

Pulsar Glitches and Superfluid Dynamics of Neutron Stars

George Alec William Howitt

ORCID: 0000-0002-2593-0189

Submitted in total fulfillment of the requirements of the degree
of Doctor of Philosophy

December 2020

School of Physics
Faculty of Science
University of Melbourne

Declaration

I hereby declare that

- This thesis, titled "Pulsar glitches and superfluid dynamics of Neutron Stars" comprises only my original work towards the degree of Doctor of Philosophy, except where indicated otherwise in the preface.
- Due acknowledgement has been made in the text to all other material used.
- The thesis is fewer than 100,000 words in length, exclusive of tables, figures, bibliography and appendices.

George Howitt

Abstract

Glitches are sudden jumps in the spin frequency of pulsars that occur at random times and over several orders of magnitude in size. One popular theory holds that glitches are caused by differential rotation between the neutron star crust and a superfluid component in the interior. This implies that glitches can be used to study the dense nuclear matter inside neutron stars, testing a regime of physics inaccessible to present-day laboratory experiments or other electromagnetic observations. In the vortex avalanche model of pulsar glitches, the angular momentum of the superfluid is determined by the configuration of an array of quantized vortices which pin to impurities in the interior. As the star spins down, individual vortices unpin and move outwards, triggering unpinning in other nearby vortices leading to an avalanche. The angular momentum lost by the vortices as they move outward is transferred to the crust, which causes it to spin up. This thesis tests the superfluid model of pulsar glitches using hydrodynamic simulations of glitch recovery, statistical analysis of astronomical glitch data and N -body point vortex simulations of avalanches in an idealised neutron star model.

Following a glitch, pulsars often have a lengthy recovery period of weeks to months, which is a key piece of evidence for the presence of a superfluid phase inside neutron stars. We perform simulations of glitch recovery using a pseudo-spectral Navier-Stokes-like solver. We model a neutron star as a two component superfluid consisting of a viscous proton-electron plasma and an inviscid superfluid neutron condensate in a spherical Couette geometry. The two fluid components are coupled through mutual friction. We prepare the system in a state of differential rotation between the core and the outer crust and examine the response of the outer boundary to glitches induced by instantaneously changing the angular velocity of the boundaries and by recoupling the fluid interior. We find that with strong mutual friction the characteristic glitch recovery time scale decreases by as much as a factor of three. For glitches originating in the fluid, strong mutual friction decreases the maximum spin up by a factor of up to five. These effects may be partially responsible for the diversity of glitch recoveries observed in the pulsar popu-

lation, which vary from rapid, complete recoveries to slow partial recoveries or step changes in the spin frequency with no observed change in the spin down rate. The strength of mutual friction depends on where in the star a superfluid phase exists, so these results may allow future observations to constrain the star’s internal structure.

The vortex avalanche model of glitches is similar to other non-equilibrium systems in nature such as sandpiles, earthquakes and solar flares. Such systems are often studied in the paradigm of self-organized criticality (SOC), a hallmark of which is scale invariant size probability density functions (PDFs) and exponential waiting time PDFs. We perform a statistical study of the size and waiting time PDFs for the five most prolific glitching pulsars using the non-parametric kernel density estimator. This work complements previous parametric studies of glitches, which typically fit the PDFs to known distributions such as power laws, Gaussian or log-normal distributions. Two objects, PSR J1740–3015 and the Crab pulsar PSR J0534+2200, appear to have exponential waiting time PDFs and scale-invariant size PDFs over several decades. Two other objects, PSR J0537–6910 and the Vela pulsar PSR J0835–4510 have quasiperiodic size and waiting time PDFs, typical of fast-driven SOC systems. One object, PSR J1341–6220, appears to exhibit hybrid behaviour. We test the ability of the kernel density estimator to resolve multimodality in synthetic data drawn from a composite Gaussian distribution (which is qualitatively similar to the waiting time PDF in J1341–6220), and find that the small number of glitches observed in this object ($N = 23$) makes confirmation of multimodality in the PDFs difficult.

In order to study the non-equilibrium dynamics of the vortex avalanche theory of glitches, it is useful to have an idealised model which can make falsifiable predictions about the properties of glitches, such as PDFs of sizes and waiting times, correlations between these observables, and the dependence of the PDFs on the internal parameters of the model. Previous work with Gross-Pitaevskii simulations has been unable to study systems larger than ~ 100 vortices, far from the $\sim 10^{18}$ vortices in neutron stars. With such small numbers of vortices, it is impossible to resolve the dynamics over multiple orders of magnitude. We develop a mathematical model for simulating the motion of point vortices in two dimensions under the influence of deceleration, dissipation, and pinning. We describe a numerical solver for this model, which uses an N -body method to compute the vortex velocities and an adaptive time step scheme for time evolution. We present the results of numerical experiments validating the method, including stability of a vortex ring and dissipative formation of an Abrikosov array. We then perform simulations of 1000-5000 vortices with a wide range of values for the strengths of dissipation and pinning, pinning site density and deceleration of

the container. Vortex avalanches occur routinely in the N -body simulations, when chains of unpinning events are triggered collectively by vortex-vortex repulsion, consistent with previous, smaller scale studies using the Gross-Pitaevskii equation. The PDFs of the avalanche sizes and waiting times are consistent with both exponential and lognormal distributions. We find weak cross-correlations between glitch sizes and waiting times. These correlations are consistent with astronomical data and meta-models of pulsar glitch activity as a state-dependent Poisson process or a Brownian stress-accumulation process, and inconsistent with another popular alternative, a threshold-triggered stress-release model with a single, global stress reservoir. The spatial distribution of the effective stress within the simulation volume is analysed before and after a glitch. We find that stress is distributed homogeneously throughout the system and remains near the critical threshold both before and after glitches. This implies that there is no ‘memory’ in the system; glitches do not significantly reduce the global or local stress. This is consistent with the lack of strong cross-and-auto correlations between glitch sizes and waiting times in the simulations and pulsar data.

Preface

While the majority of this thesis is the work of the author, several of the chapters are closely based on collaborative publications. Details of my contributions to each publication are presented here.

- Chapter 1 is a literature review of pulsar glitches and superfluid hydrodynamics. All of the text, tables and original figures are my own, produced in consultation with Andrew Melatos. Figure 1.1 is reproduced from Kaspi & Kramer (2016) with attribution. Publication details of all works discussed in the review are included in the bibliography.
- Chapter 2 is closely based on the following publication: G. Howitt, B. Haskell, A. Melatos, *Hydrodynamic Simulations of Pulsar Glitches*, published in Monthly Notices of the Royal Astronomical Society, Volume 460, Issue 2, p.1201-1213, August 2016.

This work describes simulations performed using a superfluid hydrodynamics code developed by Carlos Peralta [based in turn on an earlier code by Bagchi & Balachandar (2002)] with modifications made by myself. A detailed description of the original code is in Peralta et al. (2006). The main modification to the code made by me was to improve the time-stepping algorithm for the evolution of the angular velocity of the outer boundary from a first-order Euler method to fourth-order Runge-Kutta.

The simulations were performed by myself with input from Brynmor Haskell and Andrew Melatos. The text was written by myself, apart from several paragraphs in section 2.2 around equation (2.2) which were written by Haskell. The text was edited and revised based on discussion with Haskell and Melatos. All original figures were produced by myself, and Figure 2.1 is reproduced from Peralta (2006) with attribution. Preparatory work for this publication, including familiarisation with and discussion of the code and preliminary (mainly validation-style) simulations, was included in my MSc thesis at the University of Melbourne.

- Chapter 3 is closely based on the following publication:

G. Howitt, A. Melatos, A. Delaigle, *Nonparametric Estimation of the Size and Waiting Time Distributions of Pulsar Glitches*, published in The Astrophysical Journal, Volume 867, Issue 1, article id. 60, November 2018.

This work presents estimates of glitch size and waiting time probability density functions (PDFs), based on data from the Jodrell Bank glitch catalogue (Espinoza et al., 2014).

The first several drafts of this paper were prepared by Andrew Melatos prior to the beginning of my PhD, with myself initially listed as second author. The published version of the paper contains some common material with the initial drafts. However it has been modified and added to extensively by me during my PhD. I am the first author of the published version to reflect my role as the primary contributor. As the authorship history of this publication is somewhat complicated, I will outline the changes to the paper from the initial drafts and the contribution from each author in detail.

The text in the introduction discussing the background information on pulsar glitches and the motivation for the work was initially written by Melatos. The published version replaces parts of Melatos' brief original text with a more extensive discussion of previous work on the topic (second paragraph) and a more in-depth discussion of the methodology used in the paper (3rd paragraph), both written by me.

Section 3.2.1 discussing the kernel density estimator was initially written by Melatos and was edited heavily by me in consultation with Aurore Delaigle for accuracy and clarity and to correct technical mistakes in Melatos' original version. Section 3.2.2 was initially written by Melatos and edited by me to better reflect the final content of the publication. Section 3.2.3 was written by me. The first part of section 3.2.4 discussing the exponential and power law distributions was written by Andrew Melatos based on Monte Carlo simulations performed by me and figures that I produced. The latter part of the section discussing multimodality was written by me.

Section 3.3 presenting the results of the kernel density estimates of glitch size and waiting time PDFs was initially written by Melatos based on data analysis performed by me. Between the early drafts and the final publication, our approach to the data analysis changed significantly. Initially, we used several different bandwidth-selection algorithms and discussed their effect on the resulting PDF estimates.

This analysis was excluded from the published version and replaced with the discussion on error estimation in section 3.2.3 and the discussion on over/under-smoothing discussed in the final paragraph of section 3.2.4. Also, more glitches were detected in some of the pulsars discussed in the publication. These changes were incorporated into the text by me. The superseded results from the early drafts were included in a chapter of my MSc thesis at the University of Melbourne.

The discussion of the results in section 3.4 was initially written by Melatos and edited by me after the final data analysis was complete to account for the changes in the results.

All of the data analysis was performed by me based on analysis scripts written by me. All figures were produced by me at all stages of the writing process. Melatos and Delaigle did not perform any of the analysis or produce any figures. Modifications to the paper following the referee report and the response to the referee were written by me. These include the discussion of errors (section 3.2.3) and multimodality, the quantitative analysis in section 3.2.4 and a more formal definition of the KDE and the Gaussian kernel in section 3.2.1.

Delaigle provided advice on the use and applicability of the kernel density estimator and interpretation of the results, as well as comments on drafts and the referee response.

- Chapter 4 is closely based on the following publication:

G. Howitt, A. Melatos, B. Haskell, *Simulating pulsar glitches: an N-body solver for superfluid vortex motion in two dimensions*, published in Monthly Notices of the Royal Astronomical Society, Volume 498, Issue 1, pp.320-331, August 2020.

This work discusses a mathematical model for the motion of superfluid point vortices, a numerical method for solving for the vortex motion and the results of simulations that demonstrate vortex avalanches.

The model was formulated mathematically and physically by me in collaboration with Andrew Melatos and Brynmor Haskell. The numerical solver was written entirely by me. The text, as well as all tables and figures, was written by me, with edits and revisions suggested by Melatos and Haskell.

- Chapter 5 summarises the work presented in Chapters 2 – 4, as well as future directions for the work in the thesis. The future directions include preliminary results on three new projects, described in sections

5.2.1, 5.2.2 and 5.2.3. The calculations and 11 figures were done entirely by me. The 3D vortex code described in section 5.2.2 was written by me. The text was written by me in consultation with Andrew Melatos.

During my PhD I also contributed to three other publications which are not included in this thesis. They are listed below.

- G. Howitt, S. Stevenson, A. Vigna-Gómez, S. Justham, N. Ivanova, T. E. Woods, C. J. Neijssel, I. Mandel, *Luminous Red Novae: population models and future prospects*, published in Monthly Notices of the Royal Astronomical Society, Volume 492, Issue 3, p.3229-3240, March 2020

This work studied the possible connection between common envelope evolution in binaries and a class of optical transients called Luminous Red Novae (LRNe). We used the binary population synthesis code COMPAS to evolve a population of isolated stellar binaries and ran the output of the simulations through a set of models for the duration and luminosity of LRNe based on properties of binary systems that undergo common envelope evolution. We predicted the expected detection rates and brightness and duration of LRNe populations that will be detected with the Vera C. Rubin Observatory (LSST) under each of our model assumptions. The project was conceived by Ilya Mandel, Natasha Ivanova and Stephen Justham, with preliminary work undertaken by Simon Stevenson. I ran the simulations with the assistance of Stevenson, Alejandro Vigna-Gómez and Coenraad Neijssel. I performed all of the data analysis and produced all of the figures by myself, with advice from Stevenson, Vigna-Gómez, Neijssel and Mandel. I wrote the majority of the text, with contributions from Stevenson, Vigna-Gómez and Mandel, who also provided comments and revisions throughout the writing process. Justham, Ivanova and Woods provided several of the models used in the final analysis and also provided comments on the manuscript throughout the writing process.

- A. Melatos, G. Howitt, W. Fulgenzi, *Size-waiting-time Correlations in Pulsar Glitches*, published in The Astrophysical Journal, Volume 863, Issue 2, article id. 196, August 2018

This work looked at statistical correlations between pulsar glitch sizes and waiting times, in the context of a state-dependent Poisson process model for glitches. Few statistically-significant correlations are found, though this may be due to small-number statistics, and we make predictions based on the state-dependent Poisson process model that correlations will be seen in several pulsars as more glitches are detected.

The text was written by Andrew Melatos, and I provided comments and revisions throughout the writing process. Will Fulgenzi advised on technical aspects of the theory, which borrows from Fulgenzi, Melatos & Hughes (2017), including the calculations in the appendix. I performed all of the data analysis on the pulsar glitch data. I produced figures 1, 3 and 5, and Tables 1 and 2.

- A. Vigna-Gómez, M. MacLeod, C. J. Neijssel, F. S. Broekgaarden, S. Justham, G. Howitt, S. E. de Mink, S. Vinciguerra, I. Mandel, *Common envelope episodes that lead to double neutron star formation*, published in Publications of the Astronomical Society of Australia, Volume 37, article id. e038, 2020

This work studied the stellar progenitors of the galactic double neutron star population in order to match stellar and binary evolution models with observations. Using the binary population synthesis code COMPAS, we created a population of binary neutron stars with similar properties to those observed in the Milky Way, and examined the properties of the donors in the binaries at the onset of common envelope formation. We identified three significant classes of binary systems that eventually evolve to form double neutron stars: giant donors with fully-convective envelopes, cool donors with partially-convective envelopes and hot donors with radiative envelopes.

This project was led by Alejandro Vigna-Gómez, who wrote the majority of the manuscript and performed the simulations and data analysis. My contribution was limited to section 4.2 discussing luminous red novae, where I provided the numerical data based on simulations I performed, and provided comments on the text of that section.

Parts of this research were supported by

- the University of Melbourne through a Research Training Program Scholarship, The Haasz Family Fund, The Cook W E.C.H.H. Fund and The Dr Alan Kenneth Head Travelling Scholarships.
- the Australian Research Council through ARC Discovery Project DP110103347, ARC LIEF LE140100121, and the Centre of Excellence for Gravitational Wave Discovery, project number CE170100004.
- The National Science Centre Poland (NCN) through grant SONATA BIS 2015/18/E/ST9/00577, P.I.: B. Haskell

Acknowledgements

This PhD has been a long and sometimes tortuous journey, one which wouldn't have been possible to complete without the support of many people.

Firstly, I would like to thank my supervisor, Professor Andrew Melatos. I have been working with Andrew since I started my MSc in 2013, and I am hugely grateful to have such an intelligent and generous mentor. Andrew gave me a great deal of freedom to steer this project in the directions I saw fit, and his rigorous oversight has improved my work immeasurably. As well as his academic supervision of the project, Andrew's boundless intellectual curiosity and strong sense of academic and personal ethics has been a massive inspiration.

I am also grateful to my supervisory panel, Associate Professor Andy Martin and Associate Professor Michele Trenti, both of whom provided excellent oversight and support along the way.

I would also like to thank Dr Bynmor Haskell, who has been a wonderful collaborator and mentor throughout my PhD, as well as a generous host when I visited CAMK in Warsaw in late 2016.

I am very grateful to Professor Ilya Mandel and the COMPAS team, in particular Dr Alejandro Vigna-Gómez, Dr Simon Stevenson and Dr Coenraad Neijssel, who hosted me for several months at the University of Birmingham in 2017 and allowed me to work with them on a very interesting side-project to my main thesis work.

Thanks to everyone in the Astro group at the University of Melbourne. Working in such a friendly and intellectually stimulating environment made the otherwise esoteric and isolating work of a PhD fun and interesting. Finishing my thesis under lockdown has made me much more appreciative of the value of a great network of peers and collaborators.

Lastly I would like to thank my friends, family and my partner Lucy, without whose love, support and indulgence I certainly wouldn't have gotten this far.

Contents

1	Introduction	25
1.1	Pulsars	25
1.2	Pulsar timing	28
1.3	Timing irregularities	29
1.3.1	Timing noise	29
1.3.2	Glitches	30
1.3.3	Glitch statistics	31
1.3.4	Rise time and recovery	38
1.4	Superfluidity and its role in neutron stars	40
1.4.1	Observational evidence for superfluidity in neutron stars	41
1.4.2	Superfluid hydrodynamics	41
1.4.3	Mutual Friction and the HVBK model	43
1.4.4	GPE and vortex filament model	44
1.5	Physics of pulsar glitches	48
1.5.1	Crust-core coupling and mutual friction	49
1.5.2	Vortex pinning	50
1.5.3	Superfluid turbulence and hydrodynamic instabilities .	50
1.5.4	Crustquakes	51
1.5.5	Magnetic effects	51
1.5.6	Avalanche dynamics and self-organized criticality . . .	52
1.6	Thesis outline	52
2	Hydrodynamic simulations of pulsar glitch recovery	54
2.1	Introduction	54
2.2	Equations of motion	56
2.3	Numerical method	58
2.3.1	Pseudospectral solver	58
2.3.2	Initial and boundary conditions	59
2.3.3	Model assumptions and limitations	61
2.4	Steady pre-glitch spin down	63
2.4.1	Set up	63

2.4.2	Output and initial state	65
2.5	Post-glitch recovery	72
2.5.1	Activating the glitch	72
2.5.2	Fitting the glitches	73
2.5.3	Crust glitches	73
2.5.4	Bulk glitches	76
2.5.5	Inner glitches	78
2.6	Discussion	80
2.6.1	Crust vs inner glitches	80
2.6.2	Bulk glitches and mutual friction	81
2.7	Conclusions	83
3	Nonparametric Estimation of the Size and Waiting Time Dis-	
	tributions of Pulsar Glitches	85
3.1	Introduction	85
3.2	Kernel density estimator	87
3.2.1	Definition	87
3.2.2	Positive definite variables	89
3.2.3	Error estimation with nonparametric methods	90
3.2.4	Validation	90
3.3	Probability density functions	94
3.3.1	Waiting times	94
3.3.2	Sizes	97
3.4	Conclusion	100
4	Simulating pulsar glitches: an N-body solver for superfluid	
	vortex motion in two dimensions	102
4.1	Introduction	102
4.2	Vortex dynamics	104
4.2.1	Equations of motion	104
4.2.2	Vorticity-induced velocity	105
4.2.3	Boundary conditions and image vortices	106
4.2.4	Rotating frame	107
4.2.5	Pinning	107
4.2.6	Dissipation	107
4.3	Numerical method	108
4.3.1	Feedback	108
4.3.2	Dimensionless coordinates	109
4.3.3	Glitch finding	110
4.4	Validation	110
4.4.1	Stability of a vortex ring	110

4.4.2	Dissipation and Abrikosov array formation	111
4.5	An Astrophysical Example: Neutron Star Glitches	117
4.5.1	Set up	117
4.5.2	Avalanche dynamics	118
4.5.3	Sizes and waiting times	122
4.5.4	Model variations	125
4.5.5	Size-waiting time correlations	127
4.6	Conclusion	130
5	Summary and Future Work	132
5.1	Summary	132
5.1.1	Hydrodynamic simulations of pulsar glitch recovery . .	132
5.1.2	Nonparametric Estimation of the Size and Waiting Time Distributions of Pulsar Glitches	134
5.1.3	Simulations of pulsar glitches with N -body vortex code	135
5.2	Future work	136
5.2.1	Anti-glitches and the vortex avalanche model	137
5.2.2	N -body vortex method in three dimensions	140
5.2.3	Vortex annihilation and Onsager vortex formation . . .	149

List of Tables

1.1	Summary of literature on glitch completeness and what aspects are discussed in each paper.	32
1.2	Summary of numerical techniques for studying quantum turbulence with the vortex filament method.	47
1.3	Neutron star physics testable with studies of pulsar glitches. .	49
2.1	Fitted parameters to equation (2.14) with $N = 2$ for the crust glitches in Figure 2.7 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$	75
2.2	Fitted parameters to equation (2.13) with $N = 2$ for the bulk glitches in Figure 2.8 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$	76
2.3	Fitted parameters to equation (2.14) with $N = 2$ for the inner glitches in Figure 2.9 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$	78
2.4	We compare different methods for estimating the permanent change in the crust angular velocity residual, $\Delta\Omega_p$, following crust and inner glitches with strong and weak mutual friction. The third column estimates $\Delta\Omega_p$ as the value of $\Delta\Omega$ after 500 time units of simulation post-glitch, $\Delta\Omega(t = 500)$. The fourth column estimates $\Delta\Omega_p$ using the formula $\Delta\Omega_p^{\text{est}} = \Delta\Omega_{\text{max}} I_{\text{crust,core}}/I_{\text{total}}$. Finally we estimate $\Delta\Omega_p$ by fitting the data to glitch recovery models similar to equation (2.14) with five parameters (fifth column) and six parameters (sixth column).	81
4.1	Summary of input parameters for the vortex avalanche simulations. Values used in the simulation are discussed in the text in Sections 4.5.1 and 4.5.4.	118

4.2	Effect of varying simulation parameters (first column) on vortex avalanche statistics (columns two and three): mean waiting time between glitches $\langle \Delta t / T_0 \rangle$, and mean glitch size $\langle \Delta \Omega_C \rangle$. Symbols $+$ ($-$) indicate that $\langle \Delta t \rangle$ or $\langle \Delta \Omega \rangle$ increases (decreases) relative to the default value in Section 4.5.1 as the parameter in column one increases; “.” indicates no noticeable effect.	126
4.3	Pearson correlation coefficients r and Spearman rank coefficients ρ for forward ($+$) and backward ($-$) size-waiting time correlations in each of our model variations.	129
4.4	Spearman autocorrelation coefficient ρ and p-values for glitch waiting times Δt and sizes $\Delta \Omega$ in each of our model variations [cf. Table 1 in Carlin & Melatos (2019)].	130

List of Figures

1.1	$P - \dot{P}$ diagram [from Kaspi & Kramer (2016)]. Red dots indicate pulsars. Blue circles indicate pulsars with binary companions. Stars indicate pulsars found in supernova remnants.	27
1.2	Histogram of number of glitches per pulsar. Data taken from http://www.jb.man.ac.uk/pulsar/glitches.html (Espinoza et al., 2011).	32
1.3	Cumulative number of pulsar glitches detected (blue curve) and number of glitching pulsars discovered (orange dashed curve) versus time. Dash-dotted black line is a linear fit to the glitch discovery $N \propto t$ for $52500 < \text{MJD} < 56000$. Dotted black curve is a quadratic fit to the glitch discovery $N \propto t^2$ for $45000 < \text{MJD} < 56000$. Data taken from the Jodrell Bank glitch catalogue http://www.jb.man.ac.uk/pulsar/glitches.html (Espinoza et al., 2011).	35
1.4	Histogram of minimum and maximum glitch size in the 80 pulsars with > 1 detected glitch. Data taken from the Jodrell Bank glitch catalogue http://www.jb.man.ac.uk/pulsar/glitches.html (Espinoza et al., 2011).	37
2.1	Illustration of the model neutron star's geometry. A spherical surface of radius R_1 is embedded inside a larger spherical surface of radius R_2 with the same origin. Both spheres may rotate about independent axes, with respective angular velocity vectors $\boldsymbol{\Omega}_1$ and $\boldsymbol{\Omega}_2$, though in this chapter we only consider rotation about a common axis. The Cartesian coordinates are defined so that the inner sphere rotates around the z axis, and the orthonormal vectors $\hat{\mathbf{e}}_{r,\theta,\phi}$ form the basis of the spherical coordinate system. Taken from Peralta (2006).	59

- 2.2 Response of the crust and core angular velocities, $\Omega_1(t)$ and $\Omega_2(t)$, and their time derivatives, $\dot{\Omega}_1(t)$, and $\dot{\Omega}_2(t)$, to a constant external torque on the crust, for $0 \leq t \leq 500$, with strong (left column) and weak (right column) mutual friction, with $\Omega_1(t=0) = \Omega_2(t=0)$. In this figure, Ω_1 and Ω_2 are measured in a frame rotating with angular velocity $\Omega_f = \Omega_1(t=0) = \Omega_2(t=0)$ and are also normalized with respect Ω_f . Time is normalized with respect to Ω_f^{-1} . The top row is the crust angular velocity, $\Omega_1(t)$, (dashed curves), and the core angular velocity, $\Omega_2(t)$, (solid curves). The bottom row is the crust angular velocity derivative, $\dot{\Omega}_1(t)$, (dashed curves), and the core angular velocity derivative $\dot{\Omega}_2(t)$, (solid curves). In both mutual friction regimes the two boundaries are decelerating at approximately the same rate after $t = 500$, a state we refer to as spin-down equilibrium. 66
- 2.3 Contours of the azimuthal velocity lag, v_{pn}^ϕ , plotted at $t = 5 \approx \tau_{\text{MF}} < \tau_{\text{Ek}} < \tau_{\text{visc}}$ in the y - z plane, with strong mutual friction. Velocity is expressed in dimensionless units normalized with respect to $R_1\Omega_f^{-1}$. v_{pn}^ϕ is non-zero near $r = R_2$, but decreases monotonically to zero as $r \rightarrow R_1$. $|v_{\text{pn}}^\phi|$ increases from zero at the poles to a maximum at $\theta \approx 30^\circ$ in latitude. The flow is axisymmetric but not columnar, as the 18 contours have similar curvature to the boundaries throughout most of the domain. 67
- 2.4 Contours of the azimuthal component of the proton-neutron velocity lag, v_{pn}^ϕ , in the meridional plane, for strong mutual friction (left column), and weak mutual friction (right column), at $t = 5, 25, 100$, and 500 (top to bottom row respectively). Velocity is expressed in dimensionless units normalized with respect to $R_1\Omega_f^{-1}$. Red colours show positive lag into the page, blue colours show lag out of the page. Rotation is about the positive z -axis. 69

2.5	Snapshots of the θ and ϕ -averaged azimuthal velocity lag, $\langle v_{\text{pn}}^\phi \rangle_{\theta, \phi}$, versus cylindrical radius r , at $t = 5$ (top left), 25 (top right), 100 (bottom left), and 500 (bottom right). Time is expressed in dimensionless units normalized with respect to the angular velocity of the rotating frame Ω_f^{-1} , and velocity is expressed in dimensionless units normalized with respect to $R_1 \Omega_f^{-1}$. The red curve (circles) and the blue curve (crosses) correspond to strong and weak mutual friction respectively. The boundary condition enforces $v_{\text{pn}}^\phi = 0$ at $r = R_1$ and $r = R_2$. When mutual friction is strong, v_{pn}^ϕ changes sign near $r = R_1$. When mutual friction is weak, the viscous component spins down faster than the condensate throughout the domain.	70
2.6	In-plane meridional streamlines of the protons (top) and neutrons (bottom) with strong mutual friction at $t = 500$. Both components show meridional circulation, but the flows are distinct. In the protons the circulation is centered at $\theta \approx 30^\circ$ and $r \approx 0.5(R_1 + R_2)$, while for the neutrons the circulation cells are centered at $\theta \approx 20^\circ$ and $r \approx R_1$. This difference in flow patterns leads to the result shown in Figure 2.4, where v_{pn}^ϕ changes sign in the region $R_1 < r < R_2$ for strong mutual friction.	71
2.7	Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 100 time units following a crust glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.14) with $N = 2$ (see Table 2.1 for values). Mutual friction has little effect on the recovery time-scales.	74
2.8	Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 100 time units following a bulk glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.13) with $N = 2$ (see Table 2 for values). With strong mutual friction, the size of the glitch is significantly reduced, and there is an oscillatory component to the recovery for $t \lesssim 30$	77

2.9	Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 500 time units following an inner glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.14) with $N = 2$ (see Table 2.3 for values). In contrast to crust and bulk glitches, there is no relaxation following the spin up, $\Delta\Omega_{\max} = \Delta\Omega(t = 500)$. We see no significant difference in the recovery timescales between strong and weak mutual friction.	79
3.1	Convergence of the kernel density estimator $\hat{p}(x)$ [equation (3.1)] for an exponential PDF, $p(x) = e^{-x}$ (top row), and $\hat{p}[\log_{10}(x)]$ for the power law PDF, $p(x) = 0.2x^{-1.2}$ with $x \geq 1$ (bottom row). The solid black curve in the top row corresponds to the exponential PDF, The solid, colored curves in the left column correspond to $\hat{p}(x)$ and $\hat{p}[\log_{10}(x)]$ for realizations with $N = 25$ (blue), 100 (red), and 1000 (green). The tick marks in the panels in the left column indicate the abscissae of the data in the $N = 25$ (blue) realizations. The dashed black line in the bottom row shows the power law PDF. The right-hand column shows 10 instances of $\hat{p}(x)$ and $\hat{p}[\log_{10}(x)]$ with $N = 25$ (gold curves). Equation (3.1) is applied to the reflected data in the top row and the reflected base-10 logarithm of the data in the bottom row.	92
3.2	Kernel density estimates of a two-component Gaussian PDF $p(x) = 0.6\mathcal{N}(\mu = 4, \sigma = 1) + 0.4\mathcal{N}(\mu = 8, \sigma = 1)$, with $N = 25$ data points. The solid black curve shows the kernel density estimate with bandwidth h_0 chosen by the normal reference rule, the dotted grey curve is the kernel density estimate with bandwidth $h_0/2$ and the dot-dashed grey curve is with bandwidth $2h_0$. The dashed green curve shows the underlying PDF from which the data used to construct the estimator, shown as black tick marks along the x -axis, are drawn.	93

3.3	Kernel density estimates of the waiting time PDFs $\hat{p}(\Delta t)$ for the five most active glitchers, plotted on linear axes. Solid curves correspond to the nonparametric kernel density estimator with a normal reference bandwidth $h = h_0$. In the left-hand column, red dashed curves are parametric Poisson fits with $\lambda = \langle \Delta t \rangle^{-1}$ (maximum likelihood) computed directly by averaging the data. Grey dotted curves are Poisson fits with $\lambda = \lambda_{\pm}$ from the K-S analysis in Melatos, Peralta & Wyithe (2008). In the right-hand column, the red dashed curves are the kernel density estimator with bandwidth $h = h_0/2$, the blue dotted curve is with $h = 2h_0$. Tick marks on the horizontal axis indicate the abscissae of the raw data.	96
3.4	Kernel density estimates of the fractional size PDFs $\hat{p}(s)$, for the five most active glitchers, plotted on log-log axes. The solid black curve is $\hat{p}(s)$ using a normal reference bandwidth, $h = h_0$. In the left-hand column, the dashed red lines show the parametric power-law fit with a estimated by maximum likelihood. The grey dotted lines are power-law fits with $a = a_{\pm}$ from the K-S analysis in Melatos, Peralta & Wyithe (2008). In the right-hand column, the red dashed curves are $\hat{p}(s)$ with $h = h_0/2$, and the blue dotted curves are $\hat{p}(s)$ with $h = 2h_0$. Tick marks on the horizontal axis indicate the abscissae of the raw data.	98
4.1	Vortex-averaged angular velocity of N vortices, normalised by ω_2 , versus time, normalised by τ_N , for $2 \leq N \leq 10$. Top panel: vortex positions are initialised at the vertices of a regular polygon. Bottom panel: vortices are given a random radial displacement of $\pm 5 \times 10^{-2} R$ relative to the centroid of the initial polygon. Note: the time axis starts at $t/\tau_N = 0$ (100) in the top (bottom) panels.	112
4.2	Vortex positions (blue circles) and velocities (indicated by blue arrows) of the $N = 10$ ring initialised in a regular polygon at $t = 3\tau_{10}$ (top panel) and at $t = 50\tau_{10}$ before any radial perturbation at $t = 100\tau_{10}$ (bottom panel).	113

4.3	Convergence of the vortex array from an arbitrary initial configuration to a triangular array. Top panel: Hamiltonian $\mathcal{H}$ [equation (4.7)] versus time for 100 vortices initialised with random positions drawn from a uniform spatial distribution and evolved according to equation (4.1) (ignoring pinning and image vortices) for $0 \leq t \lesssim 10^4 T_0$, where T_0 is the rotation period of the container. Each simulation is initialised in the same configuration, with a different value of the drag angle $\phi \in \langle 0.01, 0.02, 0.05, 0.1, 0.2, 0.5 \rangle$ (see legend). Second panel: as for the top panel, but with the vortices initially situated equidistantly around a ring of radius R . Third panel: $\mathcal{H}$ versus t for 10 random uniform-density initial configurations with $\phi = 0.1$. Bottom panel: Vortex positions at the end of one of the simulations in the third panel.	115
4.4	Container angular velocity Ω_C , normalized by its initial value, as a function of time over 10^4 rotation periods. This simulation uses the default parameters listed in Section 4.5.1.	119
4.5	Vortex motion during an avalanche. Top five panels are snapshots at (top to bottom) $t/T_0 = 3079, 3086, 3093, 3104.5, 3116$. Grey dots show stationary vortices; red dots show vortices that have moved in the previous time step, black (fading to gray) tracers show the positions of the moving vortices for the 20 previous time steps. Bottom panel: container angular velocity Ω_C versus time t . Snapshots from the top five panels are marked with blue crosses.	120
4.6	Top panel: PDF of stress before (blue curve) and after (orange curve) the glitch, smoothed using a kernel density estimator. Effective stress, parametrized by the speed v_{stress} before (middle panel) and after (bottom panel) a vortex avalanche. Red dots indicate vortices with $v_{\text{stress}} > 0.5$	123
4.7	Kernel density estimates of the waiting time PDF from three simulations with the default parameters in Section 4.5.1. Black curve: aggregated data. Coloured curves: individual simulations. Grey, dashed line: maximum likelihood exponential fit. Grey, dot-dashed curve: maximum likelihood log-normal fit.	124
4.8	Kernel density estimates of the size PDF from three simulations with the default parameters in Section 4.5.1. Black curve: aggregated data. Coloured curves: individual simulations. Grey, dashed curve: maximum likelihood exponential fit. Grey, dot-dashed curve: maximum likelihood log-normal fit. Grey, dotted line: power law fit for $\Delta\Omega_C/\Omega_0 > 2 \times 10^{-4}$	124

4.9	Vortex motion during an avalanche in the weak dissipation simulation. Top four panels are snapshots at (top to bottom) $t/T_0 = 5487.5, 5510, 5520, 5570$. Grey dots show stationary vortices; red dots show vortices that have moved in the previous time step, black (fading to gray) tracers show the positions of the moving vortices for the 20 previous time steps. Bottom panel: container angular velocity Ω_C versus time t . Snapshots from the top four panels are marked with blue crosses.	128
5.1	Container angular velocity Ω_C , normalized by its initial value, as a function of time over 10^4 rotation periods.	139
5.2	Kernel density estimates of the glitch size and waiting time PDFs for spin up (blue curve) and spin down (orange dashed curve) simulations.	141
5.3	Snapshots from simulation of leapfrogging vortex rings. Initially, the rings have unit radius $R = 1$ and are separated by 0.1 in the z -direction. Each ring is discretized into 100 points. Time increasing from top to bottom (arbitrary units).	143
5.4	Radius R (averaged over all vortex points) of leapfrogging vortex rings as a function of time t over approximately five cycles. The orange curve shows R versus t for the ring which is initially above, blue curve shows same for the ring which is initially below.	144
5.5	Initial angular velocity of rectilinear N -filament array, $\omega_N(t = 0)$, averaged over all vortex points in filaments and all filaments in array, normalised by the initial value of the 2-filament array, $\omega_2(t = 0)$. Blue dots indicate data, black curve is a power law fit $\omega_N = \omega_2 N^{1.1}$	145
5.6	10-filament vortex array configurations at $t = 0$ (top panel) and $t = 0.13/T_2(t = 0)$ (bottom panel).	146
5.7	Top panel: angular velocity $\langle \omega_N \rangle$ (averaged over each point in filament and each filament in array) versus time for $N = 2, 3, 5, 9$ filament arrays. Bottom panel: radial displacement $\Delta R = R - 1$ versus time for $N = 2, 3, 5, 9$ filament arrays. R shown is maximum value of R for all vortex points in the array.	147
5.8	Oscillation frequency ν_N versus initial azimuthal velocity $\omega_{N,i}$ for $2 \leq N \leq 7$ vortex filament arrays. Blue dots indicate data, black curve shows a quadratic fit $\nu_n \propto -\omega_N^2$ to the data.	148
5.9	Number of vortices N vs time as a function of the annihilation threshold σ	150

5.10	Vortex configurations during Onsager vortex formation. Blue dots indicate positions of vortices with counter-clockwise circulation, orange dots indicate positions of vortices with clockwise circulation of equal magnitude. Annihilation threshold σ increasing from top to bottom. Left column shows configurations at beginning of simulations, right column shows same after $t = 1000T$	152
5.11	Number of vortices N , ensemble-averaged over five simulations, versus time with $\phi = 0$ (blue curve) and $\phi = 10^{-3}$ (orange curve). Grey dashed line is a fit to $N \propto t^a$ for the $\phi = 0$ (or $T = 0$) case, grey dotted line is same for $\phi = 10^{-2}$ ($T > 0$).	153

Chapter 1

Introduction

1.1 Pulsars

Neutron stars are dense stellar remnants formed in the death throes of massive stars. Stars support themselves against gravitational collapse by thermal pressure gradients created by thermonuclear reactions in their cores. In the latter stages of the lives of massive stars, lighter elements become exhausted in the core leading to gravitational contraction, which then ignites nuclear burning of heavier elements. Once iron begins forming, however, fusion of heavier elements is endothermic, and so gravitational collapse proceeds until protons and electrons begin undergoing inverse beta decay, forming neutrons. Neutrons, being fermions, cannot occupy the same quantum state, so as they become more densely packed they exert quantum degeneracy pressure upon one another, resisting further gravitational collapse (Shapiro & Teukolsky, 2008).

