extreme long-period end of the spider population. Its orbital period exceeds that of known globular-cluster redbacks and is comparable to that of the proposed “Huntsman” systems, such as PSR J1417–4402. Further multiwavelength observations will be required to determine whether the companion is non-degenerate or evolved, and to establish whether J1848–0129C represents a long-period spider system in a globular cluster or an unusually wide

detached configuration.

The orbital parameters of J1848–0129D have also been determined. Using detections from MeerKAT and FAST, we obtained a Keplerian orbital solution that coherently folds all available measurements and is consistent with a circular orbit. The resulting mass function,  $f(M) = 0.166 M_\odot$ , implies, assuming  $M_p = 1.4 M_\odot$ , a minimum companion mass  $M_{c,\min} \simeq 0.98 M_\odot$  for an edge-on configuration and  $M_c \simeq 1.20 M_\odot$  for a median inclination of  $i = 60^\circ$ . While this mass range overlaps the lowest measured neutron-star masses, the absence of measurable eccentricity disfavors a double-neutron-star interpretation and instead points to a mildly recycled pulsar with a massive CO/ONeMg white dwarf companion formed via stable mass transfer. Continued phase-connected timing will refine the orbital parameters and may enable detection of post-Keplerian effects.

J1848–0129E and J1848–0129F are fainter binary millisecond pulsars for which the current detection set is insufficient to obtain coherent orbital solutions. The Glimpse-C01 pulsars are also being pursued in ongoing timing efforts: a timing analysis of J1848–0129A is in preparation (McCarver et al. in prep.), J1848–0129B (Baoda et al. in prep.), and timing work on the newly discovered systems, including J1848–0129C and J1848–0129D, is also in preparation (Nag et al. in prep.).

One notable feature of the pulsars currently known in Glimpse-C01 is that a significant fraction exhibit spin periods between 10 and 20 ms, indicating that they are only mildly recycled. In the case of J1848–0129C, this can be understood in terms of its relatively massive white-dwarf companion, and its spin period is consistent with those observed in systems with similarly massive companions in the Galactic field. The low eccentricity does not indicate a dynamical exchange origin for this system. This raises the question of whether such mildly recycled systems with relatively massive companions are over-abundant in this cluster, which could point to an unusual evolutionary history. However, observational biases must be considered: the high dispersion measure and potential pulse scattering may hinder the detection of faster-spinning pulsars. If so, the intrinsic pulsar population of Glimpse-C01 may be more similar to that of other globular clusters with comparable encounter rates, implying that a substantial population of faster millisecond pulsars remains to be discovered.

GLIMPSE-C01 is one of the most heavily obscured globular clusters in the Milky Way, with a visual extinction of  $A_V \sim 18$  magnitudes [? ], indicating that the cluster lies behind a large column of intervening material. It is also reported to exhibit strong spatial variations in optical extinction (differential reddening; Cadelano et al., in prep.). Among clusters hosting identified pulsars, it is currently the most heavily

extincted system known. The combination of a high mean DM and a large internal DM dispersion may therefore be consistent with substantial small-scale structure in the intervening medium along neighbouring sightlines. This raises the possibility of a connection between fluctuations in ionized gas traced by DM and the complex dust distribution probed at optical and near-infrared wavelengths (Pallanca et al., in prep.), although a quantitative investigation of this link lies beyond the scope of the present work.

The dispersion-measure behaviour observed toward Glimpse-C01 has direct implications for survey strategy at high foreground DMs. Two pulsars detected toward the cluster, J1848–0129D and J1848–0129F, exhibit substantial offsets from the median cluster DM. A comparison across globular clusters with  $\geq 5$  known pulsars shows that Glimpse-C01 lies at the high-DM extreme and follows the empirical trend in which larger median DMs are associated with larger intra-cluster DM dispersions (see Fig. 4.4). Practically, this motivates adopting broader or adaptive DM search ranges for high-DM clusters, since narrowly padded DM windows become increasingly likely to miss genuine cluster pulsars in complex lines of sight.

Finally, the S-band and L-band flux-density measurements reported here provide a uniform high-frequency reference set for re-detected pulsars and a basis for illustrative spectral-index constraints in selected non-detections.

The combined results demonstrate that MeerKAT S-band tied-array observations are capable of recovering faint millisecond pulsars in dense cluster environments and open a complementary discovery space for highly dispersed systems that are challenging to detect at lower frequencies.

## Acknowledgements

This chapter is based on work presented in Nag et al. (2026, submitted). The author acknowledges the support of the South African Radio Astronomy Observatory (SARAO) for MeerKAT operations, the MeerTime and TRAPUM collaborations for access to data products, and the University of Cagliari and INAF–Osservatorio Astronomico di Cagliari for computational resources.

This work was supported in part by the Italian Research Center on High-Performance Computing, Big Data and Quantum Computing (ICSC), funded by the European Union – NextGenerationEU – within the National Recovery and Resilience Plan (PNRR), Mission 4, Component 2.

Full acknowledgements, including funding details and institutional contributions, are provided in the submitted manuscript.

# Chapter 5

## PULSEJET\_beta:

# Pulsar Searching through Efficient Jerk Templates

## 5.1 Going Beyond Acceleration Searches

As discussed in Section 3.6.6, the effect of binary motion on a pulsar signal can be understood most directly through its phase evolution. In the presence of line-of-sight motion, the pulse phase is no longer described by a simple linear function of time, but instead acquires higher-order corrections. Expanding the phase as in Eq. 3.58, one obtains

$$\Phi(t) = 2\pi f_0 \left( t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 + \dots \right), \quad (5.1)$$

where  $a_{\parallel}$  and  $j_{\parallel}$  represent the line-of-sight acceleration and jerk, respectively.

The constant-acceleration approximation corresponds to truncating this expansion at quadratic order in time. Under this assumption, the signal remains locally well-described by a linear frequency evolution, and both time-domain and frequency-domain search techniques can recover the signal efficiently. This regime applies when the orbital motion varies slowly over the observation, such that higher-order terms remain negligible.

However, in more compact binary systems, the orbital motion evolves significantly over the duration of the observation. In such cases, the cubic term in Eq. 3.58 becomes non-negligible, and the signal can no longer be accurately described by a constant acceleration. The phase evolution acquires curvature, and any method that assumes a purely quadratic phase model will become progressively mismatched to

the true signal.

A useful way to characterise the breakdown of the constant-acceleration approximation is to compare the relative contributions of the acceleration and jerk terms over the observation timescale. From Eq. 3.58, the ratio of the cubic to quadratic phase contributions scales as

$$\frac{\Phi_{\text{jerk}}}{\Phi_{\text{acc}}} \sim \frac{j_{\parallel} T_{\text{obs}}}{a_{\parallel}}. \quad (5.2)$$

When this ratio approaches unity, the jerk term contributes significantly to the phase evolution, and the constant-acceleration approximation breaks down. In binary systems, this condition is naturally related to the orbital timescale, leading to the well-known requirement that  $T_{\text{obs}}$  must be small (less than or equal to 10 times) compared to the orbital period for acceleration searches to remain valid [34, 35].

The implications of this breakdown depend on how the search is performed. In frequency-domain approaches, the signal distortion manifests as a non-linear evolution of the apparent frequency, leading to a mismatch with templates that assume linear drift. In time-domain approaches, the same effect corresponds to an incorrect demodulation of the time series when only acceleration is accounted for. In both cases, the underlying issue is identical: the assumed signal model does not capture the true phase evolution.

It is therefore more appropriate to view acceleration and jerk not as features of a particular search method, but as successive terms in a unified description of the signal phase. Both frequency-domain and time-domain techniques attempt to account for these terms, but they do so in different ways: either by modelling the spectral response after the Fourier transform, or by correcting the time series before the transform is applied.

For the majority of pulsar searches, particularly in blind surveys, the constant-acceleration approximation remains adequate and computationally efficient. More advanced treatments, including jerk searches, are therefore best viewed as targeted extensions designed to recover signals from systems that lie beyond the reach of standard acceleration searches.

## 5.2 Time-Domain Resampling Formalism

Having established that the limitations of acceleration searches arise from an incomplete modelling of the signal phase, we now focus on the time-domain approach, in which the correction is applied directly to the time series prior to the Fourier transform. A qualitative description of this method was given in Section 3.6.6; here

we develop the formalism in a more explicit and quantitative manner.

The starting point is the phase model introduced in Eq. 3.58, which may be written as

$$\Phi(t) = 2\pi f_0 \left( t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 + \dots \right). \quad (5.3)$$

This expression describes the phase evolution of a pulsar signal observed in the presence of line-of-sight motion, truncated at the jerk term. The essential idea of time-domain resampling is to construct a new time coordinate  $\tau$  such that the phase becomes linear in  $\tau$ , i.e.

$$\Phi(t) = 2\pi f_0 \tau, \quad (5.4)$$

as already introduced in Eq. 3.59.

By equating the two expressions, one obtains the transformation

$$\tau = t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 + \dots, \quad (5.5)$$

which is identical to Eq. 3.60. This transformation defines the mapping between the observed time coordinate  $t$  and the corrected time coordinate  $\tau$ , in which the pulsar appears as an isolated periodic source.

In practice, the observed signal is discretely sampled at intervals  $t_{\text{samp}}$ , yielding a time series  $s(t_i)$  with  $t_i = i t_{\text{samp}}$ . The goal is to construct a new time series  $s^{\text{rs}}(\tau_j)$  sampled uniformly in  $\tau$ , where  $\tau_j = j t_{\text{samp}}$ . This requires evaluating the inverse mapping  $t(\tau)$  and determining which sample in the original time series corresponds to each resampled point.

Since the mapping between  $t$  and  $\tau$  is generally non-linear, the values of  $t(\tau_j)$  will not coincide exactly with the original sampling grid. As a result, the resampled time series must be constructed using either nearest-neighbour assignment or interpolation between adjacent samples. In the simplest implementation, one identifies the index  $i$  such that  $t_i$  is closest to  $t(\tau_j)$  and sets

$$s_j^{\text{rs}} = s_i. \quad (5.6)$$

More accurate implementations employ linear or higher-order interpolation schemes to reduce discretisation errors [103].

Once the resampled time series has been constructed, it can be analysed using a standard Fourier transform. If the trial parameters  $(a_{\parallel}, j_{\parallel})$  are sufficiently close to the true values, the phase modulation induced by the binary motion is effectively removed, and the signal power is concentrated into a narrow set of Fourier bins corresponding to the intrinsic spin frequency and its harmonics. In this sense, time-

domain resampling performs a coherent demodulation of the signal prior to the FFT.

The key computational implication of this approach is that the entire procedure—resampling followed by FFT—must be repeated for each trial set of parameters. If one searches over a grid in acceleration and jerk, the total number of trials scales as

$$N_{\text{trial}} \sim N_a \times N_j, \quad (5.7)$$

where  $N_a$  and  $N_j$  are the number of trial acceleration and jerk values, respectively. This scaling leads to a rapid increase in computational cost as higher-order terms are included.

Despite this cost, the time-domain formulation provides a direct way to apply a chosen phase model to the time series before the FFT is computed. This makes it conceptually convenient to implement polynomial corrections, such as constant acceleration or acceleration plus jerk, and it can also be extended to more detailed orbital prescriptions. This generality is not unique to the time-domain approach: frequency-domain methods can also search for orbital modulation, including through phase-modulation or sideband techniques implemented in PRESTO’s binary-search tools [161]. The distinction made here is therefore not that full orbital information is exclusive to time-domain methods, but that PULSEJET beta explores a time-domain implementation of the acceleration–jerk polynomial correction.

### 5.3 The PULSEJET\_beta Pipeline: Motivation and Design

Modern pulsar search pipelines broadly fall into two categories: frequency-domain approaches, such as those implemented in PRESTO, and time-domain resampling approaches, as implemented in software such as `Peasoup`<sup>1</sup>. While both methods are rooted in the same physical description of binary motion, they differ in how the signal model is applied in practice.

`Peasoup` is a GPU-accelerated pulsar search code designed to perform coherent time-domain acceleration searches. It operates by resampling the input time series according to trial values of the line-of-sight acceleration and subsequently applying a Fourier transform to recover periodic signals. In contrast to frequency-domain matched filtering, this approach directly corrects the phase evolution of the signal prior to the FFT, enabling fully coherent recovery when the assumed model is

---

<sup>1</sup><https://github.com/ewanbarr/peasoup>

accurate.

Time-domain resampling techniques have been used in pulsar searches for decades [34, 103]. In their modern GPU-accelerated form, they are implemented in the PEASOUP search code and have been used in several major pulsar surveys and search efforts. For example, PEASOUP was used in the HTRU-S medlat/hilat reprocessing that reported 23 new pulsars [162], in the SUPERB survey real-time periodicity pipeline [163], in the TRAPUM L-band survey for pulsars in *Fermi*-LAT sources [164], in the MMGPS-L survey pipeline description [165], and in the GPU-accelerated reprocessing of archival PMPS data that yielded 37 new pulsars [166]. These examples show that PEASOUP is not merely a conceptual implementation of time-domain resampling, but an established and practically validated search code in modern pulsar astronomy.

Despite this, existing implementations of time-domain searches have been largely limited to the constant-acceleration approximation. While higher-order corrections such as jerk can, in principle, be incorporated into the phase model, there has been no widely adopted pipeline that performs a systematic jerk search in the time domain for blind pulsar surveys. This limitation motivates the development of the PULSEJET\_beta pipeline <sup>2</sup>, which extends the resampling formalism to include jerk and enables a coherent search over both acceleration and its first derivative.

The development of PULSEJET\_beta is driven by two primary scientific motivations which are discussed in the next two subsections.

### 5.3.1 Accessing Extreme Regions of Parameter Space

The primary scientific motivation for extending the Peasoup framework to a jerk-capable time-domain pipeline is not that standard acceleration searches are ineffective in general, but rather that they become progressively sub-optimal in a very specific and astrophysically important region of parameter space. That region corresponds to compact binaries in which the line-of-sight motion varies rapidly over the observation and the higher-order terms in the signal phase can no longer be neglected. These are precisely the systems that are often regarded as the “holy grail” targets of radio pulsar searches: extremely compact double neutron star systems, and eventually pulsar–black hole binaries, whose discovery would have exceptional value for relativistic gravity, binary evolution, and gravitational-wave astrophysics [77, 167, 168].

The point is therefore not that acceleration searches fail everywhere, nor that time-domain jerk searches should replace them in all circumstances. On the con-

---

<sup>2</sup>[https://github.com/Rouhin1997/PULSEJET\\_beta.git](https://github.com/Rouhin1997/PULSEJET_beta.git)

trary, standard Fourier-domain acceleration and jerk searches are highly efficient in their native coordinates because they tile the Fourier response directly in terms of frequency, frequency drift, and frequency-drift derivative. The motivation for `PULSEJET_beta` is instead that a time-domain search over physical acceleration and jerk provides a complementary way of specifying the search region. In particular, it allows the bounds of the search to be chosen directly in terms of the line-of-sight acceleration and jerk expected for a target class of compact binaries.

To make this comparison quantitative, it is useful to examine how fixed Fourier-domain search limits map into physical acceleration and jerk as a function of pulsar spin period. In the Fourier-domain acceleration search implemented in `PRESTO`, the basic measurable quantity is the drift of the signal through Fourier bins over an observation of duration  $T_{\text{obs}}$ . For a signal observed at spin frequency  $f_0 = 1/P$ , the total bin drift under constant acceleration is

$$z \approx \dot{f} T_{\text{obs}}^2 = \frac{a_{\parallel}}{cP} T_{\text{obs}}^2, \quad (5.8)$$

where  $a_{\parallel}$  is the line-of-sight acceleration. When harmonic summing is used, the same physical acceleration produces a proportionally larger drift at the  $h$ -th harmonic. Therefore, for a maximum search value  $z_{\text{max}}$ , the corresponding physical acceleration boundary may be written as

$$|a_{\parallel}| \leq \frac{z_{\text{max}} cP}{h T_{\text{obs}}^2}. \quad (5.9)$$

This relation shows that a fixed Fourier-domain choice of  $z_{\text{max}}$  corresponds to a spin-period-dependent boundary in physical acceleration. This should not be interpreted as an intrinsic inefficiency of the Fourier-domain method; rather, it is the result of expressing the search in Fourier-bin drift coordinates rather than directly in physical acceleration units.

An analogous mapping applies to jerk. If the acceleration varies over the observation, the quadratic term in the frequency evolution introduces a drift-rate derivative in Fourier space, usually denoted by  $w$  [106]. Using

$$\ddot{f} = \frac{\dot{j}_{\parallel}}{cP}, \quad (5.10)$$

the corresponding Fourier-space curvature parameter scales as

$$w \approx \ddot{f} T_{\text{obs}}^3 = \frac{\dot{j}_{\parallel}}{cP} T_{\text{obs}}^3. \quad (5.11)$$

Including harmonic summing gives the corresponding physical jerk boundary

$$|j_{\parallel}| \leq \frac{w_{\max} c P}{h T_{\text{obs}}^3}. \quad (5.12)$$

Thus, for fixed values of  $z_{\max}$  and  $w_{\max}$ , the region covered in physical  $(a_{\parallel}, j_{\parallel})$  space scales with the pulsar spin period. In a period–acceleration or period–jerk plot this naturally appears as a wedge. The wedge is therefore best understood as the mapping of uniform Fourier-domain limits into physical units, not as a fundamental limitation of Fourier-domain searching.

The time-domain resampling approach is parameterised differently. Once the phase model is specified, one may choose directly the physical search ranges

$$|a_{\parallel}| \leq a_{\max}^{(\text{td})}, \quad |j_{\parallel}| \leq j_{\max}^{(\text{td})}, \quad (5.13)$$

and resample the time series over that band. In other words, the search boundary is imposed directly in physical parameter space rather than indirectly through Fourier-bin drift. For fixed choices of  $a_{\max}^{(\text{td})}$  and  $j_{\max}^{(\text{td})}$ , the adopted time-domain search region therefore appears as a rectangular band in physical acceleration and jerk, independent of spin period. This is useful when the intended target region is defined by physical binary parameters rather than by fixed Fourier drift limits.