The existence of compact stellar remnants supported by neutron degeneracy pressure was theorised soon after the discovery of the neutron, as a similar process for electrons was understood to be responsible for the hydrostatic equilibrium of white dwarves (Landau, 1932). It was also quickly realised that neutron stars might be born in supernovae (Baade & Zwicky, 1934). There was theoretical interest in neutron stars over the following decades, but it was widely believed that observational detection of neutron stars would be challenging, if not impossible, as they exist as luminous X-ray sources for only a short time after formation (Chiu, 1964; Morton, 1964; Cameron, 1965; Bahcall & Wolf, 1965).

In the late 1960s, radio astronomers at Cambridge discovered a radio source that flashed every 1.337 s. This source originated outside of the solar system, and its high frequency ruled out a stellar binary (Hewish et al.,

1968). The subsequent discovery of more of these pulsating radio sources, or “pulsars”, at higher frequencies (Pilkington et al., 1968) also ruled out rotating white dwarves – centrifugal forces would cause a white dwarf rotating so quickly to break apart. Beamed radiation from rotating neutron stars became the prime candidate for explaining these pulsars (Gold, 1968). Further evidence for the rotating neutron star theory came from the discovery of pulsars in the Crab and Vela supernova remnants (Large, Vaughan & Mills, 1968; Radhakrishnan et al., 1969; Cocke, Disney & Taylor, 1969; Warner, Nather & Macfarlane, 1969).

The origin of the emission from pulsars is an active topic of research [see e.g. Melrose & Rafat (2017) and references therein for a discussion of proposed mechanisms]. The basic picture holds that the neutron star has a strong magnetic field with a significant dipole component whose axis is offset from the rotation axis of the star. Misalignment between the magnetic and rotation axes causes beamed emission from the magnetic poles, which sweeps past the Earth every time the star rotates, leading to the periodic pulse signal (Lyne & Graham-Smith, 2006).

Most pulsars discovered to date have been detected at radio frequencies, however, there are pulsars that have been observed across the electromagnetic spectrum, e.g. optical (Manchester et al., 1978; Pavlov, Stringfellow & Cordova, 1996; Zampieri et al., 2014) X-rays (Trümper & Becker, 1998; Rots, Jahoda & Lyne, 2004), and gamma-rays (Abdo et al., 2013). Pulsars are commonly characterised by their spin period P and the spin period derivative $\dot{P}$. From electrodynamical considerations, P and $\dot{P}$ can be used to estimate the age and magnetic field strength at the surface of the neutron star (Irvine, 1978). Pulsars are divided into three (sometimes more) categories based on their position in a $P - \dot{P}$ diagram [see Figure 1.1, taken from Kaspi & Kramer (2016)]. These are

- Slow pulsars, which comprise the majority of the observed pulsar population, having typical periods of $100 \text{ ms} \lesssim P \lesssim 1 \text{ s}$ and period derivatives $10^{-16} \lesssim \dot{P} \lesssim 10^{-12}$. These tend to be less than 10^6 years old, after which they become too faint to be detected (Chen & Ruderman, 1993).
- Millisecond pulsars, which have short periods $P \lesssim 10 \text{ ms}$ and $10^{-18} \lesssim \dot{P} \lesssim 10^{-20}$, and tend to be older ($10^8 - 10^9$ years old). Millisecond pulsars are almost always found in binaries, and their fast spin frequencies are thought to be due to accretion of angular momentum from a binary companion (Bhattacharya & van den Heuvel, 1991; Edwards & Bailes, 2001).




Figure 1.1: $P - \dot{P}$ diagram [from Kaspi & Kramer (2016)]. Red dots indicate pulsars. Blue circles indicate pulsars with binary companions. Stars indicate pulsars found in supernova remnants.

- Magnetars, which are young, slow pulsars with high spin-down rates ($P \gtrsim 1$ s, $\dot{P} \gtrsim 10^{-12}$) and strong magnetic fields ($B \sim 10^{14}$) most commonly detected in X-rays and soft γ -rays (Kaspi & Beloborodov, 2017). Magnetars are thought to be newly-born neutron stars whose spin-down is powered by decay of their magnetic fields, rather than electromagnetic braking (Mereghetti, Pons & Melatos, 2015).

1.2 Pulsar timing

While the first pulsar was discovered by direct observations of single pulses (Hewish et al., 1968), the vast majority of pulsars discovered to date [≈ 2800 (Manchester et al., 2005)¹] are too faint for each pulse to be seen. The shape and strength of individual pulses also varies significantly, so pulsar astronomers average the flux from 10 to 1000 consecutive pulses in order to produce a pulse profile, which remains stable on long time-scales (Helfand, Manchester & Taylor, 1975; Cordes & Shannon, 2010; Liu et al., 2011). Modern pulsar discovery campaigns use lengthy observations at fixed locations and search this data in the frequency domain for periodic signals (Manchester et al., 1996). The number of known pulsars has increased dramatically in recent decades thanks to the advent of multi-beam receivers, which allow for simultaneous measurement of large, disparate patches of the sky with high sensitivity (Manchester et al., 2001). Other innovations include the use of graphical processing units (GPUs), which allow for the efficient processing of large data volumes across a larger range of frequencies than previously feasible (Jankowski et al., 2019). Current and future radio observatories such as LOFAR (Sanidas et al., 2019), UTMOST (Jankowski et al., 2019), MeerKAT (Bailes et al., 2018), and the Square Kilometre Array (SKA) (Levin et al., 2018) will make large strides in both coverage and sensitivity, leading to the discovery of more pulsars and allowing for more frequent monitoring of the known pulsar population.

Pulsar timing experiments work by comparing the observed times of arrival (ToAs) of pulses at an observatory to a phenomenological model of expected ToAs. At the observatory, many (~ 1000) consecutive pulses are measured, then averaged in time, frequency and sometimes polarisation in order to create a stable pulse profile. The epoch of the pulse arrival is determined in the reference frame of the observatory, then timing models include corrections for the motion of the Earth relative to the solar system barycentre, the proper motion of the pulsar, and the orbital properties of

¹Also see online catalogue at <https://www.atnf.csiro.au/research/pulsar/psrcat/>.

the pulsar, e.g. effects due to binary companions or orbiting planets (Backer & Hellings, 1986). The residuals between observed and predicted ToAs are minimised using software such as `TEMP02` (Edwards, Hobbs & Manchester, 2006), and can be reduced to the order of ~ 10 ns for some pulsars (Arzoumanian et al., 2018; Parthasarathy et al., 2021). Such high-precision timing allows for decades-long phase coherence in measurements of pulsars that are monitored with week-to-month cadences. Highly accurate pulsar timing on such long time-scales allows for novel applications such as the development of time standards based on pulsars (Hobbs et al., 2020), and pulsar timing arrays: gravitational wave detectors based on measuring correlated residuals in a network of stable pulsars, to detect nHz gravitational waves from a stochastic background of merging supermassive black holes in the early Universe (Manchester et al., 2013; Rosado, Sesana & Gair, 2015; Perera et al., 2019).

1.3 Timing irregularities

Almost all pulsars are observed to be slowing down, due to the loss of angular momentum from a ‘wind’ of electromagnetic particles at the light cylinder (Goldreich & Julian, 1969; Ostriker & Gunn, 1969). Pulsar timing models tend to assume that the spin-down is steady, so that the rotational phase can be modelled using a Taylor series

$$\phi(t) = \phi(t_0) + \nu(t - t_0) + \frac{1}{2}\dot{\nu}(t - t_0)^2 + \frac{1}{2}\ddot{\nu}(t - t_0)^3 + \mathcal{O}[\ddot{\nu}t - t_0)^4], \quad (1.1)$$

where ν is the rotational frequency. Improving the timing model requires accurately measuring more frequency derivatives. Measuring $\ddot{\nu}$ and higher derivatives requires phase-connected timing over years (Johnston & Galloway, 1999), and in practice often yields contradictory results (Zhang & Xie, 2012). On these time-scales, most pulsars exhibit stochastic timing irregularities, which are grouped into two categories: timing noise and pulsar glitches.

1.3.1 Timing noise

Timing noise is time-correlated ‘red’ noise in the residuals that appears on a scale of weeks to months (D’Alessandro et al., 1995; Hobbs et al., 2004). The diverse phenomenology of timing noise in different pulsars makes consistent technical characterisation difficult, and there is no consensus on its physical origin (Cordes & Greenstein, 1981). It is likely that the term ‘timing noise’ encapsulates a host of different physical processes.

An early and enduring characterisation of timing noise is a random walk in one of the orbital parameters – either phase, frequency or frequency derivative (Boynton et al., 1972). Several analyses [e.g. Cordes & Downs (1985); D’Alessandro et al. (1995); Hobbs et al. (2004)] found that random walks alone are insufficient to describe timing noise in all pulsars, however, and suggest that other processes such as unresolved jumps in the frequency (Johnston & Galloway, 1999) or precession (Kerr et al., 2016) may be involved. While timing noise is present in all pulsars, it appears to be strongest in younger pulsars (Arzoumanian et al., 1994; Shannon & Cordes, 2010; Namkham, Jaroen-jittichai & Johnston, 2019; Lower et al., 2020)

Several theories for the physical origin of timing noise are superfluid vortex creep (Alpar, Nandkumar & Pines, 1986), superfluid turbulence (Melatos & Link, 2014), magnetospheric changes in the spin-down torque (Cheng, 1987; Price et al., 2012; Onuchukwu & Chukwude, 2016), and free precession (Jones, 2012).

1.3.2 Glitches

Glitches are sudden increases of the pulse frequency up to one part in 10^5 . The first glitch ever recorded was in the Vela pulsar in 1969 (Radhakrishnan & Manchester, 1969; Reichley & Downs, 1969), just two years after the discovery of pulsars, and since then around 10% of all known pulsars have been observed to glitch at least once (Espinoza et al., 2011).

Most glitches are observed as a secular drift in the residuals of pulsar timing experiments between two observations that can be resolved by including a step change in the frequency, often followed by a subsequent transient increase in the pulsar’s spin-down rate (Wang et al., 2000). The epoch and size of the glitch are determined by least-squares fitting to a phase-connected timing model between the observations [Though see recent work by Palfreyman et al. (2018), where a glitch in the Vela pulsar was observed in real time, and also Ashton et al. (2019), who used Bayesian parameter estimation to track the frequency evolution of the pulsar during the glitch].

The origin of glitches is an active topic of research, and its investigation forms the basis of this thesis. Generally speaking, it is believed that glitches are caused by a stochastic recoupling between the crust and other components of the neutron star (Haskell & Melatos, 2015). Two of the more enduring models are the ‘starquake’ model (Ruderman, 1969; Baym et al., 1969; Middleditch et al., 2006), and the superfluid vortex model (Anderson & Itoh, 1975; Alpar, 1977; Melatos, Peralta & Wyithe, 2008). The starquake model is based on the premise that newly-born neutron stars are rapidly rotating, so when the crust forms it is an oblate spheroid. As the star spins

down, the liquid interior becomes less oblate, causing stress to build up in the crust. At some point, the increasing stress causes catastrophic failure and reconfiguration of the crust. The new, more spherical crust has a lower moment of inertia, so conservation of angular momentum dictates that the spin frequency must increase, causing a glitch. The superfluid vortex model posits that neutron stars contain a superfluid, which forms an array of quantised vortices that pin to impurities in the interior. As the crust spins down, the pinned vortices remain in place, causing a lag to build up between the two components of the star. When vortices eventually unpin, they transfer angular momentum back to the crust, spinning it up and causing a glitch. The superfluid vortex model for glitches is discussed in greater detail in Section 1.4.

1.3.3 Glitch statistics

At the time of writing, the Jodrell Bank Centre for Astrophysics catalogue of pulsar glitches lists 555 glitches recorded in 189 pulsars over the last 50 years (Espinoza et al., 2011)². Of the glitching pulsars, 109 have only been observed to glitch once or twice, and the most recorded glitches in a single object is 45 in PSR J0537–6910 (Antonopoulou et al., 2018; Ferdman et al., 2018). Figure 1.2 shows a histogram of the number of detected glitches per pulsar.

It is unlikely that any glitch catalogue is complete. Many pulsars are observed infrequently, so a unique phase-connected timing solution cannot always be constructed between ToAs. If a glitch occurs in between two such observations, the epoch of the glitch cannot be determined, so it is impossible to distinguish between a single glitch and multiple, smaller glitches (Janssen & Stappers, 2006; Yu & Liu, 2017). Another issue is the detectability of small glitches, which may not be resolvable above the timing noise, or, if they occur shortly after a large glitch, may be absorbed into the fitted parameters of the larger glitch’s recovery (Wong, Backer & Lyne, 2001; Espinoza et al., 2014).

There is also uncertainty in the classification of small glitches. Some authors distinguish between regular (or ‘macro’) glitches and ‘micro-glitches’ [e.g. Cordes & Downs (1985); Chukwude & Urama (2010); Onuchukwu & Chukwude (2016)]. The remainder of this section contains a critical review of some of the literature on these issues, and Table 1.1 summarises which of the above issues is discussed in each of the papers. Chukwude & Urama (2010) make the distinction that regular glitches have amplitudes $\Delta\nu/\nu \gtrsim 10^{-9}$ and

²See <http://www.jb.man.ac.uk/pulsar/glitches.html> for an up-to-date list. The Australia Telescope National Facility also maintains a glitch catalogue at <https://www.atnf.csiro.au/research/pulsar/psrcat/glitchTbl.html>.




Figure 1.2: Histogram of number of glitches per pulsar. Data taken from <http://www.jb.man.ac.uk/pulsar/glitches.html> (Espinoza et al., 2011).

Reference	Time-resolution	Detectability	Classification
Janssen & Stappers (2006)	✓	✓	✓
Chukwude & Urama (2010)			
Espinoza et al. (2014)		✓	
Espinoza et al. (2021)		✓	
Yu & Liu (2017)	✓	✓	

Table 1.1: Summary of literature on glitch completeness and what aspects are discussed in each paper.

associated increases in the spin-down rate ($\Delta\dot{\nu} < 0$), while micro-glitches are smaller in amplitude ($|\Delta\nu/\nu| \lesssim 10^{-10}$), and can have either positive or negative changes in ν and $\dot{\nu}$. Chukwude & Urama (2010) and Onuchukwu & Chukwude (2016) performed an analysis on ≈ 15 yr of timing data for 26 radio pulsars with the Hartebeesthoek Radio Astronomy Observatory in South Africa, identifying 266 micro-glitches. However, it is not clear that macro-glitches are always accompanied with increases in the spin-down rate (Espinoza et al., 2011). It is also unclear if the distinction between these two classes of events is strictly adhered to in the literature. As an example, Janssen & Stappers (2006) found four small ($10^{-11} \lesssim \Delta\nu/\nu \lesssim 10^{-10}$) glitches in PSR J0358+5413 in six-years of well-sampled observing data. These glitches are recorded in the ATNF pulsar glitch catalogue (Manchester et al., 2005), and in Espinoza et al. (2011), but not the online Jodrell Bank glitch catalogue³.

Several authors have examined the issue of completeness in individual pulsars or in particular datasets. Janssen & Stappers (2006), in their analysis of glitches in PSR J0358+5413, performed Monte Carlo simulations to test their glitch detection method, and determined that they were able to detect glitches down to a threshold of $\Delta\nu/\nu \sim 10^{-11}$, about the size of the smallest glitch they detected in that pulsar. Janssen & Stappers (2006) also looked at the probability density function (PDF) of glitch size in the entire glitch population (Manchester et al., 2005), and found consistency with a power law PDF for $\Delta\nu/\nu \gtrsim 10^{-9}$, but when considering the entire population ($\Delta\nu/\nu \gtrsim 10^{-11}$), a KS test returns a p -value of $\approx 10^{-3}$ for consistency with a power law distribution. Janssen & Stappers (2006) suggested that the inconsistency with the power law PDF at small $\Delta\nu$ is due to the difficulty of detecting glitches with $\Delta\nu/\nu \lesssim 10^{-9}$, rather than a physical mechanism restricting the minimum glitch size. Espinoza et al. (2014) performed an analysis of 29 years of timing data for the Crab pulsar (PSR J0534+2200). They constructed a glitch detection algorithm that identified any changes in the timing solution between ToAs where there was a positive change in the spin frequency $\Delta\nu$ and a negative change in the frequency derivative $\Delta\dot{\nu}$. Using this algorithm, they identified all previously detected glitches, as well as several hundred candidate events with smaller $\Delta\nu$ than the previously detected glitches. They found a gap of around half a decade in $\Delta\nu$ between the largest of the glitch candidates and the smallest previously detected glitch. Espinoza et al. (2014) also found a similar number of ‘anti-glitch’ candidates

³Based on the data from Janssen & Stappers (2006), one expects ≈ 10 more glitches to have been detected in this pulsar, however, at the time of writing, no new glitches have been reported since April 2004.

(having negative $\Delta\nu$ and positive $\Delta\nu$) with similar $|\Delta\nu|$ to the glitch candidates, while the previously detected glitches are well-described by a power law size PDF, this is not the case when the glitch candidates are included. For these reasons, Espinoza et al. (2014) classified the glitch and anti-glitch candidates as timing noise, and used the gap between the population of timing noise jumps and the previously detected glitches as evidence of a minimum size for glitches in the Crab pulsar. A similar analysis for glitches in the Vela pulsar, however, found no such evidence of a minimum glitch size (Espinoza et al., 2021).

Yu & Liu (2017) performed a statistical analysis on the completeness of the data set from Yu et al. (2013), which discovered 107 glitches in 36 pulsars from ~ 20 yr of timing data from the Parkes radio telescope. Yu & Liu (2017) considered three potential causes for undetected glitches in the data:

1. at least three ToAs are needed to detect a glitch, so glitches that occurred between the first or final two ToAs in the time span of the observations are undetectable.
2. multiple glitches may have occurred between ToAs, and can't be distinguished from a single glitch.
3. Small glitches may not be distinguishable from timing noise.

Yu & Liu (2017) performed Monte Carlo simulations to construct synthetic timing series for each pulsar, added glitches to this data and searched the synthetic data for glitches using a nested sampling method to determine the glitch parameters. They used the results of these simulations to compare the expected PDF of glitch sizes to the observed size PDF, finding that the two PDFs mostly matched, except below $\Delta\nu/\nu \lesssim 10^{-7}$, where the PDF of simulated glitches was greater than the observed PDF, suggesting that perhaps not all small glitches were detected in the data from Yu et al. (2013).

Incompleteness in glitch catalogues may arise from insufficient monitoring of known glitching pulsars. We can look for incompleteness due to ‘missing’ glitches by comparing the rate of discovery of glitching pulsars over time to the detection rate of pulsar glitches. Two things are expected in such a comparison if the glitch catalogue is complete. Firstly, the rate at which glitches are detected should increase as more glitching pulsars are discovered, as previously discovered pulsars continue to glitch. Secondly, the rate at which glitches are detected should continue to increase, albeit slower, even if the rate at which new glitching pulsars are discovered decreases, for the same reason as above. Because glitches are a stochastic process, both effects are only expected to be observable on time scales longer than the average time




Figure 1.3: Cumulative number of pulsar glitches detected (blue curve) and number of glitching pulsars discovered (orange dashed curve) versus time. Dash-dotted black line is a linear fit to the glitch discovery $N \propto t$ for $52500 < \text{MJD} < 56000$. Dotted black curve is a quadratic fit to the glitch discovery $N \propto t^2$ for $45000 < \text{MJD} < 56000$. Data taken from the Jodrell Bank glitch catalogue <http://www.jb.man.ac.uk/pulsar/glitches.html> (Espinoza et al., 2011).

between observations for the most frequently glitching pulsars. In Figure 1.3, we show the cumulative number of detected pulsar glitches and also the cumulative number of glitching pulsars discovered as a function of time over the past 50 years. The data is taken from the discovery epochs recorded on the Jodrell Bank glitch catalogue, <http://www.jb.man.ac.uk/pulsar/glitches.html>, (Espinoza et al., 2011). We take the date of glitching pulsar discovery to be the epoch of the first detected glitch, because many non-glitching pulsars are not regularly monitored, so there may be an artificial lag between the discovery of a pulsar and the first glitch detection in that pulsar.

Figure 1.3 shows that the rate of glitch detection increased steadily during $\text{MJD} \lesssim 56000$. In the past 10 years, however, the glitch detection rate has decreased, from $\approx 20 \text{ glitches yr}^{-1}$ during $50000 \lesssim \text{MJD} \lesssim 56000$ to $\approx 10 \text{ yr}^{-1}$ during $56000 \lesssim \text{MJD} \lesssim 59000$. This suggests that a significant

number of glitches have been missed since $\text{MJD} > 56000$, possibly due to reduced follow up observations of pulsars discovered in the Parkes multi-beam survey after 2010⁴. We can estimate bounds on the number of missing glitches from Figure 1.3 in two ways. As a first assumption, we take the catalogue to be complete for $\text{MJD} \lesssim 56000$. To estimate a lower bound, we assume that no new glitching pulsars are discovered for $\text{MJD} \gtrsim 56000$, but known glitching pulsars continue to be monitored frequently enough so that all glitches are detected. The rate of glitch discovery after $\text{MJD} 56000$ should then be roughly constant, so that the number of glitches increases linearly with time. The dash-dotted line in Figure 1.3 shows a linear fit to the number of glitches based on the data from $52500 < \text{MJD} < 56000$. Based on this fit, the shortfall at the present day is ≈ 100 glitches. To estimate an upper bound, we assume the rate of discovery of glitching pulsars discovered is constant for $\text{MJD} > 45000$, so that the number of glitches detected increases quadratically with time. The dotted line in Figure 1.3 shows a quadratic fit to the glitch data from $45000 < \text{MJD} < 56000$, which suggests a shortfall of ≈ 200 glitches at the present day. Based on this crude analysis, we estimate that ~ 100 - 200 glitches have not been detected in the past ≈ 10 yr due to insufficient monitoring of known glitching pulsars.

Separately to the above discussion, we can also test for incompleteness by looking at whether small glitches are likely to have gone undetected, and whether there is a global minimum glitch size by looking at the PDFs of minimum and maximum glitch sizes in glitching pulsars. One out of many possibilities is that the minimum size is set by nuclear physics parameters tied to the stellar mass, which varies weakly across the pulsar population (Montoli, Antonelli & Pizzochero, 2020). If this is the case, the PDF of minimum glitch size should be peaked near this value. Figure 1.4 shows a histogram of minimum and maximum glitch sizes in the 80 pulsars with more than one recorded glitch.

Figure 1.4 shows that the maximum glitch PDF cuts off at $\Delta\nu/\nu \approx 10^{-5}$, and $\approx 50\%$ of all glitching pulsars have a maximum glitch size of $10^{-6} \lesssim \Delta\nu/\nu \lesssim 10^{-5}$. The minimum glitch size PDF peaks at $\Delta\nu/\nu \approx 10^{-9}$, but there is no apparent cut-off, $\approx 30\%$ of glitching pulsars have minimum glitch sizes $\Delta\nu/\nu < 10^{-9}$, and the smallest recorded glitch has $\Delta\nu/\nu = 2.5 \times 10^{-12}$ (McKee et al., 2016)⁵.

Figure 1.4 suggests there is a physical limit to the maximum glitch size. This is in keeping with theoretical predictions that predict that the maximum

⁴I thank the reviewer for bringing this fact to my attention.

⁵It is worth pointing out that this smallest glitch occurred in a millisecond pulsar, and thus may not be comparable to the vast majority of other glitches, which occur in slow pulsars.




Figure 1.4: Histogram of minimum and maximum glitch size in the 80 pulsars with > 1 detected glitch. Data taken from the Jodrell Bank glitch catalogue <http://www.jb.man.ac.uk/pulsar/glitches.html> (Espinoza et al., 2011).

lag between the superfluid and the crust that can be supported by pinning is $\sim 10^{-2} \text{ rad s}^{-1}$ (Haskell, Pizzochero & Sidery, 2012). There is little evidence in Figure 1.4 for a global minimum glitch size. This is consistent with the findings of Janssen & Stappers (2006); Yu & Liu (2017) discussed above, and may indicate a substantial population of undetected glitches with $\Delta\nu/\nu \lesssim 10^{-9}$. The claim from Espinoza et al. (2014) of a minimum glitch size in the Crab pulsar may be correct, but it is worth noting that the Crab is a particularly well-monitored object, and a similarly fastidious analysis is not possible for the majority of glitching pulsars.

Fortunately, current and upcoming radio instruments, such as UTMOST, MeerKAT, and SKA will be able to time many pulsars much more frequently and with greater sensitivity, so that in coming years many more glitches will be detected with greater completeness (Jankowski et al., 2019; Bailes et al., 2018; Levin et al., 2018). As well as improved instruments, improved glitch-finding algorithms are being developed that will ensure consistent classification of timing irregularities in pulsars (Melatos et al., 2020).

1.3.4 Rise time and recovery

Pulsar glitches are characterised by near-instantaneous rise times and lengthy recoveries. Recent observation of a glitch in the Vela pulsar which occurred as it was being monitored constrained the rise time to under 12 seconds with 90% confidence (Palfreyman et al., 2018; Ashton et al., 2019), a finding consistent with previous observations of Vela glitches (Dodson, McCulloch & Lewis, 2002; Dodson, Lewis & McCulloch, 2007).

A rapid spin-up time scale is consistent with both the crustquake and vortex avalanche models. In the crustquake model, the relevant time scale is the sound-crossing time for shear waves in the crust. In a neutron star crust the sound speed $c_s \approx 10^8 \text{ m s}^{-1}$, so the sound-crossing time is $\lesssim 10^{-3} \text{ s}$. In a vortex avalanche, vortices move on average one Feynman distance (or the typical separation between vortices), defined as $a = \sqrt{\kappa/2\pi\nu}$, where $\kappa = h/2m_n$ is the quantum of circulation and m_n is the mass of a neutron. If the vortex velocity is $R\Delta\Omega$, where $\Delta\Omega$ is the angular velocity lag between the crust and superfluid components, the time scale for this motion is $t_{\text{vortex}} \approx 10^{-6}(\nu/10 \text{ Hz})^{-1/2}(R/10^4 \text{ m})^{-1}(\Delta\Omega/0.5 \text{ rads})^{-1} \text{ s}$. The time-scale for the crust to spin up following a vortex avalanche, however, is the mutual friction time-scale, $t_{\text{MF}} = 1/(2\Omega B)$, where B is a constant parametrising the strength of the coupling between the crust and the superfluid (see Section 1.4.3 below). The upper limit on the rise time of the 2016 Vela glitch from Palfreyman et al. (2018); Ashton et al. (2019) can be used to constrain the value of B in Vela, which is determined by the mechanism of the crust-

superfluid coupling (Newton, Berger & Haskell, 2015; Graber, Cumming & Andersson, 2018; Sourie & Chamel, 2020).

Palfreyman et al. (2018) also observed a null in the pulse train of Vela coincident with the glitch. Bransgrove, Beloborodov & Levin (2020) proposed that this may have been due to a crustquake disturbing the pulsar’s magnetosphere and interrupting the radio emission.

Across the population of glitching pulsars, the response of the spin frequency following a glitch is diverse. Often, the instantaneous rise is followed by a temporary increase in the spin-down rate. This is often modelled as an exponential recovery (Shemar & Lyne, 1996; Wong, Backer & Lyne, 2001), so that the frequency residuals following the glitch are given by

$$\Delta\nu(t) = \Delta\nu_p + \Delta\dot{\nu}_p t + \sum_{i=1}^N \Delta\nu_i \exp(-t/\tau_i), \quad (1.2)$$

where $\Delta\nu_p$ is the permanent change in the spin frequency, $\Delta\dot{\nu}_p$ is the permanent change in the spin-down rate, and the $\Delta\nu_i$ terms are transient frequency jumps [positive or negative (van Eysden & Melatos, 2010; Basu et al., 2020)] which decay on time scales τ_i . In an analysis of 107 glitches in 36 pulsars, Yu et al. (2013) observed exponential recoveries in 27 glitches, on time scales $10 \text{ d} \lesssim \tau \lesssim 300 \text{ d}$. Yu et al. (2013) also point out that long recovery time scales may sometimes be mistaken for a permanent change in the spin-down rate, especially during observations soon after glitches. On the other hand, some glitches are observed as a simple step change in the frequency (Hobbs, Lyne & Kramer, 2010; Espinoza et al., 2011), which may be a time resolution issue due to exponential recoveries whose time scale is shorter than the gaps between observations (Dang et al., 2020). A recent analysis of 21 years of timing data in the Vela pulsar, perhaps the best-monitored glitching pulsar, used a Bayesian framework to characterise timing irregularities so that phase coherence was maintained across the entire span of the data (Shannon et al., 2016). Of the models considered, Shannon et al. (2016) found most support for glitches where there is a step change in the spin frequency, followed by a short ($\approx 25 \text{ d}$) and long ($\approx 1000 \text{ d}$) time scale exponential decay in the spin frequency, and no permanent or transient changes to the frequency derivative.

An enhanced spin-down rate following a glitch is expected in the vortex avalanche model. Glitches reduce the lag between the superfluid interior and the crust, decoupling some fraction of the superfluid from the spin-down torque (Baym, Pethick & Pines, 1969; Alpar, Langer & Sauls, 1984; Pines & Alpar, 1985). Haskell & Antonopoulou (2014), using a model from Alpar & Baykal (2006), showed that the recovery time scale may be determined by

where in the star vortices unpin, as the time for the superfluid to recouple to the crust differs in different regions of the interior.

An interesting phenomenon seen in some Crab pulsar glitches is ‘overshoot’, where following the exponential recovery the spin frequency then rises exponentially on a ~ 100 day time scale to return to the long-term spin-down rate (Wong, Backer & Lyne, 2001; Shaw et al., 2018). This behaviour is modelled as negative $\Delta\nu$ terms in equation (1.2) van Eysden & Melatos (2010) reproduced this effect, as well as the more conventional recovery of glitches in the Vela pulsar, using a two-component hydrodynamic model of the superfluid interior. van Eysden & Melatos (2010) found that the form of the recovery depends on the shear viscosity, the mutual friction strength (see Section 1.4.3 below) and the relative fraction of superfluid and charged components in the interior. Other authors have described ‘delayed spin up’ in some Crab glitches (Basu et al., 2020; Ge et al., 2020), where negative $\Delta\nu_i$ terms with $\tau_i \sim 1$ d cause a plateau in the spin frequency before the exponential recovery takes over. It is unclear whether these phenomena are related, and higher resolution timing, ideally pulse-to-pulse monitoring of the Crab pulsar during a glitch is needed to distinguish these effects. Ashton et al. (2019) found no evidence for a gradual rise in the 2016 Vela glitch.

So-called ‘slow glitches’ have been observed in several pulsars, where the spin frequency slowly increases over several months before relaxing back to its original state over a slightly longer time scale (Shabanova, 2005, 2010). The process repeats periodically, and the peak value of the glitch amplitude is the same each time. While curious, the lengthy rise time, periodic recurrence and uniform amplitude of these slow glitches strongly indicate that they are a separate, though possibly related, category of event and will not be discussed further in this thesis.

1.4 Superfluidity and its role in neutron stars

Superfluidity is a state of a liquid characterised by flow without viscosity. First discovered in liquid helium (Kapitza, 1938; Allen & Misener, 1938), it was proposed that a superfluid state may exist inside neutron stars even before their discovery (Migdal, 1959). Protons and electrons co-rotate with the rigid crust due to electromagnetic interactions. In the absence of superfluidity, the neutrons co-rotate with the protons through viscous interactions. If the neutrons are in a superfluid state, however, the interaction time-scale between the neutrons and the charged components of the neutron star is increased by many orders of magnitude (Baym, Pethick & Pines, 1969). A major consequence of superfluidity in a neutron star is that as the star spins

down, the superfluid component spins faster than the charged, viscous fluid components coupled to the crust. Differential rotation provides a significant energy and angular momentum reservoir that may be the cause of timing noise and glitches. This section reviews the observational evidence for superfluidity in neutron stars and gives an overview of models of superfluid hydrodynamics relevant to the superfluid vortex avalanche model of pulsar glitches.

1.4.1 Observational evidence for superfluidity in neutron stars

Following the first glitch observed in the Vela pulsar (Radhakrishnan & Manchester, 1969; Reichley & Downs, 1969), Baym et al. (1969) suggested that a superfluid in the interior could explain why the glitch was accompanied by an increase in the spin-down rate that persisted for ~ 1 year following the glitch. As discussed in Section 1.3.4 above, glitch recoveries are used extensively in theoretical studies of superfluidity in neutron stars, e.g. Alpar, Langer & Sauls (1984); Pines & Alpar (1985); van Eysden & Melatos (2010); Haskell & Antonopoulou (2014); Newton, Berger & Haskell (2015); Howitt, Haskell & Melatos (2016).

Another piece of evidence for superfluidity in neutron stars comes from measurements of X-ray cooling of neutron stars. Heinke & Ho (2010) measured the cooling rate of the neutron star in the Cassiopeia A supernova remnant, finding that between 2000-2010 it was higher than expected from theoretical cooling models calibrated to the current surface of the temperature of the star. Shternin et al. (2011); Page et al. (2011) proposed that the rapid cooling is due to a recent transition to a superfluid phase in the interior of the star, which causes increased cooling due to neutrino emission from the formation of Cooper pairs. More recent work (Posselt et al., 2013; Posselt & Pavlov, 2018) has called the original cooling results into question, however, finding that they depend on the model used for filter contamination on the Chandra X-ray detectors. The most recent observations (Wijngaarden et al., 2019), using 18 years of Chandra data, find a cooling rate lower than that reported in Heinke & Ho (2010), but still consistent with a recent transition to superfluidity.

1.4.2 Superfluid hydrodynamics

The key hallmark of superfluids is that flow occurs without viscosity. Landau (1941) suggested that this is because, at low temperatures, the minimum

energy of excitations is non-negligible. Below some critical flow speed, no dissipative excitations can be generated by the fluid. In the Landau model, a superfluid is described as two interpenetrating components: a viscous ‘normal’ component, and an inviscid ‘superfluid’ component. They are governed by the equations

$$\frac{\partial \mathbf{v}_s}{\partial t} + (\mathbf{v}_s \cdot \nabla) \mathbf{v}_s = -\frac{\rho_s}{\rho} \nabla p, \quad (1.3)$$

$$\frac{\partial \mathbf{v}_n}{\partial t} + (\mathbf{v}_n \cdot \nabla) \mathbf{v}_n = -\frac{\rho_n}{\rho} \nabla p + \nu_n \nabla^2 \mathbf{v}_n, \quad (1.4)$$

where $\mathbf{v}_s$ and $\mathbf{v}_n$ are the velocities of the superfluid and normal components respectively, ρ_s and ρ_n are the relative densities of the superfluid and normal components and $\rho = \rho_s + \rho_n$ is the total density, p is the pressure, and ν_n is the kinematic viscosity of the normal fluid. The relative densities are determined by the temperature. At 0 K, the fluid is in a pure superfluid state $\rho_s = \rho$. The normal fluid fraction increases with temperature up to a ‘transition temperature’ where $\rho_n = \rho$.

Because the superfluid component has no shear viscosity (or any other viscosity) it must be irrotational. However, experiments on cylinders of rotating superfluid helium (helium II) found that, above an angular velocity of $\approx 35 \text{ cm s}^{-1}$, the meniscus height is the same as that for a fluid undergoing solid body rotation (Osborne, 1950). The resolution to this paradox was realised by Onsager (1949) and expanded upon by Feynman (1955), who found that, while the bulk superfluid must be irrotational, the superfluid may be punctured by regions of zero density about which the fluid phase changes by 2π and whose circulation $\oint \mathbf{v}_s \cdot d\mathbf{l}$ is quantised in units of $\kappa = h/m$, where h is the Planck constant and m is the mass of the constituent bosons. These ‘vortices’ allow the superfluid to mimic solid-body rotation, at least on length scales larger than the average inter-vortex separation. The number of vortices N_v is determined by the angular velocity Ω by the Feynman condition

$$N_v \kappa = \oint \mathbf{v}_s \cdot d\mathbf{l} = 2\pi \Omega R^2 \quad (1.5)$$

The quantized circulation was first observed by Vinen (1961), by measuring the vibrational modes of a wire threaded through a cylindrical container of helium II. Later experiments using high-speed cameras and fluorescent tracer particles which are trapped by vortices were able to image vortices directly (Yarmchuk, Gordon & Packard, 1979; Yarmchuk & Packard, 1982).

1.4.3 Mutual Friction and the HVBK model

Superfluid vortices act as quasiparticles that interact with the normal fluid component, creating an effective drag force that couples the two components. Hall & Vinen (1956) developed a phenomenological description of this ‘mutual friction’ in helium II. Mutual friction arises from excitations in the normal fluid scattering off vortex lines, and acts both along and transverse to the velocity lag between the superfluid and viscous components. Curved vortex lines also experience a restoring tension force, which was incorporated into the model of Hall & Vinen (1956) by Bekarevich & Khalatnikov (1961). Incorporating mutual friction and vortex tension into equations (1.3) & (1.4) leads to the Hall-Vinen-Bekarevich-Khalatnikov (HVBK) equations (Barenghi & Jones, 1988; Peralta et al., 2005)

$$\frac{d\mathbf{v}_s}{dt} + (\mathbf{v}_s \cdot \nabla)\mathbf{v}_s = -\frac{\nabla p_s}{\rho} + \nu_s \mathbf{T} - \frac{\rho_n}{\rho} \mathbf{F} - \frac{\rho_s \nu_s}{\rho} \nabla |\boldsymbol{\omega}_s|, \quad (1.6)$$

$$\frac{d\mathbf{v}_n}{dt} + (\mathbf{v}_n \cdot \nabla)\mathbf{v}_n = -\frac{\nabla p_n}{\rho} + \nu_n \nabla^2 \mathbf{v}_n + \frac{\rho_s}{\rho} \mathbf{F} - \frac{\rho_s \nu_s}{\rho} \nabla |\boldsymbol{\omega}_s|, \quad (1.7)$$

where the effective pressures p_s and p_n are defined by $\nabla p_s = \nabla p - \rho_n \nabla (\mathbf{v}_n - \mathbf{v}_s)^2/2$ and $\nabla p_n = \nabla p + \rho_s \nabla (\mathbf{v}_n - \mathbf{v}_s)^2/2$, $\nu_s = \kappa/4\pi \log(b_0/a_0)$ is the stiffness parameter, where b_0 is the average inter-vortex spacing and a_0 is the vortex core radius, $\mathbf{T} = \boldsymbol{\omega}_s \times (\nabla \hat{\boldsymbol{\omega}}_s)$ is the vortex tension force, $\mathbf{F}$ is the mutual friction force, defined below, and $\boldsymbol{\omega}_s = \nabla \times \mathbf{v}_s$ is the superfluid vorticity⁶. For a rectilinear vortex array, the mutual friction is

$$\mathbf{F} = \frac{1}{2} B \hat{\boldsymbol{\omega}}_s \times [\boldsymbol{\omega}_s \times (\mathbf{v}_n) - \mathbf{v}_s - \nu_s \mathbf{T}] + \frac{1}{2} B' [\boldsymbol{\omega}_s \times (\mathbf{v}_n - \mathbf{v}_s) - \nu_s \mathbf{T}], \quad (1.8)$$

where B and B' are dimensionless coefficients that depend on the nature of the interaction between vortices and the normal fluid [see Barenghi, Donnelly & Vinen (1983) for a review on the determination of these parameters in helium II]. In neutron stars, the dominant cause of mutual friction is the scattering of electrons from magnetized vortex cores (Mendell, 1991). The magnetization is partly due to the magnetic moment of the neutrons (Sauls, Stein & Serene, 1982), but is primarily due to a process called entrainment (Andreev & Bashkin, 1976), whereby the velocity of the neutrons induces a proton current (Alpar, Langer & Sauls, 1984).

Equation (1.8) describes mutual friction for a regular vortex array undergoing laminar flow, but this is not necessarily the case in a neutron star

⁶Formally, $\boldsymbol{\omega}_s = 0$ everywhere except the centre of vortices. The hydrodynamic equations (1.6) & (1.7) are therefore only applicable on length scales where a fluid element contains many vortices and the vorticity can be treated as a continuous variable.

(Greenstein, 1970). Differential rotation of the superfluid and viscous components may trigger the Donnelly-Glaberson instability (Glaberson, Johnson & Ostermeier, 1974), leading to the formation of a vortex tangle (Peralta et al., 2006; Khomenko, Antonelli & Haskell, 2019). Laboratory experiments on turbulent counterflow in superfluid helium (Gorter & Mellink, 1949; Vinen, 1957) led to a phenomenological model for the mutual friction of an isotropic vortex tangle, known as the Gorter-Mellink mutual friction

$$\mathbf{F} = A_{\text{GM}} \frac{\rho_n \rho_s (\mathbf{v}_n - \mathbf{v}_s)^2}{\rho} (\mathbf{v}_n - \mathbf{v}_s) \quad (1.9)$$

The assumption of an isotropic vortex tangle in equation (1.9) may be inappropriate in the neutron star context. (Andersson, Sidery & Comer, 2007) discussed that a vortex tangle in a rotating neutron star may be polarised, and considered an extension to the Gorter-Mellink model based on laboratory experiments with helium II. Another approach, however, is to move to a hydrodynamical regime where the individual vortex dynamics are resolvable.

1.4.4 GPE and vortex filament model

Equations (1.6) and (1.7) are valid only in the limit where the shortest length-scale considered is much greater than the inter-vortex spacing, otherwise the continuous superfluid vorticity $\boldsymbol{\omega}_s$ is not well-defined. As discussed in Section 1.4.3 above, in some situations, such as superfluid turbulence or vortex pinning (of particular relevance to the vortex avalanche model of pulsar glitches), it is worthwhile to study vortices motion directly. Two popular methods of studying superfluid vortex dynamics are the Gross-Pitaevskii equation (GPE) and the vortex filament model.

The GPE is a quantum-mechanical description of weakly-interacting Bose gases at low temperature, in which quantized line vortices emerge as elementary excitations (Gross, 1961; Pitaevskii, 1961). It is a nonlinear Schrodinger equation for the order parameter $\psi(x, t)$, where $\rho_s \propto |\psi|^2$. Analytical studies by Fetter (1965) demonstrated that such vortices obey the Feynman condition. The GPE is most often used to study dilute Bose-Einstein condensates (BECs) in terrestrial contexts [see e.g. (Fetter & Svidzinsky, 2001; Barenghi & Parker, 2016)]. Several authors have also used the GPE in the neutron star case (Warszawski & Melatos, 2011, 2012; Melatos, Douglass & Simula, 2015; Lönnborn, Melatos & Haskell, 2019). While promising, a drawback of GPE simulations is computational expense, neutron stars contain $\sim 10^{18}$ vortices, but simulating systems of more than $\sim 10^2$ vortices is prohibitively resource-intensive.

In the limit where the inter-vortex separation is much greater than the vortex core size, quantized vortices can be treated with classical hydrodynamics (Bustamante & Nazarenko, 2015). This leads to the vortex filament model (Schwarz, 1985), in which vortices move with the local fluid velocity, which can be determined from the vorticity distribution by solving the Biot-Savart equation

$$\mathbf{v}_s(\mathbf{x}, t) = \frac{1}{4\pi} \int d^3\mathbf{x}' \frac{\boldsymbol{\omega}_s(\mathbf{x}') \times (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|^3}, \quad (1.10)$$

which is often computationally more tractable than full GPE simulations.

Equation (1.10) describes the superfluid velocity at an arbitrary point in space $\mathbf{x}$ given an arbitrary vorticity distribution $\boldsymbol{\omega}_s(\mathbf{x})$. In the vortex filament model, the vorticity is confined to line vortices parametrised by an arclength ξ . Any point along a vortex filament is specified by the vector $\mathbf{s}(\xi, t)$, relative to an arbitrary origin. The vorticity from a configuration of vortex filaments is given by

$$\boldsymbol{\omega}_s(\mathbf{x}, t) = \kappa \oint_{\mathcal{L}} d\xi \mathbf{s}'(\xi, t) \delta[\mathbf{x} - \mathbf{s}(\xi, t)], \quad (1.11)$$

where $\mathcal{L}$ denotes that the integral is over all filaments and $\mathbf{s}' = \partial\mathbf{s}/\partial\xi$. Rather than solve for the velocity at all points in the fluid, in the vortex filament model we evaluate (1.10) only at a set of discrete vortex points $\mathbf{s}_i$ along each filament, so we can rewrite (1.10)

$$\frac{\partial\mathbf{s}_i}{\partial t} = \mathbf{v}_s(\mathbf{s}_i, t) \quad (1.12)$$

$$= \frac{\kappa}{4\pi} \oint_{\mathcal{L}} d\xi \frac{\mathbf{s}'(\xi, t) \times [\mathbf{s}_i - \mathbf{s}(\xi, t)]}{|\mathbf{s}_i - \mathbf{s}(\xi, t)|^3}. \quad (1.13)$$

Equation (1.12) is the statement of the Helmholtz vortex theorem that vortex lines in an inviscid, unbounded fluid move with the local fluid velocity, while (1.13) comes from putting (1.11) into (1.10).

Two significant challenges in evaluating equation (1.13) numerically are that the integrand is divergent in the limit $\mathbf{s} \rightarrow \mathbf{s}_i$, and that the non-local nature of the integral means that the computational cost scales as $\mathcal{O}(N)^2$, where N is the number of vortex points. To deal with the first issue, following Schwarz (1985), the right-hand side of equation (1.13) is usually split into a local and non-local contribution,

$$\mathbf{v}_s(\mathbf{s}_i, t) = \frac{\kappa}{4\pi} \ln \left[\frac{(l_i l_{i+1})^{1/2}}{a_0} \right] \mathbf{s}'_i \times \mathbf{s}''_i + \frac{\kappa}{4\pi} \oint_{\mathcal{L}'} d\xi \frac{\mathbf{s}' \times (\mathbf{s}_i - \mathbf{s})}{|\mathbf{s}_i - \mathbf{s}|^3}, \quad (1.14)$$

where l_i is the arclength between $\mathbf{s}_{i-1}$ and $\mathbf{s}_i$ and l_{i+1} is the arclength between $\mathbf{s}_i$ and $\mathbf{s}_{i+1}$ respectively and the integral over $\mathcal{L}'$ excludes the region

between $\mathbf{s}_{i-1}$ and $\mathbf{s}_{i+1}$. A common solution to the issue of computational cost is to ignore the non-local second term in (1.14), called the local induction approximation (LIA) (Arms & Hama, 1965; Ricca, 1996; Strong & Carr, 2012; Horner & Van Gorder, 2019), which reduces the scaling to $\mathcal{O}(N)$.

A significant portion of this thesis is based on the application of the vortex filament method to the vortex avalanche model of pulsar glitches. However, I only consider the case of vortices in two dimensions, which simplifies the dynamics and eases the computational burden considerably. Several important features of vortex dynamics, including turbulence, vortex reconnection, and vortex tension do not apply in the two-dimensional case. These effects are likely important for the vortex avalanche model, and an extension of the model I describe in Chapter 4 to three dimensions would be a valuable continuation of this work and is discussed in Section 5.2.2.

Development of numerical vortex filament codes has primarily been motivated by the study of superfluid turbulence⁷. A thorough description of superfluid turbulence is outside the scope of this thesis [though see Nemirovskii (2013); Tsubota, Kobayashi & Takeuchi (2013); Barenghi & Parker (2016), and Mongiovì, Jou & Sciacca (2018)] for recent reviews, which also include discussion of numerical and theoretical studies of quantum turbulence using the GPE], but several theoretical and phenomenological relations are worth describing briefly, as the ability to reproduce these results is a key benchmark in assessing the performance of numerical schemes. The canonical scenario for turbulence in helium II is thermal counterflow (Gorter & Mellink, 1949). In a counterflow experiment, a cylindrical container of helium II, closed at one end and attached to a bath of helium at the other, is heated from the closed end. The temperature gradient drives a flow of the normal fluid component towards the bath, while the superfluid component flows in the opposite direction. Above some critical value of either the heat flux or the counterflow velocity $\mathbf{v}_n - \mathbf{v}_s$, the flow becomes turbulent, due to the interaction of the superfluid vortices and the normal component. A statistical measure useful in discussing superfluid turbulence is the vortex line density, $L = \Omega^{-1} \oint_{\mathcal{C}} d\xi$. In a homogenous counterflow experiment, the evolution of L is governed by Vinen’s equation (Vinen, 1957)

$$\frac{dL}{dt} = \frac{\chi_1 B \rho_n}{2\rho} (\mathbf{v}_n - \mathbf{v}_s)^2 L^{3/2} - \chi_2 \frac{\kappa}{2\pi} L^2, \quad (1.15)$$

where χ_1 is a constant, B is the same temperature-dependent coefficient in equation (1.8) and χ_2 is a temperature-dependent coefficient. The first term

⁷The term ‘quantum turbulence’, which also refers to turbulence in systems such as Bose-Einstein condensates as well as superfluids such as helium II is often used interchangeably with superfluid turbulence

Numerical method	References
LIA	Schwarz (1985, 1988)
LIA + Navier-Stokes	Idowu et al. (2000, 2001)
Full Biot-Savart	Adachi, Fujiyama & Tsubota (2010); Kondaurava & Nemirovskii (2012); Andryushchenko, Kondaurava & Nemirovskii (2016)
Biot-Savart + Navier-Stokes	Yui, Tsubota & Kobayashi (2018); Yui et al. (2020)
Tree code	Baggaley & Barenghi (2011, 2012)

Table 1.2: Summary of numerical techniques for studying quantum turbulence with the vortex filament method.

in (1.15) describes the generation of vortices through mutual friction and the second term describes the decay of vortices. In the steady state, the vortex line density is

$$L = \gamma^2(\mathbf{v}_n - \mathbf{v}_s)^2, \quad (1.16)$$

where $\gamma = \pi B \rho_n \chi_1 / \kappa \rho \chi_2$, a result which has been verified by laboratory experiments (Vinen & Niemela, 2002). Another important test of numerical schemes is the fluctuation spectrum of L , which has been shown to obey a $f^{-5/3}$ frequency dependence in experiments with turbulent superfluid ^4He and ^3He .

The vortex filament method originated in the study of classical hydrodynamics [e.g. Arms & Hama (1965)], but the first application of the vortex filament method to study superfluid vortices was by Schwarz (1985), who outlined the equations of motion (1.14) for quantized vortex filaments, and described a numerical scheme for solving them. Table 1.2 summarises the different numerical techniques that have been used to study quantum turbulence using the vortex filament method, and some references for each technique. Schwarz (1985) also includes numerical prescriptions for the interaction of vortices with planar and spherical boundaries, a reconnection procedure for vortices which intersect with either themselves or another vortex filament, and the interaction of vortices with a viscous normal fluid [usually assumed to be homogenous, though see Idowu et al. (2001); Yui, Tsubota & Kobayashi (2018), who performed simulations of vortex filaments coupled to a normal fluid obeying the Navier-Stokes equation]. In a follow-up paper, Schwarz (1988) performed simulations of counterflow which agreed well with equation

(1.16). These simulations used the LIA, however, and required an artificial mixing procedure to produce a homogenous vortex tangle. Adachi, Fujiyama & Tsubota (2010) repeated the simulations of Schwarz (1988) using both the LIA and the non-local Biot-Savart equation (1.14). They encountered the same issues with the LIA, but simulations with the full Biot-Savart equation were able to produce a homogeneous vortex tangle and reproduce equation (1.16). Other numerical studies using the full Biot-Savart equation have looked at the decay of vortex tangles Kondaurova & Nemirovskii (2012); Kondaurova et al. (2014), reconnection dynamics Andryushchenko, Kondaurova & Nemirovskii (2016), and vortex stretching and structure formation (Hänninen & Baggaley, 2014).

In order to ease the computational burden of computing (1.13), Baggaley & Barenghi (2011, 2012) developed a code that uses a tree algorithm (Barnes & Hut, 1986), a technique commonly used in N -body cosmological simulations [e.g. Springel (2005)]. Using a tree algorithm, the contribution to the velocity at a single point along a vortex filament is calculated by averaging the contribution of the rest of the points in a hierarchical fashion, the averaging becoming more coarse for more distant points. This technique reduces the computational complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$, with a small penalty in accuracy. Baggaley & Barenghi (2011) found that for a simulation with 2000 vortex points, the velocity calculated with a tree algorithm differed from a calculation with the full Biot-Savart law by $\lesssim 0.25\%$, and, using the tree algorithm, they were able to increase their simulations from $\sim 10^3$ to $\sim 10^4$ vortex points. Using the tree code, Baggaley & Barenghi (2011) examined the fluctuation spectrum of the vortex line density L in the steady state of a vortex tangle, and their results matched the $f^{-5/3}$ spectrum found in experiments on ^4He and ^3He .

1.5 Physics of pulsar glitches

Pulsar glitches are one of the few observable phenomena that can test the physics of the neutron star crust and interior. Table 1.3 outlines the different aspects of neutron star physics that can be examined with studies of pulsar glitches, and the relevant observables which can test each. The table is not exhaustive, but is focused towards those topics that are examined later in this thesis. This section briefly reviews each row in Table 1.3 with relevant references.

Physics	Rise times	Recoveries	Long-term statistics
Crust-core coupling; Mutual friction	✓	✓	✓
Vortex pinning		✓	✓
Superfluid turbulence; Hydrodynamic instabilities	✓	✓	
Crustquakes; breaking strain			✓
Magnetic effects; Superconductivity		✓	✓
Avalanche dynamics; self-organized criticality			✓

Table 1.3: Neutron star physics testable with studies of pulsar glitches.

1.5.1 Crust-core coupling and mutual friction

Neutron stars are thought to be stratified. The outermost layer is a crystalline solid consisting of elements such as iron and nickel. Deeper inside, pressure causes protons and electrons to form neutrons through inverse beta decay, leading to more neutron-rich isotopes in the crystal lattice. Eventually, the lattice structure dissolves into a mixture of a neutron fluid and a proton-electron plasma (Shapiro & Teukolsky, 2008; Chamel & Haensel, 2008). The proton-electron plasma is electromagnetically coupled to the crust, but if the neutrons are superfluid, their spin-down lags the crust, leading to differential rotation between the crust and superfluid (Baym et al., 1969). The maximum size of the lag depends on the strength of mutual friction between the superfluid and viscous components, which is determined by where a superfluid phase transition occurs in the star [Alpar, Langer & Sauls (1984); Andersson et al. (2012), also see discussion in Section 1.3.4]. It can be estimated from measurements of glitch rise times and recoveries (van Eysden & Melatos, 2010; Haskell & Antonopoulou, 2014; Graber, Cumming & Andersson, 2018). Long-term measurements of glitch activity can also be used to measure the fraction of the moment of inertia that is decoupled from the crust, a proxy for the lag between superfluid and viscous components

(Link, Epstein & Lattimer, 1999; Pizzochero et al., 2017; Montoli, Antonelli & Pizzochero, 2020).

1.5.2 Vortex pinning

Superfluid vortices resist spin-down by ‘pinning’ to nuclear lattice sites in the crust (Anderson & Itoh, 1975; Alpar, 1977). As the angular momentum of the superfluid is determined by the spatial distribution of vortices (Fetter, 1965), where vortices pin affects the maximum differential rotation and hence the maximum glitch size (Haskell, Pizzochero & Sidery, 2012; Hooker, Newton & Li, 2015; Seveso et al., 2016; Gügercinoğlu & Alpar, 2019). Glitches may be triggered by unpinned vortices causing knock-on unpinning of neighbouring vortices in an avalanche process (Melatos, Peralta & Wyithe, 2008; Warszawski, Melatos & Berloff, 2012; Warszawski & Melatos, 2013). The potential for vortices to cause knock-on also depends on where pinning occurs (Haskell & Melatos, 2016), which may be observable from glitch size distributions (Cheng et al., 1988; Link & Epstein, 1991) and glitch recoveries (Alpar, Cheng & Pines, 1989).

1.5.3 Superfluid turbulence and hydrodynamic instabilities

It is an open question whether superfluid vortices in a neutron star form a straight vortex array or a vortex tangle (Melatos & Peralta, 2007; Andersson, Sidery & Comer, 2007). Mastrano & Melatos (2005) argued that the Reynolds number is high enough that a build-up of differential rotation can excite a Kelvin-Helmholtz instability at the crust-core interface, leading to superfluid turbulence and potentially providing a trigger mechanism for glitches. Peralta et al. (2006) showed that a vortex tangle formed by differential rotation relaxes back to a laminar vortex array following a sudden spin-up of the crust, i.e. a glitch, and suggested that the transition between the two states may be observable in the recovery from glitches. Similarly, Mongiovi, Russo & Sciacca (2017) modelled the evolution of the vortex line density in a decelerating neutron star and showed that a transition from a straight vortex array to a vortex tangle acts to bring the crust and the superfluid into co-rotation, successfully fitting their model to glitches in the Vela pulsar. Haskell, Antonopoulou & Barenghi (2020) also found evidence for a vortex tangle by comparing hydrodynamical simulations with and without turbulence to measurements of glitch recoveries in Vela and PSR J0537–6910.

1.5.4 Crustquakes

In the crustquake model of glitches, differential rotation of the crust and superfluid leads to catastrophic failure and reconfiguration of the crust. The new, less oblate crust spins up in order to conserve angular momentum (Ruderman, 1969; Alpar et al., 1996). The crustquake model requires that the stress built up between glitches is enough to overcome the breaking strain of the crust. However, recent molecular dynamics simulations suggest that the breaking strain is large, $\sigma_c \approx 10^{-1}$, (Horowitz & Kadau, 2009). For older pulsars such as Vela, this suggests a typical time between glitches $\sim 10^5$ yr, much longer than the 2 – 3 year observed (Baym & Pines, 1971; Haskell & Melatos, 2015; Hooker, Newton & Li, 2015). The crustquake model may still be viable for younger, more rapidly spinning-down pulsars such as the Crab and J0537–6910 (Middleditch et al., 2006; Akbal & Alpar, 2018; Ferdman et al., 2018; Antonopoulou et al., 2018). Longer-term monitoring of glitching pulsars, as well as improved modelling of crust failure mechanisms will further test the viability and limitations of the model (Franco, Link & Epstein, 2000; Jones, 2003; Giliberti et al., 2020).

1.5.5 Magnetic effects

While not considered explicitly in this thesis, magnetic effects likely play a significant role in superfluid vortex dynamics in neutron stars (Jones, 1991; Graber et al., 2015). If the protons in the core form a type II superconductor vortices can be pinned by an array of magnetic flux tubes (Jones, 2006; Glampedakis & Andersson, 2011). Turbulence is more likely for pinning to flux tubes than nuclear lattice sites (Link, 2012a,b), particularly if the magnetic field axis is misaligned with the rotation axis (Drummond & Melatos, 2017, 2018). (Glampedakis & Andersson, 2011) suggested that pinning to flux tubes can occur early in the life of a neutron star, but becomes energetically unfavourable on a time scale much shorter than the spin-down time scale of typical radio pulsars. Haskell, Pizzochero & Sidery (2012); Haskell, Pizzochero & Seveso (2013) found that strong pinning of vortices to flux tubes reduces the angular momentum reservoir of the superfluid, and were unable to fit the size and recovery time scales of glitches in Vela with models that included pinning in the core.

1.5.6 Avalanche dynamics and self-organized criticality

The vortex avalanche model of pulsar glitches has many similarities with systems exhibiting self-organized criticality (SOC) (Jensen, 1998), where microscopic threshold-driven dynamics determine global properties of the system. Some frequently-glitching pulsars appear to have exponential waiting time distributions and power-law size distributions (Morley & García-Pelayo, 1993; Melatos, Peralta & Wyithe, 2008; Espinoza et al., 2011; Howitt, Melatos & Delaigle, 2018), a typical characteristic of SOC systems such as sandpiles (Rosendahl, Vekić & Kelley, 1993; Lőrincz & Wijngaarden, 2007), earthquakes (Sornette & Sornette, 1989; Bak et al., 2002), solar flares (Wheatland, 2000; Watkins et al., 2015) and flux tube avalanches in superconductors (Field et al., 1995; Aegerter, 1998). Similar avalanche behaviour is seen in Gross-Pitaevskii simulations of ~ 100 vortices (Warszawski & Melatos, 2011, 2013).

1.6 Thesis outline

The remainder of this thesis is structured as follows.

- Chapter 2 describes simulations of pulsar glitch recovery using a hydrodynamics code that solves the HVBK equations (1.7) and (1.6) in a spherical Couette geometry. Simulations are performed in the strong and weak mutual friction regimes using a variety of glitch mechanisms, and the results are analyzed in the context of observational studies of pulsar glitch recoveries and previous analytical results and numerical simulations.
- Chapter 3 is a statistical study of the glitch size and waiting time probability density functions (PDFs) for the five most prolific glitching pulsars. We use a kernel density estimator to construct the size and waiting time PDFs based on data from the Jodrell Bank glitch catalogue (Espinoza et al., 2014) and discuss the results in the context of the vortex avalanche model for pulsar glitches.
- Chapter 4 describes an N -body code for vortex motion in two dimensions in the presence of pinning, dissipation and a decelerating boundary. We describe the numerical method, present tests of the code based on and perform simulations vortex avalanches in an array of 2000 pinned vortices in a decelerating container. We examine the size

and waiting time PDFs of vortex avalanches in simulations with a wide range of input parameters and compare the results to previous GPE simulations with ~ 100 vortices and observations of glitching pulsars.

- Chapter 5 summarises the results of the previous chapters and discusses future work to be done with the N -body vortex code in a variety of contexts, including preliminary results.

Chapter 2

Hydrodynamic simulations of pulsar glitch recovery

2.1 Introduction

Radio pulsar glitches are sporadic jumps in the spin frequency of pulsars which occur against a background of steady electromagnetic spin down. During a glitch, the spin frequency increases by as much as one part in 10^5 (Manchester & Hobbs, 2011), over a time-scale unresolvable by radio timing experiments (Dodson, Lewis & McCulloch, 2007; Palfreyman et al., 2018). Glitches are usually followed by a recovery period, during which some (though not necessarily all) of the increase in spin frequency is reversed, and the pulsar returns to a state of steady electromagnetic spin down (Wong, Backer & Lyne, 2001; Yu et al., 2013). The impermanent part of the glitch is often fitted by multiple decaying exponentials, which can have characteristic timescales ranging from minutes to weeks for a single glitch (Dodson, McCulloch & Lewis, 2002). Often, glitches are also followed by a change in the spin-down rate, which can persist for as long as the time between glitches (Yu et al., 2013).

Glitches provide a window into the interior of neutron stars and inform theoretical models of bulk nuclear matter (Link, Epstein & Lattimer, 1999; van Eysden & Melatos, 2010). Since the work of Baym, Pethick & Pines (1969) and Anderson & Itoh (1975), it has been widely believed that glitches originate in the superfluid interior of a neutron star; see Haskell & Melatos (2015) for a recent review of glitch models. The temperature is likely low enough for neutrons to form Cooper pairs and hence an inviscid Bose-Einstein condensate in certain regions of the interior (Baym, Bethe & Pethick, 1971). Direct evidence for superfluidity has come from recent observations of the

young neutron star in the supernova remnant Cassiopeia A, whose current temperature is too high if its current cooling rate has been maintained since birth. A recent transition to a superfluid state in the interior, leading to enhanced cooling from neutrino emission, seems to be implied (Heinke & Ho, 2010; Shternin et al., 2011; Page et al., 2011; Elshamouty et al., 2013).

A superfluid in a rotating container forms an array of vortices with quantized circulation, whose configuration determines the angular velocity of the superfluid as a whole (Tilley & Tilley, 1990). As the pulsar spins down due to electromagnetic braking, there is a hydrodynamical lift force, called a Magnus force, which pushes the vortices out of the star (Seveso, Pizzochero & Haskell, 2012). In the absence of obstructions, the Magnus force keeps the superfluid neutrons in corotation with the charge-neutral electron-proton fluid, which in turn is kept in corotation with the crust by the strong magnetic field. As the superfluid vortices pass through the crystalline crust, however, it is energetically favourable for vortex cores to overlap with crustal ions, which means that they ‘pin’ to lattice sites in the crust and decouple from the smoothly decelerating proton-electron fluid (Alpar, 1977; Anderson et al., 1982; Seveso et al., 2016). Vortices also pin to magnetic flux tubes in the core (Mendell, 1991; Link, 2012a), due to the superconducting nature of the protons in that region (Migdal, 1959; Page, Geppert & Weber, 2006). When a single vortex unpins, it can knock-on surrounding vortices before it repins, causing them to unpin as well — a vortex avalanche (Cheng et al., 1988; Warszawski & Melatos, 2013). The dynamics of such avalanches are well-studied in terrestrial systems, such as sandpiles and forest fires (Bak, Tang & Wiesenfeld, 1987; Turcotte, 1999), and also in astrophysical contexts, such as solar flares [see Watkins et al. (2015) for a review]. Recent studies of the statistical distributions of glitch sizes and waiting times between glitches find consistency between pulsar glitch data and the superfluid vortex avalanche model (Melatos, Peralta & Wyithe, 2008; Warszawski & Melatos, 2011, 2013; Melatos, Douglass & Simula, 2015).

The dynamics of individual vortices is highly complex and has been studied extensively for terrestrial superfluids such as ^4He and Bose-Einstein condensates (Donnelly, 1991; Fetter, 2009). Such systems can be modelled microscopically using the Gross-Pitaevskii equation (GPE). While the GPE approach is feasible for small systems, it does not scale well to neutron stars, which contain $\gtrsim 10^{15}$ vortices. Moreover, the GPE describes weakly interacting systems such as Bose-Einstein condensates, rather than strongly interacting fermionic systems such as neutron stars. On the other hand, the average inter-vortex spacing is much smaller than the neutron star radius, so a neutron star superfluid can be described in the continuum, or hydrodynamic, limit. Each fluid element contains many vortices, so the rapid spatial

variation of the superfluid velocity field around the vortices is smoothed, and the superfluid vorticity can be defined as a continuous quantity (Hall & Vinen, 1956; Hills & Roberts, 1977; Andersson, Sidery & Comer, 2006).

In this chapter, we simulate the response of a model neutron star to glitches hydrodynamically. We start by describing the equations of motion for a two-component superfluid (Section 2.2), then describe the system geometry and numerical method (Section 2.3). We prepare the system in a state of differential rotation (Section 2.4), then investigate the response of the system to glitches induced in three different ways in Section 2.5. Firstly, we spin up the ‘crust’ (Section 2.5.3); secondly, we recouple the two fluid components in the bulk (Section 2.5.4); and finally, we spin up the core (Section 2.5.5). We study all three types of glitches in regimes where the coupling between the two fluid components, called mutual friction, is either strong or weak. The response of the spin frequency to a glitch is dramatically different for each of the three glitch types. We find that mutual friction plays an important role in those glitches which originate in the bulk. The observed behaviour is discussed critically in Section 2.6.

2.2 Equations of motion

We model the neutron star as a system of two coupled fluids: a neutron condensate, labelled n, and a charge-neutral fluid of protons and electrons, labelled p. For the purpose of studying glitch relaxation this is an adequate description (Sidery, Passamonti & Andersson, 2010; Haskell, Andersson & Comer, 2012; Haskell & Antonopoulou, 2014), as electrons can be considered locked to the protons on length-scales larger than the electron screening length and time-scales longer than the inverse of the plasma frequency (Mendell, 1991; Andersson, Haskell & Samuelsson, 2011). We have the usual conservation laws for the number densities n_x ,

$$\partial_t n_x + \nabla_i (n_x v_x^i) = 0, \quad (2.1)$$

where $x = n, p$ indexes the constituent, i labels Cartesian components, and we adopt the Einstein convention of summing over repeated indices. The two momentum equations can be written as (Prix, 2004)

$$\begin{aligned} & (\partial_t + n_x^j \nabla_j) (v_i^x + \varepsilon_x v_i^{yx}) + \nabla_i (\tilde{\mu}_x + \Phi) + \varepsilon_x v_{yx}^j \nabla_i v_j^x \\ & = (f_i + \nabla^j D_{ij}^x) / \rho_x, \end{aligned} \quad (2.2)$$

where ε_x is the entrainment coefficient, $\tilde{\mu}_x = \mu_x / m_x$ is the chemical potential μ_x scaled by the mass m_x of each component, $v_i^{yx} = v_i^y - v_i^x$ is the relative

flow velocity, and Φ is the gravitational potential. Viscous terms are encoded in the tensor D_{ij}^x ; they are more numerous than in a single Newtonian fluid, given the additional degrees of freedom of a multi-fluid system (Andersson, Sidery & Comer, 2006; Haskell, Andersson & Comer, 2012). The forces f_i on the right-hand side represent various other interactions between the fluids, such as the Lorentz force or, as we see below, the mutual friction force.

To simplify the problem we take both fluids to be incompressible and consider only the shear viscosity acting on the electron-proton ‘normal’ fluid (Haskell, Andersson & Comer, 2012). We also neglect the effect of entrainment ($\epsilon_x = 0$). In the crust the latter can actually be quite a poor approximation, as entrainment coefficients can be large (Chamel, 2012; Andersson et al., 2012; Chamel, 2013; Newton, Berger & Haskell, 2015). However, in the outer core of the star it is a good approximation (Carter, Chamel & Haensel, 2006).

Even with these simplifications, the problem is challenging numerically. The flow at every point is characterised by four quantities: the velocity fields of the proton-electron fluid $\mathbf{v}_p$ and the superfluid neutrons $\mathbf{v}_n$ and their respective chemical potentials, $\tilde{\mu}_p$ and $\tilde{\mu}_n$. The gravitational potential Φ is taken to be constant and absorbed in the chemical potential terms. In the isothermal, incompressible, constant density regime, the flow is described by the dimensionless equations⁸

$$\frac{d\mathbf{v}_p}{dt} + (\mathbf{v}_p \cdot \nabla)\mathbf{v}_p = -\nabla\tilde{\mu}_p + \frac{1}{Re}\nabla^2\mathbf{v}_p + \frac{1}{\rho_p}\mathbf{F}, \quad (2.3)$$

$$\frac{d\mathbf{v}_n}{dt} + (\mathbf{v}_n \cdot \nabla)\mathbf{v}_n = -\nabla\tilde{\mu}_n - \frac{1}{\rho_n}\mathbf{F}, \quad (2.4)$$

$$\nabla \cdot \mathbf{v}_p = \nabla \cdot \mathbf{v}_n = 0, \quad (2.5)$$

where ρ_p and ρ_n are the relative densities of protons and neutrons (normalised so that $\rho_p + \rho_n = 1$), and Re is the Reynolds number. The mutual friction force, $\mathbf{F}$, is a term that describes the interaction between the proton and neutron fluids, which arises primarily from the scattering of electrons off vortex cores (Alpar, Langer & Sauls, 1984). It is a combination of the Magnus and drag forces of the form (Hall & Vinen, 1956)

$$\mathbf{F} = \rho_n B \hat{\boldsymbol{\omega}} \times (\boldsymbol{\omega} \times \mathbf{v}_{pn}) + \rho_n B' \boldsymbol{\omega} \times \mathbf{v}_{pn}, \quad (2.6)$$

where $\boldsymbol{\omega} = \nabla \times \mathbf{v}_n$ is the superfluid vorticity, $\hat{\boldsymbol{\omega}}$ is the vorticity unit vector and $\mathbf{v}_{pn} = \mathbf{v}_p - \mathbf{v}_n$ is the velocity lag between the two fluids, written as

⁸By neglecting the entrainment and all viscous terms but the shear viscosity of the ‘normal’ fluid, the second-rank tensor equation (2.2) is reduced to a vector equation. From here on, all equations are written using conventional vector notation.

$v_i^{yx} = v_i^y - v_i^x$ in equation (2.2). The parameters B and B' are related ($B' \approx B^2$) dimensionless constants (Andersson & Comer, 2006). Vortex lines have a tendency to resist bending, which results in a tension force which can be included in (2.6) (Hills & Roberts, 1977; van Eysden, 2015). For simplicity, we ignore the effect of vortex tension, c.f. Peralta et al. (2008). We also neglect vortex tangles and their effect on (2.6) (Gorter & Mellink, 1949; Peralta et al., 2006; Andersson, Sidery & Comer, 2007).

Equations (2.3)–(2.5), although derived to describe a condensate coupled to a viscous fluid in a neutron star, are formally identical to the Hall-Vinen-Bekarevich-Khalatnikov (HVBK) equations generally used to describe a (laboratory) condensate coupled to its thermal excitations (Hall & Vinen, 1956; Chandler & Baym, 1986), thus allowing us to make use of numerical schemes developed for the HVBK formalism (Henderson & Barenghi, 1995; Peralta et al., 2008). In this context we note that the superfluid pairing gaps are density dependent in a neutron star, so there may be regions in which the superfluid transition temperature is small, and thermal excitations are important. In the present analysis, given the many other simplifying assumptions, we neglect this effect.

2.3 Numerical method

2.3.1 Pseudospectral solver

To solve the equations of motion (2.3)–(2.5), a pseudospectral collocation method is used to discretize the spatial coordinates, and a fractional timestep algorithm is used to advance the solution in time (Canuto et al., 1993). An explicit algorithm (third-order Adams-Bashforth) is used to solve the nonlinear terms, while the diffusion terms are solved using an implicit Crank-Nicolson scheme to maintain numerical stability (Boyd, 2013). The spectral solver is the same as that used by Peralta et al. (2005, 2008), who in turn followed Bagchi & Balachandar (2002). The radial coordinate r is expanded as a series of Chebyshev polynomials, and the angular coordinates θ and ϕ are expanded as Fourier series with parity correction at the coordinate singularity at the poles. For a detailed description of the solver, see Section 3 and the Appendix of Peralta et al. (2008). The computational domain and boundary conditions are described in Section 2.3.2 below. The equations of motion in this work are slightly different to those in Peralta et al. (2005, 2008), as noted in Section 2.2.




Figure 2.1: Illustration of the model neutron star’s geometry. A spherical surface of radius R_1 is embedded inside a larger spherical surface of radius R_2 with the same origin. Both spheres may rotate about independent axes, with respective angular velocity vectors Ω_1 and Ω_2 , though in this chapter we only consider rotation about a common axis. The Cartesian coordinates are defined so that the inner sphere rotates around the z axis, and the orthonormal vectors $\hat{\mathbf{e}}_{r,\theta,\phi}$ form the basis of the spherical coordinate system. Taken from Peralta (2006).

2.3.2 Initial and boundary conditions

We solve equations (2.3)–(2.5) within a spherical Couette geometry, consisting of two concentric spherical boundaries rotating about a common axis. Figure 2.1 shows a schematic of the computational domain.

The inner sphere radius is R_1 , the outer sphere radius is R_2 , and the respective angular speeds are Ω_1 and Ω_2 . In spherical coordinates the computational domain is

$$R_1 \leq r \leq R_2 \quad (2.7)$$

$$-\pi/2 \leq \theta \leq \pi/2 \quad (2.8)$$

$$0 \leq \phi \leq 2\pi. \quad (2.9)$$

The domain is discretised in space according to a Gauss-Lobatto quadrature scheme [Canuto et al. (1993); also see Appendix A of Peralta et al. (2008)], and we typically have $N_r \times N_\theta \times N_\phi = 121 \times 120 \times 4$ discretisation points. The coarse discretisation in ϕ is appropriate, since the common rotation axis of the crust and core prevents non-axisymmetric flow states, a result which was verified in Peralta et al. (2005, 2008).

In what follows, we express all variables in dimensionless form. We normalize times with respect to $\Omega_1(t = 0)^{-1}$, so that one rotation period at

$\Omega_1(t=0)$ corresponds to 2π time units. Lengths are normalized with respect to R_1 , and velocities with respect to $R_1\Omega_1(t=0)^{-1}$. The mass normalization, which affects quantities like the density, moment of inertia, and torque, is discussed below equation (2.13) in Section 2.4.1.