This distinction is illustrated in Figure 5.0, where the PRESTO Fourier-domain limits are mapped into physical units and compared with a representative time-domain band for  $T_{\text{obs}} = 600$  s. In the acceleration case, the transition period at which the Fourier-domain boundary first contains the chosen time-domain acceleration band follows directly from Eq. 5.9:

$$P_{\text{cross},a} = \frac{h a_{\max}^{(\text{td})} T_{\text{obs}}^2}{z_{\max} c}. \quad (5.14)$$

Similarly, for jerk,


$$P_{\text{cross},j} = \frac{h j_{\max}^{(\text{td})} T_{\text{obs}}^3}{w_{\max} c}. \quad (5.15)$$

These crossing periods are marked in the figures. They quantify where a particular choice of Fourier-domain search limits and a particular choice of time-domain physical bounds cover the same region. Below  $P_{\text{cross}}$ , the selected Fourier-domain limits correspond to a narrower physical interval than the chosen time-domain band; above  $P_{\text{cross}}$ , they contain that band.


For the representative values adopted here, the acceleration comparison shows how fixed  $z_{\max}$  limits map into period-dependent physical acceleration coverage.

The effect is more visible in the jerk case. For example, with  $w_{\max} = 4000$ ,  $j_{\max}^{(\text{td})} = 20 \text{ m s}^{-3}$ , and  $T_{\text{obs}} = 600 \text{ s}$ , the crossing period moves from  $\sim 14.4 \text{ ms}$  for  $h = 4$  to  $\sim 28.8 \text{ ms}$  for  $h = 8$  and  $\sim 57.6 \text{ ms}$  for  $h = 16$ . These numbers should be interpreted only as a comparison between two specific choices of search parameterisation and computational budget. They do not imply that Fourier-domain methods are intrinsically unable to cover the same physical region: increasing  $z_{\max}$  or  $w_{\max}$  would extend the Fourier-domain coverage, at the cost of additional computation. Conversely, achieving full sensitivity to very fast spin periods with time-domain resampling also requires a sufficiently dense set of acceleration–jerk trials, which can be computationally expensive.


The practical distinction is therefore not one of fundamental accessibility, but one of implementation and computational trade-off. Fourier-domain methods tile the signal response efficiently in their natural coordinates  $(f, z, w)$  using matched filters, whereas time-domain resampling applies each trial correction to the time series and then performs a one-dimensional FFT. This latter structure is attractive for GPU implementation because the resampling trials can be parallelised and the subsequent harmonic summing is performed on ordinary one-dimensional spectra. However, the total number of floating-point operations for equivalent coverage may be larger than in a highly optimised Fourier-domain matched-filter implementation. The aim of `PULSEJET_beta` is therefore to explore whether GPU acceleration and metric-informed template placement can make this time-domain strategy competitive for targeted searches over physically motivated acceleration–jerk regions.




$$h = 4$$




$$h = 4$$




$$h = 8$$



$$h = 8$$



$$h = 16$$



$$h = 16$$

Figure 5.0: Comparison between the parameter-space coverage of the Fourier-domain PRESTO search (blue wedge) and a representative time-domain search band (orange), shown for acceleration (left column) and jerk (right column). The three rows correspond to harmonic sums of  $h = 4, 8,$  and  $16$ . Increasing  $h$  improves sensitivity to narrow pulse profiles through harmonic summing, but simultaneously makes the Fourier-domain search boundary more restrictive, particularly at short spin periods. The vertical dotted lines indicate the crossing period below which the Fourier-domain wedge no longer fully covers the chosen time-domain band. This effect is most pronounced for jerk, where the loss of coverage occurs precisely in the region occupied by rapidly spinning pulsars in extremely compact binaries.

This is the central scientific rationale for `PULSEJET_beta`. The pipeline is designed specifically to interrogate that small but high-value region of parameter space where compact orbital motion produces large jerk and where the standard acceleration-search formalism is no longer the most natural or sensitive way to proceed. As Ewan Barr put it during the development of this work:

“We are looking for those holy grail systems like pulsar–black hole binaries and DNS systems before merger and stuff, otherwise what even is the point of writing codes like this? Otherwise it is just madness.”

That statement captures the philosophy behind the present pipeline. The purpose is not to replace existing survey search tools, but to extend the reachable search space towards the most extreme compact-binary systems, where the scientific return of a discovery would far outweigh the additional computational cost. Once that motivation is established, the need for an efficient GPU implementation follows naturally: if the goal is to search coherently in this difficult regime, then the computational burden of the time-domain jerk search must be brought under control.

### 5.3.2 Computational Cost of Jerk Searches

The second motivation arises from the computational cost associated with extending the phase model to include higher-order terms. In the time-domain framework, each trial set of parameters  $(a_{\parallel}, j_{\parallel})$  requires the construction of a new resampled time series followed by a Fourier transform. As a result, the computational cost scales with the total number of templates in the search space.

The inclusion of the jerk term introduces a quadratic dependence in the resampling transformation (Eq. 3.60), which significantly increases the complexity of the mapping between input and output time indices. This effect is illustrated in Figures 5.1 and 5.2, which show the deformation of the resampling index and its local increments for increasing values of jerk.

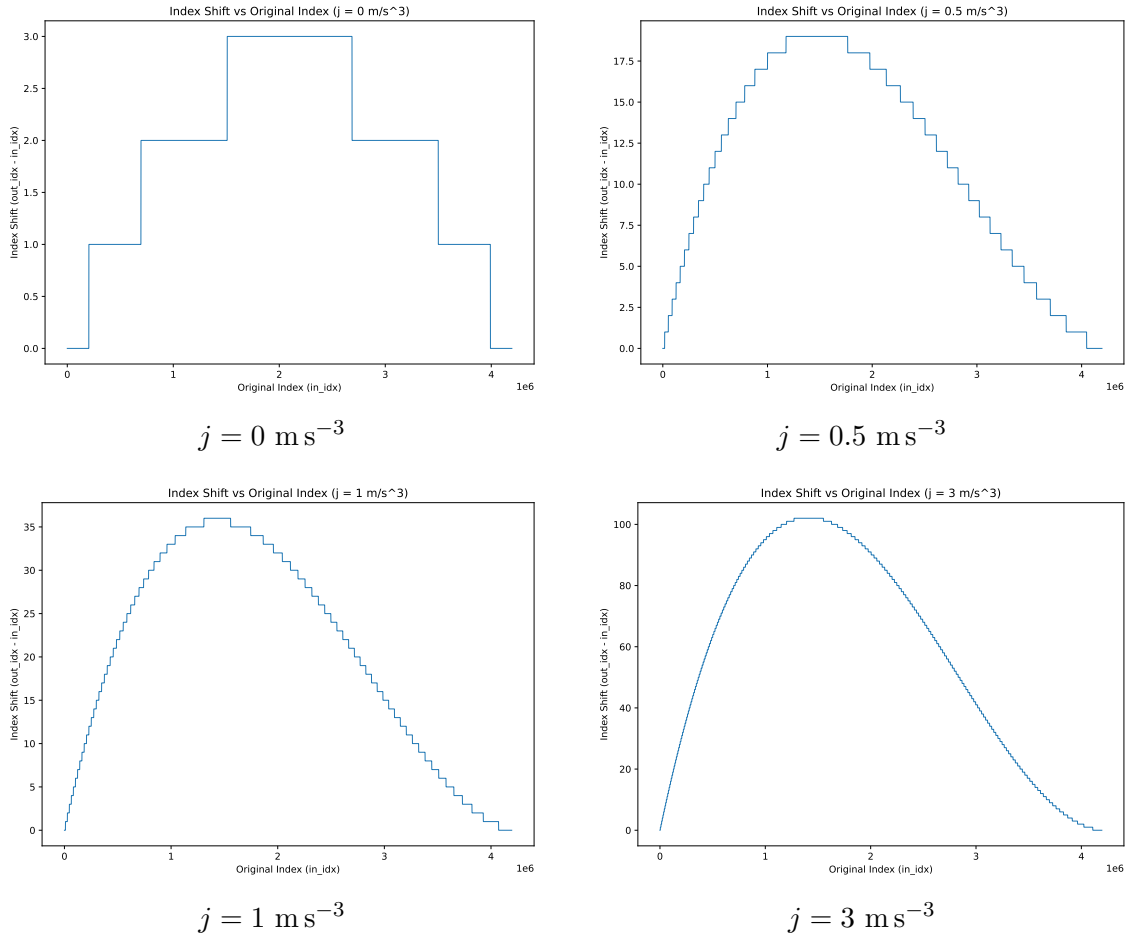


Figure 5.1: Evolution of the resampling index shift,  $(\text{out\_idx} - \text{in\_idx})$ , for increasing trial jerk for the same value of acceleration. In the pure-acceleration case ( $j = 0$ ), the mapping remains only weakly distorted, whereas progressively larger jerk values introduce an increasingly strong non-linear deformation of the output index as a function of input index. This reflects the quadratic contribution of the jerk term to the time-domain resampling map. As  $j$  increases, the resampler must account for progressively larger departures from the original regular sampling pattern, implying a larger correction to the time series before the FFT is performed.

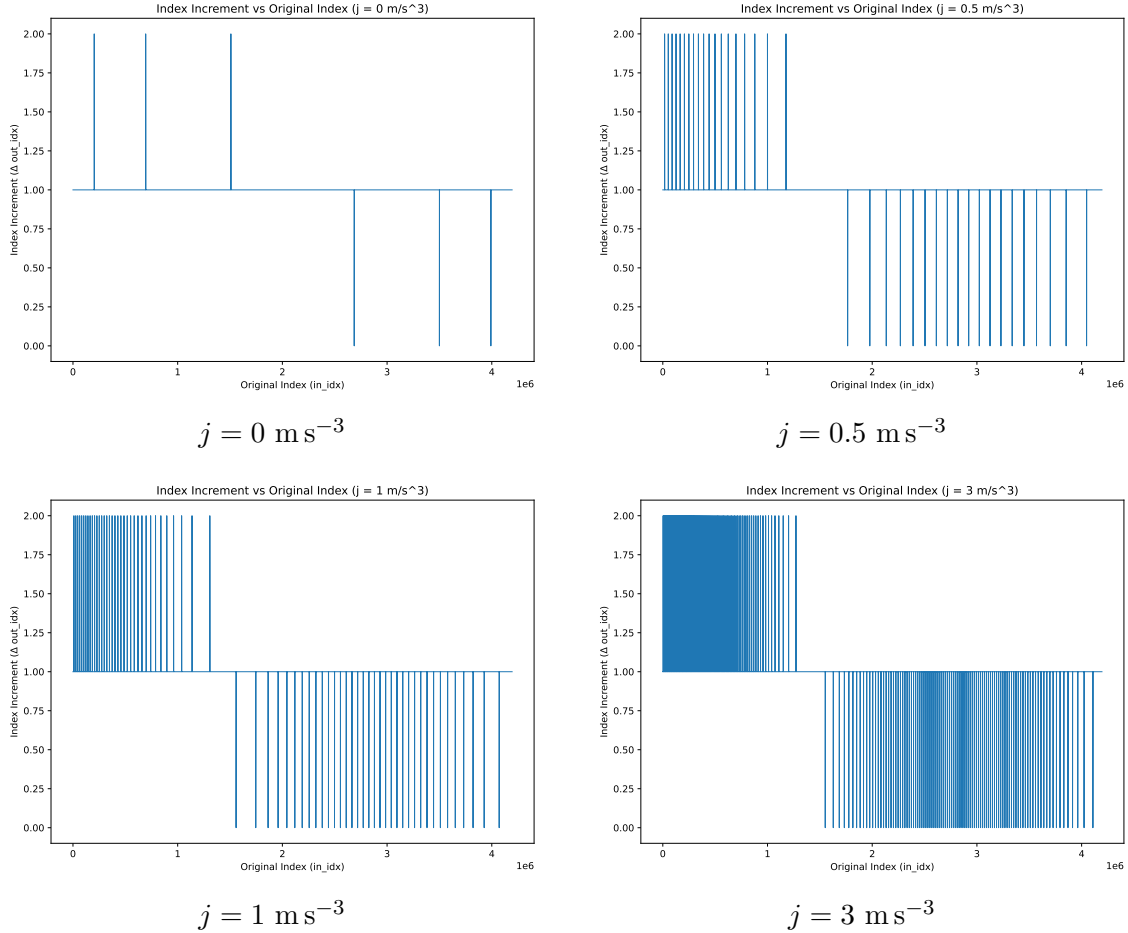


Figure 5.2: Local increment of the resampling index,  $\Delta \text{out\_idx}$ , for increasing trial jerk keeping the value of acceleration constant. While the  $j = 0$  case remains comparatively regular, the inclusion of jerk introduces a position-dependent variation in the local stepping of the output index. This means that neighbouring output samples no longer map back to the input array with a uniform stride. Instead, the access pattern becomes increasingly irregular as the jerk increases, which is precisely the behaviour that drives up the computational cost of the resampling stage. In practice, this leads to less regular memory access and motivates the use of massively parallel GPU implementations when extending time-domain searches beyond constant acceleration.

As the jerk increases, the mapping between input and output indices becomes increasingly non-linear, leading to irregular memory access patterns and reduced computational efficiency. When extended over a large template bank, this results in a rapid increase in total computational cost, making a naive CPU implementation impractical for large-scale searches.

This motivates the adoption of GPU-based implementations, where the parallel nature of the resampling and FFT operations can be exploited to achieve the required performance. The PULSEJET\_beta pipeline is therefore designed as a GPU-accelerated extension of the Peasoup framework, enabling efficient exploration of a two-dimensional parameter space in acceleration and jerk.

In the following sections, we describe the implementation of the jerk resampling algorithm, the construction of the template bank, and the performance characteristics of the pipeline.

## 5.4 Astrophysical Determination of the Acceleration and Jerk Bounds

Before constructing a template bank in the  $(a_{\parallel}, j_{\parallel})$  plane, one must first decide which region of that plane the search should cover at all. This is a scientific choice before it is a computational one. The adopted bounds should be large enough to include the classes of binaries that the search is intended to detect, but not so large that computational effort is wasted on clearly implausible regions of parameter space.

In the present work, the guiding requirement is sensitivity to very compact binary systems in which the pulsar can experience large line-of-sight acceleration and large line-of-sight jerk during a single observation. The relevant search boundaries are therefore set not from an abstract Fourier-space criterion, but from physically motivated orbital models. We ask the following question: for a binary system with plausible masses, inclination, orbital period, and eccentricity, what are the largest values of  $|a_{\parallel}|$  and  $|j_{\parallel}|$  that can occur over an orbit?

### 5.4.1 Keplerian description of the line-of-sight motion

Consider a pulsar of mass  $M_p$  in a binary with companion mass  $M_c$ , orbital period  $P_b$ , eccentricity  $e$ , longitude of periastron  $\omega$ , and inclination  $i$ . Let the pulsar orbit the system barycentre with semi-major axis  $a_p$ . Kepler's third law gives the relative

semi-major axis of the binary,

$$a^3 = \frac{G(M_p + M_c)P_b^2}{4\pi^2}, \quad (5.16)$$

and the pulsar semi-major axis is

$$a_p = a \frac{M_c}{M_p + M_c}. \quad (5.17)$$

The projected semi-major axis is therefore

$$a_p \sin i = \left( \frac{GP_b^2}{4\pi^2} \right)^{1/3} \frac{M_c \sin i}{(M_p + M_c)^{2/3}}. \quad (5.18)$$

Equivalently, in pulsar-timing notation one often writes

$$x \equiv \frac{a_p \sin i}{c}, \quad (5.19)$$

but for the present derivation it is convenient to retain the dimensional quantity  $a_p \sin i$  explicitly.

Define the mean orbital angular frequency

$$n \equiv \frac{2\pi}{P_b}. \quad (5.20)$$

The orbital phase is most conveniently parametrised through the mean anomaly

$$M(t) = n(t - T_0), \quad (5.21)$$

where  $T_0$  is the epoch of periastron passage. For an eccentric orbit,  $M$  is not itself the geometric orbital angle. Instead, one introduces the eccentric anomaly  $E$ , defined implicitly by Kepler's equation,

$$M = E - e \sin E. \quad (5.22)$$

This transcendental equation has no closed-form solution in elementary functions and must therefore be solved iteratively. In practice, we use Newton–Raphson iteration,

$$E_{k+1} = E_k + \frac{M - (E_k - e \sin E_k)}{1 - e \cos E_k}, \quad (5.23)$$

with a suitable initial guess, until convergence is achieved. Once  $E$  is known, the

true anomaly  $\nu$  follows from

$$\tan \frac{\nu}{2} = \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2}. \quad (5.24)$$

The true anomaly  $\nu$  is the physical orbital angle measured from periastron and is the most natural variable for evaluating the instantaneous line-of-sight motion [2].

### 5.4.2 Line-of-sight velocity

For a Keplerian orbit, the pulsar line-of-sight velocity can be written as

$$v_{\parallel}(t) = \frac{n a_p \sin i}{\sqrt{1-e^2}} [\cos(\omega + \nu) + e \cos \omega]. \quad (5.25)$$

This is the quantity that enters the Doppler relation for the observed spin frequency,

$$f_{\text{obs}}(t) = f_0 \left( 1 + \frac{v_{\parallel}(t)}{c} \right), \quad (5.26)$$

up to the adopted sign convention for  $v_{\parallel}$ . The important point for the present purpose is that once  $v_{\parallel}(t)$  is known, the corresponding acceleration and jerk follow by differentiation.

### 5.4.3 Line-of-sight acceleration

The line-of-sight acceleration is

$$a_{\parallel}(t) = \frac{dv_{\parallel}}{dt}. \quad (5.27)$$

To evaluate this derivative, one uses the standard Keplerian identity

$$\dot{\nu} = n \frac{(1 + e \cos \nu)^2}{(1 - e^2)^{3/2}}. \quad (5.28)$$

Differentiating Eq. 5.25 then gives

$$a_{\parallel}(t) = -\frac{n^2 a_p \sin i}{(1 - e^2)^2} (1 + e \cos \nu)^2 \sin(\omega + \nu). \quad (5.29)$$

This expression shows explicitly how the maximum line-of-sight acceleration depends on the projected orbit size  $a_p \sin i$ , the orbital timescale through  $n^2$ , and the eccentricity-dependent factor  $(1 + e \cos \nu)^2 / (1 - e^2)^2$ .