A spherical Couette geometry is adopted primarily for numerical reasons, i.e. to avoid the coordinate singularity at $r=0$ and stabilize the evolution; flows with $R_2 - R_1 \geq 0.5R_1$ are notoriously unstable (Benton & Clark Jr, 1974; Yavorskaya et al., 1980; Nakabayashi, Zheng & Tsuchida, 2002). The geometry is justified physically as an idealized model of either the outer core region of a neutron star, in which case the outer boundary is the crust/core interface, or the inner crust region, in which case the outer boundary is the neutron drip point. In either case the inner boundary represents some phase separatrix, below which the fluids are more tightly coupled by (say) a rapid increase in the strength of mutual friction with decreasing radius (Alpar, Langer & Sauls, 1984). Nonetheless, the presence of a rigid inner boundary is an artificial constraint in our model system. In our simulations we typically set the dimensionless gap width to be $\delta = (R_2 - R_1)/R_1 = 0.2$. The region between R_1 and R_2 is filled with a two-component fluid with component densities $\rho_p = \rho_n = 0.5$ and a Reynolds number of $Re = 500$ for the proton fluid. Both of these numbers are artificial, most estimates suggest $\rho_n \approx 0.9$ (Lattimer & Prakash, 2004), and Re may be as high as 10^{11} (Mastrano & Melatos, 2005; Melatos & Peralta, 2007), but these values are chosen for numerical reasons.

At $t=0$, the inner and outer boundaries are corotating at $\Omega(t=0) = 1$. The simulation runs in a reference frame rotating at $\Omega_f = \Omega(t=0)$ for all t , which introduces Coriolis terms, $2\mathbf{\Omega}(t=0) \times \mathbf{v}_{p,n}$, on the right-hand sides of equations (3) and (4). The initial conditions imply $\mathbf{v}_p = \mathbf{v}_n = 0$ everywhere at $t=0$.

For both the proton and neutron velocity fields, we impose no-penetration and no-slip boundary conditions. Mathematically, these boundary conditions can be expressed as

$$\mathbf{v}_{p,n}(R_{1,2}, \theta, \phi) = R_{1,2}\mathbf{\Omega}_{1,2} \times \hat{\mathbf{r}}, \quad (2.10)$$

where $\mathbf{\Omega}_{1,2}$ is the angular velocity of the boundary, and $\hat{\mathbf{r}}$ is the radial unit vector. The use of these boundary conditions is motivated primarily by numerical considerations. However, there is some physical justification as well. At $r = R_2$, no-penetration reflects the inability of either fluid to flow through the solid crust. No-penetration at $r = R_1$ is more artificial. Viscous processes as well as magnetic coupling keep the proton fluid near both boundaries in corotation, implying no-slip for $\mathbf{v}_p$. The boundary condition for the superfluid neutrons depends on the interaction between individual vortex lines and

the surface. Following Khalatnikov & Hohenberg (1965), the relative motion between a vortex line with tangential velocity $\mathbf{v}_L$ and a boundary moving with velocity $\mathbf{u}$ can be written as

$$\mathbf{v}_L - \mathbf{u} = c_1 \hat{\boldsymbol{\omega}}_{\mathbf{n}} \times (\mathbf{n} \times \hat{\boldsymbol{\omega}}_{\mathbf{s}}) + c_2 \mathbf{n} \times \hat{\boldsymbol{\omega}}_{\mathbf{n}}, \quad (2.11)$$

where $\mathbf{n}$ is the unit normal to the boundary. The coefficients $c_{1,2}$ parametrize the amount of slip. The two extremal cases are perfect sliding ($c_1 = c_2 \rightarrow \infty$) and no-slip ($c_1 = c_2 = 0$). One is usually obliged to set c_1 and c_2 on empirical grounds, even for well-studied and controlled situations involving liquid helium, let alone under the uncertain conditions present in a neutron star. Peralta et al. (2005, 2008) investigated boundary conditions extensively with this solver and encountered numerical difficulties for choices other than no-slip [c.f. Reisenegger (1993) and van Eysden (2015) for a discussion of alternatives in analytic calculations]. As mentioned above, we consider the the viscous and inviscid components to be locked together for $r < R_1$ and postulate that the fluid immediately adjacent to this region is also strongly coupled, i.e. $c_1 = c_2 = 0$ at $r = R_1$.

2.3.3 Model assumptions and limitations

Necessarily, the above model involves several simplifying assumptions. Some of these stem from a lack of theoretical consensus about aspects of the physics, while others are made in order to make a difficult numerical problem more tractable. We introduce the assumptions as they arise in previous sections. Here we summarize them together and discuss why they have been made, how they affect the applicability of our results to real pulsars, and how future work might refine the model.

- *Equations of motion.* The incompressibility condition (2.5) is justified, because the sound speed in a neutron star is much greater than the flow speeds. By contrast, the assumption of constant, uniform density does break down in the outer core modelled here. Work is currently under way to adapt the Navier-Stokes solver on which our two-fluid solver is based to work with non-uniform densities (K. Poon et al, in prep.) The absence of entrainment is valid as long as the simulation volume represents the outer core, where entrainment is weak (Carter, Chamel & Haensel, 2006). However, the inclusion of an entrainment term in equations (2.3) and (2.4) is relatively straightforward and represents a promising direction for future work. Vortex tension is another effect we have neglected. van Eysden & Melatos (2013) showed that in the crust of a neutron star the effect of vortex tension is small and confined to a

boundary layer, though more recent work (van Eysden, 2015) suggests that vortex tension adds an oscillatory component to glitch recovery. As with entrainment, it is relatively straightforward to include vortex tension in the solver as a natural next step.

- *Spherical Couette geometry and boundary conditions.* In Section 2.3.2 we explain the reasons for using a spherical Couette geometry with no-slip boundaries at R_1 and R_2 . This is an obvious limitation, as it restricts the model to a specific region of the star and imposes an artificial boundary condition at the inner boundary, yet it is unavoidable: the solver in its present form (Section 2.3.1) is unstable numerically when applied to a complete sphere. The appropriate boundary conditions for a superfluid in a rotating spherical container remain unclear and depend on the configuration of the vortex array and its interaction with the boundary (Khalatnikov & Hohenberg, 1965; Henderson, Barenghi & Jones, 1995; Peralta et al., 2008). Theoretical work by Campbell & Krasnov (1982) suggests that the vortex-boundary interaction is important in a spin-down context, as the rate of vortex nucleation is greater for rough boundaries than for smooth boundaries, so a rough boundary may decrease spin-down rate of the condensate. Future laboratory experiments with superfluid helium may shed some light on appropriate boundary conditions, but there is no guarantee that the results of such an experiment would be applicable to a fermion condensate in a neutron star. Experiments with liquid helium demonstrate that vorticity is transported erratically instead of smoothly across a two-phase boundary, such as at $r = R_1$, in response to interfacial Kelvin-Helmholtz instabilities, an effect which we do not include (Blaauwgeers et al., 2002; Mastrano & Melatos, 2005). The use of a solid boundary at R_2 is also unphysical, because the crust-core interface is not sharply delineated, and the solid and superfluid components interact non-trivially through magnetic and elastic coupling. Andersson, Haskell & Samuelsson (2011) developed a hydrodynamic formalism for investigating the crust-core coupling, which may be useful in the future for determining appropriate boundary conditions in an astrophysical context. Numerical studies of vortex motion are another avenue through which this issue might be investigated in the future (Schwarz, 1985; Adachi, Fujiyama & Tsubota, 2010; Baggaley & Barenghi, 2012).
- *Astrophysical parameters.* Certain parameters of the model are poorly constrained, because bulk nuclear matter in the low-temperature, high-density regime of neutron stars cannot be studied easily in laboratory

experiments. For example, the relative moments of inertia of the crust and the superfluid depend on the precise equation of state for bulk nuclear matter, and inferred values from measurements of pulsar glitches only weakly constrain these (Sections 2.4.1 and 2.6.2). Where possible we use the best current estimates and cite relevant literature, e.g. the values of the mutual friction parameters B and B' (Section 2.2). In other places, we use artificial values for numerical stability or for illustrative purposes, all of which are explained in the relevant sections, e.g. the value of Re (Section 2.4.1). Future work will explore the parameter space more widely, in preparation for new observational studies of pulsar glitches and gravitational waves (Weber et al., 2007; van Eysden & Melatos, 2010).

2.4 Steady pre-glitch spin down

Before investigating the post-glitch response in Section 2.5, we discuss how to set up the simulation to achieve a realistic initial state leading up to the glitch.

2.4.1 Set up

Many models of glitches posit that, as a neutron star spins down, the deceleration of the condensate lags that of the viscous component and the crust, due to vortex pinning. As the lag builds up, the pinned vortices act as a reservoir of angular momentum, which can recouple with the viscous fluid and crust spasmodically, causing glitches (Anderson & Itoh, 1975; Warszawski & Melatos, 2011; Haskell & Melatos, 2015). To simulate glitches in this paradigm, we need to prepare the system in a state where there is a lag between the neutron and proton velocities prior to a glitch. There is no analytic solution for high-Reynolds-number spherical Couette flow in such a state. Instead, we begin with the inner and outer boundaries co-rotating and spin down the outer sphere by imposing an external torque, analagous to the star’s electromagnetic braking torque (Melatos, 1997). As the external torque spins down the outer boundary, the neutron and proton fluids near the no-slip boundary also spin down, and the deceleration is communicated to the interior by Ekman pumping (Melatos, 2012; van Eysden & Melatos, 2014; van Eysden, 2015).

In response to the external torque, opposing viscous torques N_1 and N_2 are induced by the gradients of $\mathbf{v}_p$ at $r = R_1$ and R_2 (Landau & Lifshitz, 1959). At each time step, the viscous torque is calculated by integrating

the non-vanishing terms of the stress tensor over the surfaces $r = R_1, R_2$. Assuming axisymmetric flow one finds

$$N_{1,2} = \frac{2\pi R_{1,2}}{Re} \int d\theta \sin \theta \left(\left. \frac{\partial v_p^\phi}{\partial r} \right|_{r=R_{1,2}} - \frac{v_p^\phi}{R_{1,2}} \right). \quad (2.12)$$

The angular velocities of the boundaries, $\Omega_{1,2}$, evolve according to

$$I_{1,2} \frac{d\Omega_{1,2}}{dt} = N_{1,2} + N_{\text{ext}}, \quad (2.13)$$

where I_1 and I_2 are the moments of inertia of the inner core and the crust, and N_{ext} is the external torque. Equation (2.13) is solved at each timestep using a third-order Adams-Bashforth algorithm, so as to maintain the same time-accuracy as the pseudospectral solver. The moment of inertia is expressed in dimensionless units, normalized with respect to $(\rho_p + \rho_n)R_1^5$. This means that the moment of inertia of the fluid in the interior, notionally regarded as a rigid body, is $I_3 = 2.5$. We choose the moment of inertia of the crust, I_2 so that $I_2/I_3 = 4$, and the moment of inertia of the core, I_1 , to satisfy $I_2/I_3 = 4$ and $I_1/I_2 = 10$. These numbers give the fractional moment of inertia due to the condensate as 1.1%, which is similar to the value of this ratio in glitching pulsars (Link, Epstein & Lattimer, 1999; Andersson et al., 2012). The external torque, N_{ext} is chosen so that a decoupled crust with moment of inertia I_2 decelerates at $\dot{\Omega} = N_{\text{ext}}/I_2 = -10^{-3}$. This is artificially high; most pulsars have $|\dot{\Omega}| \sim 10^{-15}$ (Manchester et al., 2005). However, it is necessary so that an appreciable lag builds up between the crust and other components in the duration of a typical simulation.

Three important time-scales in this system are the viscous time-scale, $\tau_{\text{visc}} = Re/\Omega$, the Ekman time-scale, $\tau_{\text{Ek}} = \sqrt{Re}/\Omega$, and the mutual friction time-scale, $\tau_{\text{MF}} = 1/(2\Omega B)$. In a neutron star, we are interested in two regimes: the viscous-dominant regime, $\tau_{\text{Ek}} < \tau_{\text{visc}} < \tau_{\text{MF}}$, and the mutual friction-dominant regime, $\tau_{\text{MF}} < \tau_{\text{Ek}} < \tau_{\text{visc}}$, which we refer to as weak and strong mutual friction respectively for the remainder of this chapter. For the strong mutual friction case, we take $B = 0.1$, and for weak mutual friction we have $B = 10^{-4}$, which are typical values for Kelvin wave damping and electron scattering respectively (Alpar, Langer & Sauls, 1984; Andersson & Comer, 2006; Haskell, Pizzochero & Sidery, 2012). We typically have $Re = 500$, which is significantly lower than realistic neutron star values of $Re \sim 10^{11}$ (Mastrano & Melatos, 2005; Melatos & Peralta, 2007). In the interest of computational tractability we choose lower values of Re to avoid turbulence in the viscous component (Peralta et al., 2005), which would require increased spatial resolution, and to reduce the important time-scales τ_{Ek} and τ_{visc} .

Starting from a state of co-rotation, we impose a constant torque on the outer boundary and evolve the system until the inner and outer boundaries are spinning down at approximately the same rate, i.e. we have $\dot{\Omega}_1 \approx \dot{\Omega}_2$. We refer to this state as ‘spin-down equilibrium’ and use it as a starting point for our simulations of glitches. Note that $\dot{\Omega}_1 \approx \dot{\Omega}_2$ does not imply $\dot{\mathbf{v}}_{pn} \approx 0$. The majority of the angular momentum in the model pulsar is in the heaviest region, $r < R_1$, which couples to the crust and proton fluid on either the viscous or Ekman timescales (van Eysden & Melatos, 2014). To reach spin-down equilibrium, we need $t > \tau_{\text{visc}} > \tau_{\text{Ek}}$, however, it is still possible for the lag between the proton and neutron fluids to grow, as the maximum value of $\mathbf{v}_{pn}$ is determined by mutual friction, and is not reached until we have $t > \tau_{\text{MF}}$. Evolving the system until one has $\dot{\mathbf{v}}_{pn} = 0$ may seem a more natural initial condition for a glitch, but in practice τ_{MF} can be prohibitively long for numerical experiments. Moreover, it is unclear whether neutron stars ever reach such a state in reality, as the glitch trigger mechanism is unknown, hydrodynamical instabilities could set in long before the velocity lag reaches equilibrium (Andersson, Comer & Prix, 2004), and stratification acts to maintain a shear (Melatos, 2012). Similar issues arise in laboratory experiments with liquid helium (van Eysden & Melatos, 2011).

2.4.2 Output and initial state

Figure 2.2 shows the evolution of $\Omega_1(t)$, $\Omega_2(t)$, $\dot{\Omega}_1(t)$, and $\dot{\Omega}_2(t)$, for $0 \leq t \leq 500$. From the bottom two panels, we can see that the system reaches steady deceleration by $t = 500$: we obtain $\dot{\Omega}_1 \approx \dot{\Omega}_2 \approx -9.5 \times 10^{-5}$, and $|\dot{\Omega}_1 - \dot{\Omega}_2|/|\dot{\Omega}_1| \lesssim 5\%$, and the angular velocity lag between the crust and core $\Omega_2(t = 500) - \Omega_1(t = 500)$ is -0.52 with strong mutual friction and -0.61 with weak mutual friction.

The deceleration is faster with weak mutual friction, as shown in the top two panels. This is because the neutron and proton components couple only after $t > \tau_{\text{MF}} = 500$ (weak), so for $t < 500$ the external torque is only spinning down the viscous component, the core and the crust. With strong mutual friction, the inviscid component couples to the viscous component after $t > \tau_{\text{MF}} = 5$ (strong), so the moment of inertia of the coupled system is greater than with weak mutual friction, and the magnitude of its deceleration is reduced.

For a more detailed view of the flow, we look in figure 2.3 at contours of v_{pn}^ϕ in the strong mutual friction regime in a slice through the y - z plane, at $t = 5 \approx \tau_{\text{MF}} < \tau_{\text{Ek}} \ll \tau_{\text{visc}}$. It is clear that the flow is axisymmetric but not columnar; the contours of v_{pn}^ϕ are curved. An expanded view of the flow in the top-right quadrant of the y - z plane appears in Figure 2.4, at




Figure 2.2: Response of the crust and core angular velocities, $\Omega_1(t)$ and $\Omega_2(t)$, and their time derivatives, $\dot{\Omega}_1(t)$, and $\dot{\Omega}_2(t)$, to a constant external torque on the crust, for $0 \leq t \leq 500$, with strong (left column) and weak (right column) mutual friction, with $\Omega_1(t=0) = \Omega_2(t=0)$. In this figure, Ω_1 and Ω_2 are measured in a frame rotating with angular velocity $\Omega_f = \Omega_1(t=0) = \Omega_2(t=0)$ and are also normalized with respect Ω_f . Time is normalized with respect to Ω_f^{-1} . The top row is the crust angular velocity, $\Omega_1(t)$, (dashed curves), and the core angular velocity, $\Omega_2(t)$, (solid curves). The bottom row is the crust angular velocity derivative, $\dot{\Omega}_1(t)$, (dashed curves), and the core angular velocity derivative $\dot{\Omega}_2(t)$, (solid curves). In both mutual friction regimes the two boundaries are decelerating at approximately the same rate after $t = 500$, a state we refer to as spin-down equilibrium.




Figure 2.3: Contours of the azimuthal velocity lag, v_{pn}^ϕ , plotted at $t = 5 \approx \tau_{\text{MF}} < \tau_{\text{Ek}} < \tau_{\text{visc}}$ in the y - z plane, with strong mutual friction. Velocity is expressed in dimensionless units normalized with respect to $R_1 \Omega_f^{-1}$. v_{pn}^ϕ is non-zero near $r = R_2$, but decreases monotonically to zero as $r \rightarrow R_1$. $|v_{\text{pn}}^\phi|$ increases from zero at the poles to a maximum at $\theta \approx 30^\circ$ in latitude. The flow is axisymmetric but not columnar, as the 18 contours have similar curvature to the boundaries throughout most of the domain.

$t = 5, 25, 100$, and 500 . The contours in the remainder of the domain can be inferred from the symmetries in Figure 2.4. For strong mutual friction (left column), there is a noticeable difference between the flow at $t = 5$ (top row) and at later times. At $t = 5$, v_{pn}^ϕ has a maximum at $r \approx R_2$, at a latitude of $\theta \approx 30^\circ$. At fixed radius, v_{pn}^ϕ decreases slowly with latitude: $v_{\text{pn}}^\phi(r \approx R_2, \theta = 0) \approx 0.75v_{\text{pn,max}}^\phi$. For $t \geq 25$ (second row and below), $v_{\text{pn,max}}^\phi$ is still near $r = R_2$ and $\theta = 30^\circ$, however, v_{pn}^ϕ falls off more steeply with decreasing latitude, with $v_{\text{pn}}^\phi(r \approx R_2, \theta = 0) \approx 0.25v_{\text{pn,max}}^\phi$ at $t = 25$ and $\approx 0.1v_{\text{pn,max}}^\phi$ at $t = 500$. For weak mutual friction, the flow pattern does not change significantly between $t = 5$ and $t = 100$, apart from a steady increase in $|v_{\text{pn}}^\phi|$. Another notable difference between the two regimes is that v_{pn}^ϕ decreases monotonically from $v_{\text{pn,max}}^\phi$ at $r = R_2$ to zero at $r = R_1$ for weak mutual friction, while it decreases to $-v_{\text{pn,max}}^\phi$ at $r = R_1$ for strong mutual friction. To investigate the sign-reversal of v_{pn}^ϕ more closely, we first average v_{pn}^ϕ over the angular coordinates, θ and ϕ , and examine the radial profile of the averaged velocity, $\langle v_{\text{pn}}^\phi \rangle_{\theta,\phi}$, in Figure 2.5. At $r = R_1$ and R_2 , the two fluids are locked together by the no-slip boundary condition.

Away from the boundaries, an interesting feature develops with strong mutual friction: $\langle v_{\text{pn}}^\phi \rangle_{\theta,\phi}$ undergoes a sign change. Naïvely, one expects the inviscid neutron condensate to lag the protons, because Ekman pumping acts on the protons to bring them into corotation with the boundaries after $t \approx \tau_{\text{Ek}}$, while there is no equivalent process for the condensate. When mutual friction is weak, this is indeed the case. For $\tau_{\text{MF}} < \tau_{\text{Ek}}$, however, the protons drag the neutrons along as Ekman pumping proceeds. Equations (2.3) and (2.4) differ in form, so different flow patterns develop in the two components [van Eysden & Melatos (2013); in particular see Figures 1 – 6 in the latter chapter]. In Figure 2.6, we show the streamlines of the protons (top) and neutrons (bottom) at $t = 500$ with strong mutual friction. It is clear that the flow states are distinct. For the proton component, the meridional circulation is centred at $\theta \approx 30^\circ$ at $r \approx 0.5(R_1 + R_2)$. The streamlines are parallel to the boundaries at $r \approx R_1$ and $r \approx R_2$ and elsewhere they bend inwards. For the neutron component, the streamlines are less symmetric. The primary cell is centered at $\theta \approx 20^\circ$ and is closer to $r = R_1$ than $r = R_2$. The flow is roughly columnar, with a Stewartson layer partially forming at $r \sin \theta \approx R_1$ (Peralta & Melatos, 2009). Subtracting the flow patterns in the two panels of Figure 2.6 (see bottom left panel of Figure 2.4), we find that the neutrons spin down faster than the protons at $r \approx R_1$.




Figure 2.4: Contours of the azimuthal component of the proton-neutron velocity lag, v_{pn}^ϕ , in the meridional plane, for strong mutual friction (left column), and weak mutual friction (right column), at $t = 5, 25, 100$, and 500 (top to bottom row respectively). Velocity is expressed in dimensionless units normalized with respect to $R_1 \Omega_f^{-1}$. Red colours show positive lag into the page, blue colours show lag out of the page. Rotation is about the positive z -axis.




Figure 2.5: Snapshots of the θ and ϕ -averaged azimuthal velocity lag, $\langle v_{\text{pn}}^\phi \rangle_{\theta, \phi}$, versus cylindrical radius r , at $t = 5$ (top left), 25 (top right), 100 (bottom left), and 500 (bottom right). Time is expressed in dimensionless units normalized with respect to the angular velocity of the rotating frame Ω_f^{-1} , and velocity is expressed in dimensionless units normalized with respect to $R_1 \Omega_f^{-1}$. The red curve (circles) and the blue curve (crosses) correspond to strong and weak mutual friction respectively. The boundary condition enforces $v_{\text{pn}}^\phi = 0$ at $r = R_1$ and $r = R_2$. When mutual friction is strong, v_{pn}^ϕ changes sign near $r = R_1$. When mutual friction is weak, the viscous component spins down faster than the condensate throughout the domain.




Figure 2.6: In-plane meridional streamlines of the protons (top) and neutrons (bottom) with strong mutual friction at $t = 500$. Both components show meridional circulation, but the flows are distinct. In the protons the circulation is centered at $\theta \approx 30^\circ$ and $r \approx 0.5(R_1 + R_2)$, while for the neutrons the circulation cells are centered at $\theta \approx 20^\circ$ and $r \approx R_1$. This difference in flow patterns leads to the result shown in Figure 2.4, where v_{pn}^ϕ changes sign in the region $R_1 < r < R_2$ for strong mutual friction.

2.5 Post-glitch recovery

2.5.1 Activating the glitch

We simulate glitches by taking, as initial conditions, the velocity fields and boundary conditions of the system after 500 time units of steady spin down and modifying them impulsively in one of three ways. Firstly, we spin up the outer boundary instantaneously, which we call a ‘crust glitch’. Secondly, we recouple the proton and neutron fluids instantaneously, so that v_{pn}^ϕ is reduced to zero everywhere in a step, which we call a ‘bulk glitch’. Thirdly, we spin up the core instantaneously, which we call an ‘inner glitch’.

We can think of a crust glitch as the response to a crustquake, caused by shear stresses (Ruderman, 1969; Middleditch et al., 2006), or crust cracking due to a build-up of superfluid vortices at the crust (Alpar et al., 1996). Similarly, inner glitches represent a violent event occurring within the core, at $r < R_1$, e.g. an avalanche of superfluid vortices pinned to magnetic flux tubes (Link, 2012a). Some models of core matter, such as color superconducting condensates, predict a high shear modulus (Mannarelli, Rajagopal & Sharma, 2007) so that seismic disturbances (‘corequakes’) may have observable effects on the crust (Ruderman, 1976; Ruderman, Zhu & Chen, 1998).

Bulk glitches correspond to trigger mechanisms distributed throughout the superfluid itself, in the region $R_1 \leq r \leq R_2$. When $|\mathbf{v}_{\text{pn}}|$ grows too large, a hydrodynamical instability similar to the two-stream instability can occur, which acts to reduce the velocity lag between the two components (Andersson, Comer & Prix, 2004; Mastrano & Melatos, 2005; Glampedakis & Andersson, 2009). Another potential trigger is a vortex avalanche, driven by the shear between the bulk fluid and pinned vortices (Warszawski, Melatos & Berloff, 2012).

Following a glitch, we evolve the system for a further 500 time units. We also run a control simulation, where steady spindown continues uninterrupted for a further 500 time units, i.e. no glitch. We compare $\Omega_2(t)$ with and without a glitch and construct ‘residuals’ $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, where $\Omega_{2,g}(t)$ and $\Omega_{2,ng}(t)$ are the values of Ω_2 with and without a glitch respectively. The residual has two advantages. First, in reality, glitches are detected as deviations in pulse times-of-arrival from a non-glitch spin-down model (Hobbs, Edwards & Manchester, 2006; Espinoza et al., 2011). Second, the residuals separate the glitch evolution from that due to the external torque, which remains constant throughout the post-glitch simulation (the glitch recovery time-scale is much shorter than the spin-down time-scale). Note that $\Omega_1(t)$ also evolves, consistent with equation (2.13).

2.5.2 Fitting the glitches

Generally, pulsar glitches are followed by a recovery phase, during which some or all of the increase in spin frequency is reversed. It is common to fit glitches to a function of the form [cf. (1.2)]

$$\Omega(t) = \Omega_0(t) + \Delta\Omega_p + \sum_{i=1}^N \Delta\Omega_n e^{-t/\tau_n}, \quad (2.14)$$

where $\Delta\Omega_p$ is the permanent change in the angular velocity and $\Delta\Omega_n$ are transient changes in the angular velocity of either sign that decay on a characteristic timescale τ_n [see e.g. Shemar & Lyne (1996); Wong, Backer & Lyne (2001)]. Further time derivatives of Ω may be included in equation (2.14), but are normally neglected. In Section 2.6.1, we discuss the effect of including a permanent change in the angular velocity derivative, $\Delta\dot{\Omega}_p$. As many as four distinct transient timescales can be fitted for some glitches (Dodson, Lewis & McCulloch, 2007), ranging from tens of seconds to hundreds of days. One-dimensional superfluid simulations in Haskell, Andersson & Comer (2012) show that the recovery occurs on a combination of time-scales and is not well-approximated by a single exponential, a result which also holds for two-dimensional Ekman pumping (van Eysden & Melatos, 2010). A multiple exponential fit, though, is well-motivated by theoretical work done by van Eysden & Melatos (2010, 2014), who solved analytically for the rotational evolution of a rotating vessel filled with a helium-II-like superfluid, described by the equations of motion of Chandler & Baym (1986), following an impulsive acceleration. They found that, in the limits $\tau_{\text{MF}} \ll \tau_{\text{Ek}}$, $B' \ll 1$, $I_3/I_2 \gg 1$, the recovery involves two exponential timescales, which depend on the relative densities of the viscous and inviscid components, the relative strength of viscous and mutual friction forces, and the inertia of the container.

For our simulated glitches, we fit the residual $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$ with five parameters: τ_1 , τ_2 , $\Delta\Omega_1$, $\Delta\Omega_2$, and $\Delta\Omega_p$. The fit is extracted using the Levenberg-Marquardt least-squares algorithm, implemented in the statistics software R with the package `minpack.lm` (Elzhov et al., 2013).

2.5.3 Crust glitches

To induce a crust glitch, we instantaneously increase Ω_2 by an amount $|\Omega_2 - \Omega_1|/2$ at $t = 500$. The jump is chosen arbitrarily; we find that its size does not alter the dynamics of the recovery qualitatively. Figure 7 shows the evolution of $\Delta\Omega$ following the crust glitch for strong (top panel) and




Figure 2.7: Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 100 time units following a crust glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.14) with $N = 2$ (see Table 2.1 for values). Mutual friction has little effect on the recovery time-scales.

Parameter	Weak MF	Strong MF
$\Delta\Omega_1$	2.55×10^{-3}	4.15×10^{-3}
τ_1	1.53	5.15
$\Delta\Omega_2$	2.42×10^{-2}	1.86×10^{-2}
τ_2	66.4	62.0
$\Delta\Omega_p$	2.96×10^{-3}	2.46×10^{-3}

Table 2.1: Fitted parameters to equation (2.14) with $N = 2$ for the crust glitches in Figure 2.7 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$.

weak (bottom panel) mutual friction. The parameters of the associated dual exponential fits are shown in Table 2.1. Immediately following the glitch $\Delta\Omega$ decreases monotonically. This happens because the no-slip boundary condition causes $v_p^\phi(r = R_2)$ to also increase, so the viscous torque, equation (2.12), is reduced (relative to the no-glitch simulation), and the external spin-down term in equation (2.13) dominates. Both strong and weak mutual friction recover similarly, with $\tau_2 \approx 64$. The fitted values of τ_1 are discrepant by a factor of ≈ 3 between the mutual friction regimes, however, in both cases $\Delta\Omega_1 \ll \Delta\Omega_2$, so doesn't affect the shape of the recovery profile significantly. $\Delta\Omega_p$ is $\approx 10\%$ of $\Delta\Omega_{\max}$ with strong and weak mutual friction. There is no qualitative difference between the recovery profile in the two mutual friction regimes for crust glitches. In the limits $\tau_{\text{MF}} \ll \tau_{\text{Ek}}$, $B' \ll 1$, $I_3/I_2 \gg 1$, van Eysden & Melatos (2010) derived simple analytic expressions for the time-scales of a double-exponential recovery,

$$\tau_1 \approx \tau_{\text{MF}}, \quad (2.15)$$

$$\tau_2 \approx \frac{\tau_{\text{Ek}}}{\rho_p(1 + I_3/I_2)} \quad (2.16)$$

(or vice versa). For $B = 0.1$ (strong mutual friction), equations (2.15) and (2.16) give $\tau_1 = 5$, $\tau_2 = 36$, compared to $\tau_1 = 5.15$, $\tau_2 = 62.0$ obtained from fitting to simulations. This is a decent level of agreement, though it should be noted that equations (2.15) and (2.16) are not exactly the expressions as written in van Eysden & Melatos (2010), who considered a spherical container ($R_1 = 0$) filled with a HVBK superfluid (Chandler & Baym, 1986), rather than a spherical Couette geometry filled with fluid obeying equations (2.3) – (2.5). More specifically, in van Eysden & Melatos (2010), the quantity I_3/I_2 is the ratio of the moment of inertia of the interior of the sphere to that of the container, while in our analysis we exclude the moment of inertia of the core. We are also in a different parameter regime to van Eysden & Melatos (2010): we have $\tau_{\text{Ek}} \approx 5\tau_{\text{MF}}$ for strong mutual friction, and $I_3/I_2 \approx 0.25$.

Parameter	Weak MF	Strong MF
$\Delta\Omega_1$	-4.27×10^{-3}	-2.61^{-4}
τ_1	1.60	0.201
$\Delta\Omega_2$	4.40×10^{-3}	2.26×10^{-4}
τ_2	68.0	59.2
$\Delta\Omega_p$	6.56×10^{-4}	1.21×10^{-6}

Table 2.2: Fitted parameters to equation (2.13) with $N = 2$ for the bulk glitches in Figure 2.8 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$.

2.5.4 Bulk glitches

To induce a bulk glitch, we instantaneously recouple the proton and neutron fluids, so that the velocity lag is reduced to zero, in a way that conserves total angular momentum. Generally, the recoupling algorithm depends on the relative moments of inertia of each fluid, but since we have $I_n = I_p$, we simply set $v_p^\phi \mapsto v_p^\phi - v_{pn}^\phi/2$, and $v_n^\phi \mapsto v_n^\phi + v_{pn}^\phi/2$. Figure 2.8 displays the residual $\Delta\Omega$ in the 100 time units of simulation time following a bulk glitch with strong and weak mutual friction. The recovery is fitted to equation (2.13), and the fitted parameters are quoted in table 2.2. For both strong and weak mutual friction, the crust spins up immediately after recoupling. Unlike crust glitches, where the spin up is instantaneous by construction, there is a noticeable rise time in $\Delta\Omega$, which is reflected in the fitting algorithm returning $\Delta\Omega_1 < 0$ in Table 2.2. The reason that $\Delta\Omega$ increases initially is that spinning up the protons increases $\partial v_p^\phi / \partial r|_{r=R_2}$, N_2 [via equation (2.12)], and hence $\Delta\Omega$ [via equation (2.13)]. As the crust spins up, however, the no-slip boundary condition increases $v_p^\phi(r = R_2)$, which acts to reduce the viscous torque and halt the spin up after $t \approx 5$. $\Delta\Omega$ reaches a maximum at similar times for both strong and mutual friction regimes ($t = 5.0$ and $t = 6.3$ respectively), which corresponds to approximately one rotation period of the star $P = 2\pi/\Omega_2 = 7.0$. The spin up time is also similar to τ_{MF} for strong mutual friction, as discussed further in Section 2.6.2. For $t \gtrsim 5$, $\Delta\Omega$ decreases on a time-scale of $\tau_2 \approx 65$. This is because the spin up of the crust following the glitch reduces the angular velocity lag between the crust and core and upsets spin-down equilibrium. In order to restore spin-down equilibrium, the crust spins down faster than the core for ≈ 300 time units $\sim \tau_{visc}$, after which spin-down equilibrium is restored. This is similar to what we see during the set up of the initial condition. In Figure 2.2, we find $|\dot{\Omega}_2| > |\dot{\Omega}_1|$ for $t \lesssim 300$, and $|\dot{\Omega}_2| \approx |\dot{\Omega}_1|$ for $t \gtrsim 300$.

The evolution of $\Delta\Omega$ is different in the strong and weak mutual friction




Figure 2.8: Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 100 time units following a bulk glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.13) with $N = 2$ (see Table 2 for values). With strong mutual friction, the size of the glitch is significantly reduced, and there is an oscillatory component to the recovery for $t \lesssim 30$.

Parameter	Weak MF	Strong MF
$\Delta\Omega_1$	1.46×10^{-3}	2.09×10^{-3}
τ_1	0.98	1.80
$\Delta\Omega_2$	-3.90×10^{-2}	-3.73×10^{-2}
τ_2	70.3	63.5
$\Delta\Omega_p$	3.78×10^{-2}	3.55×10^{-2}

Table 2.3: Fitted parameters to equation (2.14) with $N = 2$ for the inner glitches in Figure 2.9 with strong ($B = 0.1$) and weak ($B = 10^{-4}$) mutual friction (MF) at $Re = 500$.

regimes in several ways. Firstly, the peak size of the glitch is larger for weak than strong mutual friction, with $\Delta\Omega_{\max}^{\text{weak}} \approx 20\Delta\Omega_{\max}^{\text{strong}}$. Secondly, while weak mutual friction produces a fast, monotonic rise followed by a slower, monotonic recovery, there is a noticeable oscillatory component in both the rise and recovery for $t \leq 30 \approx \tau_{\text{Ek}}$ with strong mutual friction. The oscillation period is ≈ 5 time units, similar to τ_{MF} , which may explain the absence of oscillation in the glitch with weak mutual friction: any oscillations which occur on a time-scale of $\tau_{\text{MF,weak}} = 500$ are invisible because $\Delta\Omega(t)$ decays exponentially on a time scale of $\tau_{2,\text{weak}} = 68$ (table 2.2). Thirdly, the fitted rise time τ_1 is ~ 10 times faster with strong mutual friction than with weak, as shown in Table 2.2. However, looking at Figure 2.8, this is an artifact of the fit; $\Delta\Omega$ actually peaks at $t = 5.0$ and $t = 6.3$ for strong and weak mutual friction respectively, which corresponds to approximately one rotation period. Finally, weak mutual friction causes a permanent change in the spin frequency of the crust, viz. $\Delta\Omega^{\text{weak}}(t = 500)/\Delta\Omega_{\max}^{\text{weak}} = 0.14$, c.f. $\Delta\Omega^{\text{strong}}(t = 500)/\Delta\Omega_{\max}^{\text{strong}} = 0.003$. Though this result is not obvious in Figure 2.8, which truncates the recovery at $t = 100$ for clearer presentation of the oscillation, it can be seen clearly in the fitted values of $\Delta\Omega_p$ in Table 2.2, viz. $\Delta\Omega_p^{\text{weak}}/\Delta\Omega_{\max}^{\text{weak}} = 0.14$ and $\Delta\Omega_p^{\text{strong}}/\Delta\Omega_{\max}^{\text{strong}} = 0.005$.

2.5.5 Inner glitches

Inner glitches resemble crust glitches, except that Ω_1 jumps initially. The results below describe a glitch where Ω_1 is instantaneously increased to $\Omega_1(t = 0) = 1$, although in general we find that the size of the jump does not qualitatively change the recovery dynamics. Figure 2.9 plots the residual $\Delta\Omega$ for 500 time units following the glitch. The fitted parameters of a dual exponential fit are shown in Table 2.3. The recovery following an inner glitch is noticeably different from crust and bulk glitches. The glitch algorithm adds




Figure 2.9: Angular velocity residuals, $\Delta\Omega(t) = \Omega_{2,g}(t) - \Omega_{2,ng}(t)$, for 500 time units following an inner glitch. The red diamonds show the data from the simulation. The black dashed curve shows a curve of best fit to a dual exponential recovery given by equation (2.14) with $N = 2$ (see Table 2.3 for values). In contrast to crust and bulk glitches, there is no relaxation following the spin up, $\Delta\Omega_{\text{max}} = \Delta\Omega(t = 500)$. We see no significant difference in the recovery timescales between strong and weak mutual friction.

angular momentum to the system by increasing Ω_1 , which causes $v_p^\phi(r = R_1)$ and hence N_1 to also increase, increasing the spin-down torque on the core and upsetting spin-down equilibrium. The system returns to spin-down equilibrium after $t \approx \tau_{\text{visc}}$ by redistributing the added angular momentum from the glitch between $r = R_1$ and $r = R_2$, which also causes the crust to spin up. Figure 2.9 shows how, immediately following the glitch, $\Delta\Omega$ rises monotonically and asymptotes towards a maximum value after $t \sim 300 \sim \tau_{\text{visc}}$. There is no relaxation.

2.6 Discussion

We discuss the above results in two groups. Firstly, we consider the crust and inner glitches, which can be grouped naturally, since both involve impulsive acceleration of a boundary and a net increase of the angular momentum of the system (Section 2.6.1). Secondly, we consider the angular momentum-conserving bulk glitches (Section 2.6.2).

2.6.1 Crust vs inner glitches

A simple way to model the observed behaviour is by considering an idealised system consisting of four rigidly rotating components: the crust, the proton fluid, the neutron fluid, and the inner core, all of which may couple to or decouple from one another.

To describe crust glitches in this picture, we consider the inner core and the crust as being coupled prior to the glitch (see the discussion of spin-down equilibrium in Section 2.4.2 and the bottom panels of Figure 2.2). The inertia of the neutrons and protons is small compared to the crust and core. The glitch is induced by suddenly spinning up the crust, upsetting spin-down equilibrium. A restoring force quickly spins down the crust and increases the angular velocity of the fluid and the core. Eventually, spin-down equilibrium is restored, and the permanent increase in the spin frequency of the crust can be estimated by multiplying the initial glitch size $\Delta\Omega_{\text{max}} = \Delta\Omega(t = t_{\text{glitch}})$ by the ratio of the moments of inertia of the crust and the whole system, $I_{\text{crust}}/I_{\text{total}}$. For inner glitches the argument is the same, but the relevant ratio is $I_{\text{core}}/I_{\text{total}}$, and $\Delta\Omega_{\text{max}}$ is the amount by which the inner boundary spins up.

In Table 2.4 we compare the estimated permanent spin-up, $\Delta\Omega_p^{\text{est}} = \Delta\Omega_{\text{max}} I_{\text{crust,core}}/I_{\text{total}}$ to the measured value of $\Delta\Omega(t = 500)$, as well as the fitted value $\Delta\Omega_p$ for crust and inner glitches with two different spin-down models described below. For the crust glitches, Table 2.4 shows that the

Location	MF	$\Delta\Omega(t = 500)$ ($\times 10^{-3}$)	$\Delta\Omega_p^{\text{est}}$ ($\times 10^{-3}$)	$\Delta\Omega_p^{(5)}$ ($\times 10^{-3}$)	$\Delta\Omega_p^{(6)}$ ($\times 10^{-3}$)
Crust	Strong	2.48	2.31	2.46	2.42
Crust	Weak	2.99	2.67	2.60	2.90
Inner	Strong	3.58	3.44	3.55	3.44
Inner	Weak	3.81	3.54	3.78	3.67

Table 2.4: We compare different methods for estimating the permanent change in the crust angular velocity residual, $\Delta\Omega_p$, following crust and inner glitches with strong and weak mutual friction. The third column estimates $\Delta\Omega_p$ as the value of $\Delta\Omega$ after 500 time units of simulation post-glitch, $\Delta\Omega(t = 500)$. The fourth column estimates $\Delta\Omega_p$ using the formula $\Delta\Omega_p^{\text{est}} = \Delta\Omega_{\text{max}} I_{\text{crust,core}} / I_{\text{total}}$. Finally we estimate $\Delta\Omega_p$ by fitting the data to glitch recovery models similar to equation (2.14) with five parameters (fifth column) and six parameters (sixth column).

estimated values agree approximately, though in both mutual friction regimes we obtain $\Delta\Omega(t = 500) \approx 1.1\Delta\Omega_p^{\text{est}}$. The discrepancy arises because the system is not completely coupled, even after 500 time units, so that the outer boundary is decelerating faster under the external torque than it would in equilibrium. In Table 2.4 we also consider the effect of adding a permanent change in the spin-down rate to our glitch recovery model, as such a change is often reported in the literature [e.g. Wang et al. (2000)]. $\Delta\Omega_p^{(5)}$ is the fitted value implied by equation (2.14), while $\Delta\Omega_p^{(6)}$ also includes a sixth parameter, the permanent change in the spin-down rate $\Delta\dot{\Omega}_p t$ in equation (2.14). Both values agree to within 10%, both in comparison to each other and also to the predicted and observed values, and the sum-of-squares errors returned by the fitting algorithm for each model are similar.

2.6.2 Bulk glitches and mutual friction

An interesting result from the bulk glitch simulations is the large size difference between glitches with strong and weak mutual friction. Figure 2.5 shows that $|v_{\text{pn}}^\phi|$ is greater for weak mutual friction than strong, but the difference is at most a factor of ≈ 4 , not the factor of ≈ 20 observed in glitch sizes. To understand this discrepancy, we treat both fluid components and the boundaries as rigid bodies, so that the angular momentum of each is given by $I_x \Omega_x$, where the index $x = \text{p, n or b}$ denotes the protons, neutrons and boundary respectively. Conservation of angular momentum

before and after the glitch implies

$$I_p \Omega_p^i + I_n \Omega_n^i + I_b \Omega_b^i = I_p \Omega_p^f + I_n \Omega_n^f + I_b \Omega_b^f, \quad (2.17)$$

where the superscripts i and f denote the initial (pre-glitch) and final (post-glitch) values respectively. Rearranging (2.17), we get

$$\Delta \Omega_b = -\frac{I_p}{I_b} \Delta \Omega_p - \frac{I_p}{I_b} \Delta \Omega_n, \quad (2.18)$$

where $\Delta \Omega_x = \Omega_x^f - \Omega_x^i$. If the proton fluid couples to the boundary much faster than the neutrons, so that the boundary spins up before the neutrons recouple, then the effect of the neutrons can be neglected, and we can estimate the maximum glitch size as

$$\Delta \Omega_{\max}^{\text{est}} = \frac{\rho_p I_3 \langle v_{\text{pn}}^\phi \rangle_{\theta, \phi}^{\max}}{2I_2}. \quad (2.19)$$

For weak mutual friction, equation (2.19) yields $\Delta \Omega_{\max}^{\text{est}} = 3.8 \times 10^{-3}$, compared to $\Delta \Omega_{\max} = 4.6 \times 10^{-3}$ from the simulation. For strong mutual friction, the prediction is worse, $\Delta \Omega_{\max}^{\text{est}} = 1.2 \times 10^{-3}$, versus $\Delta \Omega_{\max} = 2.3 \times 10^{-4}$ from simulation. That the glitch size is overestimated with strong mutual friction is unsurprising, since setting $\Delta \Omega_n = 0$ in equation (2.19) implicitly assumes that the spin-up time is much greater than τ_{MF} , whereas actually the spin-up time is $\approx$ one rotation period $\approx \tau_{\text{MF}}^{\text{strong}}$ (Section 2.5.4). In this regime, mutual friction suppresses spin up, limiting the size of a glitch. The oscillations with period $\approx \tau_{\text{MF}}^{\text{strong}}$ seen in the top panel of Figure 2.8 may be related to the reduction in glitch size with strong mutual friction.

Another interesting result from the bulk glitch simulations is that weak mutual friction produces a permanent change in $\Delta \Omega$, while strong mutual friction restores the pre-glitch trend after 500 time units. During spin down, the neutron condensate is not affected by the spin-down torque initially, and only spins down with the crust and viscous component after a time $\sim \tau_{\text{MF}}$. At $t \lesssim \tau_{\text{MF}}$, the condensate spins down slower than the crust and viscous components, building up a ‘reservoir’ of angular momentum, $J_{\text{res}} = I_n |v_{\text{pn}}^\phi| / R_2$. With weak mutual friction, this reservoir grows at all times during the simulation, while with strong mutual friction it saturates.

When the glitch occurs, angular momentum is transferred from the ‘reservoir’ to the protons, which then spin up the crust. As discussed in Section 2.5.3 spinning up the crust also increases the spin-down rate, so $\Delta \Omega$ decreases for $t \gtrsim 5$, until spin-down equilibrium is restored after $t \approx \tau_{\text{visc}} = 500$. Following a bulk glitch, the angular momentum of the overall system is unchanged, but the ‘observable’ $\Delta \Omega$ is the crust angular velocity, and the crust

only couples to the neutrons when mutual friction is strong. When mutual friction is weak, the neutrons are decoupled from the crust, so the decrease in v_n^ϕ accompanying the increase in v_p^ϕ cannot be seen in $\Delta\Omega$, and the result of the glitch is an increase in the angular momentum of the observable components in the system.

This result has interesting implications in the context of the healing parameter, usually defined as (Wang et al., 2000)

$$Q = 1 - \frac{\Delta\Omega_p}{\Delta\Omega_{\max}}, \quad (2.20)$$

so that $Q \rightarrow 0$ for a glitch that recovers fully, and $Q \rightarrow 1$ for a glitch with no recovery. Taking $\Delta\Omega_p = \Delta\Omega(t = 500)$, we find $Q = 0.143$ with weak mutual friction and $Q = 0.003$ with strong mutual friction. Generally, Q is thought to be a measure of the relative moments of inertia of the inviscid and viscous components locked to the crust, but our results above indicate that mutual friction is important. Specifically, our results affect the standard argument, which holds that if some fraction of the superfluid with moment of inertia I_{res} spins down slower than the crust and viscous components, whose total moment of inertia is denoted I_c , then the ratio I_{res}/I_c provides an upper bound on the angular momentum which is ‘recovered’ by glitches, i.e. if $\dot{\Omega}_{\text{sd}}$ is the average spin-down rate of a pulsar and $\dot{\Omega}_{\text{glitch}}$ is the average spin-up rate from glitches then $I_{\text{res}}/I_c \geq \dot{\Omega}_{\text{glitch}}/\dot{\Omega}_{\text{sd}}$. Often [e.g. Link, Epstein & Lattimer (1999); Andersson et al. (2012)], I_{res} estimated in this way is taken to be an indication of the superfluid fraction in the inner crust, and is used to put constraints on the nuclear equation of state. However, our simulations show that mutual friction affects the size of the reservoir more than ρ_n/ρ_p .

2.7 Conclusions

In this chapter, we perform numerical simulations of pulsar glitches by solving the two-fluid equations of motion for the neutrons and protons in spherical Couette geometry. Glitches are simulated in three ways: firstly, by an instantaneous increase in the angular velocity of the crust; secondly, by an instantaneous increase in the angular velocity of the inner core; and thirdly, by instantaneously locking together the neutron and proton velocities throughout the interior. All three experiments start from a realistic initial state which is set up from initially co-rotating boundaries, after which the outer boundary is spun down by a constant torque on the crust for $\approx 10^2$ rotation periods. In all cases we observe that the angular velocity of the crust increases, but we find that the response of the crust angular velocity following a glitch varies

depending on the way the glitch is activated. Glitches that originate in the crust exhibit an instantaneous angular velocity jump, followed by an exponential relaxation towards the pre-glitch trend. Glitches that originate in the core exhibit a permanent angular velocity increase building up over ≈ 50 rotation periods with no subsequent relaxation. Glitches that are caused by a sudden recoupling of the two fluid components in the bulk have a rise time that is similar to the rotation period of the star, followed by an exponential relaxation. Glitch sizes and the smoothness and completeness of the recovery are affected by the strength of mutual friction. The finding that different glitch activation mechanisms produce different kinds of recoveries may help to explain some of the diversity seen in the population of glitches (Wong, Backer & Lyne, 2001; Espinoza et al., 2011; Haskell & Antonopoulou, 2014).

Our results demonstrate the importance of mutual friction in bulk glitches. When mutual friction is strong, the lag between the two fluid components can reverse between $r = R_1$ and $r = R_2$, so that the neutron condensate spins down faster than the protons in some parts of the star. Mutual friction also affects the size of bulk glitches: with stronger mutual friction glitches are smaller, and the recovery is non-monotonic, exhibiting oscillations with a period similar to the mutual friction time-scale which are damped after an Ekman time. As well as these interesting effects, we also find that mutual friction controls the healing parameter Q , a result which has implications for using pulsar glitches to test models of nuclear matter (Link, Epstein & Lattimer, 1999; Andersson et al., 2012; Newton, Berger & Haskell, 2015).

In order to preserve numerical stability and work with available computational resources, we necessarily make a number of simplifying assumptions. In particular, the choice of superfluid boundary conditions is an open question in the literature, which will be refined in light of future experimental and theoretical work. We also aim to improve the model in the future by including entrainment and vortex tension in the equations of motion and adapting the two-fluid solver to allow for non-uniform densities.

Chapter 3

Nonparametric Estimation of the Size and Waiting Time Distributions of Pulsar Glitches

3.1 Introduction

Rotational glitches are impulsive, irregularly spaced spin-up events observed in pulsars. The discovery of glitches came soon after the birth of pulsar astronomy itself: the first glitch was detected in the Vela pulsar in 1968 (Radhakrishnan & Manchester, 1969; Reichley & Downs, 1969), and in 1969 the first glitch was discovered in the Crab pulsar (Nelson et al., 1970; Lyne, Pritchard & Graham-Smith, 1993). Glitches have been discovered through large-scale monitoring programs with multibeam receivers at the Parkes and Jodrell Bank Observatories (Espinoza et al., 2011; Yu et al., 2013). At the time of writing, 504 (430) events have been detected in 187 (143) objects⁹, amounting to $\sim 10\%$ of the known pulsar population. The glitch catalogues record the instantaneous fractional change in the pulse frequency at the time of the glitch, known as the glitch size, $\Delta\nu/\nu$, and the epoch when each glitch occurs. From the epochs, we calculate the waiting time, Δt , as the difference between the epochs of successive glitches.

Glitches are thought to be caused by sudden readjustments within neutron stars; see Haskell & Melatos (2015) for a recent review of theoretical

⁹Up-to-date catalogues are kept by the Jodrell Bank Centre for Astrophysics at <http://www.jb.man.ac.uk/pulsar/glitches.html> and the Australia Telescope National Facility (ATNF) at <http://www.atnf.csiro.au/research/pulsar/psrcat/glitchTbl.html>. Numbers quoted in the text without (with) parentheses refer to the Jodrell Bank (ATNF) data. Numbers are current as at 2018 May 28.

glitch models. Studying the statistical properties of glitches may therefore lead to new insights into the physics of dense nuclear matter. Previous statistical studies can be split into two categories: those that examine the population of glitching pulsars in aggregate, and those that examine the properties of individual pulsars. In the first category, we mention the work of Morley & García-Pelayo (1993), who used size data to fit a power law to the energy released during glitches, and Lyne, Shemar & Smith (2000), who looked at a sample of 32 glitches in 15 pulsars reported in Shemar & Lyne (1996) and found a correlation between glitch activity and spin-down rate. More recently, Fuentes et al. (2017) studied a sample of 384 glitches in 141 pulsars and found a correlation between glitch activity and the spin-down luminosity of the pulsar. Fuentes et al. (2017) also examined glitch sizes in aggregate and found evidence for a multimodal size probability density function (PDF) by fitting a mixed Gaussian model to the histogram of the data. A similar analysis (with similar results) was performed by Konar & Arjunwadkar (2014). In the second category, we mention the work of Melatos, Peralta & Wyithe (2008), who constructed the empirical cumulative density functions (CDFs) of glitch sizes and waiting times for the nine pulsars with the most recorded glitches and tested for consistency with avalanche models of glitch activity (Warszawski & Melatos, 2011; Warszawski, Melatos & Berloff, 2012; Warszawski & Melatos, 2013). Melatos, Peralta & Wyithe (2008) found that two pulsars, PSR J0537–6910 and PSR J0835–4510, differ from the rest of the pulsars studied in both their size and waiting time PDFs, a finding also reported in Espinoza et al. (2011). Onuchukwu & Chukwude (2016) performed a similar analysis on ‘microglitches’ (jumps in pulsar frequency with $|\Delta\nu/\nu| \sim 10^{-10}$) in 20 pulsars using data from the Hartebeesthoek radio telescope and again interpreted the results in the context of avalanche processes. Studies of the most active glitching pulsar, PSR J0537–6910, reveal a strong linear correlation between size and waiting time to the following glitch (Middleditch et al., 2006; Ferdman et al., 2018; Antonopoulou et al., 2018). Similarly, Shaw et al. (2018) claimed to find a correlation in PSR J0534+2200 between size and waiting time since the previous glitch. Eya, Urama & Chukwude (2017) found that the size PDF in 12 pulsars is well fitted by a normal distribution.

Statistical glitch studies usually posit functional forms for the size and waiting time distributions, e.g. finite mixture models (Konar & Arjunwadkar, 2014), or normal distributions (Eya, Urama & Chukwude, 2017). These methods assume that a set of global parameters define the PDF across its domain. By contrast, a nonparametric estimator makes no assumptions about the global form of the PDF but instead estimates the local probability den-

sity around each data point (Hall, 1992)¹⁰. Nonparametric estimators are widely used in a host of scientific applications (Silverman, 1986). Recently, nonparametric estimation has been used as an independent way to verify the discovery of two interesting new features in the Crab pulsar: a resolved minimum glitch size (Espinoza et al., 2014), and an 11-year episode of accelerated glitch activity (Lyne et al., 2015). In another recent application of nonparametric estimation in pulsar glitch research, Ashton, Prix & Jones (2017) estimated the size distribution of all glitches to assess whether glitches will impact the discovery of gravitational waves from rotating neutron stars. In this paper, we study the individual glitch size and waiting time PDFs for the five pulsars with the highest number of recorded glitches using the kernel density estimator (Wand & Jones, 1995).

We review the algorithm briefly, and evaluate its performance, in Section 3.2. We then apply the algorithm to construct waiting time and size PDF estimates for the most active glitchers in Sections 3.3.1 and 3.3.2 respectively, and the case for multimodality in certain objects is examined critically. We compare the results with previous parametric studies and astrophysical models in Section 3.4. The results strengthen the empirical basis for theoretical work, e.g. by firming up the identification of distinct classes of glitching pulsars. They also offer a guide to designing the next generation of glitch monitoring campaigns with phased radio arrays like LOFAR (Kramer & Stappers, 2010), UTMOST (Caleb et al., 2016) and the Square Kilometer Array (Levin et al., 2018).

3.2 Kernel density estimator

3.2.1 Definition

We begin by defining the nonparametric kernel density estimator (Wand & Jones, 1995) we use to estimate the PDFs of glitch sizes and waiting times. Let $x_1, \dots, x_N$ be N independent, identically distributed samples of a random variable x , with underlying PDF $p(x)$. The kernel density estimator $\hat{p}(x)$ of $p(x)$ is given by

$$\hat{p}(x) = \frac{1}{N} \sum_{i=1}^N K\left(\frac{x - x_i}{h}\right), \quad (3.1)$$

where $K(x)$ is a symmetric, positive definite kernel function $K(x)$ with normalization $\int dy K(y) = 1$. The output of the kernel density estimator is

¹⁰The astrophysicist reader may recognize this as the fundamental idea behind smoothed particle hydrodynamics (Monaghan, 2005).

usually insensitive to the exact shape of $K(x)$. Truncated polynomials and smooth functions multiplied by a Gaussian are common choices (Wand & Jones, 1995), and these kernels produce an equivalent kernel density estimate under an appropriate rescaling of the bandwidth (Marron & Nolan, 1988). An exception is sharply peaked distributions, such as atomic spectra, where a different class of kernels known as ‘infinite order kernels’ are more appropriate (Davis, 1975, 1977). In this paper, we use a Gaussian kernel exclusively; the kernel corresponding to the datum x_i is

$$K_i(x) = \frac{1}{h\sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{x - x_i}{h} \right)^2 \right]. \quad (3.2)$$

The value of the estimator at x is a weighted tally of the observations x_i in a neighborhood $|x - x_i| \lesssim h$ of x ; or, equivalently, it is the unweighted sum of N identical copies of the kernel function, centered at $x_1, \dots, x_N$. Either way, each data point is spread across several bins to give a smoothly differentiable PDF. Kernel density estimation has achieved broad acceptance in many applications because it represents a more optimal trade-off between bias and variance than a bin-centred histogram. We also note that a bin-centred histogram is a kernel density estimator with a rectangular function as the kernel; in this work we use a Gaussian kernel so that the estimator inherits the useful properties of continuity and differentiability.

The key challenge when applying equation (3.1) is to select the bandwidth h in a way that ensures good practical performance. The ensemble-averaged, x -integrated bias $\int dx \langle \hat{p}(x) - p(x) \rangle$ and variance $\int dx \langle [\hat{p}(x) - p(x)]^2 \rangle$ increase and decrease respectively, as h increases in the limit $N \rightarrow \infty$. Hence optimizing h involves a compromise between bias and variance. One approach is to let h vary with x according to the local density of data points (see footnote 10) but it is ill-suited to small glitch samples. An alternative, which has proven its worth in many applications, is to choose a single, global h , that minimizes the asymptotic mean integrated square error (Wand & Jones, 1995),

$$\text{AMISE} = \frac{1}{Nh} \int dx [K(x)]^2 + \frac{h^4}{4} \left[\int dx x^2 K(x) \right]^2 \int dx [p''(x)]^2. \quad (3.3)$$

As $p''(x)$ is unknown a priori, it must be approximated. Many techniques have been developed to estimate $p''(x)$ in order to select h to minimise the AMISE [a summary of several of the more common techniques can be found in Wand & Jones (1995)]; in this work we use the normal reference bandwidth¹¹,

¹¹This is, in effect, a second-order parametric assumption. The normal reference band-

which assumes that $p(x)$ is a Gaussian with variance σ^2 . In this case, the h that minimises the AMISE can be written analytically,

$$h = \left\{ \frac{8\pi^{1/2} \int dx [K(x)^2]}{3N[\int dx x^2 K(x)^2]} \right\}^{1/5} \sigma. \quad (3.4)$$

The normal reference bandwidth rule replaces the true variance in equation (3.4) with some estimate $\hat{\sigma}^2$ (Silverman, 1986; Scott, 1992). We take $\hat{\sigma}$ to be the minimum of the sample’s interquartile range and standard deviation. It is important to realise that the AMISE is, as the name implies, an asymptotic measure. An h chosen to minimise the AMISE is just an approximation to the bandwidth that truly optimizes the trade-off between bias and variance in the $N \leq 35$ regime of pulsar glitch statistics. In Section 3.2.4 we explore the effect of bandwidth selection on synthetic data drawn from a bimodal distribution. In Section 3.3 we examine how the estimated PDFs change qualitatively with h .

3.2.2 Positive definite variables

If a PDF is defined on a finite domain and is non-zero at an endpoint, the kernel density estimator overspills the boundary. Relative to an estimator that takes the boundary into account, the local bias $\langle \hat{p}(x) - p(x) \rangle$ in the vicinity of the boundary is greater than the x -integrated bias $\int dx \langle \hat{p}(x) - p(x) \rangle$. When equation (3.1) is applied to a positive definite random variable, such as glitch size or waiting time, probability leaks spuriously into the region $x < 0$. Leakage is significant, when a sizable fraction of the data points satisfy $x_i \lesssim h$. Renormalization does not fix the problem, as it introduces extra bias near $x = 0$. Many methods have been developed to counteract leakage. Here we describe two procedures which are robust and straightforward to implement: (i) reflect the data about $x = 0$ and apply (3.1) with a symmetric kernel;¹² or (ii) transform the data (e.g. by taking their logarithm). Taking the logarithm compresses the domain of the size variable (which can otherwise extend over 4 dex in an individual object), so that it can be modelled usefully by a single, global h chosen according to (3.4). We apply the reflection method to the

width tends to produce an ‘oversmoothed’ estimate, sacrificing bias in order to reduce variance, and is more appropriate in the small- N regime than other techniques such as Sheather-Jones plug-in (Sheather & Jones, 1991) and smoothed cross-validation (Hall, Marron & Park, 1992).

¹²One can also reflect about a right-hand boundary, but this is irrelevant for glitches, whose observed sizes and waiting times are much smaller than the maxima that radio timing experiments can detect.

waiting time data in Section 3.3.1 and the logarithmic transform method to the size data in Section 3.3.2.

3.2.3 Error estimation with nonparametric methods

An obvious question when using the kernel density estimator is how accurate are the estimates produced? When using a parametric estimator, one can construct confidence intervals by varying the parameters within a range and testing the likelihood of excluding the resulting fit, e.g. with a K-S test as in Figures 3.2 and 3.3. With the kernel density estimator, a similar approach is hindered by two issues: the set of estimated parameters is large (formally the number of points where the curve is estimated), and bias is one of the main contributors to the inaccuracy of the estimate (cf. parametric estimators, where the variance dominates) (Hall, 1992; Chaudhuri & Marron, 1999). One approach to constructing confidence intervals that may seem appealing is to create data replicates by resampling and construct $\hat{p}$ for the replicated data sets. The simplest resampling ‘bootstrapping’ method is inappropriate, because it produces a zero-bias confidence interval on the estimate $\hat{p}$ rather than on the true PDF (Hall, 1992). Rather than seeking to construct a (potentially misleading) quantitative measure of the goodness of fit, we instead take a qualitative approach, which aims to test whether particular features of the estimated distribution, e.g. monotonicity and multimodality, are robust features that appear for a wide range of bandwidth choices (Chaudhuri & Marron, 1999; Marron & Chung, 2001). This concept is discussed further in Section 3.2.4 and Section 3.3 below.

3.2.4 Validation

Before analyzing actual data, we run some tests on synthetic data drawn from the exponential distribution $p(x) = e^{-x}$, which acts as a proxy for the glitch waiting time distribution, and the power law distribution $p(x) = 0.2x^{-1.2}$, $x \geq 1$ (Melatos, Peralta & Wyithe, 2008). The tests are not a substitute for definitive convergence studies presented elsewhere (Silverman, 1986) but give some sense of the reliability of the estimator and, importantly, the sorts of artifacts (e.g. wiggles, plateaus) that arise from noise. To construct kernel density estimates, we use the statistical software R (R Core Team, 2014), and the `bkde` function included in the package `KernSmooth` (Wand, 2014).

Figure 3.1 illustrates qualitatively how the estimator (3.1) performs with N , when the underlying PDF is an exponential (top row) or a power law

(bottom row). For the exponential distribution, the reflection boundary correction described in Section 3.2.2 is used. For the power law distribution, the kernel density estimate is applied to the logarithm¹³ of x . Applying the estimate to $\log_{10}(x)$ reduces fluctuations (bumps and large gaps) that would otherwise appear, so that a global bandwidth gives good performance. As the domain is restricted to $x \geq 1$, and the probability density is significant at $x = 1$, we apply the reflection method to the logarithmic data at $\log_{10} x = 0$. In the top left panel, the colored curves show $\hat{p}(x)$ for three individual realizations with $N = 25, 100$, or 1000 for the exponential distribution (solid black curve). The bottom left panel shows $\hat{p}[\log_{10}(x)]$ for three realizations with $N = 25, 100$, or 1000 , for the power law distribution (solid black line). In the right column, each solid gold curve shows $\hat{p}(x)$ and $\hat{p}[\log_{10}(x)]$ for one realisation with $N = 25$, in order to give a sense of the scatter in the estimator. We have checked many realizations and the results consistently exhibit the following properties. (i) Wiggles appear in $\hat{p}(x)$ for $N \leq 100$. They are noise artifacts, which should not be interpreted as multimodality when analysing real data in Section 3.3. (ii) Broadly speaking, the estimator performs similarly on the power law and the exponential. (iii) The estimator plateaus at $x = 0$, when the data are reflected (top row), because K being symmetric implies $\hat{p}'(0) = 0$. Hence $\hat{p}(x)$ systematically underestimates $p(x)$ near $x = 0$, if the underlying PDF is cuspy there, although the bias decreases in a controlled fashion as N increases¹⁴. In the top-right panel, $\langle \hat{p}(x=0)/p(x=0) \rangle = 0.71$, with a standard deviation of 0.18; in the bottom-right panel, $\langle \hat{p}[\log_{10}(x)]/p[\log_{10}(x)] \rangle = 0.72$, with a standard deviation of 0.2.

Another issue worth addressing is the performance of the kernel density estimator for multimodal PDFs, which have been proposed for glitch sizes in previous aggregated studies (Konar & Arjunwadkar, 2014; Ashton, Prix & Jones, 2017; Fuentes et al., 2017) and in analyses of individual pulsars (Melatos, Peralta & Wyithe, 2008). We aim to determine whether the kernel density estimator can reliably reproduce multiple peaks in a PDF as h varies. As a test, we consider a bimodal PDF obtained by summing two Gaussians, viz. $p(x) = 0.6\mathcal{N}(\mu = 4, \sigma = 1) + 0.4\mathcal{N}(\mu = 8, \sigma = 1)$, where $\mathcal{N}(\mu, \sigma)$ is the normal distribution with mean μ and standard deviation σ . We draw a variate of dimension $N = 25$ from $p(x)$, and construct the kernel density estimate with three choices of bandwidth: the normal reference bandwidth,

¹³If a variable x is distributed according to a power law, i.e. $p(x) \propto x^{-a}$, then the PDF of the logarithmic variable $y = \log_{10}(x)$ obeys $q(y) \propto (10^y)^{1-a}$.

¹⁴In the absence of reflection, leakage can underestimate $p(x = 0)$ by up to a factor of two, independent of N . One can recognise this by noting that the positive and negative points in the reflected data contribute equally to $\hat{p}(x = 0)$.




Figure 3.1: Convergence of the kernel density estimator $\hat{p}(x)$ [equation (3.1)] for an exponential PDF, $p(x) = e^{-x}$ (top row), and $\hat{p}[\log_{10}(x)]$ for the power law PDF, $p(x) = 0.2x^{-1.2}$ with $x \geq 1$ (bottom row). The solid black curve in the top row corresponds to the exponential PDF, The solid, colored curves in the left column correspond to $\hat{p}(x)$ and $\hat{p}[\log_{10}(x)]$ for realizations with $N = 25$ (blue), 100 (red), and 1000 (green). The tick marks in the panels in the left column indicate the abscissae of the data in the $N = 25$ (blue) realizations. The dashed black line in the bottom row shows the power law PDF. The right-hand column shows 10 instances of $\hat{p}(x)$ and $\hat{p}[\log_{10}(x)]$ with $N = 25$ (gold curves). Equation (3.1) is applied to the reflected data in the top row and the reflected base-10 logarithm of the data in the bottom row.




Figure 3.2: Kernel density estimates of a two-component Gaussian PDF $p(x) = 0.6\mathcal{N}(\mu = 4, \sigma = 1) + 0.4\mathcal{N}(\mu = 8, \sigma = 1)$, with $N = 25$ data points. The solid black curve shows the kernel density estimate with bandwidth h_0 chosen by the normal reference rule, the dotted grey curve is the kernel density estimate with bandwidth $h_0/2$ and the dot-dashed grey curve is with bandwidth $2h_0$. The dashed green curve shows the underlying PDF from which the data used to construct the estimator, shown as black tick marks along the x -axis, are drawn.

h_0 , as well as $h_0/2$, and $2h_0$. The results are shown in Figure 3.2. Figure 3.2 illustrates the characteristic property (Silverman, 1986) of the kernel density estimator with the normal reference bandwidth: the features of the underlying PDF are recovered in the estimator, which has local maxima at $x \approx 4$ and $x \approx 8$, but the estimator is somewhat ‘oversmoothed’, with $p(x)$ at the peaks underestimated by $\approx 30\%$ and $p(x)$ at the minimum overestimated by a factor of ≈ 2 . The estimator with $h = h_0/2$ tends closer to the underlying PDF, however, the location of the peaks is less accurate. The estimator with $h = 2h_0$ is excessively oversmoothed, as it completely fails to capture the bimodality in $p(x)$, instead producing a single, broad peak centered about $x \approx 5$. As Figure 3.1 shows, it is possible for the kernel density estimator to produce local maxima even when the underlying PDF is monotonic, so we tend to prefer the normal reference bandwidth with its tendency to slightly oversmooth. However, in our analysis of glitch data in Section 3.3 we use the same set of bandwidths as in Figure 3.2 in order to illustrate the uncertainty

inherent to the method (Chaudhuri & Marron, 1999).

3.3 Probability density functions

According to the Jodrell Bank catalogue at 2018 May 28, 504 glitches have been detected in 187 pulsars¹⁵. Yu & Liu (2017) performed a detectability study of glitches in the Yu et al. (2013) data set and concluded that all detectable glitches in these data had been identified. Espinoza et al. (2014) performed a detectability study of glitches in the Crab pulsar with Jodrell Bank data and reached the same conclusion; see also Janssen & Stappers (2006). Other than these two studies, however, no other authors have presented results on the completeness of glitch catalogues. We stress that both the Jodrell Bank and the ATNF glitch catalogues may not be complete, and our analysis in this paper may be based on an incomplete data set. In addition to the glitches in the Jodrell Bank catalogue, approximately 20 additional glitches have been reported in PSR J0537–6910 in two papers by Ferdman et al. (2018); Antonopoulou et al. (2018). In this section, we construct size and waiting time PDFs for the five objects that have glitched most prolifically: PSR J1740–3015 ($N = 35$), PSR J0534+2200 (26), PSR J0537–6910 (42),¹⁶ PSR J1341–6220 (23), and PSR J0835–4510 (21). Experience across many scientific applications teaches that one needs $N \gtrsim 20$ to get meaningful nonparametric results (Silverman, 1986), as verified by Figure 3.1. The sample sizes for the next most active objects, PSR J0631+1036 ($N = 15$) and PSR J1801–2304 ($N = 13$), are too small to be analysed nonparametrically [cf. Melatos, Peralta & Wyithe (2008)]; their estimated PDFs are too bumpy.

3.3.1 Waiting times

Figure 3.3 displays the waiting time PDFs for the five pulsars selected above. In each panel, the nonparametric kernel density estimator $\hat{p}(\Delta t)$ (units: yr^{-1}) from equation (3.1), using the normal reference bandwidth, $h = h_0$, is graphed on a linear scale as a solid curve. In the left column of 3.3, we displays fits to the Poisson PDF, $p(\Delta t) = \lambda e^{-\lambda \Delta t}$. The red dashed curve corresponds to the maximum likelihood estimate $\lambda = \langle \Delta t \rangle^{-1}$. The grey dotted

¹⁵We work with the larger Jodrell Bank catalogue henceforth and cross-check against the ATNF catalogue.

¹⁶There is tension in the number of glitches and the parameters of glitches for PSR J0537–6910 between Ferdman et al. (2018); Antonopoulou et al. (2018) and earlier work by Middleditch et al. (2006). We cross-check all three references and include the glitches common to two out of three sources in our analysis.

curves correspond to the values of λ which demarcate the interval (λ_-, λ_+) , where the null hypothesis that the data are drawn from a Poisson distribution is rejected by a K-S test with less than 68% confidence (Melatos, Peralta & Wyithe, 2008). In the right-hand column, we display two additional kernel density estimators using different bandwidths. The dashed red curve is with $h = h_0/2$, the dotted blue curve is with $h = 2h_0$.

Figure 3.3 reveals three important things. First, broadly speaking, $\hat{p}(\Delta t)$ decreases with Δt for PSR J0534+2200 and PSR J1740–3015, notwithstanding the bumps in the body and tail of the PDF estimate and the low- Δt plateau, which are both peculiarities of the kernel density estimator, as implied by Figure 3.1. It is apparent by eye that $\hat{p}(\Delta t)$ is broadly consistent with the parametric Poisson PDF posited previously (Wong, Backer & Lyne, 2001; Melatos, Peralta & Wyithe, 2008; Espinoza et al., 2011); the dashed curve in the left-hand column hugs the solid curve and falls between (and nearly coincides with) the dotted fits from Melatos, Peralta & Wyithe (2008), even after adding three events since 2008.

Second, $\hat{p}(\Delta t)$ rises to a maximum in the other objects. This is interesting, since PSR J0537–6910 and PSR J0835–4510 are traditionally classified as quasiperiodic glitchers (Middleditch et al., 2006; Melatos, Peralta & Wyithe, 2008), but PSR J1341–6220 was previously classified as Poisson-like (Melatos, Peralta & Wyithe, 2008). The peak in $\hat{p}(\Delta t)$ in the latter object is broad; in fact, the dashed Poisson curve is still a fair fit to the solid curve, albeit displaced downwards. However, $\hat{p}(0)$ is less than 30% of $\hat{p}(\Delta t)$ at the peak for all three non-Poisson-like objects, and the maximum in $\hat{p}$ is at $\Delta t > 0$ using the most oversmoothing bandwidth $h = 2h_0$; suggesting the existence of a maximum at $\Delta t > 0$ instead of $\Delta t = 0$.

Third, there is some evidence of bimodality in PSR J0835–4510 and PSR J1341–6220, both of which show local minima in $\hat{p}(\Delta t)$ with $h = h_0$ and $h = h_0/2$. Quasiperiodic glitchers have been modelled with a two-component PDF in previous work (Melatos, Peralta & Wyithe, 2008; Konar & Arjunwadkar, 2014; Ashton, Prix & Jones, 2017; Fuentes et al., 2017). In particular, it is sometimes said that the ‘big’ glitches in PSR J0835–4510 are spaced regularly, but the ‘small’ ones are not. The short- Δt component, if present, arguably dominates the long- Δt component in PSR J1341–6220, whereas the situation is the other way around in PSR J0835–4510. Looking at $\hat{p}(\Delta t)$ for PSR J0534+2200 and PSR J1740–3015, with $h = h_0/2$, however, shows that this choice of bandwidth significantly ‘undersmooths’ the PDF, producing multiple peaks centred about one or two events. On balance, we do not see convincing evidence for bimodality in $\hat{p}(\Delta t)$ in Figure 3.3. The gentle bumps in Figure 3.3 look much like the small- N features identified in the validation experiments in Figure 3.1 due to sparse sampling of the




Figure 3.3: Kernel density estimates of the waiting time PDFs $\hat{p}(\Delta t)$ for the five most active glitchers, plotted on linear axes. Solid curves correspond to the nonparametric kernel density estimator with a normal reference bandwidth $h = h_0$. In the left-hand column, red dashed curves are parametric Poisson fits with $\lambda = \langle \Delta t \rangle^{-1}$ (maximum likelihood) computed directly by averaging the data. Grey dotted curves are Poisson fits with $\lambda = \lambda_{\pm}$ from the K-S analysis in Melatos, Peralta & Wyithe (2008). In the right-hand column, the red dashed curves are the kernel density estimator with bandwidth $h = h_0/2$, the blue dotted curve is with $h = 2h_0$. Tick marks on the horizontal axis indicate the abscissae of the raw data.

PDF. More data are required to rule out bimodality definitively, but there are certainly no strong grounds for claiming its existence on the basis of a nonparametric analysis at present.