For a circular orbit ( $e = 0$ ), Eq. 5.29 reduces to

$$a_{\parallel}(t) = -n^2 a_p \sin i \sin(\omega + \nu), \quad (5.30)$$

so that the maximum absolute acceleration is simply

$$|a_{\parallel}|_{\max} = n^2 a_p \sin i = \left(\frac{2\pi}{P_b}\right)^2 a_p \sin i. \quad (5.31)$$

This already makes the principal scaling transparent: for fixed masses and inclination, the maximum acceleration grows rapidly as  $P_b^{-2}$ .

#### 5.4.4 Line-of-sight jerk

The line-of-sight jerk is the time derivative of the acceleration,

$$j_{\parallel}(t) = \frac{da_{\parallel}}{dt}. \quad (5.32)$$

Differentiating Eq. 5.29 and again using Eq. 5.28 yields

$$j_{\parallel}(t) = -\frac{n^3 a_p \sin i}{(1 - e^2)^{7/2}} (1 + e \cos \nu)^3 [\cos(\omega + \nu) + e \cos \omega - 3e \sin(\omega + \nu) \sin \nu]. \quad (5.33)$$

As for the acceleration, the circular-orbit limit is especially simple. Setting  $e = 0$  gives

$$j_{\parallel}(t) = -n^3 a_p \sin i \cos(\omega + \nu), \quad (5.34)$$

so that

$$|j_{\parallel}|_{\max} = n^3 a_p \sin i = \left(\frac{2\pi}{P_b}\right)^3 a_p \sin i. \quad (5.35)$$

Thus, for fixed masses and inclination, the maximum jerk grows even more steeply than the acceleration, scaling as  $P_b^{-3}$ .

#### 5.4.5 Determining the search bounds

Equations 5.29 and 5.33 provide the full orbital expressions for the line-of-sight acceleration and jerk. In the circular case, the maxima are given analytically by Eqs. 5.31 and 5.35. For eccentric systems, however, the location of the extrema depends on  $\omega$ ,  $e$ , and  $\nu$  in a non-trivial way, and it is generally more robust to determine the extrema numerically over a full orbit:

$$a_{\parallel, \max} = \max_t |a_{\parallel}(t)|, \quad j_{\parallel, \max} = \max_t |j_{\parallel}(t)|. \quad (5.36)$$

This is the strategy adopted here. For each trial astrophysical binary configuration, we sample one full orbit, solve Kepler’s equation iteratively at each time step, convert to the true anomaly, and evaluate Eqs. 5.29 and 5.33. The search boundaries are then chosen so as to encompass the maximum absolute values of  $a_{\parallel}$  and  $j_{\parallel}$  expected for the class of systems to which the pipeline is intended to remain sensitive.

The scientific purpose of this exercise is straightforward. If the search is to remain sensitive to extremely compact binaries, then its  $(a_{\parallel}, j_{\parallel})$  coverage must be set by the physical maxima that such systems can produce. Only after those physical boundaries are defined does it become meaningful to discretise that region into a practical template bank, which is the subject of the next section.

## 5.5 Metric Construction and Stochastic Population of the $(a_{\parallel}, j_{\parallel})$ Template Bank

Once physically motivated bounds on the maximum line-of-sight acceleration and jerk have been established (Section 5.4), the next task is to populate that region of parameter space with a finite set of search templates. The purpose of the template bank is to ensure that for any true signal whose parameters lie inside the chosen search box,

$$-a_{\parallel,\max} \leq a_{\parallel} \leq a_{\parallel,\max}, \quad -j_{\parallel,\max} \leq j_{\parallel} \leq j_{\parallel,\max}, \quad (5.37)$$

there exists at least one nearby template whose phase model is sufficiently close that the loss in coherent signal power remains below a specified tolerance.

This section describes that construction in full detail for the acceleration–jerk model adopted in `PULSEJET.beta`. The procedure follows the general logic of metric-based template placement developed in coherent searches for pulsars and gravitational waves, in which the local mismatch between neighbouring parameter-space points is approximated by a quadratic form defined by a metric tensor [116–118]. The overall strategy is:

1. choose a phase model for the signal;
2. compute the metric on the full parameter space;
3. project out the spin-frequency direction, since the FFT already covers it;
4. use the determinant of the projected metric to define the local template density;

5. estimate the proper volume of the chosen search region;
6. populate that region stochastically according to the required density;
7. validate the resulting bank by explicitly checking the nearest-template mismatch distribution.

In the present implementation, all of these steps are carried out for the polynomial time-domain phase model in acceleration and jerk that underlies the resampling formalism developed earlier in this chapter. The final result is therefore not an arbitrary cloud of points in the  $(a_{\parallel}, j_{\parallel})$  plane, but a bank constructed from the geometry of the signal manifold itself.

### 5.5.1 Choice of phase model

The starting point is the jerk-corrected phase model. In the earlier discussion of time-domain resampling, the signal phase was written in the form

$$\Phi(t) = 2\pi f_0 \left( t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 + \dots \right), \quad (5.38)$$

which is convenient when the resampling correction is expressed as a mapping of the observed time coordinate. For the template-bank construction, however, it is algebraically more convenient to write the template phase in the equivalent form

$$\Phi(t; f, a, j) = 2\pi f \left( t - \frac{a t^2}{2c} - \frac{j t^3}{6c} \right), \quad (5.39)$$

where  $f$  is the template spin frequency and  $(a, j)$  denote the template acceleration and jerk parameters. The change of sign simply reflects whether one chooses to write the modulation in the signal itself or in the demodulating template. The two viewpoints are equivalent provided the sign convention is used consistently.

Equation 5.39 is the model implemented in the template-bank code. It is a truncated polynomial phase model with three parameters,

$$\boldsymbol{\lambda} = (f, a, j). \quad (5.40)$$

The role of the template bank is to cover the  $(a, j)$  subspace for a specified maximum signal frequency, here denoted by  $f_{\max}$ , at which the search is intended to remain sensitive. This is a crucial design parameter of the bank. In the phase model used here, the derivatives with respect to acceleration and jerk scale linearly with the signal frequency, and therefore the metric components in the projected  $(a, j)$  subspace

scale approximately as  $f^2$ . Consequently, the local template density, proportional to  $\sqrt{\det \gamma}$ , scales approximately as  $f^2$  for fixed  $T_{\text{obs}}$ , acceleration–jerk bounds, and mismatch tolerance. Increasing  $f_{\text{max}}$  therefore improves the ability of the bank to preserve sensitivity to short spin periods and higher harmonics, but it also increases the number of templates and hence the computational cost.

For a pulsar with fundamental spin frequency  $f_0 = 1/P$ , the  $h$ -th harmonic appears at frequency  $hf_0$ . A conservative choice of bank frequency is therefore

$$f_{\text{max}} = \min \left( f_{\text{Nyq}}, \frac{h_{\text{max}}}{P_{\text{min}}} \right), \quad (5.41)$$

where  $f_{\text{Nyq}}$  is the Nyquist frequency of the sampled time series,  $h_{\text{max}}$  is the highest harmonic one aims to recover with controlled mismatch, and  $P_{\text{min}}$  is the shortest spin period targeted by the search. A bank generated at  $f_{\text{max}}$  is conservative for lower-frequency signals and lower harmonics, because their phase mismatch for the same  $(\Delta a, \Delta j)$  offsets is smaller. Thus,  $f_{\text{max}}$  should be treated as one of the primary bank-generation parameters, together with  $T_{\text{obs}}$ ,  $a_{\text{max}}$ ,  $j_{\text{max}}$ , the target mismatch, and the desired coverage.

## 5.5.2 Mismatch and local metric approximation

Suppose a signal is present with true parameters  $\boldsymbol{\lambda}$ , but is filtered with a nearby template at  $\boldsymbol{\lambda} + \Delta\boldsymbol{\lambda}$ . The fractional loss in coherent detection statistic is called the mismatch, conventionally denoted by  $m$ . In the small-offset limit, the mismatch can be expanded to second order in the coordinate offsets,

$$m(\boldsymbol{\lambda}, \Delta\boldsymbol{\lambda}) \approx g_{\alpha\beta} \Delta\lambda^\alpha \Delta\lambda^\beta, \quad (5.42)$$

where repeated indices imply summation and  $g_{\alpha\beta}$  is the metric tensor on the signal manifold [118, 119]. For a phase model search, the metric can be written as

$$g_{\alpha\beta} = \langle \partial_\alpha \Phi \partial_\beta \Phi \rangle - \langle \partial_\alpha \Phi \rangle \langle \partial_\beta \Phi \rangle, \quad (5.43)$$

where  $\partial_\alpha \equiv \partial/\partial\lambda^\alpha$  and

$$\langle X \rangle \equiv \frac{1}{T_{\text{obs}}} \int_0^{T_{\text{obs}}} X(t) dt \quad (5.44)$$

denotes a time average over the observation.

Equation 5.43 is the central mathematical object of the template-bank construction. It measures how rapidly the signal phase changes under infinitesimal displacements in parameter space. Directions in which the phase changes rapidly require

denser template placement; directions in which it changes slowly require fewer templates.

### 5.5.3 Derivatives of the acceleration–jerk phase model

For the phase model in Eq. 5.39, the derivatives with respect to the three parameters are

$$\frac{\partial\Phi}{\partial f} = 2\pi \left( t - \frac{at^2}{2c} - \frac{jt^3}{6c} \right), \quad (5.45)$$

$$\frac{\partial\Phi}{\partial a} = -\frac{\pi f}{c} t^2, \quad (5.46)$$

$$\frac{\partial\Phi}{\partial j} = -\frac{\pi f}{3c} t^3. \quad (5.47)$$

Substituting these into Eq. 5.43 yields the full  $3 \times 3$  metric on the parameter space  $(f, a, j)$ . After carrying out the time averages analytically over the interval  $0 \leq t \leq T_{\text{obs}}$ , one obtains

$$g_{\alpha\beta} = \begin{pmatrix} g_{ff} & g_{fa} & g_{fj} \\ g_{af} & g_{aa} & g_{aj} \\ g_{jf} & g_{ja} & g_{jj} \end{pmatrix}, \quad (5.48)$$

with

$$g_{ff} = \frac{\pi^2 T_{\text{obs}}^2}{5040 c^2} \left( 45 T_{\text{obs}}^4 j^2 + 280 T_{\text{obs}}^3 a j + 448 T_{\text{obs}}^2 a^2 - 504 T_{\text{obs}}^2 c j - 1680 T_{\text{obs}} a c + 1680 c^2 \right), \quad (5.49)$$

$$g_{fa} = g_{af} = \frac{\pi^2 T_{\text{obs}}^3 f}{180 c^2} (5 T_{\text{obs}}^2 j + 16 T_{\text{obs}} a - 30 c), \quad (5.50)$$

$$g_{fj} = g_{jf} = \frac{\pi^2 T_{\text{obs}}^4 f}{5040 c^2} (45 T_{\text{obs}}^2 j + 140 T_{\text{obs}} a - 252 c), \quad (5.51)$$

$$g_{aa} = \frac{4\pi^2 T_{\text{obs}}^4 f^2}{45 c^2}, \quad (5.52)$$

$$g_{aj} = g_{ja} = \frac{\pi^2 T_{\text{obs}}^5 f^2}{36 c^2}, \quad (5.53)$$

$$g_{jj} = \frac{\pi^2 T_{\text{obs}}^6 f^2}{112 c^2}. \quad (5.54)$$

At this stage the metric still includes the spin-frequency dimension. This is not yet the metric that should be used for the template bank, because the search does

not place explicit templates in  $f$ . Instead, the FFT already provides an efficient and effectively regular search over frequency.

#### 5.5.4 Projection onto the $(a, j)$ subspace

The metric relevant to the template bank must therefore be the metric on the  $(a, j)$  subspace after projecting out the frequency direction. Following the standard construction used by Knispel for orbital template banks [116, 169], the projected metric is obtained through the Schur complement,

$$\gamma_{ij} = g_{ij} - \frac{g_{if}g_{jf}}{g_{ff}}, \quad i, j \in \{a, j\}. \quad (5.55)$$

This projected metric measures mismatch in the acceleration–jerk plane once the frequency has been allowed to adjust optimally through the FFT.

It is convenient to define the common denominator

$$D(a, j) = 45T_{\text{obs}}^4 j^2 + 280T_{\text{obs}}^3 a j + 448T_{\text{obs}}^2 a^2 - 504T_{\text{obs}}^2 c j - 1680T_{\text{obs}} a c + 1680c^2. \quad (5.56)$$

The projected metric elements then become

$$\gamma_{aa} = \frac{\pi^2 T_{\text{obs}}^4 f^2}{45c^2 D} (5T_{\text{obs}}^4 j^2 + 84T_{\text{obs}}^2 c j + 420c^2), \quad (5.57)$$

$$\gamma_{aj} = \gamma_{ja} = \frac{\pi^2 T_{\text{obs}}^5 f^2}{90c^2 D} (-10T_{\text{obs}}^3 a j + 45T_{\text{obs}}^2 c j - 84T_{\text{obs}} a c + 420c^2), \quad (5.58)$$

$$\gamma_{jj} = \frac{\pi^2 T_{\text{obs}}^6 f^2}{45c^2 D} (5T_{\text{obs}}^2 a^2 - 45T_{\text{obs}} a c + 108c^2). \quad (5.59)$$

The local mismatch between a signal at  $(a_0, j_0)$  and a nearby template at  $(a, j)$  is then approximated by

$$m(a_0, j_0; a, j) \approx \gamma_{aa} \Delta a^2 + 2\gamma_{aj} \Delta a \Delta j + \gamma_{jj} \Delta j^2, \quad (5.60)$$

where

$$\Delta a = a - a_0, \quad \Delta j = j - j_0, \quad (5.61)$$

and the metric coefficients are evaluated at the signal location  $(a_0, j_0)$ .

This quadratic form is the object used throughout the template-bank code to estimate how well any point in the target region is covered by the nearest template.

### 5.5.5 Determinant and local template density

The determinant of the projected metric plays a special role because it defines the local proper volume element on the signal manifold,

$$dV_{\text{proper}} = \sqrt{\det \gamma(a, j)} da dj. \quad (5.62)$$

For the present phase model, the determinant simplifies to

$$\det \gamma(a, j) = \frac{\pi^4 T_{\text{obs}}^{10} f^4}{2700 c^2 D(a, j)}. \quad (5.63)$$

Hence

$$\sqrt{\det \gamma(a, j)} = \frac{\pi^2 T_{\text{obs}}^5 f^2}{\sqrt{2700} c \sqrt{D(a, j)}}. \quad (5.64)$$

This quantity determines the required template density. Where  $\sqrt{\det \gamma}$  is large, the signal manifold is locally stretched and neighbouring templates must be closer together in coordinate space to maintain the same mismatch. Where  $\sqrt{\det \gamma}$  is smaller, templates may be placed more sparsely. In this sense, the determinant is the bridge between the phase model and the actual spatial density of templates.

### 5.5.6 Proper volume of the target region

The total proper volume of the chosen search box is

$$V = \int_{-a_{\text{max}}}^{a_{\text{max}}} \int_{-j_{\text{max}}}^{j_{\text{max}}} \sqrt{\det \gamma(a, j)} da dj. \quad (5.65)$$

In general this integral is not particularly convenient to evaluate analytically, so in the implementation it is estimated by Monte Carlo integration. Random points are drawn uniformly over the target box,  $\sqrt{\det \gamma}$  is evaluated at those points, and the mean value is multiplied by the coordinate-space area of the box,

$$V \approx \left\langle \sqrt{\det \gamma} \right\rangle_{\text{MC}} (2a_{\text{max}})(2j_{\text{max}}). \quad (5.66)$$

This step is crucial because the total proper volume tells us, in a metric sense, how large the target region really is. Two rectangular boxes with the same coordinate-space dimensions can require very different numbers of templates if the metric density differs substantially across them.

### 5.5.7 From proper volume to template count

The next step is to decide how many templates are required [116, 118]. In the code, this is done by specifying two quantities:

1. a target mismatch  $\mu$ , which is the maximum tolerated nearest-template mismatch;
2. a desired coverage fraction  $C$ , which is the fraction of the target region that should lie within mismatch  $\mu$  of at least one template.

For a stochastic covering in  $n$  dimensions, the volume of the mismatch ball is proportional to  $\mu^{n/2}V_n$ , where  $V_n$  is the volume of the unit  $n$ -ball. Since the present bank is two-dimensional,

$$n = 2, \quad V_2 = \pi. \quad (5.67)$$

The code therefore uses the standard stochastic-covering estimate

$$N_{\text{temp}} \approx \frac{\ln(1 - C)}{\ln(1 - \mu \pi / V)}, \quad (5.68)$$

rounded to the nearest integer.

This formula should not be read as an exact theorem for the final realised bank; rather, it provides a practical estimate of how many random templates are needed so that the bank has the desired statistical coverage of the target region. In practice, the realised coverage is always checked empirically afterwards.

### 5.5.8 Sampling the bank from the metric density

Having determined the target number of templates, the bank must now be populated. The key design choice here is that templates are not placed uniformly in coordinate space. Instead, they are sampled with probability density

$$p(a, j) = \frac{\sqrt{\det \gamma(a, j)}}{V}, \quad (5.69)$$

within the target region.

This choice ensures that regions requiring a higher local template density receive more samples automatically. The code implements this through Markov-chain Monte Carlo. Two possible samplers are provided:

1. an independence Metropolis–Hastings sampler for comparatively small banks;

2. an `emcee` ensemble sampler for larger banks.