3.3.2 Sizes

Figure 3.4 displays the estimated PDF of the fractional size logarithm, $s = \log_{10}(\Delta\nu/\nu)$. In the left column, the nonparametric estimate $\hat{p}(s)$ from (1) using a normal reference bandwidth, $h = h_0$, is graphed as a solid curve. Also displayed are parametric fits to the power law distribution $p(\Delta\nu) = (1 - a)(\Delta\nu_{\max}^{1-a} - \Delta\nu_{\min}^{1-a})^{-1}\Delta\nu^{-a}$. The exponent a is estimated by maximum likelihood (red dashed line) and by calculating the end-points $a = a_{\pm}$ (grey dotted lines) of $a_- \leq a \leq a_+$, where the K-S probability exceeds 32% as in Figure 3.2; see Melatos, Peralta & Wyithe (2008) for details. There is an art to choosing $\Delta\nu_{\min}$ and $\Delta\nu_{\max}$, as discussed in Section 4.3 in Melatos, Peralta & Wyithe (2008). The smallest glitch observed is likely to be a reasonable estimate of $\Delta\nu_{\min}$, because $p(\Delta\nu)$ rises steeply as $\Delta\nu \rightarrow 0$, but this has not been quantified systematically except by Janssen & Stappers (2006), who simulated microglitch detection in a noisy time series and found $\Delta\nu_{\min} = 10^{-10}\nu$ for recent Jodrell Bank observations; see also Espinoza et al. (2014), where it is claimed that the lower cut-off in $p(\Delta\nu)$ can be resolved. Below we set $\Delta\nu_{\min}$ and $\Delta\nu_{\max}$ to the observed minimum and maximum respectively; the results are insensitive to this choice, as demonstrated previously (Melatos, Peralta & Wyithe, 2008). If the glitch size distribution truly follows a power law, then it is bounded from below and the estimator $\hat{p}(s)$ should include a reflection boundary correction as in Section 3.2.4 and Figure 3.1. However, we have neglected to perform this boundary correction, since we don't know where (if anywhere) the boundary is. In the right column of Figure 3.4, we show $\hat{p}(s)$ with $h = h_0$ as a black solid curve, and also $\hat{p}(s)$ with $h = h_0/2$ (dashed red curve) and $h = 2h_0$ (dotted blue curve).

Figure 3.4 elaborates the classification stemming from Figure 3.3 in two interesting ways. First, we find that $\hat{p}(s)$ generally decreases with s , with a prominent peak at $\approx s_{\min}$, for the Poisson-like duo, PSR J0534+2200 and PSR J1740–3015, whose waiting time PDFs also decrease ¹⁷. In contrast, PSR J0537–6910 and PSR J0835–4510 are not monotonic, with $\hat{p}(s_{\min})$ less than 20% of $\hat{p}(s)$ at the peak and the bulk of the probability confined within $\lesssim 1$ dex. PSR J1341–6220, previously classified as Poisson-like by Melatos, Peralta & Wyithe (2008) and now a quasiperiodic candidate on

¹⁷The drop in the PDF at $s < s_{\min}$ is an artifact; it traces the leftward Gaussian tail of the leftmost kernel.




Figure 3.4: Kernel density estimates of the fractional size PDFs $\hat{p}(s)$, for the five most active glitches, plotted on log-log axes. The solid black curve is $\hat{p}(s)$ using a normal reference bandwidth, $h = h_0$. In the left-hand column, the dashed red lines show the parametric power-law fit with a estimated by maximum likelihood. The grey dotted lines are power-law fits with $a = a_{\pm}$ from the K-S analysis in Melatos, Peralta & Wyithe (2008). In the right-hand column, the red dashed curves are $\hat{p}(s)$ with $h = h_0/2$, and the blue dotted curves are $\hat{p}(s)$ with $h = 2h_0$. Tick marks on the horizontal axis indicate the abscissae of the raw data.

the basis of Figure 3.1, displays interesting intermediate behavior: $\hat{p}(s)$ does not decrease monotonically but nor is it narrowly peaked, staying nearly flat over ≈ 2 dex. This may mean one of two important things: (i) the broad waiting time plateau in Figure 3.1 actually extends to $\Delta t = 0$ and would emerge more clearly given more data, ultimately putting PSR J1341–6220 in the Poisson class; or (ii) the correspondence between power-law size and exponential waiting time statistics is not one-to-one (Melatos, Peralta & Wyithe, 2008).

Second, it is apparent by eye that in the two cases where $\hat{p}(s)$ decreases monotonically (PSR J0534+2200 and PSR J1740–3015), $\hat{p}(s)$ is broadly consistent with a power law. The dashed line is consistent with the solid curve and is bracketed by the dotted fits from Melatos, Peralta & Wyithe (2008), even after adding three events since 2008. Interestingly, the Poisson-like duo have $1.11 \leq a \leq 1.22$ for the best-fitting dashed line, whereas PSR J1341–6220 has $a = 0.92$ ¹⁸. This division is significant physically, because the mean is dominated by different extremes of the distribution in the two cases, with $\langle \Delta \nu \rangle = (a - 1)(2 - a)^{-1}(\Delta \nu_{\max}/\Delta \nu_{\min})^{1-a}\Delta \nu_{\max}$ for $1 < a < 2$ and $\langle \Delta \nu \rangle = (a - 1)(2 - a)^{-1}\Delta \nu_{\max}$ for $a < 1$ (assuming $\Delta \nu_{\min} \ll \Delta \nu_{\max}$).

It is sometimes claimed in the literature that the quasiperiodic glitchers harbor two distinct event populations (Melatos, Peralta & Wyithe, 2008). We find no compelling evidence to support this claim in Figure 3.4. With $h = h_0$, there are weak hints of two peaks at $s \approx -7.7$ and $s \approx -6.5$ in PSR J0537–6910, the larger events being more frequent. The hypothetical populations are well separated relative to h , but the lesser group contains just four out of 40 events, which are probably random outliers; the validation experiments in Figure 3.1 routinely produce a few outliers in any given realization, which show up as a small, low- s bump for $N \leq 25$. Likewise, in PSR J1341–6220, the broad plateau in $\hat{p}(s)$ contains a hint of two bumps at $s \approx -7.4$ and $s \approx -6.1$, but again similar noise artifacts are seen in Figure 3.1. As in Figure 3.3, using $h = h_0/2$ produces multiple spurious peaks in every object, suggesting this bandwidth is undersmoothing significantly, while $h = 2h_0$ produces a similarly-shaped estimate with a single peak that is broadly consistent with the result for $h = h_0$. More data may alter this picture, but for now the nonparametric case for bimodality in glitch sizes is weak in the five objects analysed here.

¹⁸Recall from footnote 4 that the PDF of the logarithm of a power law-distributed variable with power law index a is a power law with index $1 - a$, and hence has a positive slope for $a < 1$.

3.4 Conclusion

As the number of detected radio pulsar glitches grows, it is ever more feasible to disaggregate the data and construct size and waiting time distributions for individual objects. Previous statistical analyses have assumed theoretically inspired parametric PDFs (e.g. power law, exponential, finite mixture) and tested for K-S inconsistency (Melatos, Peralta & Wyithe, 2008; Espinoza et al., 2011; Konar & Arjunwadkar, 2014; Onuchukwu & Chukwude, 2016). In this paper, we estimate the PDFs nonparametrically using the kernel density estimator, a powerful tool which converges optimally in many circumstances and has proved its mettle in many applications, where data sets are modest in size (Wand & Jones, 1995).

Our analysis yields three main results. (i) It confirms the existence of two classes of glitcher: Poisson-like objects with monotonically decreasing waiting time and size PDFs, consistent with exponential and power-law functional forms respectively, and objects with peaked waiting time and size PDFs, which trigger quasiperiodically. (ii) It suggests that one object which was previously classified as Poisson-like, PSR J1341–6220, shows evidence of hybrid behaviour: $\hat{p}(\Delta t)$ peaks at $\Delta t > 0$, but the peaks in $\hat{p}(\Delta t)$ and $\hat{p}(s)$ are broader than for PSR J0537–6910 and PSR J0835–4510, and $\hat{p}(\Delta t)$ as a whole remains marginally consistent with an exponential distribution in a K-S sense. (iii) One sees weak hints of bimodality in $\hat{p}(\Delta t)$ and $\hat{p}(s)$, consistent with previous parametric modeling of quasiperiodic glitchers (Melatos, Peralta & Wyithe, 2008) and a modest excess of large glitches in aggregate statistics (Lyne, Shemar & Smith, 2000). However, Monte-Carlo experiments (Figure 3.1) and a sensitivity study addressing bandwidth selection in Section 3.3.1 and Section 3.3.2 present an inconclusive picture. The putative bimodality is as likely to be a noise artifact as not, given the data at hand.

Do the above results shed new light on glitch physics? This complicated question will be addressed in depth in a forthcoming theoretical paper. Here we venture to make one, model-independent point without favoring any of the microphysical mechanisms referenced in Section 3.1. It is thought that glitches are driven by the electromagnetic torque, which spins down the stellar crust differentially with respect to internal components. Stresses of various kinds (e.g. elastic forces in the context of crust quakes, Magnus forces in the context of superfluid vortex dynamics) build up globally yet inhomogeneously, until a glitch is triggered, whereupon they relax locally via chains (‘avalanches’) of threshold-activated, stick-slip events (e.g. crust cracking, or vortex unpinning). The avalanches must be cooperative phenomena mediated by a knock-on process, otherwise the central limit theorem predicts size and waiting time PDFs dramatically narrower than those observed,

given the large number of interacting elements involved (e.g. $\gtrsim 10^{10}$ vortices) (Warszawski, Melatos & Berloff, 2012; Warszawski & Melatos, 2013; Fulgenzi, Melatos & Hughes, 2017). Avalanche systems of this sort, like sand piles, tend towards a self-organized critical state and operate in two regimes: (i) slow driving, where successive events are triggered in spatially distinct regions and are mutually independent, i.e. Poisson-like, with scale-invariant size distributions, and (ii) fast driving, where successive events involve the whole system (e.g. all the vortices in a superfluid, or all the grains in a sand pile) and recur quasiperiodically with roughly equal sizes (Jensen, 1998; Melatos, Peralta & Wyithe, 2008). In this sense, therefore, the nonparametric PDF estimates in Figures 3.3 and 3.4 are broadly consistent with theoretical expectations, irrespective of the microphysics. However, three additional implications flow from the new results in this paper. First, with the hybrid behaviour of PSR J1341–6220 defying easy categorization, it makes sense to look harder for independent ways to classify radio pulsars as slowly/rapidly driven in the sense above¹⁹. Second, among the quasiperiodic candidates, $\hat{p}(s)$ is considerably narrower in PSR J0835–4510 (≈ 0.5 dex) than in the others (≈ 2.5 dex), and $\hat{p}(\Delta t)$ is relatively broad in PSR J1341–6220 (indeed almost exponential), raising the interesting possibility that the nexus between power-law size and exponential waiting time statistics, characteristic of self-organized criticality, may sometimes be broken. Third, if multimodality is uncovered in any object in the future, as more data are collected, it will argue for the existence of an additional, non-scale-invariant trigger in the glitch mechanism, at least in some pulsars. How this relates (if at all) to the violation of scale-invariance implied by the minimum glitch size resolved in PSR J0534+2200 (Espinoza et al., 2014) is an interesting question for future work.

¹⁹Preliminary analysis of the next most active glitchers, PSR J0631+1036 ($N = 15$) and PSR J1801–2304 ($N = 13$), suggests that the former is Poisson-like, and the latter is quasiperiodic with a broad $\hat{p}(\Delta t)$ like PSR J1341–6220. However, N is too small to interpret the results reliably. More data are needed and will be forthcoming soon, as both objects are active, with $\langle \Delta t \rangle = 1.1$ yr and 2.3 yr respectively.

Chapter 4

Simulating pulsar glitches: an N -body solver for superfluid vortex motion in two dimensions

4.1 Introduction

A superfluid supports rotation through the formation of an array of quantized vortex lines (Onsager, 1949; Feynman, 1955). In a container with cylindrical symmetry, the array is rectilinear, if the angular velocity is less than the critical value, where instabilities set in (Glaberson, Johnson & Ostermeier, 1974). Over length-scales much larger than the inter-vortex separation, the fluid mimics solid-body rotation, the motion of individual vortices can be averaged over volume, and key aspects of the flow can be described hydrodynamically, e.g. via a multi-fluid system (Andersson & Comer, 2006). However, there are some macroscopic properties of the flow whose treatment requires tracking the motion of individual vortices, e.g. transport coefficients in a vortex tangle (Peralta, 2006; Andersson, Sidery & Comer, 2007), and far-from-equilibrium phenomena such as vortex pinning (Alpar, 1977; Haskell & Melatos, 2016; Drummond & Melatos, 2018).

One system where vortex pinning is important is the interior of a neutron star. Due to their extreme density ($\sim 10^{14}$ g cm $^{-3}$) and relatively low temperature ($k_{\text{B}}T \approx 10^6$ eV), neutrons inside a neutron star are believed to condense into a superfluid phase (Baym, Bethe & Pethick, 1971; Pines & Alpar, 1985). The superfluid neutrons couple loosely to other components of the star, such as the solid crust and an interpenetrating fluid of charged

particles. As the crust brakes electromagnetically, the superfluid vortices migrate outwards. However, nuclear lattice sites and magnetic flux tubes ‘pin’ the vortices by providing an attractive force, which resists the Magnus force which brings the neutrons and the crust to co-rotation. Hence the neutrons lag the crust. As the lag grows, so does the Magnus force, until the vortices unpin. If enough vortices unpin, the resulting back-reaction on the crust is observable as an instantaneous increase in the frequency of the pulsar, i.e. a pulsar glitch (Anderson & Itoh, 1975; Haskell, Pizzochero & Sidery, 2012).

Another application of vortex pinning with close connections to pulsar glitches is laboratory experiments on magnetic flux tube avalanches in type II superconductors subjected to changing magnetic fields. In these systems, magnetic flux is distributed as a triangular array of quantized flux tubes (Abrikosov, 1957; Bean, 1964; Fetter, Hohenberg & Pincus, 1966). As the field ramps down, flux tubes are expelled from the superconductor. It is observed that pinning of the flux tubes causes the expulsion to occur in bursts involving up to thousands of flux tubes (Field et al., 1995; Altshuler & Johansen, 2004).

Gross-Pitaevskii simulations of small, idealized systems containing $\sim 10^2$ vortices and $\sim 10^4$ pinning sites demonstrate that unpinning occurs collectively (Warszawski & Melatos, 2011; Warszawski, Melatos & Berloff, 2012; Warszawski & Melatos, 2013; Melatos, Douglass & Simula, 2015). Analytic studies have found that collective unpinning is sensitive to the strength of pinning and to stellar parameters such as mass and temperature, which affect how far vortices can move before re-pinning (Haskell & Melatos, 2016). One-dimensional hydrodynamic simulations reveal that accumulation of vorticity in one region of a neutron star can lead to glitch-like travelling waves which reduce differential rotation (Khomenko & Haskell, 2018). The latter simulations rely on assumptions of how pinned vortices are distributed within a neutron star which will be improved by a better understanding of the microscopic pinning dynamics, an important motivation for this paper. Collective unpinning is hard to study theoretically in the many-vortex regime, due to the computational expense of scaling up Gross-Pitaevskii simulations (neutron stars have $\sim 10^{18}$ vortices and $\sim 10^{50}$ pinning sites, for example). While existing simulations show clear evidence of unpinning knock-on through short-range interactions and long-range acoustic processes (Warszawski, Melatos & Berloff, 2012), it is difficult to reliably determine the many-vortex size and waiting time probability distribution functions (PDFs) from simulations of small systems. More broadly, it is unclear how the knock-on behaviour scales up to larger systems. Similar comments apply to experiments with type II superconductors.

The ‘vortex avalanche’ model resembles systems exhibiting self-organized

criticality (SOC) (Bassler & Paczuski, 1998; Melatos, Peralta & Wyithe, 2008), a paradigm that has found applications in numerous fields of study; see Watkins et al. (2015) for a review. Power-law size and exponential waiting-time PDFs are characteristic of SOC. They are observed in the glitch histories of two pulsars, PSR J0534+2200 and PSR J1740–3015 (see Chapter 2). Two other pulsars, PSR J0537–6910 and PSR J0835–4510 exhibit quasiperiodicity in their glitch size and waiting time PDFs, which is also predicted for fast-driven SOC systems (Jensen, 1998). In experimental studies on superconducting flux tube avalanches, the PDF of burst sizes is a power law, consistent with the predictions of SOC (Field et al., 1995).

In this chapter, we describe a two-dimensional, N -body vortex filament code that can handle $N \sim 10^4$ vortices, given reasonable computational resources. In Section 4.2 we describe the mathematical framework of the solver. In Section 4.3 we describe the algorithm and its implementation. In Section 4.4 we validate the numerical method through a systematic set of numerical tests. In Section 4.5 we illustrate the astrophysical applications of the solver by performing numerical experiments on a pinned, decelerating superfluid to produce SOC-like vortex avalanches and investigate them in the context of pulsar glitches.

4.2 Vortex dynamics

We study the motion of a system of N point-like rotational vortices in a two-dimensional geometry with cylindrical symmetry. On scales larger than the inter-vortex separation, the quantum mechanical structure of the vortices can be ignored, and the vortices move according to classical hydrodynamics. In the absence of external influences such as lift forces and pinning, the velocity $d\mathbf{x}/dt$ of a quantized vortex at position $\mathbf{x}(t)$ is the same as the bulk fluid velocity at $\mathbf{x}(t)$ induced by all the other vortices. (In practice, we calculate the bulk velocity from the vorticity; see Section 4.2.2). The convective motion is supplemented by other effects, such as the interaction of vortices with boundaries, impurities, and other, viscous fluid components.

4.2.1 Equations of motion

In a reference frame that co-rotates with the container, the position of a vortex in Cartesian coordinates, (x_i, y_i) , evolves according to

$$\frac{d}{dt} \begin{pmatrix} x_i \\ y_i \end{pmatrix} = \mathcal{R}_\phi \begin{pmatrix} v_{i,x} \\ v_{i,y} \end{pmatrix}, \quad (4.1)$$

with

$$v_{i,x} = -\sum_{j \neq i} \frac{\kappa y_{ij}}{r_{ij}^2} + \sum_{j=1}^N \frac{\kappa y_{ij,\text{image}}}{r_{ij,\text{image}}^2} + \omega y_i - \sum_k \frac{\partial V(\mathbf{x}_i - \mathbf{x}_k)}{\partial y_i} \quad (4.2)$$

$$v_{i,y} = \sum_{j \neq i} \frac{\kappa x_{ij}}{r_{ij}^2} - \sum_{j=1}^N \frac{\kappa x_{ij,\text{image}}}{r_{ij,\text{image}}^2} - \omega x_i + \sum_k \frac{\partial V(\mathbf{x}_i - \mathbf{x}_k)}{\partial x_i}. \quad (4.3)$$

In (4.2) and (4.3), we define $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j = (x_{ij}, y_{ij})$ to be the displacement between vortices at $\mathbf{x}_i$ and $\mathbf{x}_j$, with $r_{ij} = |\mathbf{x}_{ij}|$. Similarly, $\mathbf{x}_{ij,\text{image}} = \mathbf{x}_i - \mathbf{x}_{j,\text{image}} = (x_{ij,\text{image}}, y_{ij,\text{image}})$ is the displacement between a vortex at $\mathbf{x}_i$ and the image vortex of a vortex at $\mathbf{x}_j$, with $r_{ij,\text{image}} = |\mathbf{x}_{ij,\text{image}}|$ (see Section 4.2.3). Furthermore, $2\pi\kappa$ is the quantum circulation, ω is the angular velocity of the inertial frame, and $V(\mathbf{x}_i - \mathbf{x}_k)$ is the pinning potential at $\mathbf{x}_i$ due to a pinning site located at $\mathbf{x}_k$, with $1 \leq k \leq N_{\text{pin}} \neq N$ in general. The first terms in equations (4.2) and (4.3) are the components of the fluid velocity at $\mathbf{x}_i$ induced by the other vortices $1 \leq j \neq i \leq N$. The second terms describe the motion due to image vortices, the third term is due to a rotating reference frame, the fourth term describes the velocity induced by pinning sites, and the rotation matrix $\mathcal{R}_\phi$ in equation (4.1) describes the effect of dissipation. We describe the origin and detailed form of each of these terms below.

4.2.2 Vorticity-induced velocity

The vorticity of a fluid is defined as the curl of the velocity field, i.e.,

$$\boldsymbol{\omega}(\mathbf{x}) = \nabla \times \mathbf{u}(\mathbf{x}). \quad (4.4)$$

Given $\boldsymbol{\omega}(\mathbf{x})$, the velocity can be calculated using the Biot-Savart law,

$$\mathbf{u}(\mathbf{x}) = \frac{1}{4\pi} \int d^3\mathbf{x}' \frac{\boldsymbol{\omega}(\mathbf{x}') \times (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|^3}. \quad (4.5)$$

In this paper, we consider infinitely long, rigid vortices aligned parallel to each other. This describes a two-dimensional system exactly but ignores important three-dimensional phenomena, such as vortex tangles, vortex tension, and reconnection (Saffman, 1995). A preliminary treatment of a three-dimensional code using the same method described in this chapter is included in Section 5.2.2.

For a system of N vortices moving only under vortex-induced motion in an unbounded fluid, equation (4.5) reduces to

$$\frac{d\mathbf{x}_i}{dt} = \sum_{j \neq i} \kappa \frac{\hat{\mathbf{z}} \times \mathbf{x}_{ij}}{r_{ij}^2}, \quad (4.6)$$

for a vortex at point $\mathbf{x}_i$. For a pair of vortices, equation (4.6) describes counter-clockwise rotation of each vortex about the centroid of the pair. The term $j = i$ is excluded in equation (4.6) because there is no centre of rotation for an unbounded fluid: in the absence of other vortices, an initially stationary isolated vortex remains stationary regardless of its initial position.

Following Lin (1941), we note that (4.6) is equivalent to Hamilton's equations of motion for the Hamiltonian

$$\mathcal{H} = \sum_{i=1}^N \sum_{j \neq i} \kappa \ln(r_{ij}), \quad (4.7)$$

where the Cartesian coordinates x_i and y_i are the conjugate variables.

4.2.3 Boundary conditions and image vortices

If the fluid is bounded by a container, the equations of motion are modified in order to enforce the boundary condition that the normal component of the induced velocity at the boundary of the container vanishes. This boundary condition can be solved using the method of images. For a single vortex at polar coordinates (r, ψ) , in a cylindrical container of radius R with $r < R$, the radial velocity at an arbitrary boundary point (R, θ) is

$$u_r = \frac{\kappa r \sin(\theta - \psi)}{R^2 + r^2 - 2Rr \cos(\theta - \psi)}. \quad (4.8)$$

The radial component of the induced velocity at the boundary due to a vortex at polar coordinates $(R^2/r, \psi)$ is the same as equation (4.8), so the boundary condition $u_r = 0$ can be satisfied by placing an image vortex with opposite circulation (i.e. $\kappa \rightarrow -\kappa$) at $(R^2/r, \psi)$.

With a circular boundary, the right-hand side of equation (4.6) includes an extra term due to the image vortices,

$$- \sum_{j=1}^N \kappa \frac{\hat{\mathbf{z}} \times \mathbf{x}_{ij,\text{image}}}{r_{ij,\text{image}}^2}, \quad (4.9)$$

where $\mathbf{x}_{ij,\text{image}}$, and $r_{ij,\text{image}}$ are defined in Section 4.2.1. The sum over the image vortices does not exclude the $j = i$ term, because the presence of the boundary imposes a fixed centre, i.e. even a single vortex offset from the centre of the container rotates about the centre.

4.2.4 Rotating frame

We run simulations in a rotating frame, whose angular velocity ω is chosen initially such that the total circulation satisfies $\oint \mathbf{v} \cdot d\mathbf{l} = R^2\omega = N\kappa$. The rotating frame enters through the third term in equations (4.2) and (4.3). The container and any pinning sites attached to the container (see Section 4.2.5) are stationary in the rotating frame. Note that the container and the superfluid do not in general have the same angular velocity in the simulations. The relation $R^2\omega = N\kappa$ is chosen such that, in equilibrium (if reached hypothetically), the vortices and container co-rotate.

4.2.5 Pinning

In a neutron star, nuclear lattice sites or magnetic flux tubes interrupt the smooth outward flow of vortices as the neutron star decelerates. This effect is called ‘pinning’ (Alpar, 1977). It arises due to quantum mechanical interactions between the vortex cores and the lattice sites or flux tubes (Link, 2009; Drummond & Melatos, 2017, 2018). In laboratory superfluids, pinning occurs at imperfections in the container (Tsakadze & Tsakadze, 1980). In this paper, we consider a grid of pinning sites at positions $\mathbf{x}_k$ in the computational domain, which are stationary in the rotating frame described in Section 4.2.4. Pinning leads to a term $\sum_k V(\mathbf{x} - \mathbf{x}_k)$ in equation (4.7), where $V(\mathbf{r}) = -V_0 f(r)$ is an attractive pinning potential, and $f(r)$ is some radially symmetric function, e.g. a Gaussian. Evaluating the equations of motion for this extended Hamiltonian gives the fourth terms in equations (4.2) and (4.3). The terms describe clockwise rotation of vortices about the pinning sites, which counteracts the counter-clockwise rotation of vortex pairs about their centroids (Acheson, 1990).

4.2.6 Dissipation

A vortex array has a tendency to minimize its free energy [see e.g. Abrikosov (1957); Fetter (1965); Campbell & Ziff (1979)]. However, the motion described by equation (4.7), is conservative; $\mathcal{H}$ is a constant of the motion. Some dissipative mechanism is necessary to reach this lowest energy state. Campbell & Ziff (1979) showed that the free energy minimum can be reached for vortices in an arbitrary initial configuration by moving the vortices along a vector parallel to the gradient of the free energy. However, these authors were interested only in finding stable equilibrium states and did not suggest a physical mechanism for reaching these states. We incorporate dissipation using the formalism developed for hydrodynamic descriptions of superfluid

helium. Phenomenological models of superfluids posit a dissipative interaction between the condensate and the viscous ‘normal’ component that is analogous to drag [e.g. Hall & Vinen (1956)]. Sedrakian (1995) showed that in this formalism the equations of motion for a vortex line in two dimensions can be written as a combination of induced motion due to the other vortices considered in aggregate, and a rotation through a ‘dissipation angle’, whose value is related to the strength of the mutual friction between the inviscid and viscous components of the fluid. In our model, we first evaluate the velocity of the vortices due to real vortices, image vortices, and pinning sites, then rotate the velocity vectors through a dissipation angle, ϕ , related to thermodynamic properties of the superfluid, by multiplying with the rotation matrix $\mathcal{R}_\phi$ in (4.1). In Section 4.4.2, we verify that, in the absence of pinning, this procedure causes a vortex array initialized in an arbitrary configuration to eventually form an Abrikosov array which co-rotates with the frame of the simulation (Abrikosov, 1957; Campbell & Ziff, 1979).

4.3 Numerical method

We write a Python code that solves Equations (4.1)–(4.3) using a fourth-order explicit adaptive Runge-Kutta Cash-Karp (RKCK) method (Press et al., 1992). A vortex simulation begins by initializing the vortex locations, calculating the initial vortex velocities, then stepping the system forward in time.

4.3.1 Feedback

In order to study avalanches, e.g. in neutron star applications, we consider a superfluid which is confined by a rotating container which decelerates. The dissipation mechanism discussed in Section 4.2.6 causes the vortices to migrate outwards in order to maintain corotation. As the vortices migrate outwards, the angular momentum of the superfluid decreases. An equal and opposite angular impulse feeds back onto the container, slowing the spin-down rate. The astrophysical details of the feedback mechanism lie outside the scope of this paper. We assume here for simplicity that it is instantaneous and lossless (Fulgenzi, Melatos & Hughes, 2017; Carlin & Melatos, 2020).

Let $L_c = I_c \Omega_c$ and $L_s = I_s \Omega_s$ be the angular momentum of the container and the superfluid respectively and I_c and I_s be their respective moments of inertia. The evolution equation for Ω_c is

$$\frac{d\Omega_c}{dt} = N_{\text{ext}} - I_{\text{rel}} \frac{d\Omega_s}{dt}, \quad (4.10)$$

with $I_{\text{rel}} = I_s/I_c$, where N_{ext} is the external spin-down torque divided by I_c . For a rotating superfluid, the angular momentum within radius R is (Fetter, 1965)

$$L_s = k \sum_{i=1}^N (R^2 - r_i^2), \quad (4.11)$$

where k is a constant with units of kg s^{-1} , and r_i is the radial coordinate of the i -th vortex. If we assume axisymmetry, equation (4.11) reduces to $L_s = kN(R^2 - \langle r^2 \rangle)$, where $\langle r^2 \rangle$ denotes the average over the vortices of the square of the radial coordinate. For a uniform vortex array we obtain $\langle r^2 \rangle = R^2/2$, $L_s = kNR^2/2$.

The spin-down and feedback procedure works as follows. At each time step:

- (1) compute Ω_s from equation (4.11);
- (2) increment Ω_c by $N_{\text{ext}}\Delta t$;
- (3) update the vortex positions at the new time step according to equation (4.1);
- (4) compute Ω_s from equation (4.11);
- (5) decrement Ω_c by $I_{\text{rel}}\Delta\Omega_s$.

Updating Ω_c only at the beginning and end of the time step does not strictly adhere to the fifth-order RKCK scheme used to solve equations (4.2) and (4.3). However, it is a good approximation if we assume that the time-scale over which vortices adjust to changes in the angular velocity of the container is much faster than other relevant time-scales.

4.3.2 Dimensionless coordinates

We run our simulations in a dimensionless coordinate system where $\kappa = 1$ and the fundamental length unit is 1. All other quantities, such as the time and velocity, are defined through these quantities. Our system of equations does not include any mass terms, as the vortex ‘particles’ are features in the velocity field and intrinsically massless. Other quantities involving mass, such as angular momenta and moments of inertia, only enter our equations of motion in dimensionless units such as I_{rel} .

4.3.3 Glitch finding

The main observables in pulsar glitch studies are the glitch sizes and inter-glitch waiting times. Pulsar radio emission is thought to be phase-locked to the rigid stellar crust, so both observables relate directly to the evolution of $\Omega_c(t)$. The procedure we use in order to extract sizes and waiting times from our simulations is as follows. The angular velocity of the container, Ω_c , is recorded at each time step. When the sign of $d\Omega_c/dt$, the rate of change in Ω_c between successive time steps, is positive, we flag the epoch as t_g . We then compute the cumulative increase in Ω_c between t_g and the next epoch when $d\Omega_c/dt$ becomes negative, t_{end} . The cumulative increase in Ω_c between t_g and t_{end} is recorded as the size of the glitch at epoch t_g . The waiting times are computed as the difference between successive epochs t_g . We do not compute a waiting time for the first glitch, because the simulation takes a while to build up enough stress to trigger the first vortex avalanche, i.e., it is not yet stationary statistically.

We note in passing that the above method of glitch finding broadly mirrors how astronomical glitches are detected. Most pulsars are monitored sporadically, with typical cadences of days to months. Hence pulsar glitches are almost never detected in real time [though see Palfreyman et al. (2018)]. Instead, they are identified through secular changes in timing residuals relative to a glitch-free timing model, which incorporates the astrometric and rotational evolution of the pulsar between observations (Lorimer & Kramer, 2004).

4.4 Validation

We verify that the solver described in Section 4.3 works as intended by reproducing two classical results in point vortex dynamics (Acheson, 1990): the stability of a vortex ring (Section 4.4.1), and the formation of an Abrikosov array (Section 4.4.2).

4.4.1 Stability of a vortex ring

The study of vortex motion in two dimensions goes back to the work of Helmholtz and others in the 19th Century. An important early result, first discussed by Thomson (1883) and later proved by Havelock (1931), describes the motion of N point vortices evenly spaced around a ring of radius R in an infinite medium. In the absence of dissipation, the vortices rotate uniformly about the centre of the ring, with angular velocity $\omega_N = \kappa(N-1)/(2R^2)$. For $N < 7$, the system is stable to small perturbations and maintains circular

motion indefinitely. For $N = 7$, the system is neutrally stable. For $N > 7$, the ring configuration is unstable, and a transition from circular to chaotic motion is unavoidable. Havelock (1931) performed a full linear stability analysis of this problem, including an extension to systems involving inner and outer boundaries, as well as counter-propagating vortex rings.

In order to test the solver, we reproduce three results: (i) the angular velocity of each vortex in the ring is proportional to the number of vortices, with $\omega_N = \kappa(N - 1)/(2R^2)$; (ii) rings with $N < 7$ are stable indefinitely while those with $N > 7$ result in chaotic motion; and (iii) the $N = 7$ case is metastable, transitioning from stable to chaotic motion when perturbed. We perform a suite of simulations where we solve equation (4.6) for $2 \leq N \leq 10$, and $0 \leq t \leq 200\pi/\omega_N$ (i.e. 100 rotation periods). In Figure 4.1, we show the angular velocity of the vortices, averaged over all vortices, normalised by ω_2 , versus time, in units of the rotation period of the vortex ring $\tau_N = 2\pi/\omega_N$. In the top panel, the vortices are placed initially at the vertices of a regular polygon and are evolved for $0 \leq t/\tau_N \leq 100$. In the bottom panel, we perturb the system by displacing each vortex radially by $\pm 5 \times 10^{-2}R$ (with $+$ or $-$ chosen at random for each vortex) relative to the centroid of the initial polygon and evolve the system for $100 \leq t/\tau_N \leq 200$. Figure 4.1 shows that $\langle \omega_N \rangle$ increases linearly with N , as expected. In the top panel, for $N \leq 7$, $\langle \omega_N \rangle/\omega_2$ is constant with time. For $N > 7$, the motion becomes disordered at $t \approx 10\tau_N$ and $\langle \omega_N \rangle/\omega_2$ fluctuates noisily. In the bottom panel, after perturbing the vortex positions initially, the rings with $N < 7$ and $N > 7$ are unaffected, but the metastable $N = 7$ ring comes to resemble the $N > 7$ cases.

To further illustrate the transition from ordered to disordered motion, in Figure 4.2 we show the state of the 10-vortex ring at $t = 3\tau_{10}$ and $t = 50\tau_{10}$. The top panel shows the ordered behaviour at the beginning of the simulation. The vortices are all equidistant from the centre of the circle, and their velocities are equal and tangent to the circle. In the bottom panel, the positions and velocities are randomized.

4.4.2 Dissipation and Abrikosov array formation

Abrikosov (1957) showed that flux tubes in a type II superconductor arrange themselves in a triangular array, as long as the system is unbounded. The same is true of flux tubes in a rotating superfluid such as helium II (Campbell & Ziff, 1979). When the vortices are confined to a cylindrical container, e.g. superfluid helium in a rotating bucket, the array is not exactly triangular near the boundary (Campbell & Ziff, 1979).

In order to verify that the dissipation mechanism in Section 4.2.6 is work-




Figure 4.1: Vortex-averaged angular velocity of N vortices, normalised by ω_2 , versus time, normalised by τ_N , for $2 \leq N \leq 10$. Top panel: vortex positions are initialised at the vertices of a regular polygon. Bottom panel: vortices are given a random radial displacement of $\pm 5 \times 10^{-2}R$ relative to the centroid of the initial polygon. Note: the time axis starts at $t/\tau_N = 0$ (100) in the top (bottom) panels.




Figure 4.2: Vortex positions (blue circles) and velocities (indicated by blue arrows) of the $N = 10$ ring initialised in a regular polygon at $t = 3\tau_{10}$ (top panel) and at $t = 50\tau_{10}$ before any radial perturbation at $t = 100\tau_{10}$ (bottom panel).

ing correctly, we show (i) that $\mathcal{H}$ converges for any initial vortex configuration and value of dissipation angle ϕ , and (ii) that the final value of $\mathcal{H}$ corresponds to an approximately triangular array which co-rotates with the container. We perform a suite of simulations solving equations (4.2) and (4.3) for an array of 100 vortices, choosing $\phi \in \langle 0.01, 0.02, 0.05, 0.1, 0.2, 0.5 \rangle$, to study convergence as a function of dissipation angle. The initial vortex positions are either drawn at random from a uniform spatial distribution or spaced uniformly around a ring of radius R (cf Section 4.4.1). We also perform 10 simulations with $\phi = 0.1$, with different initial vortex positions drawn at random from a uniform spatial distribution to study the ensemble statistics. We run each simulation for $t \approx 10^4 T_0$, where T_0 is the rotation period of the container, determined by the Feynman condition $T_0 = 2\pi R^2 / N\kappa$. We run the simulations in a rotating frame with $\omega = 2\pi / T_0$. While the terms involving ω in equations (4.2) and (4.3) confine vortices within $r < R$ at equilibrium, the vortices often overspill the boundary transiently at early times. When image vortices are present, the overspill vortices annihilate, and leave the simulation. Because the initial positions are randomized, the number of vortices that survive is different in each simulation, and $\mathcal{H}$ does not reliably converge to the same value each time. In the simulations described in this section, where we are mostly interested in the array configuration at equilibrium, we do not enforce the boundary condition in order to keep the vortex number constant. All the simulations described in Section 4.5 below include image vortices.

Figure 4.3 graphs $\mathcal{H}/\mathcal{H}_f$ versus t , where $\mathcal{H}_f$ is the value of $\mathcal{H}$ at the end of the simulation with $\phi = 0.5$ and random initial vortex positions. The top two panels show $\mathcal{H}$ versus t for $0.01 \leq \phi \leq 0.5$. In the top panel the initial vortex positions are drawn at random from a uniform spatial distribution. In the second panel the initial vortex positions are the vertices of a regular polygon. The third panel shows $\mathcal{H}$ versus t for 10 random uniform-density initial configurations with $\phi = 0.1$. The bottom panel shows the vortex configuration at the end of one of the simulations shown in the third panel.