In both cases, the target density is Eq. 5.69. The independence Metropolis–Hastings implementation works as follows:

1. propose a candidate point  $(a', j')$  uniformly inside a generation box;
2. evaluate  $\sqrt{\det \gamma(a', j')}$  at the proposal and at the current point;
3. accept the proposal with probability

$$\alpha = \min \left( 1, \frac{\sqrt{\det \gamma(a', j')}}{\sqrt{\det \gamma(a, j)}} \right). \quad (5.70)$$

Because the proposal distribution is uniform and independent of the current location, it cancels in the acceptance ratio.

The code does not sample directly inside the final target box, but instead inside a slightly larger padded box,

$$|a| \leq (1 + \epsilon)a_{\max}, \quad |j| \leq (1 + \epsilon)j_{\max}, \quad (5.71)$$

with padding fraction  $\epsilon$ . This is an edge-protection strategy: without padding, stochastic banks often become under-dense near the boundaries simply because no templates are permitted beyond the edge. By sampling in a slightly enlarged region and then trimming back to the target box, the interior coverage becomes more uniform.

### 5.5.9 Saving and interpreting the bank

Once the Markov-chain sampler has produced the required number of accepted points, the bank is stored as a list of template coordinates,

$$\{(a_k, j_k)\}_{k=1}^{N_{\text{temp}}}. \quad (5.72)$$

These are the trial acceleration and jerk values supplied to the search pipeline. Each point corresponds to one resampling template. In the search stage, the time series is resampled using that template, Fourier transformed, and searched for periodic power.

Thus the template bank is the direct computational realisation of the metric geometry derived above.

### 5.5.10 Empirical validation of the bank

A metric-based bank is only useful if it actually provides the intended coverage. For that reason, the code includes an explicit validation step.

A set of random test points  $(a_0, j_0)$  is drawn uniformly from the target region. For each test point, the mismatch to every template in the bank is computed using the quadratic form of Eq. 5.60,

$$m_k(a_0, j_0) = \gamma_{aa} (a_k - a_0)^2 + 2\gamma_{aj}(a_k - a_0)(j_k - j_0) + \gamma_{jj}(j_k - j_0)^2. \quad (5.73)$$

The nearest-template mismatch is then

$$m_{\min}(a_0, j_0) = \min_k m_k(a_0, j_0). \quad (5.74)$$

The empirical coverage is estimated as the fraction of test points satisfying

$$m_{\min}(a_0, j_0) \leq \mu. \quad (5.75)$$

This validation serves two purposes. First, it checks whether the realised bank coverage is close to the requested coverage used in Eq. 5.68. Second, it reveals whether particular parts of the parameter space are systematically under-covered. The code therefore examines not only the global coverage fraction, but also the full distribution of nearest-template mismatches across random test points and across a coarse  $(a, j)$  map.

### 5.5.11 Relation to exact stochastic banks and the role of the metric approximation

It is worth noting an important methodological distinction. In Knispel's Einstein@Home binary search, the final stochastic bank was based on the *exact* mismatch rather than only the quadratic metric approximation, precisely because the orbital parameter-space geometry can become non-Euclidean and the metric approximation may overestimate the mismatch in some regions [? ]. Knispel explicitly shows that isomismatch surfaces can deviate from simple ellipses, making purely lattice-based constructions problematic in certain coordinates.

In the present work, however, the situation is different in two important respects. First, the phase model is deliberately restricted to the polynomial  $(a, j)$  description rather than a full Keplerian orbital manifold. Second, the metric is not used blindly: the bank is always validated empirically through the nearest-template

mismatch tests described above. The adopted bank is therefore best described as a *metric-weighted stochastic template bank*: the metric supplies the local density and mismatch estimate, while the final bank quality is judged by direct coverage tests.

### 5.5.12 Summary of the construction

The complete template-bank methodology used in `PULSEJET_beta` may therefore be summarised as follows:

1. choose physically motivated search bounds  $(a_{\max}, j_{\max})$  from the class of binaries to which the search should remain sensitive;
2. adopt the polynomial phase model in Eq. 5.39;
3. compute the full metric on  $(f, a, j)$  using Eq. 5.43;
4. project out the frequency dimension using Eq. 5.55;
5. evaluate the determinant of the projected metric to obtain the local density function;
6. estimate the proper metric volume of the chosen target region;
7. infer the required number of templates from the target mismatch and desired coverage;
8. populate the target region stochastically with density proportional to  $\sqrt{\det \gamma}$ ;
9. validate the realised bank by explicitly computing nearest-template mismatches over random test points.

This procedure provides the mathematical and computational foundation for the acceleration-jerk template banks used for the our Jerk search pipeline.

## 5.6 Iterative GPU Resampling over the $(a_{\parallel}, j_{\parallel})$ Template Bank

Once the physically relevant search region has been identified and populated by a template bank (Section 5.5), the core operation of the pipeline is straightforward in principle: for each template point

$$\mathcal{T} = \{(a_k, j_k)\}_{k=1}^{N_{\text{temp}}}, \quad (5.76)$$

the input time series is resampled according to the corresponding trial acceleration and jerk, and the corrected series is then searched in the Fourier domain. In practice, however, this step dominates the computational cost of the entire pipeline, because it must be repeated for thousands of templates.

The central design choice in `PULSEJET_beta` is that these templates are *not* generated by stepping independently in acceleration and jerk on a fixed Cartesian grid. Instead, they are drawn from the metric-informed template bank developed in Section 5.5. This distinction is important, because the phase model induces a non-trivial correlation between acceleration and jerk. A naive fixed-step strategy ignores that correlation and therefore tends to oversample some parts of the parameter space while undersampling others. The metric-based bank, by contrast, uses the local mismatch geometry itself to determine where templates are needed.

### 5.6.1 Why a fixed-step grid in $a$ and $j$ is inefficient

A conceptually simple implementation of a jerk search would proceed by selecting a uniform step size  $\Delta a$  in acceleration and a uniform step size  $\Delta j$  in jerk, and then evaluating every pair in the Cartesian product

$$a \in \{-a_{\max}, -a_{\max} + \Delta a, \dots, a_{\max}\}, \quad j \in \{-j_{\max}, -j_{\max} + \Delta j, \dots, j_{\max}\}. \quad (5.77)$$

The resulting number of trials scales approximately as

$$N_{\text{grid}} \sim \left( \frac{2a_{\max}}{\Delta a} \right) \left( \frac{2j_{\max}}{\Delta j} \right), \quad (5.78)$$

which already grows rapidly as the target bounds are increased.

More importantly, such a grid assumes that mismatch contours in the chosen physical  $(a, j)$  coordinates are aligned with the coordinate axes and may be covered efficiently with fixed independent step sizes. That is not generally the case for the projected time-domain phase metric used here. As shown in Section 5.5, the local mismatch in the projected  $(a, j)$  subspace is

$$m \approx \gamma_{aa}\Delta a^2 + 2\gamma_{aj}\Delta a\Delta j + \gamma_{jj}\Delta j^2, \quad (5.79)$$

where the off-diagonal term  $\gamma_{aj}$  is generally non-zero. In this coordinate system, acceleration and jerk are therefore correlated from the point of view of phase mismatch: a displacement in  $a$  can, to some extent, be compensated by a displacement in  $j$ , and vice versa. The natural isomismatch contours of the projected time-domain

metric are tilted ellipses rather than axis-aligned rectangles.

This statement is coordinate-dependent. It applies to the physical  $(a, j)$  bank used in the present time-domain resampling implementation after projecting over the signal frequency. It should not be interpreted as a statement about the native coordinates of Fourier-domain acceleration or jerk searches, where the search is naturally formulated in terms of frequency and its derivatives, or equivalently in terms of  $(f, z, w)$ , and can be sampled on a regular grid in those variables. The point here is only that, once the bank is specified in physical acceleration and jerk, the local phase-mismatch geometry need not be aligned with the  $(a, j)$  coordinate axes.

A fixed-step grid does not follow this geometry [170]. In regions where the mismatch ellipses are strongly tilted or elongated, a Cartesian grid wastes templates along directions where the phase changes slowly, while still risking poor coverage along the true principal directions of the metric. The metric-based stochastic bank solves precisely this problem: it populates the  $(a, j)$  plane according to the phase-model geometry and the chosen mismatch tolerance, rather than according to an arbitrary coordinate spacing.

The practical meaning of the bank is therefore simple. Each template is nothing more than one allowed pair

$$(a_k, j_k), \quad (5.80)$$

but the collection of such pairs is chosen so that the target region is covered efficiently at the required mismatch level. The bank may contain thousands of templates, but it is still substantially more efficient than a brute-force dense rectangular sweep over acceleration and jerk.

### 5.6.2 From the continuous resampling model to the discrete index map

The time-domain correction introduced earlier is based on the polynomial phase model

$$\Phi(t) = 2\pi f_0 \left( t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3 \right), \quad (5.81)$$

or equivalently, in corrected-time form,

$$\tau = t + \frac{a_{\parallel}}{2c} t^2 + \frac{j_{\parallel}}{6c} t^3. \quad (5.82)$$

To apply this to a sampled time series, the continuous time coordinate must be converted into a discrete index map.

Let the input series consist of  $N$  uniformly sampled points,

$$t_n = n t_{\text{samp}}, \quad n = 0, \dots, N - 1. \quad (5.83)$$

In the implementation, it is convenient to work with indices directly rather than repeatedly converting back and forth between time and sample number. A centred index coordinate is introduced,

$$n_c = \frac{N}{2}, \quad \Delta n = n - n_c, \quad (5.84)$$

so that the resampling is performed relative to the midpoint of the observation. This centring is useful numerically because it makes the correction symmetric about the observation midpoint and reduces the magnitude of large polynomial excursions at one edge of the series.

For each template  $(a_k, j_k)$ , the code constructs two coefficients,

$$\alpha_k = \frac{t_{\text{samp}} a_k}{2c}, \quad \beta_k = \frac{t_{\text{samp}}^2 j_k}{6c}, \quad (5.85)$$

which encode the acceleration and jerk contributions in discrete index units. The actual index mapping used by the resampler is

$$n_{\text{out}} = n + n \alpha_k \Delta n + \beta_k (\Delta n)^3. \quad (5.86)$$

In the GPU kernel, this mapping is evaluated for each input index and rounded to the nearest integer output location.

Equation 5.86 is the discrete implementation of the polynomial time-domain correction. The first correction term is the acceleration contribution, which depends linearly on the offset from the observation centre. The second is the jerk contribution, which introduces a cubic dependence in the centred coordinate. This cubic term is the essential new ingredient relative to an acceleration-only resampler.

### 5.6.3 Nearest-neighbour resampling in the implemented kernel

The core device function in the implementation is

```
__device__ unsigned long getAcceleratedJerkedIndex(
    double accel_factor,
    double jerk_factor,
```

```
double size,
unsigned long idx)
```

which evaluates the discrete mapping in Eq. 5.86. In the code, the midpoint

$$n_c = \frac{N}{2} \quad (5.87)$$

is first computed, then the centred coordinate

$$\Delta n = n - n_c \quad (5.88)$$

is formed, and finally the mapped index is returned as the nearest integer,

$$n_{\text{out}} = \text{round} \left( n + n \alpha_k \Delta n + \beta_k (\Delta n)^3 \right). \quad (5.89)$$

The global kernel then applies this map to every sample:

```
__global__ void resample_acc_jerk_kernel(float* input_d,
                                         float* output_d,
                                         double accel_factor,
                                         double jerk_factor,
                                         size_t size)
```

so that, for each index  $n$ ,

$$s_k^{\text{rs}}[n] = \begin{cases} s[n_{\text{out}}], & 0 \leq n_{\text{out}} < N, \\ 0, & \text{otherwise.} \end{cases} \quad (5.90)$$

This is therefore a nearest-neighbour resampler: the output sample is assigned the value of the closest input sample identified by the mapping in Eq. 5.89. If the mapped index falls outside the array bounds, the output is set to zero.

The use of nearest-neighbour assignment here is deliberate. It is inexpensive, simple to parallelise, and avoids the additional arithmetic and memory traffic required by higher-order interpolation. Since the resampling must be repeated for every template in the bank, keeping the per-sample kernel lightweight is important for overall throughput. The price is that the method is formally less accurate than linear or higher-order interpolation, but the gain in computational efficiency is substantial and is well matched to a large-scale iterative search.

### 5.6.4 The iterative structure of the pipeline

With the template bank and resampler in place, the main search loop is conceptually

$$(a_1, j_1) \rightarrow \text{resample} \rightarrow \text{FFT} \rightarrow \text{candidate search}, \quad (5.91)$$

repeated for every template in the bank,

$$k = 1, \dots, N_{\text{temp}}. \quad (5.92)$$

In other words, the pipeline does not search one acceleration at a time and then sweep separately over jerk with a fixed step size. Instead, it iterates over the template bank itself:

$$\mathcal{T} = \{(a_k, j_k)\}_{k=1}^{N_{\text{temp}}}, \quad (5.93)$$

and each element of that bank represents one trial polynomial phase correction.

Operationally, the process is:

1. select the next template pair  $(a_k, j_k)$  from the bank;
2. convert it to the discrete coefficients  $\alpha_k$  and  $\beta_k$  using Eq. 5.85;
3. launch the GPU resampling kernel to construct the corrected time series  $s_k^{\text{rs}}$ ;
4. perform the FFT and subsequent periodicity search on the corrected series;
5. record any significant detections;
6. move to the next template.

This must be repeated thousands of times. The search is therefore fundamentally iterative even though each individual resampling kernel is highly parallel. It is the repeated application of the resampler over a large template bank that makes the computational cost of jerk searches so much larger than that of acceleration-only searches.

### 5.6.5 Why the jerk term changes the computational character of the resampling

In an acceleration-only resampler, the index correction is dominated by the lower-order term,

$$n_{\text{out}} \approx n + n \alpha_k \Delta n, \quad (5.94)$$

which already produces a non-trivial remapping of the time series. Once jerk is added, however, the new term

$$\beta_k(\Delta n)^3 \tag{5.95}$$

introduces a much stronger non-linearity in the index map. This has two immediate consequences.

First, the number of templates increases because the search must now cover a two-dimensional parameter space rather than a one-dimensional one. Second, even for a fixed template, the resampling map itself becomes more complicated. The cubic term causes the input-to-output index relation to vary more rapidly across the observation, especially away from the midpoint. This is precisely the behaviour illustrated by the index-shift 5.1 and index-increment 5.2 figures discussed in the previous section: increasing jerk not only changes the total amount of resampling, but also makes the local stepping of the map increasingly irregular.

Thus the computational difficulty of the jerk search is *not* merely that there are more templates. It is also that the resampling kernel for each template becomes more non-linear in its memory access pattern as jerk increases.

### 5.6.6 Why GPUs are the natural implementation target

This structure makes the search particularly well suited to GPU acceleration. For a fixed template  $(a_k, j_k)$ , the computation of the mapped index  $n_{\text{out}}$  at sample  $n$  is independent of every other sample. The resampling stage is therefore an embarrassingly parallel operation:

$$n \mapsto n_{\text{out}}(n; a_k, j_k) \tag{5.96}$$

may be evaluated independently for all  $n = 0, \dots, N - 1$ .

The GPU kernel exploits exactly this property. Each thread evaluates the index map for one or more samples, reads the corresponding input location, and writes the resampled value into the output array. Once one template has been processed, the same procedure is repeated for the next. The total search therefore becomes a sequence of massively parallel resampling operations, each followed by the corresponding FFT-based search stage.

This is the computational justification for implementing the polynomial template-bank jerk search on GPUs. The bank itself is chosen to minimise the number of trial  $(a, j)$  pairs required at fixed mismatch, while the GPU implementation minimises the cost of applying each of those trials to the data.

## 5.7 Current Status of the Pipeline and Preliminary Template-Bank Performance

At the time of writing, `PULSEJET_beta` should be regarded as a functional but still evolving search pipeline. The core GPU-based resampling framework is in place and the polynomial acceleration–jerk resampler described in Section 5.6 is fully operational. In particular, the pipeline is already capable of iterating over a bank of trial  $(a_{\parallel}, j_{\parallel})$  templates, applying the corresponding time-domain correction, and passing the resampled series to the subsequent Fourier-domain search stages.

The part of the pipeline that remains under active development is the template-bank construction itself. The present implementation uses a stochastic bank generated from the projected phase metric, with the bank populated through a Monte Carlo / Markov-chain procedure designed to follow the local metric density. This approach is theoretically well motivated and, in broad terms, does succeed in populating the target region of the  $(a_{\parallel}, j_{\parallel})$  plane. However, in the current beta version, its realised coverage is not yet as uniform or as robust as desired for production science searches.

Figure 5.3 illustrates this issue using a coarse coverage map of the  $(a, j)$  plane. At each grid point, the quantity shown is the mismatch to the nearest template in the current bank, computed using the projected metric derived in Section 5.5. In an ideal bank, one would expect this map to remain relatively smooth and to stay below the chosen mismatch threshold across the full target region, apart from modest stochastic fluctuations. Instead, the present implementation shows a patchy structure, with numerous localised regions of elevated mismatch. These appear as bright yellow or green patches in the map and indicate locations where the nearest available template is appreciably farther away, in the metric sense, than intended.

This behaviour is not merely aesthetically undesirable; it has direct consequences for search sensitivity. A locally under-covered region implies that a true signal falling in that part of parameter space would be matched less accurately by the bank, leading to a larger phase mismatch and therefore a larger loss in recovered signal-to-noise ratio. In principle, some degree of stochastic irregularity is expected in any finite random bank. In practice, however, the present level of patchiness is not ideal, especially because the most prominent under-covered regions often appear near the boundaries of the search box.

Those boundary regions are scientifically important. They correspond to the largest trial values of  $|a_{\parallel}|$  and  $|j_{\parallel}|$ , i.e. the most extreme ends of the search space. These are precisely the regions where one expects the signatures of the most compact

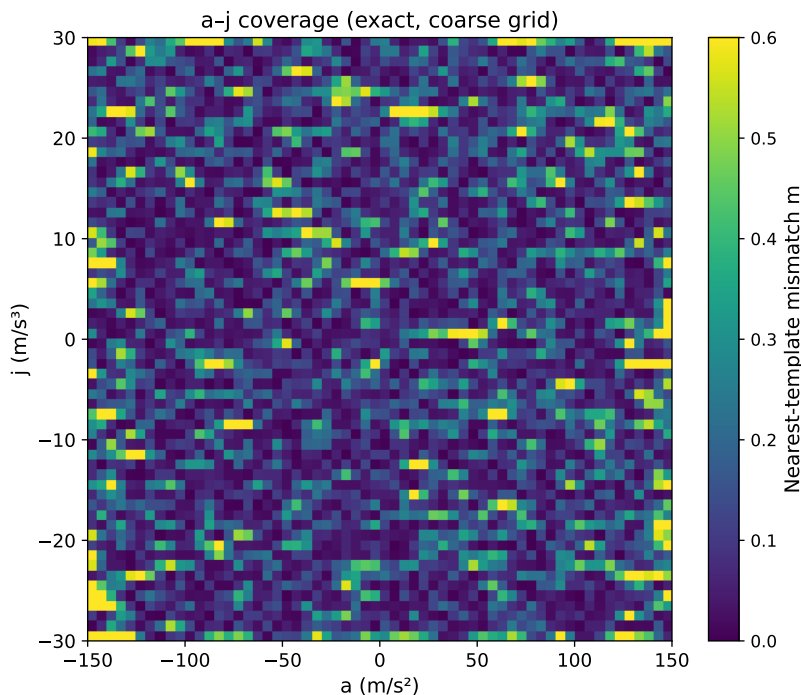


Figure 5.3: Coarse mismatch map of the current stochastic  $(a, j)$  template bank. The colour at each point shows the mismatch to the nearest template, evaluated using the projected metric on the acceleration–jerk manifold. Darker regions indicate better local coverage, while brighter patches correspond to under-covered regions where the nearest-template mismatch is higher. The patchy structure, particularly toward the boundaries of the search box, shows that the present Monte Carlo / MCMC bank generator does not yet provide sufficiently uniform coverage for a production-quality search.

and dynamically extreme binaries to lie. If the bank under-performs there, then the sensitivity loss is concentrated in exactly the part of parameter space that motivated the development of a jerk-capable time-domain search in the first place.

The present results therefore suggest a clear distinction between the status of the different pipeline components. The GPU resampling machinery itself is already functional and capable of carrying out the large number of iterative trial corrections required by the search. The current limitation lies instead in the quality of the stochastic template bank used to drive those corrections. In other words, the computational engine is ready, but the geometric organisation of the search space still requires improvement.

At present, the metric-informed Monte Carlo / MCMC bank generator provides a useful proof of concept and demonstrates that the  $(a, j)$  space can be populated in a physically motivated way. Nevertheless, the realised coverage remains spatially non-uniform, and the edge behaviour is particularly problematic. This indicates

that further work is required on the bank-generation strategy before the pipeline can be considered fully mature. Possible directions include improved edge treatments, adaptive oversampling in boundary regions, hybrid stochastic–deterministic placement methods, or more sophisticated bank-relaxation schemes that explicitly suppress local holes after the initial stochastic draw.

## 5.8 Injection Tests and Preliminary Comparison with Existing Jerk-Search Implementations

To evaluate the present beta version of `PULSEJET_beta`, a series of controlled injection tests were carried out using the `Inception` pulsar injection software<sup>3</sup>. The purpose of these tests was twofold. First, they were used to verify that the polynomial acceleration–jerk resampler and the GPU search loop behave as expected when the correct trial parameters are supplied. Second, they were used to assess the practical recovery performance of the current template-bank implementation in progressively more realistic blind-search scenarios.

In addition to these injection tests, a preliminary computational comparison was also made against `PrestoZL`<sup>4</sup>, a GPU implementation of the `PRESTO` jerk-search formalism. At the present stage of development, this comparison is intended primarily as an exploratory benchmark of computational behaviour rather than as a definitive one-to-one sensitivity comparison, since the two pipelines are based on different search strategies and the `PULSEJET_beta` template-bank construction is still under active refinement.

### 5.8.1 Validation of the resampler with exact trial parameters

The first and most direct validation step was to inject synthetic pulsars with known acceleration and jerk values and then run the resampler using exactly those same values as the trial parameters. In this idealised setting, the template mismatch is, by construction, zero at the signal location, and any residual phase drift after resampling should therefore vanish apart from numerical effects.

In these tests, all injected pulsars were recovered successfully. After applying the exact acceleration–jerk correction, the recovered signals showed no residual acceleration or jerk signature and could be folded coherently. This is an important result

---

<sup>3</sup><https://github.com/erc-compact/inception.git>

<sup>4</sup><https://github.com/zhejianglab/PrestoZL.git>

because it isolates the core resampling machinery from the template-bank problem. It shows that the polynomial jerk resampler itself, together with the GPU execution path and subsequent search stages, is functioning as intended. In other words, the pipeline is able to undo the injected orbital modulation correctly when supplied with the correct trial parameters.

This test therefore establishes the baseline correctness of the resampling engine. Any subsequent loss of performance in blind searches must then be attributed not to a failure of the underlying resampler, but to the finite coverage and quality of the template bank used to drive the search.

### 5.8.2 Illustrative comparison with an acceleration-only search

A second test was designed to examine the practical utility of introducing jerk into the time-domain correction. For this purpose, pulsars with relatively high acceleration and jerk were injected, and the resulting data were searched both with the standard acceleration-only `Peasoup` pipeline and with `PULSEJET_beta`.

Figure 5.4 shows one such example. Both pipelines recover the same injected pulsar, but the quality of the recovery differs substantially. The `PULSEJET_beta` fold shows a much clearer candidate, with a significantly stronger diagnostic response and a more tightly localised parameter-space peak. By contrast, the acceleration-only search still detects the pulsar, but with a visibly degraded candidate response. This behaviour is fully consistent with expectation: when the injected signal contains non-negligible jerk, an acceleration-only correction can at best remove only the lower-order component of the phase drift, leaving residual curvature in the signal. The jerk-enabled resampler is able to remove that additional structure and therefore recover the signal more coherently.

This comparison should not yet be interpreted as a population-level sensitivity statement, since only individual controlled cases have been examined at this stage. It does, however, provide direct qualitative evidence that the inclusion of jerk in the time-domain resampling can improve recovery for signals lying in the high-acceleration, high-jerk regime.

### 5.8.3 Blind recovery tests with small sets of injected pulsars

The next step was to move from exact-template validation to genuinely blind recovery tests. In the first such experiment, ten pulsars were injected with randomly selected acceleration and jerk values inside a chosen target range, and the pipeline was then run without prior knowledge of the true signal parameters. In this mode,

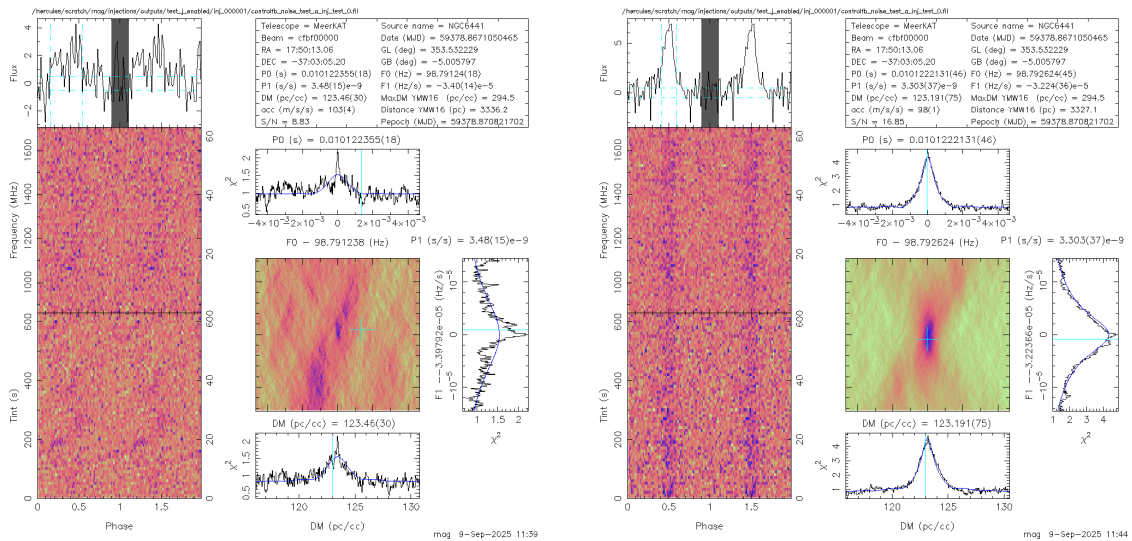


Figure 5.4: Illustrative recovery of the same injected pulsar using an acceleration-only search (left) and the jerk-enabled PULSEJET\_beta search (right). The PULSEJET\_beta candidate shows a substantially cleaner and stronger recovery, consistent with the expectation that a signal containing significant jerk is not optimally corrected by an acceleration-only search.

recovery depends critically on whether the stochastic template bank provides a sufficiently close template to each injected source.

With the current bank implementation, seven out of ten injected pulsars were recovered. Inspection of the missed signals showed that two of the failures corresponded to pulsars placed close to the boundaries of the  $(a, j)$  search region, where the nearest-template mismatch approached values of order  $m \sim 0.6$ . This behaviour is consistent with the coarse coverage map shown earlier in Fig. 5.3, which already indicated that the current bank generator tends to leave locally under-covered regions, especially near the edges of the search box.

To test whether the problem was indeed driven by edge coverage rather than by the resampling machinery itself, the same experiment was repeated after enlarging the template bank so that the injected pulsars lay more centrally within the populated region rather than near its boundaries. In that configuration, nine out of ten pulsars were recovered 5.5. This result strongly suggests that the main limitation in the present blind-search performance is not the GPU jerk-search kernel, but the realised quality of the stochastic bank, particularly its behaviour near the boundary of the search domain.

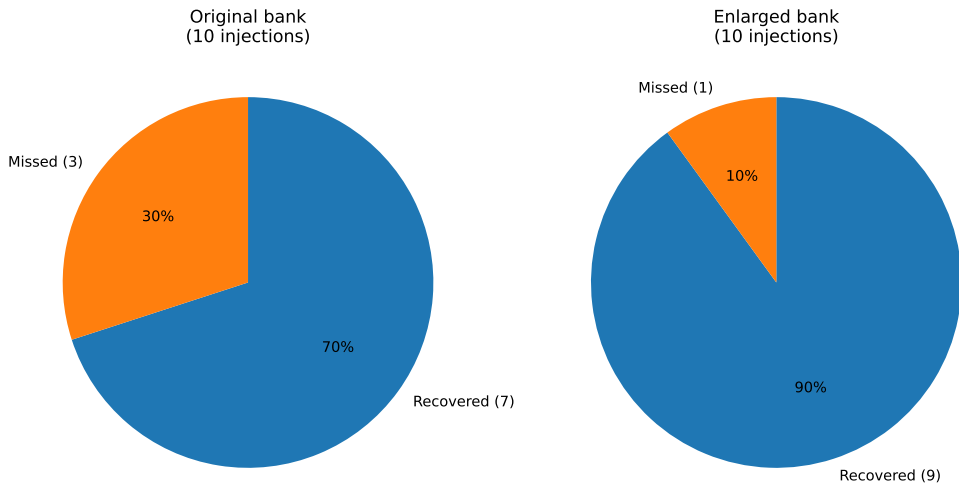


Figure 5.5: Recovery fractions for the 10-pulsar blind injection tests using the current stochastic template-bank implementation. The left panel shows the result obtained with the original template bank, for which 7 out of 10 injected pulsars were recovered. The right panel shows the result after enlarging the template bank so that the injected pulsars lay more centrally within the populated  $(a, j)$  region, yielding 9 out of 10 recoveries. The improved performance in the enlarged-bank test supports the conclusion that the dominant limitation of the present implementation is not the jerk-resampling kernel itself, but incomplete template-bank coverage, especially near the boundaries of the acceleration–jerk search space.

#### 5.8.4 Larger injection campaigns

The same general conclusion emerged from larger injection campaigns containing 30 and 100 pulsars. In these tests, the injected pulsars were drawn randomly over a range of astrophysical and observational parameters, including spin period, dispersion measure, acceleration, jerk, duty cycle, and signal-to-noise ratio. Before interpreting any non-recoveries as search failures, every injected pulsar was checked through folding to confirm that the injection itself was valid and present in the data.

For these larger tests, the current version of the pipeline recovered roughly 60–70% of the injected population. This level of performance is encouraging for a beta-stage implementation, since it demonstrates that the end-to-end search architecture is already able to recover a substantial fraction of the target signals in genuinely blind conditions. At the same time, the recovery fraction is not yet at the level required for a mature production search. The outcome is again consistent with the picture developed in the preceding sections: the dominant limitation at present is not the correctness of the jerk-resampling formalism, but the incompleteness and spatial non-uniformity of the current template bank.

During these tests, different stochastic bank-generation schemes were also ex-

plored, including an `emcee`-based approach in addition to the original Monte Carlo / MCMC random bank generator. At present, neither variant has yet delivered the level of boundary robustness and uniform practical coverage required for a final science-ready implementation. These experiments therefore define the current state of the code: the pipeline is operational and capable of recovering injected signals, but the template-bank generation remains the principal component still under active optimisation.

### 5.8.5 Template-bank trials versus a naive acceleration–jerk scan

In addition to the blind injection tests, a separate study was carried out to compare the number of search trials implied by the template-bank strategy against a naive approach based on stepping independently through acceleration and jerk with fixed step sizes. This comparison was not performed as a search on data, but as a direct accounting exercise designed to measure the reduction in the number of trial parameter combinations.

The motivation for this comparison is straightforward. A naive jerk search would evaluate every pair in a rectangular grid of trial accelerations and jerks. Such a procedure ignores the correlation between  $a$  and  $j$  encoded by the phase metric and therefore requires many more evaluations than are actually needed at a fixed mismatch tolerance. The metric-informed template bank, by contrast, seeks to populate only those regions and directions in parameter space that are required by the phase geometry itself.

Figure 5.6 shows the result of this comparison for several observation lengths. In all tested cases, the template-bank approach reduces the total number of required trials by a factor of approximately 3.4 relative to the naive scheme. This is an important validation of the metric-based construction described earlier in this chapter. Even though the present bank is not yet ideal in terms of local coverage quality, it already provides a substantial reduction in the global number of trial resamplings that must be performed.

This reduction is not merely a technical curiosity. Since each trial requires a resampling operation followed by subsequent search stages, a factor-of-few reduction in the number of templates translates directly into a factor-of-few reduction in the total computational burden of the pipeline.




Figure 5.6: Comparison between the number of trials required by a naive rectangular scan in acceleration and jerk and by the metric-informed template-bank approach, shown for several observation lengths. The right-hand panel shows the corresponding reduction factor. Even in its present beta form, the template-bank method reduces the total number of required trials by a factor of about 3.4.

### 5.8.6 Preliminary runtime comparison with PrestoZL

A further exploratory comparison was made against PrestoZL [171], which implements a GPU jerk-search approach within the PRESTO framework. Here the comparison was restricted primarily to computational behaviour, specifically the per-run jerk-search runtime and the number of candidates produced, rather than to a definitive sensitivity benchmark.

Figure 5.7 shows the outcome of one such comparison. In the tested configuration, PULSEJET\_beta completed the jerk-search run more quickly than PrestoZL, whereas the number of candidates reported by PULSEJET\_beta was substantially larger.

At present, this result should be interpreted cautiously. The two pipelines do not sample the search space in the same way. As discussed earlier with reference to Fig. 5.0, the Fourier-domain jerk-search formalism and the time-domain template-bank formalism cover the high-acceleration, high-jerk region differently. In the extreme part of the parameter space used for this test, it is plausible that PrestoZL effectively samples a smaller accessible region, while PULSEJET\_beta, being based on a time-domain search band, evaluates a larger set of physically allowed trial corrections. If so, a larger candidate count in PULSEJET\_beta may reflect broader parameter-space coverage rather than simply noisier behaviour. At the same time, it is also possible that the current candidate-collapsing and sifting stages in PULSEJET\_beta are not yet sufficiently mature, and that this contributes to the elevated candidate count. Further work on candidate consolidation is therefore clearly warranted, but it is not yet a priority compared with stabilising the template-bank




Figure 5.7: Preliminary comparison between `PrestoZL` and `PULSEJET_beta` in terms of per-run jerk-search runtime and total candidate count. In the tested configuration, `PULSEJET_beta` completed the run more quickly but produced many more candidates. At the present stage this should be interpreted only as an exploratory computational comparison, since the two pipelines probe the parameter space differently and `PULSEJET_beta` is still under active development.

performance.

### 5.8.7 Interpretation of the present results

Taken together, these tests define the present state of the project clearly.

First, the polynomial jerk-resampling GPU pipeline itself is working as intended. When supplied with the correct trial parameters, it fully removes the injected acceleration and jerk and recovers the signals coherently. Second, in targeted examples, the jerk-enabled time-domain search can recover signals more cleanly than an acceleration-only search, which is consistent with the intended scientific motivation of the pipeline. Third, in blind searches, the current stochastic template bank is already capable of recovering a substantial fraction of injected pulsars, but still fails preferentially in under-covered regions, especially near the search boundaries. Fourth, even in beta form, the metric-informed bank reduces the total number of required acceleration-jerk trials substantially relative to a naive rectangular scan. Finally, preliminary computational tests suggest that the time-domain jerk-search implementation can be competitive in runtime with an existing GPU frequency-domain jerk-search code, although a rigorous comparison of runtime and sensitivity must await a more mature and better validated version of the pipeline.

For that reason, the present section should be read as a report on preliminary performance rather than as a final statement on the search capabilities of `PULSEJET_beta`. A fair one-to-one comparison between pipelines will ultimately require controlled blind-search tests on the same data products, with matched compu-

tational budgets and well-defined injection populations. Only then will it be possible to compare both runtime and sensitivity in a quantitatively meaningful way.

# Chapter 6

## Conclusions and Future Work

### 6.1 Conclusions

This thesis has addressed two closely connected aspects of modern pulsar astronomy: the discovery and characterisation of pulsars in dense globular-cluster environments, and the development of search methods tailored to compact binary systems whose orbital dynamics challenge standard acceleration-search techniques.