Figure 4.3 verifies properties (i) and (ii) discussed in the first paragraph of Section 4.4.2. The bottom panel shows that the vortices settle to an approximately triangular array, which is stationary in the corotating frame. The top three panels verify convergence, with $\mathcal{H} \rightarrow \mathcal{H}_f$ in all simulations. In the top two panels, we see that convergence takes longer for lower values of ϕ , with $|\mathcal{H}/\mathcal{H}_f - 1| > 10^{-3}$ for $t < 1.5 \times 10^3 T$ ($\phi = 0.01$, top panel), and $|\mathcal{H}/\mathcal{H}_f - 1| > 10^{-3}$ for $t < 130T$ ($\phi = 0.5$, second panel). Across all simulations the final value of $\mathcal{H}$ differs from $\mathcal{H}_f$ by $< 0.1\%$. In the second panel, and in some of the simulations shown in the third panel, we see $\mathcal{H} < \mathcal{H}_f$ at early time. This happens when some of the vortices enter the region $r > R$




Figure 4.3: Convergence of the vortex array from an arbitrary initial configuration to a triangular array. Top panel: Hamiltonian $\mathcal{H}$ [equation (4.7)] versus time for 100 vortices initialised with random positions drawn from a uniform spatial distribution and evolved according to equation (4.1) (ignoring pinning and image vortices) for $0 \leq t \lesssim 10^4 T_0$, where T_0 is the rotation period of the container. Each simulation is initialised in the same configuration, with a different value of the drag angle $\phi \in \langle 0.01, 0.02, 0.05, 0.1, 0.2, 0.5 \rangle$ (see legend). Second panel: as for the top panel, but with the vortices initially situated equidistantly around a ring of radius R . Third panel: $\mathcal{H}$ versus t for 10 random uniform-density initial configurations with $\phi = 0.1$. Bottom panel: Vortex positions at the end of one of the simulations in the third panel.

temporarily, as discussed above.

4.5 An Astrophysical Example: Neutron Star Glitches

In this section, we report the results of a suite of simulations done with the solver described in Section 4.3. The simulations investigate the far-from-equilibrium phenomenon of vortex avalanches driven by deceleration of the vessel containing the superfluid. The results are compared with analogous numerical experiments involving Gross-Pitaevskii simulations, which reveal the microphysical knock-on mechanisms mediating the spatially correlated dynamics in vortex avalanches (Warszawski & Melatos, 2011; Warszawski, Melatos & Berloff, 2012; Warszawski & Melatos, 2013; Melatos, Douglass & Simula, 2015). The simulations form part of an idealized model of neutron star glitches triggered by vortex avalanches in a neutron superfluid coupled to a rigid stellar crust (Anderson & Itoh, 1975; Melatos, Peralta & Wyithe, 2008; Haskell & Melatos, 2015; Howitt, Haskell & Melatos, 2016).

4.5.1 Set up

We initialise an ensemble of point vortices within a circular container of radius R containing a square array of Gaussian pinning potentials, with the initial vortex positions drawn at random from a uniform spatial distribution. We evolve the ensemble according to equation (4.1) with spin down and feedback turned off, until all of the vortices are pinned. We then turn on feedback and spin down and resume evolving the system according to equations (4.1) and (4.10). The various input parameters and the meaning of each are explained in Table 4.1. The default parameters we use are $N_v = 2000$, $R = 10$, $\Delta t = 0.1T_0$ (where $T_0 = 2\pi R^2/N\kappa$ is the initial rotation period of the container), $I_{\text{rel}} = 1$, $V_0 = 2000$, $a = 0.01R$ (corresponding to a ratio of pinning sites/vortices ≈ 10), $\xi = 0.001R$, $\phi = 0.1$ radians, and $N_{\text{ext}} = -5 \times 10^{-4}\Omega_0/T_0$. We run for 2×10^5 time steps, so that, in the absence of feedback, the container’s period doubles to $2T_0$. In order to obtain a statistically useful number of glitches for analysis, we perform three simulations with the same parameters, each beginning from a different random initial vortex configuration, and aggregate the glitches. We have also performed a smaller suite of simulations with $N_v = 5000$ as a cross-check. These produced qualitatively similar results to the $N_v = 2000$ simulations described here, however, due to computational resource and time constraints we are unable to run them long enough to produce the number of glitches necessary for meaningful statistical analysis.

Parameter	Physical meaning
N_v	Number of vortices
R	Radius of container
Δt	Time step
I_{rel}	Ratio of superfluid/crust moments of inertia
V_0	Pinning strength
a	Pinning site separation
ξ	Characteristic width of pinning sites
ϕ	Dissipation strength
N_{ext}	Spin-down rate

Table 4.1: Summary of input parameters for the vortex avalanche simulations. Values used in the simulation are discussed in the text in Sections 4.5.1 and 4.5.4.

4.5.2 Avalanche dynamics

We test for avalanche behaviour by examining an ensemble of identically initialised simulations for evidence of collective vortex motion mediated by one or more knock-on mechanisms (Warszawski & Melatos, 2013). One such form of evidence is obtained by examining a movie of the vortex motion visually. Is the vortex unpinning triggered at one point and does it then spread in a connected way (avalanche), or does it occur simultaneously at multiple locations (not an avalanche)? Another form of evidence is obtained from the event statistics. Are the sizes and waiting times distributed according to a power law and exponential respectively, as in avalanche-dominated self-organised critical systems (Jensen, 1998; Melatos, Peralta & Wyithe, 2008)? Owing to computational cost, we analyse an ensemble of three simulations in this section; a fuller study will be conducted elsewhere. Identically initialised means that the macroscopic system variables ($N_v, R, \Delta t, I_{\text{rel}}, a, \xi, \phi, N_{\text{ext}}$) are the same in all three runs, but the initial vortex positions are selected randomly from a spatially uniform distribution.

Figure 4.4 shows the evolution of the container’s angular velocity, $\Omega_C(t)$, for $0 \leq t/T_0 \leq 10^4$ for one of the simulations with the default parameters. Initially, the spin down is smooth and monotonic; the vortices remain pinned, as stress builds up. At $t \approx 10^3 T_0$, the first of several small spin-up events (glitches) occurs, corresponding to collective unpinning of vortices. These continue to occur spasmodically throughout the remainder of the simulation, with 99 events detected by the glitch-finding algorithm (40 events in the time span plotted in Figure 4.4), ranging in size $\Delta\Omega_C$ over $7.1 \times 10^{-7} \leq$




Figure 4.4: Container angular velocity Ω_C , normalized by its initial value, as a function of time over 10^4 rotation periods. This simulation uses the default parameters listed in Section 4.5.1.

$\Delta\Omega_C/\Omega_0 \leq 4.5 \times 10^{-3}$. For $0 \leq t/T_0 \lesssim 10^3$, when the superfluid is effectively decoupled from the container, the spin-down rate (determined by a least-squares fit) is $-4.8 \times 10^{-4}\Omega_0/T_0$, and reduces to $-2.5 \times 10^{-4}\Omega_0/T_0$ during $10^3 \lesssim t/T_0 \lesssim 2 \times 10^4$. Once the glitches begin, they maintain the vortices near the critical unpinning threshold, repeatedly albeit spasmodically coupling the superfluid to the container, effectively doubling the inertia of the system ($I_{\text{rel}} = 1$).

Figure 4.5 depicts the motion of the vortices during the glitch that begins at $t \approx 3080T_0$, the largest glitch in the simulation. Overall the event is avalanche-like; it is triggered at a single location and spreads away in a wedge-like channel, as vortices knock-on their neighbours. We show the progression of the avalanche in five stages in the top five panels. The bottom panel shows the associated evolution of Ω_C during the avalanche, with each snapshot marked by a blue cross. The red dots in the top five panels mark the vortices that are moving at that instant. The progressively fainter black/grey circles indicate their positions at the previous 20 time steps. Light grey dots show the positions of the vortices that have not moved further than a in the previous 20 time steps. Upon studying Figure 4.5, we observe the following features.

- In the top panel, at $t/T_0 = 3079$, two nearby vortices toward the centre left unpin and begin moving radially outward. Their motion does not have a noticeable effect on Ω_C , which decreases steadily.




Figure 4.5: Vortex motion during an avalanche. Top five panels are snapshots at (top to bottom) $t/T_0 = 3079, 3086, 3093, 3104.5, 3116$. Grey dots show stationary vortices; red dots show vortices that have moved in the previous time step, black (fading to gray) tracers show the positions of the moving vortices for the 20 previous time steps. Bottom panel: container angular velocity Ω_C versus time t . Snapshots from the top five panels are marked with blue crosses.

- In the second panel, at $t/T_0 = 3086$, more vortices are on the move. Some appear to unpin by direct proximity knock-on (Warszawski, Melatos & Berloff, 2012); others unpin by themselves shortly after a vortex closer to the boundary unpins and moves outwards. Both these unpinning modes are observed in quantum mechanical, Gross-Pitaevskii simulations of smaller systems (Warszawski & Melatos, 2011; Warszawski, Melatos & Berloff, 2012; Melatos, Douglass & Simula, 2015). The collective outward motion tends to reverse the steady spin down. Two other vortices unpin in another region of the container, towards the top right. They are likely unconnected to the avalanche in progress.
- The third panel, at $t/T_0 = 3093$, shows many of the unpinned vortices leaving the container. As vortices approach the boundary, they acquire a significant counter-clockwise azimuthal velocity, due to the presence of their corresponding image vortex across the boundary. In contrast, unpinned vortices closer to the centre tend to move radially. In this panel the two other unpinned vortices in the top right re-pin close to where they first unpinned.
- The fourth panel, at $t/T_0 = 3104.5$, shows the vortex array after the majority of the unpinned vortices leave the container or re-pin. Vortices move in a slow, inward, clockwise spiral before re-pinning in the evacuated region near where the avalanche begins. Since the motion of the unpinned vortices is mostly azimuthal, rather than radial, spin up stops and Ω_C flattens out.
- In the fifth panel, at $t/T_0 = 3116$ almost all of the vortices have re-pinned and Ω_C resumes steady spin-down. During the avalanche, 68 vortices of the 1919 in the container at the beginning of the avalanche move an average distance of $\approx 0.18R$.

An important characteristic of avalanche dynamics is that avalanches relieve only a small amount of the accumulated stress, and the system remains in a marginal state close to the avalanche threshold in many places even immediately after an avalanche. To test this, we note that vortices are stationary when the pinning velocity vector $\mathbf{v}_{\text{pin}}$ is exactly equal and opposite to the vorticity-induced velocity vector (from real and image vortices) $\mathbf{v}_{\text{induced}}$. With a Gaussian pinning potential, one has

$$\mathbf{v}_{\text{pin}} = V_0 \exp \left[-\frac{(\mathbf{x} - \mathbf{x}_k)^2}{2\xi^2} \right] (\mathbf{x} - \mathbf{x}_k) \times \mathbf{e}_z \quad , \quad (4.12)$$

where $\mathbf{x}_k$ is the position of the k -th pinning centre and $\mathbf{e}_z$ is the unit vector in the z -direction. The pinning speed peaks at $\max|\mathbf{v}_{\text{pin}}| = V_0\xi e^{-1/2}$ at

$|\mathbf{x} - \mathbf{x}_k| = \xi$. A vortex leaves the region $|\mathbf{x} - \mathbf{x}_k| \leq \xi$ permanently if the absolute value of the sum of the first three terms in equations (4.2)–(4.3) exceeds $\max|\mathbf{v}_{\text{pin}}|$. We can therefore characterise the marginal stability approximately by looking at the spatial distribution of $v_{\text{stress}} = |\mathbf{v}_{\text{induced}} - \boldsymbol{\Omega}_C \times \mathbf{x}| / \max|\mathbf{v}_{\text{pin}}|$.

In Figure 4.6, we show the probability distribution function (PDF) of v_{stress} before and after the glitch shown in Figure 4.5. We also show scatter plots of the vortex positions before and after the glitch, where vortices with $v_{\text{stress}} > 0.5$ are coloured in red. Figure 4.6 conveys three important results. Firstly, stress does not accumulate exclusively in the region where the avalanche occurs. Stressed vortices occur throughout the container. Secondly, while some stress is relieved by the glitch – the median value of v_{stress} decreases from 0.47 to 0.36, and the full-width half-maximum of the PDF decreases from 0.57 to 0.50 – a lot remains afterwards. Thirdly, the remaining stress is not localised to the region outside the avalanche. Stressed vortices are more-or-less uniformly distributed throughout the container before and after the glitch, including where the avalanche occurred.

4.5.3 Sizes and waiting times

Here we examine the statistical properties of glitch sizes and waiting times in the simulations. We aggregate 314 glitches from three identically-initialized simulations with the default input parameters in Section 4.5.1 and calculate kernel density estimates of the waiting time and size PDFs. Figures 4.7 and 4.8 display the results for the waiting time and size PDFs respectively for both the aggregated data and the three individual simulations. We also include fits to an exponential distribution, $p(x) = \lambda \exp(-\lambda x)$, and a log-normal distribution, $p(x) = (2\pi x^2 \sigma^2)^{-1/2} \exp[-(\ln x - \mu)^2 / 2\sigma^2]$, where the fit parameters take their maximum likelihood values $\lambda = 1/\langle x \rangle$, $\mu = \langle \ln x \rangle$, and $\sigma^2 = \langle (\ln x - \mu)^2 \rangle$ (Howitt, Melatos & Delaigle, 2018; Fuentes, Espinoza & Reisenegger, 2019). For the size distribution, we also include a fit to a power law distribution, $p(x) = Ax^{-a}$. The fitting tool we use for the power law, `optimize.curve_fit` in the `scipy` package, is unable to produce a fit for the entire data set. However, if we consider just the tail of the data with $\Delta\Omega_C/\Omega_0 > 2 \times 10^{-4}$ (comprising $\approx 2/3$ of the glitches), we find a best-fit value of $a = 1.6$ for the power law index.

We find that both the waiting time and size distributions are well described by exponential and log-normal PDFs. An Anderson-Darling test finds consistency for both sizes and waiting times with both exponential and log-normal distributions at the 99% confidence level. Similar studies on pulsar glitch data (Melatos, Peralta & Wyithe, 2008; Howitt, Melatos & De-




Figure 4.6: Top panel: PDF of stress before (blue curve) and after (orange curve) the glitch, smoothed using a kernel density estimator. Effective stress, parametrized by the speed v_{stress} before (middle panel) and after (bottom panel) a vortex avalanche. Red dots indicate vortices with $v_{\text{stress}} > 0.5$.




Figure 4.7: Kernel density estimates of the waiting time PDF from three simulations with the default parameters in Section 4.5.1. Black curve: aggregated data. Coloured curves: individual simulations. Grey, dashed line: maximum likelihood exponential fit. Grey, dot-dashed curve: maximum likelihood log-normal fit.




Figure 4.8: Kernel density estimates of the size PDF from three simulations with the default parameters in Section 4.5.1. Black curve: aggregated data. Coloured curves: individual simulations. Grey, dashed curve: maximum likelihood exponential fit. Grey, dot-dashed curve: maximum likelihood log-normal fit. Grey, dotted line: power law fit for $\Delta\Omega_C/\Omega_0 > 2 \times 10^{-4}$.

laigle, 2018) and quantum mechanical Gross-Pitaevskii simulations of vortex avalanches (Warszawski & Melatos, 2011) often produce exponential waiting time PDFs and power-law size PDFs. However, the observational data offer support for other functional forms too [eg. log-normal, Gaussian (Fuentes, Espinoza & Reisenegger, 2019) and are based on small samples ($\lesssim 50$ events per pulsar)]. The Gross-Pitaevskii simulations are also based on small samples ($\approx 10^2$ vortices) and include acoustic knock-on (Warszawski, Melatos & Berloff, 2012), a process not present in this paper, which allows long-range unpinning. In contrast, the simulations in this paper involve more reliable statistics, with 2×10^3 vortices, and $\approx 10^2$ glitches per simulation. The applicability of the power-law PDF to only the larger glitches is interesting. A similar result was shown in observational data by Janssen & Stappers (2006), who suggested that it may be due to a population of undetected small glitches. This may also be the case in our simulations. Slow-moving avalanches of few vortices may reduce the magnitude of the spin-down rate but without changing its sign. These avalanches are not picked up by our glitch-finding algorithm, and so not be picked up by our glitch-finding algorithm.

4.5.4 Model variations

We do not attempt to fit the model in this paper to astrophysical data, because the values taken by the input parameters in a neutron star are uncertain. It is useful, however, to develop a rough sense of how the avalanche dynamics depend on the input parameters. We vary four quantities: the strength of pinning, parametrised by V_0 , the density of pinning sites, parametrised by a , the strength of dissipation, parameterised by ϕ , and the spin-down rate, parametrised by N_{ext} . We vary each parameter individually below and above the default value in Section 4.5.1. As well as the default values, we run simulations with $V_0 = 500$ and $V_0 = 2000$; $a = 0.025R$ and $a = 0.005R$ (corresponding to a ratio of pinning sites/vortices of ≈ 1 and ≈ 100 respectively); $\phi = 0.01$ rad and $\phi = 0.5$ rad; and $N_{\text{ext}} = 5 \times 10^{-4}$ and $N_{\text{ext}} = 2 \times 10^{-3}$. We perform an ensemble of three 2000-vortex simulations for each model variation and compute the total number of glitches, the mean waiting time between glitches, and the mean glitch size.

Table 4.2 shows a comparison of these quantities to the default simulation; cf. Table 9 in Warszawski & Melatos (2011). Plus (minus) signs in Table 4.2 indicate that the observable in each column increases (decreases) as the parameter increases. Dots indicate no consistent change in the observable as the parameter varies. We find that the average waiting time and glitch size both increase as the pinning strength, V_0 , increases. Increasing the spacing

Parameter	$\langle \Delta t / T_0 \rangle$	$\langle \Delta \Omega_C \rangle$
V_0	+	+
a	+	.
ϕ	.	.
N_{ext}	−	.

Table 4.2: Effect of varying simulation parameters (first column) on vortex avalanche statistics (columns two and three): mean waiting time between glitches $\langle \Delta t / T_0 \rangle$, and mean glitch size $\langle \Delta \Omega_C \rangle$. Symbols + (−) indicate that $\langle \Delta t \rangle$ or $\langle \Delta \Omega \rangle$ increases (decreases) relative to the default value in Section 4.5.1 as the parameter in column one increases; “.” indicates no noticeable effect.

between pinning sites, a , (i.e. lowering the density of pinning sites) causes the average waiting time to increase and does not change the average glitch size. Changing the drag angle, ϕ , does not consistently push the average waiting time or size in one particular direction. In fact, $\langle \Delta t \rangle$ is higher in the default case ($\phi = 0.1$) than its value with both $\phi = 0.01$ and $\phi = 0.5$. Increasing the spin-down torque, N_{ext} , results in more frequent glitches but does not change the average size. Our results extend Table 9 in Warszawski & Melatos (2011) by including ϕ . Where we can make direct comparisons, however, all of our results agree with Warszawski & Melatos (2011).

The results in Table 4.2 are interpreted physically as follows. Increasing the spin-down torque, N_{ext} , does not appear to change the dynamics of the glitches. It simply increases their frequency due to the more rapid build-up of stress. As V_0 increases, vortices withstand more stress before unpinning and hence travel further before re-pinning, leading to more knock-on and hence larger glitches. Since larger glitches release more stress, they occur less often. This also explains the longer waiting times with stronger pinning; it takes longer to build up to the critical lag threshold after a large glitch. As we increase the spacing between pinning sites, vortices are less likely to re-pin after unpinning. This leads to steady outward flow of vortices, rather than sudden, collective motion as the container decelerates, though some small avalanches still occur. Because our glitch-finding algorithm counts only unpinning events that lead to a spin-up of the container, the steady flow is not picked up, leading to a reduced number of glitches with greater waiting times between them.

The effect of changing the drag angle ϕ is complicated. When vortices unpin, the circulatory motion induced by other vortices and pinning sites

causes them to follow a spiral trajectory as they move outward. Increasing the drag angle ϕ makes the trajectory more radial, leading to less knock-on and hence smaller glitches. In contrast, when ϕ is lower, there is more knock-on, but unpinned vortices are more influenced by interactions with nearby vortices than by the dissipative radial motion. They “pinball” throughout the vortex array, with some eventually re-pinning closer to the center than where they unpin. While the overall tendency is that vortices migrate outwards, it is a more protracted process than in the default case. To better illustrate this effect, we examine the largest glitch in the weak dissipation simulation at high time resolution, cf. Figure 4.5, in Figure 4.9. The top panel in Figure 4.9 shows the first few vortices that unpin at $t = 5487.5T_0$ and begin to move toward the edge of the container. In the second panel, at $t = 5510T_0$, more vortices unpin. On average, vortices move radially outwards, but a significant number are also moving either inward or azimuthally. In the third panel, at $t = 5520T_0$, many vortices remain unpinned but the ‘swirling’ motion dominates, so the container stops spinning up (see third blue cross on bottom panel). In the fourth panel, at $t = 5570T_0$, some vortices remain unpinned some time after the container ceases spinning up. Their motion is disordered, leading to bumpy evolution of the spin frequency.

4.5.5 Size-waiting time correlations

Auto- and cross correlations between glitch sizes and waiting times are an important diagnostic of the balance between crustal spin down and vortex unpinning (Fulgenzi, Melatos & Hughes, 2017; Melatos, Howitt & Fulgenzi, 2018; Carlin & Melatos, 2019; Fuentes, Espinoza & Reisenegger, 2019). Here we look for size and waiting time correlations in our simulations across all model variations.

Firstly, we look at cross correlations. Two types of correlations are considered: forward correlations between size and the waiting time to the following glitch; and backwards correlations between size and the waiting time since the preceding glitch. In general, such correlations are not expected in systems with SOC, and are not seen in Gross-Pitaevskii simulations of glitches (Warszawski & Melatos, 2011). Only one pulsar, PSR J0537–6910 has a strong forward correlation (Middleditch et al., 2006; Ferdman et al., 2018; Antonopoulou et al., 2018). Several other pulsars appear to have weak forward correlations, though these are more uncertain due to the low number of observed glitches. No backwards correlations are observed in the pulsar population (Melatos, Howitt & Fulgenzi, 2018). The existence of size-waiting time cross correlations (or lack thereof) has been proposed as a means of falsifying meta-models of pulsar glitches as either a state-dependent Pois-




Figure 4.9: Vortex motion during an avalanche in the weak dissipation simulation. Top four panels are snapshots at (top to bottom) $t/T_0 = 5487.5, 5510, 5520, 5570$. Grey dots show stationary vortices; red dots show vortices that have moved in the previous time step, black (fading to gray) tracers show the positions of the moving vortices for the 20 previous time steps. Bottom panel: container angular velocity Ω_C versus time t . Snapshots from the top four panels are marked with blue crosses.

Model variation	N_g	r_+	r_-	ρ_+	ρ_-
Default	311	0.35	0.04	0.44	0.03
Weak pinning	360	0.25	0.07	0.24	0.03
Strong pinning	204	0.33	0.004	0.37	0.04
Low density	228	0.34	-0.1	0.22	-0.1
High density	408	0.45	0.02	0.38	0.007
Weak dissipation	392	-0.12	0.14	-0.19	-0.01
Strong dissipation	516	0.32	-0.01	0.31	-0.01
Fast spin-down	420	0.44	-0.03	0.39	-0.05
Slow spin-down	188	0.36	0.1	0.34	-0.06

Table 4.3: Pearson correlation coefficients r and Spearman rank coefficients ρ for forward (+) and backward (−) size-waiting time correlations in each of our model variations.

son process (Fulgenzi, Melatos & Hughes, 2017; Melatos, Howitt & Fulgenzi, 2018; Carlin & Melatos, 2019) or a Brownian stress-accumulation process (Carlin & Melatos, 2020). Both the Brownian and state-dependent Poisson meta-models predict positive forward correlations in certain regimes.

In Table 4.3 we show the values of the Pearson correlation coefficients r and the Spearman rank coefficients ρ for both the forward and backward size-waiting time correlations for each of our model variations. To assess significance, uncertainty is roughly the standard error $\sim r_{\pm}(N_g - 2)^{-1/2} \lesssim 0.1r_{\pm}$. All of our model variations, with the exception of weak dissipation, show a weak positive forward correlation between glitch size and waiting time. None show a significant backwards correlation.

In Table 4.4, we show the auto-correlation in size and waiting times for all model variations. None of the model variations show a significant positive autocorrelation for size or waiting time. The low density and weak dissipation cases have a statistically-significant (p-value < 0.05) weak negative autocorrelation in waiting time and the strong pinning and weak dissipation cases have a statistically-significant weak negative autocorrelation in the size. In the case of the weak dissipation size autocorrelation, the significance is marginal. We perform 18 independent significance tests, so it is likely that at least one p-value is less than 0.05 even if the null hypothesis of no autocorrelation is correct in all cases.

The results in Tables 4.3 and 4.4 are broadly consistent with both meta-models. There are no strong cross-correlations or autocorrelations in any of our models, which is consistent with what is observed in self-organized

Model variation	Waiting time	p-value	Size	p-value
	$\rho_{\Delta t}$		$\rho_{\Delta\Omega}$	
Default	0.068	0.23	-0.037	0.52
Weak pinning	-0.024	0.65	0.0010	0.85
Strong pinning	0.039	0.58	-0.18	0.011
Low density	-0.21	0.0016	0.081	0.22
High density	0.0045	0.93	-0.012	0.81
Weak dissipation	-0.19	0.00012	0.10	0.041
Strong dissipation	-0.040	0.36	0.064	0.15
Fast spin-down	0.0066	0.90	-0.083	0.12
Slow spin-down	0.024	0.75	0.071	0.33

Table 4.4: Spearman autocorrelation coefficient ρ and p-values for glitch waiting times Δt and sizes $\Delta\Omega$ in each of our model variations [cf. Table 1 in Carlin & Melatos (2019)].

critical systems (Jensen, 1998) and Gross-Pitaevskii simulations (Warszawski & Melatos, 2011), albeit counter-intuitive for a threshold-triggered stress-release process. With respect to the state-dependent Poisson model, the cross-correlations suggest that we are in the fast-driving regime [see Figure 13 in Fulgenzi, Melatos & Hughes (2017)]. This is not surprising. In order to produce statistically useful numbers of glitches, computational constraints require spin-down rates far greater than those observed in even the most rapidly decelerating pulsars. The models with negative autocorrelations in waiting time and size are inconsistent with the Brownian meta-model, and are consistent with the state-dependent Poisson model only in a restricted subset of parameter space (Carlin & Melatos, 2019, 2020).

4.6 Conclusion

Superfluid vortex avalanches have long been suggested as a mechanism for pulsar glitches, but the existence of avalanche behaviour even in simplified models of neutron stars has not been demonstrated in systems with more than ≈ 100 vortices. We have written a two-dimensional N -body solver based on the vortex filament model, including dissipation. Our results exhibit avalanche behaviour across a wide range of physical parameters for large-scale systems with $N \lesssim 5 \times 10^3$ vortices, demonstrating that vortex avalanches are ubiquitous when vortices are pinned in a decelerating container. Our results agree qualitatively with previous Gross-Pitaevskii sim-

ulations with $\approx 10^2$ vortices, despite the lack of acoustic knock-on in our classical point-vortex model. Comparing our simulations to the observed population of pulsar glitches, we find similar waiting time distributions, but our size distributions differ. We find weak cross-correlations between glitch sizes and waiting times for almost all of our model variations, a finding consistent with meta-models of pulsar glitches as a state-dependent Poisson processes or a Brownian stress-accumulation process. The correlation results are also qualitatively consistent with the observed population, in which statistically significant correlations have only been observed in one pulsar, PSR J0537 – 6910. We see weak negative autocorrelations in some of our models, which are inconsistent with the Brownian motion meta-model and restrict the parameter of the state-dependent model.

Advances in glitch modelling and detection are paving the way for falsifying specific glitch mechanisms, such as superfluid vortex avalanches. On the modelling front, large scale N -body simulations like those in this paper make specific falsifiable predictions about the long-term statistics of glitches and their individual profiles in time. Future improvements include relaxing the simplifying assumptions in the model, such as moving to three dimensions and including vortex tension (Link, 2009). Preliminary work on a three-dimensional point vortex code is shown in On the observational front, improvements include analyzing the completeness of existing datasets and enlarging the glitch sample with next generation pulsar observing campaigns with instruments such as the Square Kilometre Array. All of this can be combined with stress-release meta-models which make microphysics-agnostic predictions about the long-term glitch statistics, as exemplified by the discussion in Section 4.5.5 in this chapter.

Chapter 5

Summary and Future Work

5.1 Summary

This thesis examines the phenomenon of pulsar glitches in the context of superfluidity in neutron stars. Glitches are thought to be due to differential rotation between the crust and superfluid interior, and the thesis studies this theory from several directions:

- Hydrodynamic simulations of post-glitch coupling and relaxation of multiple stellar components (Chapter 2, summarised in Section 5.1.1).
- Quantifying the PDF of glitch sizes and waiting times minimising binning bias when event numbers are relatively small (Chapter 3, summarised in Section 5.1.2).
- Simulating superfluid vortex avalanches from first principles numerically in order to reproduce long-term glitch statistics (Chapter 4, summarised in Section 5.1.3).

5.1.1 Hydrodynamic simulations of pulsar glitch recovery

Chapter 2 describes numerical simulations of glitches using a spectral collocation solver for two-fluid hydrodynamics (Bagchi & Balachandar, 2002; Peralta et al., 2005, 2008). A neutron star is modelled as two concentric spheres with radii R_1 and R_2 and angular velocities Ω_1 and Ω_2 with $R_1 < R_2$. The region $R_1 < r < R_2$ contains a two-component superfluid consisting of a ‘normal’ component, the viscous proton-electron plasma, and an inviscid component, the superfluid neutrons. The two components interact through mutual friction, and the evolution is governed by the HVBK equations (Hall & Vinen,

1956; Chandler & Baym, 1986; Henderson, Barenghi & Jones, 1995), with no-slip boundary conditions at $r = R_1$ and $r = R_2$.

We simulate glitches by driving an initially co-rotating system into differential rotation, then triggering glitches in three ways:

- Spinning up the outer boundary.
- Spinning up the inner boundary.
- Recoupling the superfluid and viscous components of the interior.

The first trigger mechanism is analagous to a crustquake (Ruderman, 1976; Alpar et al., 1996), the second to a ‘corequake’ (Ruderman, Zhu & Chen, 1998), and the third to models which propose that glitches may be triggered by hydrodynamic instabilities (Andersson, Comer & Prix, 2004; Glampedakis & Andersson, 2009). We study the post-glitch relaxation by constructing a ‘residual’, $\Delta\Omega$, by subtracting the angular velocity of the outer boundary Ω_2 after a glitch to Ω_2 from simulations without a glitch. Simulations are performed with strong and weak mutual friction.

In the first two cases, $\Delta\Omega(t)$ is well-described by a two-component exponential recovery plus a permanent increase $\Delta\Omega_p$ [cf. equation (1.2)]. The permanent change in $\Delta\Omega(t)$ is due to the fact that these glitch mechanisms increase the total angular momentum in the system, and its magnitude closely matches a simple calculation based on relative moments of inertia. In both cases the strength of mutual friction does not significantly affect the shape of the recovery or its time scale.

In the third case, total angular momentum is conserved, and the response to the glitch differs between the simulations with strong and weak mutual friction. In both mutual friction regimes, there is a rapid spin-up followed by a more gradual recovery, both of which are well-described by exponentials. In the strong mutual friction regime there is an additional oscillatory component to $\Delta\Omega(t)$. In the strong mutual friction regime the glitch is approximately an order of magnitude less than with weak mutual friction, and $\Delta\Omega \rightarrow 0$ as $t \rightarrow \infty$, while in the weak mutual friction regime $\Delta\Omega \rightarrow 0.1\Delta\Omega_{\max}$ as $t \rightarrow \infty$. These results imply that the maximum glitch size is determined more by the strength of mutual friction than the relative moments of inertia of the superfluid and viscous components, as has been suggested by some authors (Link, Epstein & Lattimer, 1999; Andersson et al., 2012).

The different behaviour of $\Delta\Omega(t)$ from each of the glitch activation mechanisms may also explain the wide range of post-glitch recoveries seen in pulsars (Espinoza et al., 2011; Haskell & Antonopoulou, 2014).

5.1.2 Nonparametric Estimation of the Size and Waiting Time Distributions of Pulsar Glitches

Chapter 3 examines the glitch size and waiting time PDFs for the five most prolific glitching pulsars in the Jodrell Bank glitch catalogue (Espinoza et al., 2014) using kernel density estimation (Silverman, 1986). We find that two pulsars, PSR J0534+2200 (the Crab) and PSR J1740–3015, have Poisson-like waiting time PDFs and scale-invariant size PDFs. Two others, PSR J0835–4510 (Vela) and PSR J0537–6910, exhibit quasiperiodic size and waiting time PDFs. The fifth pulsar we consider, PSR J1341–6220, appears to exhibit hybrid behaviour. Because of the small number ($21 \leq N_{\text{glitch}} \leq 41$) of glitches in any pulsar in our sample, the results are not conclusive, but the use of kernel density estimation avoids the inherent bias of parametric analyses (Konar & Arjunwadkar, 2014; Eya, Urama & Chukwude, 2017; Fuentes et al., 2017) and minimises the binning bias inherent to histograms (Melatos, Peralta & Wyithe, 2008; Fuentes, Espinoza & Reisenegger, 2019). We examine the performance of the kernel density estimator (KDE) on synthetic data drawn from exponential and power law distributions, finding that for $N \approx 25$ the KDE tends to recover the overall shape of the underlying PDF, but artificial ‘wiggles’ often appear due to outlying or randomly clustered data points. We also examine the performance of the KDE on a bimodal distribution. We find that the ability of the KDE to resolve both peaks depends sensitively on the choice of kernel bandwidth.

We analyse the results in the context of the vortex avalanche model of pulsar glitches (Anderson & Itoh, 1975; Alpar, Langer & Sauls, 1984; Haskell & Melatos, 2015). In this model, glitches are considered to be driven by differential rotation between a superfluid in the interior and the decelerating crust. The angular momentum of the superfluid is determined by the configuration of an array of $\sim 10^{10}$ pinned vortices. As the crust decelerates, the vortices approach the critical unpinning threshold, and single or few-vortex unpinning events trigger unpinning avalanches which transfer angular momentum from the superfluid to the crust, causing sudden spin up. Following an avalanche, the system remains near the threshold, so that successive avalanches are uncorrelated in location in the star or time.

The vortex avalanche model has the hallmarks of systems exhibiting self-organized criticality (SOC) (Bak, Tang & Wiesenfeld, 1987; Watkins et al., 2015), where a globally driven system is kept near a critical threshold by stochastic local adjustments. In the SOC paradigm, Poisson waiting times and scale-invariance are characteristic of slowly-driven systems, implying that perhaps the Crab and PSR J1740–3015 are slowly-driven too. In contrast, quasiperiodic waiting time PDFs, as seen in Vela and PSR J0537–6910, are

typical of fast-driven systems. In pulsars, however, the meaning of fast and slow driving is not straightforward. It depends not just on the spin-down rate but also the strength of pinning and other internal parameters. Recent theoretical work (Fulgenzi, Melatos & Hughes, 2017) finds that the long-term size and waiting time statistics are controlled chiefly by a single speed-of-driving parameter, which is the spin-down rate divided by the characteristic unpinning rate when the differential rotation is half (or some fiducial fraction) its maximum value. The theory makes falsifiable predictions about the validity of the avalanche model which can be tested by observations of size-waiting time correlations and cross-correlations in glitching pulsars (Melatos, Howitt & Fulgenzi, 2018; Carlin & Melatos, 2019, 2020).

5.1.3 Simulations of pulsar glitches with N -body vortex code

Chapter 4 presents a numerical solver for quantized vortex motion in a superfluid in two dimensions in the presence of pinning sites, dissipation and a decelerating container. We develop a mathematical model for the system and describe an N -body `python` code that can solve for the motion of $\sim 10^4$ vortices. We validate the code by reproducing several classical results in vortex dynamics. Firstly, we demonstrate the instability of vortex rings with $N > 7$ vortices, and the metastability of the $N = 7$ case (Thomson, 1883; Havlock, 1931). Secondly, we show that dissipation causes an initially random vortex configuration to settle to an approximately triangular array which is stationary with respect to the rotating container (Campbell & Ziff, 1979).

We use the code to perform simulations of 2000 vortices with a wide array of input parameters, namely the pinning strength, the pinning site density, the dissipation strength and the spin-down rate. In all cases deceleration of the container drives the pinned vortices close to the unpinning threshold, leading to stochastic avalanches which spin up the container, similar to pulsar glitches. The majority of vortices remain close to the unpinning threshold throughout the simulations, even after large avalanches.

We examine the statistics of the avalanche sizes and waiting times. The shape of the size and waiting time PDFs is qualitatively similar for all parameter variations, and the effect of changing the parameters on the mean size and waiting time are the same as for Gross-Pitaevskii simulations with $N \sim 100$ (Warszawski & Melatos, 2011). For our default parameters, we find that both size and waiting time PDFs are consistent with exponential and log-normal fits, while the size PDF is also well-described by a power-law in the tail. This is broadly consistent with the behaviour seen in some glitching

pulsars such as the Crab and PSR J0537–6910 [see Chapter 2, (Melatos, Peralta & Wyithe, 2008; Fuentes, Espinoza & Reisenegger, 2019)], but it is worth emphasising that the observational data is based on $\lesssim 50$ glitches, while our simulations produce ≈ 300 avalanches.

We also test for correlations between size and waiting times across all parameter variations. Firstly, we look at the forward and backwards cross-correlations of size and waiting time, finding that all but one of our simulations produces weak but statistically-significant correlation between the size of a glitch and the waiting time to the subsequent glitch. This is in contrast to the observed population of glitching pulsars, none of which exhibit significant cross-correlations in either direction, with the notable exception of PSRJ0537–6910, which has a strong forward correlation (Middleditch et al., 2006; Ferdman et al., 2018; Antonopoulou et al., 2018; Melatos, Howitt & Fulgenzi, 2018). However, the small numbers of observed glitches make detection of the weak correlations we see in simulations difficult. The situation will become clearer as more glitches are detected. We find no significant positive auto-correlations for size or waiting time in any of our simulations, and a weak negative size auto-correlation in two of our parameter variations. Based on the theory of glitches as a state-dependent Poisson process (Fulgenzi, Melatos & Hughes, 2017), the correlations we find suggest that all of our simulations are performed in the strong-driving regime.

5.2 Future work

The N -body vortex code has potential applications in other scenarios besides the simulations of vortex avalanches in a decelerating container described in Chapter 4 and summarised in Section 5.1.3. This section describes several of these use cases for the code and preliminary work on each that has been carried out so far:

- Simulations of vortex motion in an accelerating container, leading to vortex avalanches producing spin-down of the container (Section 5.2.1).
- Extension of the N -body vortex solver to three dimensions, including simulations of interacting vortex filaments and leapfrogging of vortex rings (Section 5.2.2).
- Simulations of vortex-antivortex annihilation leading to Onsager vortex formation (Section 5.2.3).