On the observational side, the first MeerKAT S-band globular-cluster survey has demonstrated the power of high-frequency tied-array observations for probing highly dispersed and heavily scattered lines of sight. The survey resulted in the re-detection of previously known pulsars and the discovery of four new binary pulsars in Glimpse-C01, thereby establishing this cluster as one of the most interesting emerging laboratories for pulsar studies in the Galactic globular-cluster population. The survey also showed that MeerKAT S-band observations provide a complementary discovery space to lower-frequency searches, particularly in clusters where scattering and high foreground dispersion measures limit sensitivity at L-band and below.

The Glimpse-C01 discoveries are especially noteworthy. J1848–0129C exhibits a wide orbit together with reproducible orbital-phase-dependent non-detections, suggesting eclipse or absorption by extended intrabinary material. J1848–0129D has emerged as a particularly promising mildly recycled binary system with a massive companion, most naturally interpreted as a CO or ONeMg white dwarf. Together, these systems show that high-frequency globular-cluster surveys can reveal binary pulsars with unusual architectures and potentially important evolutionary histories.

On the methodological side, this thesis has developed the first version of `PULSEJET_beta`, a GPU-based time-domain jerk-search pipeline built on a metric-informed template-bank formalism in the  $(a_{\parallel}, j_{\parallel})$  plane. The scientific motivation for this development is clear: while standard Fourier-domain acceleration searches remain highly suc-

cessful and computationally efficient for large surveys, there exists a compact but astrophysically important region of parameter space in which higher-order orbital effects become significant and a coherent time-domain treatment is advantageous. `PULSEJET_beta` was designed specifically to access that regime.

The current implementation demonstrates that the polynomial jerk-resampling formalism works as intended and that the GPU execution path is capable of carrying out iterative template-driven searches efficiently. Injection tests confirm that the resampler is able to remove injected acceleration and jerk exactly when the correct trial parameters are supplied, and preliminary blind-search tests show that the pipeline can already recover a substantial fraction of injected signals. At the same time, these tests also make clear that the dominant limitation of the present beta version lies not in the resampling kernel itself, but in the practical construction of a sufficiently uniform and robust template bank near the edges of the acceleration–jerk search space.

Taken together, the results of this thesis show both the scientific potential of high-frequency globular-cluster pulsar surveys and the need for continued methodological development aimed at the most extreme compact-binary systems. In that sense, the observational and computational threads of this work are strongly complementary: the survey results motivate deeper follow-up of newly discovered binaries, while the pipeline development aims to make future searches more sensitive to precisely the kinds of systems that may otherwise remain hidden.

## 6.2 Future Work

### 6.2.1 Follow-up of the Glimpse-C01 discoveries

A major near-term priority is the continued timing follow-up of the pulsars discovered in Glimpse-C01. Ongoing timing campaigns with MeerKAT, together with awarded observing time at FAST and future follow-up observations, will extend the timing baseline and improve the orbital and spin solutions of these systems.

For J1848–0129D, the immediate goal is continued phase-connected timing. The current orbital solution already suggests a massive companion, with the measured mass function favouring a CO or ONeMg white dwarf over a double neutron star interpretation. A longer timing baseline and improved timing precision will refine the orbital parameters further and may, if the system proves sufficiently favourable, enable meaningful constraints on post-Keplerian effects. Even if such effects remain difficult to measure, the system is already a valuable target for studying the formation of mildly recycled pulsars with massive degenerate companions in dense stellar

environments. Continued monitoring may also help clarify whether Glimpse-C01 hosts an unusually large fraction of mildly recycled systems with relatively massive companions, which would have implications for the evolutionary history of the cluster.

For J1848–0129C, future work is motivated by its unusual orbital-phase-dependent detectability and the possibility that it may represent a long-period spider-like system, perhaps related to the proposed “Huntsman” subclass. The most important next step is to characterise the eclipse or absorption behaviour more precisely through denser orbital-phase coverage. This will allow the phase extent, repeatability, frequency dependence, and possible asymmetry of the non-detection window to be measured more robustly. Such observations are needed both to determine the nature of the companion and to establish whether the system hosts extended intrabinary material. If the eclipse interpretation is confirmed, J1848–0129C may provide a rare opportunity to study plasma outflows and binary interaction physics in an unusually wide system. More broadly, it could become an important test case for understanding the transition between detached binaries, redbacks, and other long-period interacting systems.

J1848–0129E and J1848–0129F also remain important targets. Both are currently limited by sparse detections, and further observations are required before coherent orbital solutions can be attempted. Additional timing detections will determine whether these systems are ordinary faint binary millisecond pulsars or whether they too display unusual orbital behaviour that was not apparent from the present data set.

## 6.2.2 Companion identification and multiwavelength follow-up

A second major direction for future work is the identification and characterisation of the binary companions through multiwavelength observations. Once sufficiently precise timing positions are available, archival and future imaging data from facilities such as the *Hubble Space Telescope* and other optical or near-infrared instruments can be searched for plausible counterpart candidates.

This is particularly important for Glimpse-C01, where the cluster is highly obscured and companion identification is observationally challenging. Nevertheless, the combination of improved pulsar localisation, infrared imaging, and colour–magnitude analysis may make it possible to constrain the evolutionary state of the companions in systems such as J1848–0129C and J1848–0129D. In the case of J1848–0129C, counterpart identification could help distinguish between a non-

degenerate or evolved companion and a detached degenerate scenario. In the case of J1848–0129D, it may provide an independent check on the inferred massive white-dwarf interpretation. More generally, multiwavelength companion studies will help place the Glimpse-C01 pulsar population in the broader context of globular-cluster binary evolution.

### 6.2.3 Further development of PULSEJET\_beta

On the methodological side, the most urgent priority is to improve the template-bank generation algorithm used by PULSEJET\_beta. The current metric-informed Monte Carlo / MCMC bank already demonstrates the feasibility of populating the  $(a_{\parallel}, j_{\parallel})$  space in a physically motivated way, but the realised coverage remains spatially non-uniform and the edge behaviour is not yet satisfactory for production blind searches.

Future work will therefore focus on investigating alternative bank-generation strategies. These include improved edge-padding and boundary-treatment schemes, adaptive oversampling of high-mismatch regions, hybrid stochastic–deterministic placement methods, and alternative sampling algorithms. The goal is to achieve sufficiently uniform coverage and sufficiently low nearest-template mismatch across the full search region, including the most extreme high-acceleration and high-jerk boundaries that are most relevant for compact-binary discovery.

Once the template-bank problem has been stabilised, the next stage will be systematic benchmarking on both injected and real data. On the injection side, this will involve controlled blind-search experiments with known pulsar populations spanning a range of spin periods, dispersion measures, duty cycles, accelerations, jerks, and signal-to-noise ratios. On real data, the goal will be to compare PULSEJET\_beta directly against existing approaches, including acceleration-only *Peasoup* searches and GPU jerk-search implementations such as *PrestoZL*, using matched data sets and controlled computational budgets.

Such comparisons should ultimately address both sensitivity and computational performance. At present, the exploratory tests reported in this thesis indicate that the time-domain jerk-search strategy is viable and can be competitive in runtime, but a definitive comparison must wait until the template-bank construction is mature enough for robust blind searches on real survey data. Once that stage is reached, the natural next step will be a dedicated methods publication describing the final implementation, its validation, and its performance relative to existing search pipelines.

## 6.3 Data Availability and Software Used

The data underlying this work are held by the author and are stored on the HPC facilities of the INAF–Osservatorio Astronomico di Cagliari. They will be made available upon reasonable request.

Table 6.1 summarises the principal software packages and pipelines used in this thesis together with the corresponding repositories or project pages.

Table 6.1: Principal software packages and pipelines used in this thesis, listed in alphabetical order.

Software pipeline	Primary use in this thesis	Repository / project page
DSPSR	Pulsar data processing and folding.	<a href="https://dspsr.sourceforge.net/">https://dspsr.sourceforge.net/</a>
Inception	Synthetic pulsar injection software used for controlled validation and recovery tests.	<a href="https://github.com/erc-compact/inception.git">https://github.com/erc-compact/inception.git</a>
PEASOUP	GPU-accelerated time-domain acceleration-search software used as the parent framework and conceptual basis for PULSEJET_beta.	<a href="https://github.com/ewanbarr/peasoup">https://github.com/ewanbarr/peasoup</a>
PRESTO	Fourier-domain pulsar searching, RFI mitigation, dedispersion, acceleration and jerk searches, candidate sifting, and related diagnostics.	<a href="https://github.com/scottransom/presto">https://github.com/scottransom/presto</a>
PrestoZL	GPU jerk-search implementation used for exploratory computational comparison with PULSEJET_beta.	<a href="https://github.com/zhejianglab/PrestoZL.git">https://github.com/zhejianglab/PrestoZL.git</a>
PULSAR_MINER	Automated PRESTO-based pulsar-search pipeline used for survey processing.	<a href="https://github.com/alex88ridolfi/PULSAR_MINER.git">https://github.com/alex88ridolfi/PULSAR_MINER.git</a>
PULSEJET_beta	GPU-based time-domain jerk-search pipeline developed in this thesis.	<a href="https://github.com/Rouhin1997/PULSEJET_beta.git">https://github.com/Rouhin1997/PULSEJET_beta.git</a>
Python scientific stack	Numerical analysis, plotting, testing, and auxiliary pipeline-development tasks.	<a href="https://www.python.org/">https://www.python.org/</a>

# Appendix A

## Flux Density Estimation from Folded Pulse Profiles

### A.1 Flux Density Estimation from Folded Pulse Profiles

We derive here the expression used to estimate the phase-averaged pulsar flux density from folded pulse profiles.

For an observation with integration time  $t_{\text{obs}}$  and usable bandwidth  $\Delta\nu$ , the rms noise of the dedispersed time series is given by the radiometer equation in terms of the system-equivalent flux density (SEFD):

$$\sigma_{\text{ts}} = \frac{\text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu t_{\text{obs}}}}, \quad (\text{A.1})$$

where  $n_{\text{pol}}$  is the number of summed orthogonal polarisations.

After folding the time series into  $N_{\text{bin}}$  pulse-phase bins, each bin contains an effective integration time of  $t_{\text{obs}}/N_{\text{bin}}$ . The rms noise per phase bin is therefore

$$\sigma_{\text{bin}} = \frac{\text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu (t_{\text{obs}}/N_{\text{bin}})}} = \frac{\text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu t_{\text{obs}}}} \sqrt{N_{\text{bin}}}. \quad (\text{A.2})$$

Let  $N_{\text{on}}$  denote the number of phase bins containing pulsed emission and  $N_{\text{off}} = N_{\text{bin}} - N_{\text{on}}$  the number of off-pulse bins used to estimate the baseline level. The integrated pulse amplitude above the baseline is

$$A = \sum_{i=1}^{N_{\text{on}}} (p_i - b), \quad (\text{A.3})$$

where  $p_i$  are the on-pulse bin values and  $b$  is the mean off-pulse baseline.

The uncertainty in the integrated pulse amplitude arises from both the noise in the on-pulse bins and the uncertainty in the baseline estimate derived from the off-pulse region. Propagating these contributions yields

$$\sigma_A = \sigma_{\text{bin}} \sqrt{N_{\text{on}} \left( 1 + \frac{N_{\text{on}}}{N_{\text{off}}} \right)}. \quad (\text{A.4})$$

The signal-to-noise ratio of the folded profile is therefore

$$\text{SNR} = \frac{A}{\sigma_A}. \quad (\text{A.5})$$

The phase-averaged flux density  $S$  is related to the integrated pulse amplitude through

$$A = S N_{\text{bin}}, \quad (\text{A.6})$$

since  $S$  represents the mean flux density over the full rotational phase.

Substituting for  $A$  and  $\sigma_A$  and rearranging yields

$$S = \frac{\text{SNR} \times \text{SEFD}}{\sqrt{n_{\text{pol}} \Delta\nu t_{\text{obs}}}} \sqrt{\frac{N_{\text{on}}}{N_{\text{off}}}}, \quad (\text{A.7})$$

which is the expression used to compute the flux densities reported in this work.

# Appendix B

## Effective SEFD and Gain Determination

### B.1 Effective SEFD and Gain Determination

This appendix describes the procedure used to determine the effective system-equivalent flux density (SEFD) for the coherently beamformed MeerKAT S-band observations, required to convert folded pulse profile signal-to-noise ratios into phase-averaged flux densities.

In the pulsar radiometer equation, the system sensitivity is expressed through the factor  $(T_{\text{rec}} + T_{\text{sky}})/G$ , which is equivalent to the SEFD. For interferometric arrays operating in tied-array beamforming mode, the effective SEFD depends on both the intrinsic per-antenna sensitivity and the number of antennas contributing coherently to the beam.

Published MeerKAT S-band performance data provide the mean SEFD as a function of frequency for a single antenna across the receiver band. The SEFD curves used in this work were obtained from official SARA O S-band capability documentation. These curves incorporate receiver temperature, sky contribution, and antenna gain, but do not include backend or processing losses. As tabulated values are not available, the curves were digitised from the published performance plots and interpolated across the observing band.

To account for the frequency dependence of the system response across the wide S-band bandwidth, we compute a band-averaged effective SEFD. Since signal-to-noise adds in quadrature across independent frequency channels, the appropriate averaging is inverse-square weighted. The effective per-antenna SEFD is therefore

given by

$$\text{SEFD}_{\text{eff,ant}} = \left( \frac{\sum_i w_i \text{SEFD}(f_i)^{-2}}{\sum_i w_i} \right)^{-1/2}, \quad (\text{B.1})$$

where  $\text{SEFD}(f_i)$  is the per-antenna SEFD at frequency  $f_i$  and  $w_i$  are per-channel weights. Given the low radio-frequency interference occupancy of the S-band data used in this work (typically  $\gtrsim 98\%$  usable bandwidth), uniform weights are adopted.

The effective SEFD for the coherently beamformed tied-array beam is obtained by scaling the per-antenna value by the number of antennas contributing to the coherent sum,

$$\text{SEFD}_{\text{eff,array}} = \frac{\text{SEFD}_{\text{eff,ant}}}{N_{\text{ant}}}. \quad (\text{B.2})$$

For observations using the full MeerKAT array, we adopt  $N_{\text{ant}}$  in the range 58–64. For beams formed using only the inner core of the array (baselines  $\lesssim 1$  km), we adopt  $N_{\text{ant}}$  between 35 and 48. This yields effective tied-array SEFD values of approximately 7.4–8.1 Jy for the full array configuration and 9.8–13.5 Jy for the core configuration.

We note that this estimate reflects only the intrinsic telescope sensitivity and does not include losses associated with digitisation, detection, or backend processing, which are often parameterised by a multiplicative factor  $\beta$  in the pulsar radiometer equation. Because signal-to-noise ratios are measured directly from the folded pulse profiles produced by the full processing chain, and no independent characterisation of backend efficiency is available, we set  $\beta = 1$  and do not apply an additional correction. Any residual processing losses are therefore implicitly absorbed into the measured signal-to-noise ratios. The derived flux densities scale linearly with any alternative choice of  $\beta$ .




Figure B.1: Mean system-equivalent flux density per antenna across the full MeerKAT S-band receiver bandwidth, digitised from SARAO performance documentation.




Figure B.2: Mean system-equivalent flux density per antenna evaluated for individual sub-bands across the observing bandwidth.

# Bibliography

- [1] T. Gold. Rotating neutron stars as the origin of the pulsating radio sources. *Nature*, 218:731–732, 1968. doi: 10.1038/218731a0.
- [2] D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Cambridge University Press, 2005. doi: 10.1017/CBO9780511544849.
- [3] A. G. Lyne and F. Graham-Smith. *Pulsar Astronomy*. Cambridge University Press, 4 edition, 2012. doi: 10.1017/CBO9781139016472.
- [4] G. Hobbs, D. R. Lorimer, A. G. Lyne, and M. Kramer. A statistical study of 233 pulsar proper motions. *Monthly Notices of the Royal Astronomical Society*, 360:974–992, 2005. doi: 10.1111/j.1365-2966.2005.09087.x.
- [5] P. Podsiadlowski, N. Langer, A. J. T. Poelarends, S. Rappaport, A. Heger, and E. Pfahl. The effects of binary evolution on the dynamics of core collapse and neutron star kicks. *The Astrophysical Journal*, 612:1044–1051, 2004. doi: 10.1086/421713.
- [6] T. M. Tauris, N. Langer, T. J. Moriya, P. Podsiadlowski, S.-C. Yoon, and S. I. Blinnikov. Ultra-stripped type ic supernovae from close binary evolution. *Monthly Notices of the Royal Astronomical Society*, 451:2123–2144, 2015. doi: 10.1093/mnras/stv990.
- [7] Ashley R. Martsen, Scott M. Ransom, Megan E. DeCesar, Paulo C. C. Freire, Jason W. T. Hessels, Anna Y. Q. Ho, Ryan S. Lynch, Ingrid H. Stairs, and Yuankun Wang. Radio Pulse Profiles and Polarization of the Terzan 5 Pulsars. , 941(1):22, December 2022. doi: 10.3847/1538-4357/aca156.
- [8] W. Baade and F. Zwicky. On super-novae. *Proceedings of the National Academy of Sciences*, 20:254–259, 1934. doi: 10.1073/pnas.20.5.254.
- [9] S. L. Shapiro and S. A. Teukolsky. *Black Holes, White Dwarfs, and Neutron Stars*. Wiley, 1983.

- [10] P. Haensel, A. Y. Potekhin, and D. G. Yakovlev. *Neutron Stars 1: Equation of State and Structure*. Springer, 2007. doi: 10.1007/978-0-387-47301-7.
- [11] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a rapidly pulsating radio source. *Nature*, 217:709–713, 1968. doi: 10.1038/217709a0.
- [12] R. N. Manchester and J. H. Taylor. *Pulsars*. W. H. Freeman, 1977.
- [13] P. Goldreich and W. H. Julian. Pulsar electrodynamics. *Astrophysical Journal*, 157:869–880, 1969. doi: 10.1086/150119.
- [14] V. S. Beskin. *MHD Flows in Compact Astrophysical Objects: Accretion, Winds and Jets*. Springer, Berlin, 2010.
- [15] F. C. Michel. *Theory of Neutron Star Magnetospheres*. University of Chicago Press, 1991.
- [16] Michael Kramer and Simon Johnston. Radio emission from beyond the light cylinder in millisecond pulsars. , 547(4):staf2258, April 2026. doi: 10.1093/mnras/staf2258.
- [17] V. Radhakrishnan and D. J. Cooke. Magnetic poles and the polarization structure of pulsar radiation. *Astrophysical Letters*, 3:225–229, 1969.
- [18] J. M. Rankin. Toward an empirical theory of pulsar emission. i. morphological taxonomy. *Astrophysical Journal*, 274:333–358, 1983. doi: 10.1086/161450.
- [19] T. M. Tauris and R. N. Manchester. On the evolution of pulsar beams. *Monthly Notices of the Royal Astronomical Society*, 298:625–636, 1998. doi: 10.1046/j.1365-8711.1998.01613.x.
- [20] D. R. et al. Lorimer. The parkes multibeam pulsar survey – vi. discovery and timing of 142 pulsars. *Monthly Notices of the Royal Astronomical Society*, 372:777–800, 2006. doi: 10.1111/j.1365-2966.2006.10887.x.
- [21] S. E. Woosley, A. Heger, and T. A. Weaver. The evolution and explosion of massive stars. *Reviews of Modern Physics*, 74:1015–1071, 2002. doi: 10.1103/RevModPhys.74.1015.
- [22] J. M. Lattimer. The nuclear equation of state and neutron star masses. *Annual Review of Nuclear and Particle Science*, 62:485–515, 2012. doi: 10.1146/annurev-nucl-102711-095018.

- [23] B. M. Gaensler and P. O. Slane. The evolution and structure of pulsar wind nebulae. *Annual Review of Astronomy and Astrophysics*, 44:17–47, 2006. doi: 10.1146/annurev.astro.44.051905.092528.
- [24] A. K. Harding. The neutron star zoo. *Frontiers of Physics*, 8:679–692, 2013. doi: 10.1007/s11467-013-0344-6.
- [25] A. Spitkovsky. Time-dependent force-free pulsar magnetospheres: Axisymmetric and oblique rotators. *Astrophysical Journal Letters*, 648:L51–L54, 2006. doi: 10.1086/507518.
- [26] V. M. Kaspi. Grand unification of neutron stars. *Proceedings of the National Academy of Sciences*, 107:7147–7152, 2010. doi: 10.1073/pnas.0911647107.
- [27] D. Bhattacharya and E. P. J. van den Heuvel. Formation and evolution of binary and millisecond radio pulsars. *Physics Reports*, 203:1–124, 1991. doi: 10.1016/0370-1573(91)90064-S.
- [28] T. M. Tauris et al. Formation and evolution of compact stellar x-ray sources. *Astrophysics and Space Science Library*, 457:1–659, 2017. doi: 10.1007/978-3-319-97616-7.
- [29] J. Frank, A. King, and D. Raine. *Accretion Power in Astrophysics*. Cambridge University Press, 3 edition, 2002.
- [30] P. Ghosh and F. K. Lamb. Accretion by rotating magnetic neutron stars. *Astrophysical Journal*, 234:296–316, 1979. doi: 10.1086/157498.
- [31] M. A. Alpar, A. F. Cheng, M. A. Ruderman, and J. Shaham. A new class of radio pulsars. *Nature*, 300:728–730, 1982. doi: 10.1038/300728a0.
- [32] G. Hobbs et al. The international pulsar timing array project. *Classical and Quantum Gravity*, 27:084013, 2010. doi: 10.1088/0264-9381/27/8/084013.
- [33] R. Blandford and S. A. Teukolsky. Arrival-time analysis for a pulsar in a binary system. *Astrophysical Journal*, 205:580–591, 1976. doi: 10.1086/154315.
- [34] H. M. Johnston and S. R. Kulkarni. A search for binary millisecond pulsars. *Astrophysical Journal*, 368:504–514, 1991. doi: 10.1086/169718.
- [35] S. M. Ransom, S. S. Eikenberry, and J. Middleditch. Fourier techniques for very long astrophysical time-series analysis. *Astronomical Journal*, 124:1788–1809, 2003. doi: 10.1086/342285.

- [36] E. S. Phinney. Pulsar recycling and evolution of binary pulsars. *Philosophical Transactions of the Royal Society of London A*, 341:39–75, 1992.
- [37] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The australia telescope national facility pulsar catalogue. *Astronomy and Astrophysics*, 433: 807–817, 2005. doi: 10.1051/0004-6361:20041485.
- [38] A. G. Lyne and F. Graham-Smith. *Pulsar Astronomy*. Cambridge University Press, 4 edition, 2012. doi: 10.1017/CBO9781139016472.
- [39] S. A. Olausen and V. M. Kaspi. The mcgill magnetar catalog. *Astrophysical Journal Supplement Series*, 212(1):6, 2014. doi: 10.1088/0067-0049/212/1/6.
- [40] D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Cambridge University Press, 2005. doi: 10.1017/CBO9780511544849.
- [41] D. B. Melrose. Coherent emission in astrophysical plasmas. *Astrophysical Journal*, 225:557–571, 1978. doi: 10.1086/156514.
- [42] D. B. Melrose. Coherent emission in pulsars. *Astrophysics and Space Science*, 291:3–14, 2004. doi: 10.1023/B:ASTR.0000044277.25821.b5.
- [43] B. J. Rickett. Radio propagation through the turbulent interstellar plasma. *Annual Review of Astronomy and Astrophysics*, 28:561–605, 1990. doi: 10.1146/annurev.aa.28.090190.003021.
- [44] D. R. Stinebring, T. V. Smirnova, T. H. Hankins, J. S. Hovis, V. M. Kaspi, J. C. Kempner, E. Myers, and D. J. Nice. Five years of pulsar flux density monitoring: Refractive scintillation and the interstellar medium. *Astrophysical Journal*, 539:300–316, 2000. doi: 10.1086/309201.
- [45] W. Sieber. Pulsar spectra. *Astronomy and Astrophysics*, 28:237–244, 1973.
- [46] O. Maron, J. Kijak, M. Kramer, and R. Wielebinski. Pulsar spectra of radio emission. *Astronomy and Astrophysics Supplement Series*, 147:195–203, 2000. doi: 10.1051/aas:2000218.
- [47] F. Jankowski, W. van Straten, E. F. Keane, M. Bailes, and E. D. Barr. Spectral properties of southern pulsars. *Monthly Notices of the Royal Astronomical Society*, 473:4436–4458, 2018. doi: 10.1093/mnras/stx2476.
- [48] M. Kramer, K. M. Xilouris, A. Jessner, D. R. Lorimer, and R. Wielebinski. The characteristics of millisecond pulsar emission. *Astrophysical Journal*, 501: 270–285, 1998. doi: 10.1086/305790.

- [49] S. D. Bates, D. R. Lorimer, and J. P. W. Verbiest. Spectral properties of radio pulsars. *Monthly Notices of the Royal Astronomical Society*, 431:1352–1358, 2013. doi: 10.1093/mnras/stt258.
- [50] J. Kijak, O. Maron, Y. Gupta, A. Jessner, and M. Kramer. Pulsar radio spectra. *Astronomy and Astrophysics*, 462:699–706, 2007. doi: 10.1051/0004-6361:20065274.
- [51] R. N. Manchester and J. H. Taylor. Polarization observations of pulsars. *Astrophysical Journal*, 201:L151–L155, 1975. doi: 10.1086/181960.
- [52] F. F. Gardner and J. B. Whiteoak. Faraday rotation in pulsar signals. *Nature*, 221:439–441, 1969. doi: 10.1038/221439a0.
- [53] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *Astrophysical Journal Letters*, 195:L51–L53, 1975. doi: 10.1086/181708.
- [54] J. H. Taylor and J. M. Weisberg. A new test of general relativity: Gravitational radiation and the binary pulsar psr b1913+16. *Astrophysical Journal*, 253:908–920, 1982. doi: 10.1086/159690.
- [55] F. Camilo, S. M. Ransom, B. M. Gaensler, D. R. Lorimer, R. N. Manchester, and N. D’Amico. Observations of 20 millisecond pulsars in 47 tucanae at 20 centimeters. *Astrophysical Journal*, 611:L25–L28, 2005. doi: 10.1086/425660.
- [56] P. C. C. Freire, A. Ridolfi, M. Kramer, C. Jordan, R. N. Manchester, and A. Possenti. On the nature and evolution of the pulsars in globular clusters. *Monthly Notices of the Royal Astronomical Society*, 430:3324–3347, 2013. doi: 10.1093/mnras/stt093.
- [57] J. M. Cordes and T. J. W. Lazio. Ne2001.i. a new model for the galactic distribution of free electrons and its fluctuations. *arXiv e-prints*, 2002.
- [58] J. P. W. Verbiest and others. The international pulsar timing array: First data release. *Monthly Notices of the Royal Astronomical Society*, 458:1267–1288, 2016. doi: 10.1093/mnras/stw347.
- [59] William E. Harris. A catalog of parameters for globular clusters in the milky way. <https://physics.mcmaster.ca/~harris/mwgc.dat>, 1996. 2010 edition (updated).
- [60] James Binney and Scott Tremaine. *Galactic Dynamics*. Princeton University Press, 2 edition, 2008.

- [61] D. C. Heggie. Binary evolution in stellar dynamics. *Monthly Notices of the Royal Astronomical Society*, 173:729–787, 1975.
- [62] F. Verbunt and P. Hut. Binary populations in globular clusters. *IAU Symposium*, 125:187–197, 1987.
- [63] D. et al. Pooley. Dynamical formation of close binary systems in globular clusters. *The Astrophysical Journal Letters*, 591:L131–L134, 2003. doi: 10.1086/377074.
- [64] F. Camilo and F. A. Rasio. Pulsars in globular clusters. *Astronomical Society of the Pacific Conference Series*, 328:147, 2005.
- [65] E. S. Phinney. Pulsars as probes of globular cluster dynamics. *ASP Conference Series*, 50:141, 1993.
- [66] Ivan R. King. The structure of star clusters. III. some simple dynamical models. *Astronomical Journal*, 71:64–75, 1966. doi: 10.1086/109857.
- [67] Lyman Spitzer. *Dynamical Evolution of Globular Clusters*. Princeton University Press, 1987.
- [68] Michela Mapelli. Dynamics of stellar systems: Collisional systems. <https://web.oapd.inaf.it/mapelli/2014colldyn1.pdf>, 2014. Lecture notes.
- [69] Douglas Heggie and Piet Hut. *The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics*. Cambridge University Press, 2003. doi: 10.1017/CBO9780511536677.
- [70] N. Ivanova, C. O. Heinke, F. A. Rasio, K. Belczynski, and J. M. Fregeau. Formation and evolution of compact binaries in globular clusters: II. binaries with neutron stars. *Monthly Notices of the Royal Astronomical Society*, 386:553–576, 2008. doi: 10.1111/j.1365-2966.2008.13064.x.
- [71] George W. Clark. X-ray sources in globular clusters. *The Astrophysical Journal Letters*, 199:L143–L145, 1975. doi: 10.1086/181867.
- [72] Fernando Camilo and Frederic A. Rasio. Pulsars in globular clusters. *Binary Radio Pulsars*, *ASP Conference Series*, 328:147, 2005.
- [73] Scott M. Ransom. Millisecond pulsars in globular clusters. *IAU Symposium*, 246:291–300, 2008. doi: 10.1017/S1743921308015817.

- [74] Paulo C. C. Freire. Pulsars in globular clusters. *IAU Symposium*, 291:243–250, 2013. doi: 10.1017/S1743921312023538.
- [75] Mallory S. E. Roberts. Surrounded by spiders! new black widows and redbacks in the galactic field. *Proceedings of the International Astronomical Union*, 291: 127–132, 2013. doi: 10.1017/S174392131202337X.
- [76] Frederic A. Rasio and Douglas C. Heggie. Binary evolution in globular clusters. *The Astrophysical Journal Letters*, 445:L133, 1995. doi: 10.1086/187904.
- [77] Ewan D. Barr, Arunima Dutta, P. C. C. Freire, Mario Cadelano, Tasha Gautam, Michael Kramer, Cristina Pallanca, Scott M. Ransom, Alessandro Riodolfi, Benjamin W. Stappers, Thomas M. Tauris, Vivek Venkatraman Krishnan, Norbert Wex, David J. Champion, Matthew Bailes, Jan Behrend, Sarah Buchner, Marisa Geyer, Weiwei Chen, Prajwal V. Padmanabh, Andrea Possenti, and Y. P. Men. A pulsar in a binary with a compact object in the mass gap between neutron stars and black holes. *Science*, 383:275–279, 2024. doi: 10.1126/science.adg3005.
- [78] Marta Burgay et al. An increased estimate of the merger rate of double neutron stars from observations of a highly relativistic system. *Nature*, 426:531–533, 2003. doi: 10.1038/nature02124.
- [79] A. G. Lyne et al. A double-pulsar system: A rare laboratory for relativistic gravity and plasma physics. *Science*, 303:1153–1157, 2004. doi: 10.1126/science.1094645.
- [80] P. C. Freire, M. Kramer, A. G. Lyne, F. Camilo, R. N. Manchester, and N. D’Amico. Detection of ionized gas in the globular cluster 47 tucanae. *The Astrophysical Journal*, 557:L105–L109, 2001. doi: 10.1086/323236.
- [81] R. T. Edwards, G. B. Hobbs, and R. N. Manchester. Tempo2, a new pulsar timing package. ii. the timing model and precision estimates. *Monthly Notices of the Royal Astronomical Society*, 372:1549–1574, 2006. doi: 10.1111/j.1365-2966.2006.10870.x.
- [82] G. B. Hobbs, R. T. Edwards, and R. N. Manchester. Tempo2, a new pulsar timing package. i. an overview. *Monthly Notices of the Royal Astronomical Society*, 369:655–672, 2006. doi: 10.1111/j.1365-2966.2006.10302.x.
- [83] I. S. Shklovskii. Possible causes of the secular increase in pulsar periods. *Soviet Astronomy*, 13:562, 1970.

- [84] P. C. C. Freire, F. Camilo, M. Kramer, D. R. Lorimer, A. G. Lyne, R. N. Manchester, and N. D’Amico. Timing of the millisecond pulsars in 47 tucanae: Proper motions, accelerations, and an improved cluster mass model. *Monthly Notices of the Royal Astronomical Society*, 340:1359–1374, 2003. doi: 10.1046/j.1365-8711.2003.06384.x.
- [85] D. C. Backer, S. R. Kulkarni, C. Heiles, M. M. Davis, and W. M. Goss. A millisecond pulsar. *Nature*, 300:615–618, 1982. doi: 10.1038/300615a0.
- [86] J. M. Cordes and T. J. W. Lazio. Ne2001.i. a new model for the galactic distribution of free electrons and its fluctuations. *arXiv e-prints*, 2002.
- [87] F. F. Gardner and J. B. Whiteoak. The rotation measures of extragalactic radio sources. *Annual Review of Astronomy and Astrophysics*, 7:245–272, 1969. doi: 10.1146/annurev.aa.07.090169.001333.
- [88] P. A. G. Scheuer. Amplitude fluctuations in radio signals propagated through the interstellar medium. *Nature*, 218:920–922, 1968. doi: 10.1038/218920a0.
- [89] Mark A. Walker, Paul B. Demorest, and Willem van Straten. Cyclic spectroscopy of the millisecond pulsar B1937+21. *The Astrophysical Journal*, 779(2):99, 2013. doi: 10.1088/0004-637X/779/2/99.
- [90] Duncan R. Lorimer. Sigproc: Pulsar signal processing programs. <http://sigproc.sourceforge.net>, 2011.
- [91] D. C. Wells, E. W. Greisen, and R. H. Harten. Fits: A flexible image transport system. *Astronomy and Astrophysics Supplement Series*, 44:363–370, 1981.
- [92] W. D. et al. Pence. Definition of the flexible image transport system (fits). *Astronomy and Astrophysics*, 524:A42, 2010.
- [93] A. W. Hotan, W. van Straten, and R. N. Manchester. Psrfits: A standard format for pulsar data storage. *Publications of the Astronomical Society of Australia*, 21:302–309, 2004.
- [94] Scott M. Ransom. Presto: Pulsar exploration and search toolkit. <https://github.com/scottransom/presto>, 2001–present.
- [95] P. Lazarus, P. C. C. Freire, B. Allen, et al. Spin period distributions of pulsars detected in radio surveys. *The Astrophysical Journal*, 812:81, 2015. doi: 10.1088/0004-637X/812/1/81.

- [96] E. van Heerden, B. W. Stappers, and E. F. Keane. The impact of red noise on pulsar search pipelines. *Monthly Notices of the Royal Astronomical Society*, 465:3966–3981, 2017. doi: 10.1093/mnras/stw3005.
- [97] P. A. Fridman and W. A. Baan. Rfi mitigation methods in radio astronomy. *Astronomy & Astrophysics*, 378:327–344, 2001. doi: 10.1051/0004-6361:20011257.
- [98] R. P. Eatough, E. F. Keane, and A. G. Lyne. An interference removal technique for radio pulsar searches. *Monthly Notices of the Royal Astronomical Society*, 395:410–415, 2009. doi: 10.1111/j.1365-2966.2009.14524.x.
- [99] Scott M. Ransom. Presto tutorial. [https://hosting.astro.cornell.edu/~shami/psrintro/papers/Ransom\\_PRESTO\\_tutorial.pdf](https://hosting.astro.cornell.edu/~shami/psrintro/papers/Ransom_PRESTO_tutorial.pdf), 2011.
- [100] Scott M. Ransom. Presto: Pulsar exploration and search toolkit. <https://github.com/scottransom/presto>, 2001–present. Astrophysics Source Code Library.
- [101] D. R. Lorimer, A. J. Faulkner, A. G. Lyne, R. N. Manchester, M. Kramer, M. A. McLaughlin, G. Hobbs, A. Possenti, I. H. Stairs, F. Camilo, M. Burgay, N. D’Amico, A. Corongiu, and F. Crawford. The parkes multibeam pulsar survey – vi. discovery and timing of 142 pulsars and a galactic population analysis. *Monthly Notices of the Royal Astronomical Society*, 372:777–800, 2006. doi: 10.1111/j.1365-2966.2006.10887.x.
- [102] M. J. Keith, A. Jameson, W. van Straten, M. Bailes, E. D. Barr, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, R. P. Eatough, S. Johnston, M. Kramer, L. Levin, A. G. Lyne, R. N. Manchester, A. Possenti, and B. W. Stappers. The high time resolution universe pulsar survey – i. system configuration and initial discoveries. *Monthly Notices of the Royal Astronomical Society*, 409: 619–627, 2010. doi: 10.1111/j.1365-2966.2010.17325.x.
- [103] Scott M. Ransom. *New Search Techniques for Binary Pulsars*. PhD thesis, Harvard University, 2001.
- [104] Scott M. Ransom. Fourier techniques for very long astrophysical time-series analysis. *The Astronomical Journal*, 124:1788–1809, 2002.
- [105] James W. Cooley and John W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301, 1965. doi: 10.1090/S0025-5718-1965-0178586-1.

- [106] Bridget C. Andersen and Scott M. Ransom. A Fourier Domain “Jerk” Search for Binary Pulsars. , 863(1):L13, August 2018. doi: 10.3847/2041-8213/aad59f.
- [107] Manjari Bagchi, Duncan R. Lorimer, and Spencer Wolfe. On the detectability of eccentric binary pulsars. , 432(2):1303–1314, June 2013. doi: 10.1093/mnras/stt559.
- [108] R. J. Lyon et al. Fifty years of pulsar candidate selection. *Monthly Notices of the Royal Astronomical Society*, 459:1104–1123, 2016.
- [109] R. P. Eatough et al. Selection of radio pulsar candidates using artificial neural networks. *Monthly Notices of the Royal Astronomical Society*, 407:2443–2450, 2010.
- [110] D. H. Staelin. Fast folding algorithm for detection of periodic pulse trains. *Proceedings of the IEEE*, 57:724–725, 1969.
- [111] E. Parent et al. A search for long-period pulsars using the fast folding algorithm. *The Astrophysical Journal*, 861:44, 2018.
- [112] V. Morello et al. Spandak: A fast folding algorithm search pipeline. *Monthly Notices of the Royal Astronomical Society*, 497:4654–4668, 2020.
- [113] M. A. McLaughlin et al. Transient radio bursts from rotating neutron stars. *Nature*, 439:817–820, 2006.
- [114] J. M. Cordes and M. A. McLaughlin. Searches for fast radio transients. *The Astrophysical Journal*, 596:1142–1154, 2003.
- [115] D. R. Lorimer et al. A bright millisecond radio burst of extragalactic origin. *Science*, 318:777, 2007.
- [116] B. Knispel, R. P. Eatough, H. Kim, E. F. Keane, B. Allen, D. Anderson, C. Aulbert, O. Bock, F. Crawford, H.-B. Eggenstein, H. Fehrmann, D. Hammer, M. Kramer, A. G. Lyne, B. Machenschalk, R. B. Miller, M. A. Papa, D. Rastawicki, J. Sarkissian, X. Siemens, and B. W. Stappers. Einstein@Home Discovery of 24 Pulsars in the Parkes Multi-beam Pulsar Survey. , 774(2):93, September 2013. doi: 10.1088/0004-637X/774/2/93.
- [117] B. Allen, B. Knispel, J. M. Cordes, J. S. Deneva, J. W. T. Hessels, D. Anderson, C. Aulbert, O. Bock, A. Brazier, S. Chatterjee, P. B. Demorest,

- H. B. Eggenstein, H. Fehrmann, E. V. Gotthelf, D. Hammer, V. M. Kaspi, M. Kramer, A. G. Lyne, B. Machenschalk, M. A. McLaughlin, C. Messenger, H. J. Pletsch, S. M. Ransom, I. H. Stairs, B. W. Stappers, N. D. R. Bhat, S. Bogdanov, F. Camilo, D. J. Champion, F. Crawford, G. Desvignes, P. C. C. Freire, G. Heald, F. A. Jenet, P. Lazarus, K. J. Lee, J. van Leeuwen, R. Lynch, M. A. Papa, R. Prix, R. Rosen, P. Scholz, X. Siemens, K. Stovall, A. Venkataraman, and W. Zhu. The Einstein@Home Search for Radio Pulsars and PSR J2007+2722 Discovery. , 773(2):91, August 2013. doi: 10.1088/0004-637X/773/2/91.
- [118] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. , 53(12):6749–6761, June 1996. doi: 10.1103/PhysRevD.53.6749.
- [119] C. Messenger and A. Patruno. A template bank to search for gravitational waves from neutron stars in binaries. *Monthly Notices of the Royal Astronomical Society*, 418:2669–2684, 2011.
- [120] Vishnu Balakrishnan, David Champion, Ewan Barr, Michael Kramer, V. Venkatraman Krishnan, Ralph P. Eatough, Rahul Sengar, and Matthew Bailes. Coherent search for binary pulsars across all Five Keplerian parameters in radio observations using the template-bank algorithm. , 511(1):1265–1284, March 2022. doi: 10.1093/mnras/stab3746.
- [121] J. H. Taylor and J. M. Weisberg. A new test of general relativity: Gravitational radiation and the binary pulsar psr 1913+16. *Astrophysical Journal*, 253:908–920, 1982.
- [122] M. Kramer et al. Strong-field gravity tests with pulsars and black holes. *Physical Review X*, 11:041050, 2021.
- [123] F. Ozel and P. Freire. Masses, radii, and the equation of state of neutron stars. *Annual Review of Astronomy and Astrophysics*, 54:401–440, 2016.
- [124] M. Bailes et al. A massive neutron star in a compact relativistic binary. *Science*, 383:275–279, 2024.
- [125] G. Agazie et al. The nanograv 15-year data set: Evidence for a gravitational-wave background. *Astrophysical Journal Letters*, 951:L9, 2023.
- [126] P. C. C. Freire, A. Ridolfi, M. Kramer, et al. Pulsars in globular clusters. *Monthly Notices of the Royal Astronomical Society*, 471:857–876, 2017.

- [127] Y. Gupta, B. Ajithkumar, H. S. Kale, et al. The upgraded gmrt: opening new windows on the radio universe. *Current Science*, 113:707–714, 2017.
- [128] G. Hobbs, R. Manchester, A. Dunning, et al. The parkes pulsar timing array project. *Publications of the Astronomical Society of Australia*, 37:e012, 2020.
- [129] R. Nan, D. Li, C. Jin, et al. The five-hundred-meter aperture spherical radio telescope (fast) project. *International Journal of Modern Physics D*, 20:989–1024, 2011.
- [130] J. Jonas and MeerKAT Team. Meerkat: The south african square kilometre array precursor. *Proceedings of Science*, MeerKAT2016:001, 2016.
- [131] M. Bailes, A. Jameson, C. Flynn, et al. The meertime pulsar timing project. *Publications of the Astronomical Society of Australia*, 37:e028, 2020.
- [132] A. Ridolfi, P. C. C. Freire, M. Kramer, et al. Trapum discovery of new pulsars in globular clusters. *Monthly Notices of the Royal Astronomical Society*, 504:1407–1421, 2021.
- [133] A. Ridolfi et al. Pulsar discoveries with meerkat. *Monthly Notices of the Royal Astronomical Society*, 513:2292–2304, 2022.
- [134] F. Abbate et al. Meerkat observations of globular cluster pulsars. *Monthly Notices of the Royal Astronomical Society*, 513:2292–2304, 2022.
- [135] N. D. R. Bhat, J. M. Cordes, F. Camilo, D. J. Nice, and D. R. Lorimer. Multifrequency observations of radio pulse broadening and constraints on interstellar electron density microstructure. *Astrophysical Journal*, 605:759–783, 2004.
- [136] S. M. Ransom, J. W. T. Hessels, I. H. Stairs, et al. Twenty-one millisecond pulsars in terzan 5. *Science*, 307:892–896, 2005.
- [137] S. Johnston, M. Kramer, and D. R. Lorimer. High-frequency pulsar searches. *Monthly Notices of the Royal Astronomical Society*, 373:L6–L10, 2006.
- [138] R. S. Lynch, S. M. Ransom, P. C. C. Freire, et al. The green bank telescope search for pulsars in globular clusters. *Astrophysical Journal Letters*, 730:L11, 2011.
- [139] J. S. Deneva, P. S. Ray, F. Camilo, et al. Discovery of pulsars in globular clusters with the gbt. *Astrophysical Journal Letters*, 807:L23, 2015.

- [140] F. Abbate, A. Ridolfi, E. D. Barr, S. Buchner, M. Burgay, D. J. Champion, W. Chen, P. C. C. Freire, T. Gautam, J. M. Grießmeier, and et al. Four pulsar discoveries in ngc 6624 by trapum using meerkat. , 513(2):2292–2301, June 2022. doi: 10.1093/mnras/stac1041.
- [141] Andrew Douglas, Prajwal V. Padmanabh, Scott M. Ransom, Alessandro Ridolfi, Paulo Freire, Vivek Venkatraman Krishnan, Ewan D. Barr, Cristina Pallanca, Mario Cadelano, Andrea Possenti, and et al. Two new black widow millisecond pulsars in m28. , 927(1):L27, March 2022.
- [142] L. Vleeschower, B. W. Stappers, M. Bailes, E. D. Barr, M. Kramer, S. Ransom, A. Ridolfi, V. Venkatraman Krishnan, A. Possenti, M. J. Keith, and et al. Discoveries and timing of pulsars in ngc 6440. , 513(1):L95–L100, June 2022.
- [143] W. Chen, P. C. C. Freire, A. Ridolfi, E. D. Barr, B. Stappers, M. Kramer, A. Possenti, S. M. Ransom, L. Levin, R. P. Breton, and et al. Meerkat discovery of 13 new pulsars in omega centauri. , 520(3):3237–3252, April 2023.
- [144] P. C. C. Freire, S. M. Ransom, S. Bégin, et al. Eight new millisecond pulsars in ngc 6440 and ngc 6441. *Astrophysical Journal*, 675:670–682, 2008.
- [145] R. S. Lynch, P. C. C. Freire, and S. M. Ransom. New pulsars discovered in globular clusters. *Astrophysical Journal*, 745:109, 2012.
- [146] S. Dai et al. Pulsar observations at s-band. *Monthly Notices of the Royal Astronomical Society*, 493:3843–3856, 2020.
- [147] Deven Bhakta, Scott M. Ransom, Megan DeCesar, and Shi Dai. Discovery and Timing of the First Millisecond Pulsar in NGC 6316. *arXiv e-prints*, art. arXiv:2603.06476, March 2026.
- [148] W. Chen, E. Barr, R. Karuppusamy, and M. Kramer. Wide-field beamformed observations with meerkat. *Journal of Astronomical Instrumentation*, 10:2150013, 2021.
- [149] Z. Yan, Z. Pan, S. M. Ransom, et al. An eclipsing black widow pulsar in ngc 6712. *Astrophysical Journal*, 921:120, 2021.
- [150] A. V. McCarver, T. J. Maccarone, S. M. Ransom, et al. A pulsar discovery in glimpse-c01. *Astrophysical Journal*, 969:30, 2024.

- [151] Y. Lian, Z. Pan, H. Zhang, S. Cao, P. C. C. Freire, L. Qian, R. P. Eatough, L. Shao, S. M. Ransom, and D. R. Lorimer. The fast globular cluster pulsar survey (gc fans). , 279(2):51, 2025. doi: 10.3847/1538-4365/ade4ba.
- [152] T. M. Tauris and G. J. Savonije. Formation of millisecond pulsars with helium white dwarf companions. , 350:928–944, 1999.
- [153] J. Strader, L. Chomiuk, C. C. Cheung, D. J. Sand, D. Donato, R. H. D. Corbet, D. Koeppe, P. G. Edwards, J. Stevens, and L. Petrov. 1fgl j1417.7–4407: A likely gamma-ray bright binary with a massive neutron star and a giant secondary. , 804:L12, 2015. doi: 10.1088/2041-8205/804/1/L12.
- [154] J. Strader, P. S. Ray, R. Urquhart, S. J. Swihart, L. Chomiuk, E. Aydi, E. C. Bellm, K. C. Dage, M. E. DeCesar, and J. S. Deneva. Psr j1947–1120: A new huntsman millisecond pulsar binary. , 980:124, 2025. doi: 10.3847/1538-4357/ada897.
- [155] L. Zhang, A. Ridolfi, H. Blumer, P. C. C. Freire, R. N. Manchester, M. McLaughlin, K. Kremer, and A. D. Cameron. Radio detection of an elusive millisecond pulsar in the globular cluster ngc 6397. , 934:L21, 2022. doi: 10.3847/2041-8213/ac81c3.
- [156] T. M. Tauris, M. Kramer, P. C. C. Freire, N. Wex, H.-T. Janka, N. Langer, P. Podsiadlowski, E. Bozzo, and S. Chaty. Formation of double neutron star systems. , 846:170, 2017. doi: 10.3847/1538-4357/aa7e89.
- [157] P. Lazarus, T. M. Tauris, B. Knispel, P. C. C. Freire, J. S. Deneva, V. M. Kaspi, B. Allen, and S. Bogdanov. Timing of a young mildly recycled pulsar with a massive white dwarf companion. , 437:1485–1494, 2014. doi: 10.1093/mnras/stt1996.
- [158] J. M. Yao, R. N. Manchester, and N. Wang. A new electron-density model for estimation of pulsar and frb distances. *ApJ*, 835:29, 2017.
- [159] P. C. C. et al. Freire. On the nature and evolution of globular cluster pulsars. *ASPC*, 328:405, 2005.
- [160] P. V. Padmanabh, S. M. Ransom, P. C. C. Freire, A. Ridolfi, J. D. Taylor, C. Choza, C. J. Clark, F. Abbate, M. Bailes, E. D. Barr, S. Buchner, M. Burgay, M. E. DeCesar, W. Chen, A. Corongiu, D. J. Champion, A. Dutta, M. Geyer, J. W. T. Hessels, M. Kramer, A. Possenti, I. H. Stairs, B. W. Stappers, V. Venkatraman Krishnan, L. Vleeschower, and L. Zhang. Discovery and

- timing of ten new millisecond pulsars in the globular cluster Terzan 5. , 686: A166, June 2024. doi: 10.1051/0004-6361/202449303.
- [161] Scott M. Ransom, James M. Cordes, and Stephen S. Eikenberry. A New Search Technique for Short Orbital Period Binary Pulsars. *The Astrophysical Journal*, 589:911–920, 2003. doi: 10.1086/374806.
- [162] V. Morello et al. Using deep learning for pulsar candidate classification. *Monthly Notices of the Royal Astronomical Society*, 483:3673–3685, 2019.
- [163] E. F. Keane, E. D. Barr, A. Jameson, et al. The survey for pulsars and extragalactic radio bursts – ii. new frb discoveries and their follow-up. *Monthly Notices of the Royal Astronomical Society*, 473(1):116–135, 2018. doi: 10.1093/mnras/stx2126.
- [164] C. J. Clark, E. D. Barr, M. C. Beuidenthout, et al. The trapum l-band survey for pulsars in fermi-lat gamma-ray sources. *Monthly Notices of the Royal Astronomical Society*, 519(4):5590–5611, 2023. doi: 10.1093/mnras/stac3791.
- [165] P. V. Padmanabh, E. D. Barr, W. Chen, et al. Mpirf–meerkat galactic plane survey – i. system set-up and discovery of 78 pulsars from the l-first results. *Monthly Notices of the Royal Astronomical Society*, 524(1):1291–1312, 2023. doi: 10.1093/mnras/stad1830.
- [166] R. Sengar, V. Balakrishnan, E. D. Barr, et al. Discovery of 37 new pulsars through gpu-accelerated reprocessing of the parkes multibeam pulsar survey. *Monthly Notices of the Royal Astronomical Society*, 522(1):1071–1089, 2023. doi: 10.1093/mnras/stad508.
- [167] K. Liu, R. P. Eatough, N. Wex, and M. Kramer. Pulsar–black hole binaries: prospects for new gravity tests with future radio telescopes. *Monthly Notices of the Royal Astronomical Society*, 445(3):3115–3132, 2014. doi: 10.1093/mnras/stu1913.
- [168] A. D. Cameron, D. J. Champion, M. Kramer, et al. An extremely relativistic double neutron star system. *Monthly Notices of the Royal Astronomical Society*, 475(1):L57–L61, 2018. doi: 10.1093/mnrasl/slx175.
- [169] L. Nieder, B. Allen, C. J. Clark, and H. J. Pletsch. Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-Ray Pulsars. , 901(2): 156, October 2020. doi: 10.3847/1538-4357/abaf53.

- [170] Holger J. Pletsch and Colin J. Clark. Optimized Blind Gamma-Ray Pulsar Searches at Fixed Computing Budget. , 795(1):75, November 2014. doi: 10.1088/0004-637X/795/1/75.
- [171] Kuang Mao, Zhaorong Tang, Qihong Pan, Pei Wang, Huaxi Chen, Scott M. Ransom, Di Li, Xuefei Tang, Qi Wang, Yi Feng, Lei Chen, Donghui Quan, and Zujie Ren. PrestoZL: A gpu-accelerated high-throughput jerk search toolkit for binary pulsars. *The Astrophysical Journal Supplement Series*, 280(1):36, 2025. doi: 10.3847/1538-4365/adf4e5.



**Funding Acknowledgement:**

This thesis was produced while attending the PhD programme in Physics at the University of Cagliari, Cycle XXXVIII, with the support of a scholarship financed by NRRP, funded by the European Union - NextGenerationEU - Mission 4 "Education and Research", Component 2 "From Research to Business" – Investment 1.4 "Strengthening research structures and creating R&D "national champions" on specific key enabling technologies".