5.2.1 Anti-glitches and the vortex avalanche model

Almost all pulsar glitches have been observed as increases in the spin frequency of decelerating pulsars. However, Archibald et al. (2013) reported an ‘anti-glitch’, i.e. a sudden *decrease* in the spin frequency, in the magnetar 1E 2259+586. The antiglitch was detected as a discontinuity in the timing model between two TOAs taken with the Swift X-ray telescope at 14 April 2012 and 28 April 2012. On 21 April 2012, the Fermi Gamma-ray Burst Monitor detected a 36ms hard X-ray burst from the same object, possibly coincident with the antiglitch. Swift monitoring subsequent to 28 April 2012 measured an increase in the 2 – 10 keV X-ray flux of a factor of ≈ 2 , which then decayed following a power law in time. These associated radiative changes are similar to other glitches in magnetars (Dib, Kaspi & Gavril, 2008), including one other spin up glitch in 1E 2259+586 (Kaspi et al., 2003), and unlike radio pulsar glitches. Following the anti-glitch, the spin-down rate of the magnetar approximately doubled for ~ 100 d before returning to its pre-glitch value following a second glitch. The sign of the second glitch is uncertain, (Archibald et al., 2013) considered two models, one with an anti-glitch followed by a glitch, the other with an anti-glitch followed by a second anti-glitch, and found no statistical preference for either, but preferred both to a model with only a single anti-glitch. More recently, a spin-up and spin-down glitch have been reported in the same object (Younes et al., 2020), both without any associated pulse profile changes or enhanced X-ray flux. As 1E 2259+586 experiences glitch-like timing irregularities of both signs, and is spinning down, it is challenging to reconcile anti-glitches in this object with the vortex avalanche model, and it is possible that they are instead due to changes in the internal magnetic field (Garcia & Ranea-Sandoval, 2015; Mastrano, Suvorov & Melatos, 2015). However, in the case of an accreting neutron star which is spinning up, the vortex avalanche model implies the possibility of anti-glitches from avalanches of vortices moving inward.

Recent observations from the *Neutron Star Interior Composition Explorer* (NICER) mission on the International Space Station (Gendreau et al., 2016) reported three spin-down glitches in the accreting pulsar NGC 300 ULX-1, (Ray et al., 2019). NGC 300 ULX-1 is a high-mass X-ray binary in the nearby galaxy NGC 300, which is visible only in outburst (Monard, 2010; Binder et al., 2011). A NICER observation campaign in 2018 was able to time this object with daily sampling for ≈ 100 d before it became too faint. During this period, Ray et al. (2019) constructed a coherent timing model over several contiguous data segments of ≈ 10 d each, finding that the pulse frequency is steadily increasing with $\dot{\nu} = 4.3 \times 10^{-10}$ Hz s $^{-1}$. Ray et al. (2019) also measured an anti-glitch with $\Delta\nu/\nu = -4.4 \times 10^{-4}$ at MJD 58243

and a second anti-glitch with $\Delta\nu/\nu = -5.5 \times 10^{-4}$ at MJD 58265. A third anti-glitch with $\Delta\nu/\nu = -7 \times 10^{-4}$ occurred at MJD 58334, however, a phase-coherent timing model could not be constructed around this event due to reduced X-ray flux. No spectral changes were detected around any of the glitches, however, this finding is not conclusive due to the inherent variability of NGC 300 ULX-1. As this pulsar is spinning up, these anti-glitches appear to be the exact inverse of most other glitches – spin-up events in spinning-down pulsars.

In order to test whether the vortex code can produce anti-glitches similar to those in NGC 300 ULX-1, we have performed simulations on pinned vortices in an accelerating container. In this scenario, the tendency of the vortices is to move toward the centre of the container. In order to maintain the Feynman condition (1.5), new vortices must be added as the container accelerates. This is implemented in the code via a simple algorithm which tracks the accumulated spin up, and inserts a new vortex near the boundary at a random angular coordinate when Ω increases by $1/2\pi R^2$. This reflects that vortex nucleation occurs near boundaries in terrestrial superfluids and BECs (Donnelly, 1991; Berloff & Barenghi, 2004; Stagg & Parker, 2017).

We perform simulations similar to those described in 4.5, comparing the evolution of vortices in both accelerating and decelerating containers. In these simulations, we have $N_v = 1000$, $R = 10$, $\Delta t = 0.1T_0$, $I_{\text{rel}} = 1$, $V_0 = 1000$, $0.01R$, $\xi = 0.0001R$, $\phi = 0.1$ rad, and $N_{\text{ext}} = \pm 10^{-3}\Omega_0/T_0$. We initialise the vortices by drawing their positions from a random spatial distribution. We then evolve the system without spin up or feedback until all vortices are pinned, and use the same initial vortex configuration for both the spin up and spin down simulations. We run the simulation for 5×10^4 time steps.

Figure 5.1 shows the evolution of Ω_C during the spin up (left panel) and spin down (right panel) simulations. The behaviour of Ω_C in Figure 5.1 is similar in both the spin up and spin down cases. The steady spin up/down is punctuated by abrupt spin down/up events of varying size at random times. However, the anti-glitches are smaller than the glitches in general. The largest anti-glitch has $|\Delta\Omega/\Omega_0| = 3 \times 10^{-5}$ while the largest glitch has $|\Delta\Omega/\Omega_0| = 7 \times 10^{-3}$. Observing the vortex motion directly, we see analogous avalanche behaviour in the spin up simulations to Figure 4.5. Single vortices unpin and move inwards, causing knock-on unpinning of other vortices which also move inwards.

We run the glitch-finding algorithm described in Section 4.3.3 on both simulations, finding 45 glitches in the spin down simulation and 51 anti-glitches in the spin up simulation. Figure 5.2 shows kernel density estimates of the size and waiting time PDFs in both simulations. For sizes, we look




Figure 5.1: Container angular velocity Ω_C , normalized by its initial value, as a function of time over 10^4 rotation periods.

at the PDF of the absolute value $|\Delta\Omega|/\Omega$ for ease of comparison. The PDFs for both simulations appear qualitatively similar, though the size PDF for the spin up simulation is steeper than for the spin down simulation. A 2-sample KS test returns a p -value of 0.51 for the waiting times and 5.0×10^{-4} for the sizes. This means that the null hypothesis that the PDFs are the same for each simulation is not rejected for the waiting time PDFs, but is strongly disfavoured for the size PDFs. Looking at the vortex motion directly, it appears that adding new vortices in the spin up simulation often causes immediate knock-on of the nearby vortices, leading to many small avalanches, not all of which are picked up by the glitch-finding algorithm as they are not large enough to change the sign of $\dot{\Omega}$. This appears to be a boundary effect, and so longer simulations of larger systems producing more glitches should be done to see whether this discrepancy in the size PDFs persists.

We do not perform a quantitative comparison between the results shown here and those in Chapter 4, due to the smaller system size and smaller glitch samples produced by these preliminary simulations. Nonetheless, we show that pinned vortex avalanches are ubiquitous in decelerating and accelerating containers, and that the vortex avalanche model has the potential to explain both pulsar glitches and the antiglitches observed in NGC 300 ULX-1. Future many-vortex simulations, as well as more observations of accelerating pulsars such as NGC 300 ULX-1, present a good opportunity for better understanding superfluid dynamics in neutron stars.

5.2.2 N -body vortex method in three dimensions

The results in this thesis demonstrate that pinned vortices in a decelerating container exhibit avalanche dynamics across a wide range of input parameters. An obvious extension to the model is to consider the motion of vortices in three dimensions. Several important effects emerge in three dimensions, namely vortex tension, reconnection and turbulence, as discussed in Section 1.4.4 and references therein.

Superfluid hydrodynamics is a mature research field within condensed matter physics, and several existing codes use the vortex filament method to study quantum turbulence, e.g. (Adachi, Fujiyama & Tsubota, 2010; Baggaley & Barenghi, 2012). However, these simulations often assume periodic boundary conditions and a fixed background normal fluid flow, though see Yui, Tsubota & Kobayashi (2018), who solved the Biot-Savart equation (1.13) coupled to a Navier-Stokes solver for the normal fluid in a square channel with hard boundaries in two dimensions and a periodic boundary condition along the initial direction of the normal fluid flow. For vortex avalanche simulations, interactions between vortices and pinning sites, the container




Figure 5.2: Kernel density estimates of the glitch size and waiting time PDFs for spin up (blue curve) and spin down (orange dashed curve) simulations.

boundary, and coupling to a time-varying normal fluid must be included.

We have begun work on adapting the solver described in Chapter 4 to three dimensions, closely following the numerical methods described in (Schwarz, 1985) and (Baggaley & Barenghi, 2012). Vortex filaments are discretized into a set of points $\mathbf{s}_i$, and we solve for the motion of the vortices using the split Biot-Savart equation (1.14)

$$\mathbf{v}_s(\mathbf{s}_i, t) = \frac{\kappa}{4\pi} \ln \left[\frac{(l_i l_{i+1})^{1/2}}{a_0} \right] \mathbf{s}'_i \times \mathbf{s}''_i + \frac{\kappa}{4\pi} \oint_{\mathcal{L}'} d\xi \frac{\mathbf{s}' \times (\mathbf{s}_i - \mathbf{s})}{|\mathbf{s}_i - \mathbf{s}|^3}. \quad (5.1)$$

We use a fourth-order Runge-Kutta scheme for the time evolution and compute the spatial derivatives using fourth-order finite difference methods. We have not yet implemented pinning, dissipation or interaction with a normal fluid into the 3D version of the code, nor have we implemented a reconnection procedure for vortex crossings.

In order to test the code, we reproduce the well-known phenomenon of leapfrogging vortex rings. In the absence of viscosity, a vortex ring is advected along the direction of its normal, with a speed inversely proportional to the radius of the ring. If two identical vortex rings are placed one above the other along the same axis of symmetry, the lower ring contracts, passes through the upper ring, and the ‘leapfrogging’ process continues ad infinitum (Acheson, 1990). We reproduce this result using the 3D point vortex code with two vortex rings of unit radius (using the same coordinate system described in 4.3.2) separated in the z -direction by $R/10$ length units. Each ring is discretized into 100 points. We show snapshots of one cycle of the leapfrogging process in Figure 5.3. Figure 5.4 shows the radius of each ring, averaged over all vortex points, as a function of time over approximately five leapfrogging cycles. Figures 5.3 and 5.4 demonstrate that the 3D vortex code is able to reproduce the classical result as expected.

As a comparison between vortex motion in two and three dimensions, we run simulations of vortex filaments arranged in a ring, analogous to the stability analysis in Section 4.4.1. We begin the simulations with $2 \leq N \leq 10$ straight vortex filaments of length 2 discretized into 20 points. The filaments are initially placed at equal angular separation about the edge of a ring of radius 1. The boundary condition is that the top and bottom vortex points of each filament are at $z = \pm 1$ throughout the simulation, as vortices must either form closed loops or attach to a boundary. We run the simulations for a length of time approximately equal to one rotation period of the 2-filament vortex array at its initial rotation speed $T_2(t = 0)$.

As in Section 4.4.1, the vortices initially rotate about the centre of the starting ring, with the rotation speed increasing with the number of filaments.




Figure 5.3: Snapshots from simulation of leapfrogging vortex rings. Initially, the rings have unit radius $R = 1$ and are separated by 0.1 in the z -direction. Each ring is discretized into 100 points. Time increasing from top to bottom (arbitrary units).




Figure 5.4: Radius R (averaged over all vortex points) of leapfrogging vortex rings as a function of time t over approximately five cycles. The orange curve shows R versus t for the ring which is initially above, blue curve shows same for the ring which is initially below.




Figure 5.5: Initial angular velocity of rectilinear N -filament array, $\omega_N(t=0)$, averaged over all vortex points in filaments and all filaments in array, normalised by the initial value of the 2-filament array, $\omega_2(t=0)$. Blue dots indicate data, black curve is a power law fit $\omega_N = \omega_2 N^{1.1}$.

However, unlike in the 2-dimensional case, the increase with N is not exactly linear, but instead appears to obey a power law $\omega_N \propto (N-1)^{1.1}$. Figure 5.5 shows the initial angular velocity $\omega_N(t=0) = v/r$, where $v = (v_x^2 + v_y^2)^{1/2}$, as a function of N , normalized by its initial value in the $N=2$ case.

Another interesting difference between the two and three-dimensional cases is that the vortex filaments do not rotate at uniform velocity. Instead, they begin to stretch and curve outward in an oscillatory fashion, likely a signature of Kelvin waves being excited along the filaments. As the vortices stretch, the rotation speed decreases. Figure 5.6 shows several snapshots of the array as the simulation progresses. Figure 5.7 shows the angular velocity of the vortex points, averaged over each point in filament and each filament in array, as a function of time, as well as the maximum radial displacement $\Delta R = R - 1$ for all vortex points in the array for $N=2$, $N=3$, $N=5$ and $N=9$ cases. Figure 5.7 shows that both the amplitude and frequency of the




Figure 5.6: 10-filament vortex array configurations at $t = 0$ (top panel) and $t = 0.13/T_2(t = 0)$ (bottom panel).




Figure 5.7: Top panel: angular velocity $\langle \omega_N \rangle$ (averaged over each point in filament and each filament in array) versus time for $N = 2, 3, 5, 9$ filament arrays. Bottom panel: radial displacement $\Delta R = R - 1$ versus time for $N = 2, 3, 5, 9$ filament arrays. R shown is maximum value of R for all vortex points in the array.




Figure 5.8: Oscillation frequency ν_N versus initial azimuthal velocity $\omega_{N,i}$ for $2 \leq N \leq 7$ vortex filament arrays. Blue dots indicate data, black curve shows a quadratic fit $\nu_n \propto -\omega_N^2$ to the data.

oscillations increase with N . For $N > 7$, after several oscillations the array becomes unstable and the filaments disintegrate, as shown by the red curves in Figure 5.7. It is unclear whether this is a natural instability similar to that seen in Section 4.4.1 or a numerical instability in the code, but the fact that it appears for the same number of vortices as in the two-dimensional case is intriguing. Another interesting effect is that while the amplitude of the oscillations increases approximately linearly for $N \leq 5$, for $N > 5$ it does not increase beyond $A = \Delta R \approx 0.25$.

Figure 5.8 shows the oscillation frequency ν_N , in units of $1/T_2(t=0)$, versus initial angular velocity for the arrays with $2 \leq N \leq 7$. ν_n is measured by hand from successive peaks and troughs from Figure 5.7. For $N > 8$, the oscillation frequency cannot be reliably measured due to the above-mentioned instability. We also include a quadratic fit $\nu_n \propto -\omega_N^2$ in Figure 5.8, which closely follows the data. This is consistent with the dispersion relation for Kelvin waves in superfluids $\nu \propto v^2$, where v is the linear velocity of the vortex

filament. (Andereck & Glaberson, 1982; Simula, Mizushima & Machida, 2008).

The results shown here for short simulations of several vortex filaments already demonstrate that three-dimensional vortex dynamics are more complex than the two-dimensional case. Further development of this code to incorporate effects such as pinning, reconnection and dissipation in larger systems is an important step for testing the vortex avalanche model of pulsar glitches.

5.2.3 Vortex annihilation and Onsager vortex formation

Onsager (1949)²⁰ predicted that macroscopic vortices in fluids can form from the turbulent dynamics of point vortices in two dimensions. Simula, Davis & Helmerson (2014) demonstrated the formation of these Onsager vortices in GPE simulations of an isolated zero-temperature Bose-Einstein condensate in 2D. In the simulations, the BEC is seeded with an equal number of opposite-signed vortices at random initial locations. As the system evolves, when two vortices come within the vortex core radius of one another they annihilate, emitting sound waves. Not all of the energy from the vortices is converted into sound waves however, so each annihilation increases the kinetic energy of the remaining vortices. Because the Hamiltonian is a function only of the inter-vortex separation, the remaining vortices form like-signed clusters in a process referred to as evaporative heating. The clusters persist on long time scales $t \sim 100\tau$, where τ is the inverse of the trap frequency of the BEC. Simula, Davis & Helmerson (2014) also performed point vortex simulations, presumably using a similar method to that described in Section 4.2, which largely agreed with the results of the GPE simulations. More recent GPE simulations also demonstrate Onsager vortex formation (Groszek et al., 2016; Han & Tsubota, 2019). Gauthier et al. (2019) performed vortex clustering experiments on a BEC of ^{87}Rb in a quasi-2D elliptical trap. In one experiment, they spun up vortices in same-signed clusters on opposite signs of the trap, while in another experiment vortex-antivortex pairs were spun up throughout the trap. In the first experiment, the vortex clusters persisted throughout the hold time of the experiment, with gradual loss of vortices mainly due to thermal dissipation. In the case where the vortices and antivortices are not initially clustered together, evaporative heating does not occur, which Gauthier et al. (2019) suggest is due to the loss of vortices from thermal dissipation.

²⁰The same paper that predicted quantized vorticity in superfluids.




Figure 5.9: Number of vortices N vs time as a function of the annihilation threshold σ .

In order to perform simulations to study Onsager vortex formation with the point vortex code, which does not include sound waves, we need to include a prescription for vortex anti-vortex annihilation. In GPE simulations, annihilations occur naturally via the emission of sound waves when vortices come within one another's core radius, but the point vortex model assumes an infinitesimal core radius. We introduce an annihilation threshold parameter, σ , which has units of length. We implement an annihilation algorithm by checking the separation between each vortex and all of the vortices of opposite sign at each time step. If the separation of any vortex anti-vortex pair is below σ , both vortices are removed from the simulation. The choice of σ affects the results of simulations significantly. Figure 5.9 shows the number of vortices N versus time for a suite of simulations where we begin with 100 vortices in total (50 of each sign) and take $\sigma \in \{10^{-4}R, 10^{-3}R, 10^{-2}R, 2 \times 10^{-2}R, 5 \times 10^{-2}R\}$. The time unit T is taken to be the equilibrium rotation period for a 100-vortex array confined to a container of radius R , as in Section 4.3. Figure 5.9 shows that the rate of vortex loss dN/dt increases with increasing σ . For

the $\sigma = 10^{-3}R$ and $\sigma = 10^{-4}R$ cases, the shape of the decay curve is different than for $\sigma \geq 10^{-2}R$. For low values of σ vortices do not come close enough to annihilate, and only leave the simulation through the boundary. For $\sigma \geq 10^{-2}R$, annihilation of vortex-antivortex pairs are the dominant way that vortices leave the simulation, and they become more common as σ increases. For $\sigma \geq 10^{-2}R$, we also see clustering of the vortices and antivortices into Onsager vortices. Figure 5.10 shows the vortex configuration at the beginning and end of the simulations with $\sigma = 10^{-2}$, $2 \times 10^{-2}R$ and $5 \times 10^{-2}R$. For the remainder of this section we take $\sigma = 10^{-2}R$.

Groszek et al. (2016) performed GPE simulations with $N \approx 100$ and examined the dynamics of annihilations in systems at zero and non-zero temperature. They found that in the zero temperature regime, vortex-antivortex pairs annihilate via three and four-vortex processes, as an isolated vortex-antivortex pair is stable in the absence of dissipation. At non-zero temperature many-body processes become less important. Thermal dissipation allows vortex-antivortex pairs to come close enough to annihilate on their own. These results are consistent with experiments of quasi-2D BECs by Kwon et al. (2014), who created a turbulent vortex array by stirring a static BEC with a repulsive optical potential and observed that the decay of vortices was consistent with one-body (i.e. drift from the container) and two-body (i.e. annihilation) processes. Baggaley & Barenghi (2018) also studied vortex decay by performing GPE simulations of vortices in an unbounded fluid and measuring the exponent k of the rate equation

$$\frac{dN}{dt} = -\Gamma N^k, \quad (5.2)$$

where Γ is a rate constant. In the limit $t \gg 0$, the solution to (5.2) is $N \propto t^{1/(1-k)}$. Baggaley & Barenghi (2018) found that their zero temperature results were consistent with $k = 4$, suggesting the four-body annihilation process described by Groszek et al. (2016). The non-zero temperature results were more consistent with $k = 2$, i.e. direct vortex-antivortex annihilation.

To test whether we similar behaviour in the point vortex code, we run a set of simulations with $N = 500$ with and without the dissipation mechanism described in Section 4.2. No dissipation, i.e. $\phi = 0$, corresponds to zero-temperature in the BEC case, while the simulations with dissipation, where we take $\phi = 10^{-2}$, correspond to non-zero temperature. For each case we run five simulations with different random initial vortex configuration and take the ensemble average of N at each time step. Figure 5.11 shows plots of N versus time, as well as fits to $N \propto t^a$. The power law fits shown in Figure 5.11 correspond to values of $k = 4.0$ in the $\phi = 0$ case and $k = 2.1$ in the $\phi = 10^{-2}$ case, consistent with the experimental results for



Figure 5.10: Vortex configurations during Onsager vortex formation. Blue dots indicate positions of vortices with counter-clockwise circulation, orange dots indicate positions of vortices with clockwise circulation of equal magnitude. Annihilation threshold σ increasing from top to bottom. Left column shows configurations at beginning of simulations, right column shows same after $t = 1000T$.




Figure 5.11: Number of vortices N , ensemble-averaged over five simulations, versus time with $\phi = 0$ (blue curve) and $\phi = 10^{-3}$ (orange curve). Grey dashed line is a fit to $N \propto t^a$ for the $\phi = 0$ (or $T = 0$) case, grey dotted line is same for $\phi = 10^{-2}$ ($T > 0$).

BECs (Kwon et al., 2014) and GPE simulations (Groszek et al., 2016; Baggaley & Barenghi, 2018). These results show that the point vortex method can achieve qualitatively similar results to more computationally-expensive methods such as GPE simulations, as well as demonstrating the versatility of the code in applications beyond pulsar glitch studies.

Bibliography

- Abdo A. A. et al., 2013, ApJS, 208, 17
- Abrikosov A., 1957, Journal of Physics and Chemistry of Solids, 2, 199
- Acheson D. J., 1990, Elementary fluid dynamics. Oxford University Press
- Adachi H., Fujiyama S., Tsubota M., 2010, Phys. Rev. B, 81, 104511
- Aegerter C. M., 1998, Phys. Rev. E, 58, 1438
- Akbal O., Alpar M. A., 2018, MNRAS, 473, 621
- Allen J. F., Misener A. D., 1938, Nature, 141, 75
- Alpar M. A., 1977, ApJ, 213, 527
- Alpar M. A., Baykal A., 2006, MNRAS, 372, 489
- Alpar M. A., Chau H. F., Cheng K. S., Pines D., 1996, ApJ, 459, 706
- Alpar M. A., Cheng K. S., Pines D., 1989, ApJ, 346, 823
- Alpar M. A., Langer S. A., Sauls J. A., 1984, ApJ, 282, 533
- Alpar M. A., Nandkumar R., Pines D., 1986, ApJ, 311, 197
- Altshuler E., Johansen T. H., 2004, Reviews of Modern Physics, 76, 471
- Andereck C. D., Glaberson W. I., 1982, Journal of Low Temperature Physics, 48, 257
- Anderson P. W., Alpar M. A., Pines D., Shaham J., 1982, Philosophical Magazine, Part A, 45, 227
- Anderson P. W., Itoh N., 1975, Nature, 256, 25
- Andersson N., Comer G. L., 2006, Classical and Quantum Gravity, 23, 5505
- Andersson N., Comer G. L., Prix R., 2004, MNRAS, 354, 101

- Andersson N., Glampedakis K., Ho W. C. G., Espinoza C. M., 2012, *Physical Review Letters*, 109, 241103
- Andersson N., Haskell B., Samuelsson L., 2011, *Monthly Notices of the Royal Astronomical Society*, 416, 118
- Andersson N., Haskell B., Samuelsson L., 2011, *MNRAS*, 416
- Andersson N., Sidery T., Comer G. L., 2006, *MNRAS*, 368, 162
- Andersson N., Sidery T., Comer G. L., 2007, *MNRAS*, 381, 747
- Andreev A. F., Bashkin E. P., 1976, *Soviet Journal of Experimental and Theoretical Physics*, 42, 164
- Andryushchenko V. A., Kondaurova L. P., Nemirovskii S. K., 2016, *Journal of Low Temperature Physics*, 185, 377
- Antonopoulou D., Espinoza C. M., Kuiper L., Andersson N., 2018, *MNRAS*, 473, 1644
- Archibald R. F. et al., 2013, *Nature*, 497, 591
- Arms R. J., Hama F. R., 1965, *Physics of Fluids*, 8, 553
- Arzoumanian Z. et al., 2018, *ApJS*, 235, 37
- Arzoumanian Z., Nice D. J., Taylor J. H., Thorsett S. E., 1994, *ApJ*, 422, 671
- Ashton G., Lasky P. D., Graber V., Palfreyman J., 2019, *Nature Astronomy*, 3, 1143
- Ashton G., Prix R., Jones D. I., 2017, *Phys. Rev. D*, 96, 063004
- Baade W., Zwicky F., 1934, *Proceedings of the National Academy of Science*, 20, 259
- Backer D. C., Hellings R. W., 1986, *ARA&A*, 24, 537
- Bagchi P., Balachandar S., 2002, *Journal of Fluid Mechanics*, 466, 365
- Baggaley A. W., Barengi C. F., 2011, *Phys. Rev. B*, 84, 020504
- Baggaley A. W., Barengi C. F., 2012, *Journal of Low Temperature Physics*, 166, 3
- Baggaley A. W., Barengi C. F., 2018, *Phys. Rev. A*, 97, 033601
- Bahcall J. N., Wolf R. A., 1965, *Phys. Rev. Lett.*, 14, 343

- Bailes M. et al., 2018, arXiv e-prints, arXiv:1803.07424
- Bak P., Christensen K., Danon L., Scanlon T., 2002, Phys. Rev. Lett., 88, 178501
- Bak P., Tang C., Wiesenfeld K., 1987, Physical Review Letters, 59, 381
- Barenghi C., Parker N., 2016, A Primer on Quantum Fluids, SpringerBriefs in Physics. Springer International Publishing
- Barenghi C. F., Donnelly R. J., Vinen W. F., 1983, Journal of Low Temperature Physics, 52, 189
- Barenghi C. F., Jones C. A., 1988, Journal of Fluid Mechanics, 197, 551
- Barnes J., Hut P., 1986, Nature, 324, 446
- Bassler K. E., Paczuski M., 1998, Phys. Rev. Lett., 81, 3761
- Basu A., Joshi B. C., Krishnakumar M. A., Bhattacharya D., Nandi R., Bandhopadhyay D., Char P., Manoharan P. K., 2020, MNRAS, 491, 3182
- Baym G., Bethe H. A., Pethick C. J., 1971, Nuclear Physics A, 175, 225
- Baym G., Pethick C., Pines D., 1969, Nature, 224, 673
- Baym G., Pethick C., Pines D., Ruderman M., 1969, Nature, 224, 872
- Baym G., Pines D., 1971, Annals of Physics, 66, 816
- Bean C. P., 1964, Reviews of Modern Physics, 36, 31
- Bekarevich I., Khalatnikov I., 1961, Sov. Phys. JETP, 13, 643
- Benton E., Clark Jr A., 1974, Annual Review of Fluid Mechanics, 6, 257
- Berloff N. G., Barenghi C. F., 2004, Phys. Rev. Lett., 93, 090401
- Bhattacharya D., van den Heuvel E. P. J., 1991, Phys. Rep., 203, 1
- Binder B., Williams B. F., Kong A. K. H., Gaetz T. J., Plucinsky P. P., Dalcanton J. J., Weisz D. R., 2011, ApJ, 739, L51
- Blaauwgeers R. et al., 2002, Physical Review Letters, 89, 155301
- Boyd J., 2013, Chebyshev and Fourier Spectral Methods: Second Revised Edition, Dover Books on Mathematics. Dover Publications
- Boynton P. E., Groth E. J., Hutchinson D. P., Nanos, G. P. J., Partridge R. B., Wilkinson D. T., 1972, ApJ, 175, 217

Bransgrove A., Beloborodov A. M., Levin Y., 2020, ApJ, 897, 173

Bustamante M. D., Nazarenko S., 2015, Phys. Rev. E, 92, 053019

Caleb M. et al., 2016, MNRAS, 458, 718

Cameron A. G. W., 1965, Nature, 205, 787

Campbell L. J., Krasnov Y. K., 1982, Journal of Low Temperature Physics, 49, 377

Campbell L. J., Ziff R. M., 1979, Phys. Rev. B, 20, 1886

Canuto C., Hussaini Y., Quarteroni A., Thomas A. J., 1993, Spectral Methods in Fluid Dynamics, Scientific Computation. Springer Berlin Heidelberg

Carlin J. B., Melatos A., 2019, MNRAS, 488, 4890

Carlin J. B., Melatos A., 2020, MNRAS

Carter B., Chamel N., Haensel P., 2006, International Journal of Modern Physics D, 15, 777

Chamel N., 2012, Phys Rev C, 85, 035801

Chamel N., 2013, Physical Review Letters, 110, 011101

Chamel N., Haensel P., 2008, Living Reviews in Relativity, 11

Chandler E., Baym G., 1986, Journal of Low Temperature Physics, 62, 119

Chaudhuri P., Marron J. S., 1999, Journal of the American Statistical Association, 94, 807

Chen K., Ruderman M., 1993, ApJ, 402, 264

Cheng K. S., 1987, ApJ, 321, 799

Cheng K. S., Pines D., Alpar M. A., Shaham J., 1988, ApJ, 330, 835

Chiu H. Y., 1964, Annals of Physics, 26, 364

Chukwude A. E., Urama J. O., 2010, MNRAS, 406, 1907

Cocke W. J., Disney M. J., Taylor D. J., 1969, Nature, 221, 525

Cordes J. M., Downs G. S., 1985, ApJS, 59, 343

Cordes J. M., Greenstein G., 1981, ApJ, 245, 1060

- Cordes J. M., Shannon R. M., 2010, arXiv e-prints, arXiv:1010.3785
- D'Alessandro F., McCulloch P. M., Hamilton P. A., Deshpande A. A., 1995, MNRAS, 277, 1033
- Dang S. J. et al., 2020, ApJ, 896, 140
- Davis K. B., 1975, The Annals of Statistics, 1025
- Davis K. B., 1977, The Annals of Statistics, 530
- Dib R., Kaspi V. M., Gavril F. P., 2008, ApJ, 673, 1044
- Dodson R., Lewis D., McCulloch P., 2007, Ap&SS, 308, 585
- Dodson R. G., McCulloch P. M., Lewis D. R., 2002, ApJ, 564, L85
- Donnelly R. J., 1991, Quantized Vortices in Helium II
- Drummond L. V., Melatos A., 2017, MNRAS, 472, 4851
- Drummond L. V., Melatos A., 2018, MNRAS, 475, 910
- Edwards R. T., Bailes M., 2001, ApJ, 553, 801
- Edwards R. T., Hobbs G. B., Manchester R. N., 2006, MNRAS, 372, 1549
- Elshamouty K. G., Heinke C. O., Sivakoff G. R., Ho W. C. G., Shternin P. S., Yakovlev D. G., Patnaude D. J., David L., 2013, ApJ, 777, 22
- Elzhov T. V., Mullen K. M., Spiess A.-N., Bolker B., 2013, minpack.lm: R interface to the Levenberg-Marquardt nonlinear least-squares algorithm found in MINPACK, plus support for bounds. R package version 1.1-8
- Espinoza C. M., Antonopoulou D., Dodson R., Stepanova M., Scherer A., 2021, A&A, 647, A25
- Espinoza C. M., Antonopoulou D., Stappers B. W., Watts A., Lyne A. G., 2014, MNRAS, 440, 2755
- Espinoza C. M., Lyne A. G., Stappers B. W., Kramer M., 2011, MNRAS, 414, 1679
- Eya I. O., Urama J. O., Chukwude A. E., 2017, ApJ, 840, 56
- Ferdman R. D., Archibald R. F., Gourgouliatos K. N., Kaspi V. M., 2018, ApJ, 852, 123
- Fetter A. L., 1965, Physical Review, 138, 429

- Fetter A. L., 2009, *Reviews of Modern Physics*, 81, 647
- Fetter A. L., Hohenberg P. C., Pincus P., 1966, *Physical Review*, 147, 140
- Fetter A. L., Svidzinsky A. A., 2001, *Journal of Physics Condensed Matter*, 13, R135
- Feynman R. P., 1955, in *Progress in Low Temperature Physics*, C.J. Gorter, ed., Vol. 1, Elsevier
- Field S., Witt J., Nori F., Ling X., 1995, *Phys. Rev. Lett.*, 74, 1206
- Franco L. M., Link B., Epstein R. I., 2000, *ApJ*, 543, 987
- Fuentes J. R., Espinoza C. M., Reisenegger A., 2019, *A&A*, 630, A115
- Fuentes J. R., Espinoza C. M., Reisenegger A., Shaw B., Stappers B. W., Lyne A. G., 2017, *A&A*, 608, A131
- Fulgenzi W., Melatos A., Hughes B. D., 2017, *MNRAS*, 470, 4307
- Garcia F., Ranea-Sandoval I. F., 2015, *MNRAS*, 449, L73
- Gauthier G. et al., 2019, *Science*, 364, 1264
- Ge M. Y. et al., 2020, *ApJ*, 896, 55
- Gendreau K. C. et al., 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9905, *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*, p. 99051H
- Giliberti E., Cambiotti G., Antonelli M., Pizzochero P. M., 2020, *MNRAS*, 491, 1064
- Glaberson W. I., Johnson W. W., Ostermeier R. M., 1974, *Phys. Rev. Lett.*, 33, 1197
- Glampedakis K., Andersson N., 2009, *Physical Review Letters*, 102, 141101
- Glampedakis K., Andersson N., 2011, *ApJ*, 740, L35
- Gold T., 1968, *Nature*, 218, 731
- Goldreich P., Julian W. H., 1969, *ApJ*, 157, 869
- Gorter C. J., Mellink J. H., 1949, *Physica*, 15, 285
- Graber V., Andersson N., Glampedakis K., Land er S. K., 2015, *MNRAS*, 453, 671

- Graber V., Cumming A., Andersson N., 2018, *ApJ*, 865, 23
- Greenstein G., 1970, *Nature*, 227, 791
- Gross E. P., 1961, *Il Nuovo Cimento*, 20, 454
- Groszek A. J., Simula T. P., Paganin D. M., Helmerson K., 2016, *Phys. Rev. A*, 93, 043614
- Gügercinoğlu E., Alpar M. A., 2019, *MNRAS*, 488, 2275
- Hall H. E., Vinen W. F., 1956, *Proceedings of the Royal Society of London Series A*, 238, 215
- Hall P., 1992, *The bootstrap and Edgeworth expansion*. Springer Science & Business Media
- Hall P., Marron J., Park B. U., 1992, *Probability theory and related fields*, 92, 1
- Han J., Tsubota M., 2019, *Phys. Rev. A*, 99, 033607
- Hänninen R., Baggaley A. W., 2014, *Proceedings of the National Academy of Science*, 111, 4667
- Haskell B., Andersson N., Comer G. L., 2012, *Physical Review D*, 86, 063002
- Haskell B., Antonopoulou D., 2014, *MNRAS*, 438, L16
- Haskell B., Antonopoulou D., Barenghi C., 2020, *arXiv e-prints*, arXiv:2007.02748
- Haskell B., Melatos A., 2015, *International Journal of Modern Physics D*, 24, 30008
- Haskell B., Melatos A., 2016, *MNRAS*, 461, 2200
- Haskell B., Pizzochero P. M., Seveso S., 2013, *ApJ*, 764, L25
- Haskell B., Pizzochero P. M., Sidery T., 2012, *MNRAS*, 420, 658
- Havelock T., 1931, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 11, 617
- Heinke C. O., Ho W. C. G., 2010, *ApJ*, 719, L167
- Helfand D. J., Manchester R. N., Taylor J. H., 1975, *ApJ*, 198, 661
- Henderson K. L., Barenghi C. F., 1995, *Journal of Low Temperature Physics*, 98, 351
- Henderson K. L., Barenghi C. F., Jones C. A., 1995, *Journal of Fluid Mechanics*, 283, 329

- Hewish A., Bell S. J., Pilkington J. D. H., Scott P. F., Collins R. A., 1968, *Nature*, 217, 709
- Hills R. N., Roberts P. H., 1977, *Archive for Rational Mechanics and Analysis*, 66, 43
- Hobbs G. et al., 2020, *MNRAS*, 491, 5951
- Hobbs G., Lyne A. G., Kramer M., 2010, *MNRAS*, 402, 1027
- Hobbs G., Lyne A. G., Kramer M., Martin C. E., Jordan C., 2004, *MNRAS*, 353, 1311
- Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655
- Hooker J., Newton W. G., Li B.-A., 2015, *MNRAS*, 449, 3559
- Horner C. L., Van Gorder R. A., 2019, *Physics of Fluids*, 31, 065103
- Horowitz C. J., Kadau K., 2009, *Phys. Rev. Lett.*, 102, 191102
- Howitt G., Haskell B., Melatos A., 2016, *MNRAS*, 460, 1201
- Howitt G., Melatos A., Delaigle A., 2018, *ApJ*, 867, 60
- Idowu O. C., Kivotides D., Barengi C. F., Samuels D. C., 2001, *Numerical Methods for Coupled Normal-Fluid and Superfluid Flows in Helium II*, Barengi C. F., Donnelly R. J., Vinen W. F., eds., Vol. 571, p. 162
- Idowu O. C., Willis A., Barengi C. F., Samuels D. C., 2000, *Phys. Rev. B*, 62, 3409
- Irvine J., 1978, *Neutron stars*, Oxford studies in physics. Clarendon Press
- Jankowski F. et al., 2019, *MNRAS*, 484, 3691
- Janssen G. H., Stappers B. W., 2006, *A&A*, 457, 611
- Jensen H. J., 1998, *Self-Organized Criticality*. Cambridge: University Press
- Johnston S., Galloway D., 1999, *MNRAS*, 306, L50
- Jones D. I., 2012, *MNRAS*, 420, 2325
- Jones P. B., 1991, *MNRAS*, 253, 279
- Jones P. B., 2003, *ApJ*, 595, 342
- Jones P. B., 2006, *MNRAS*, 365, 339

- Kapitza P., 1938, *Nature*, 141, 74
- Kaspi V. M., Beloborodov A. M., 2017, *ARA&A*, 55, 261
- Kaspi V. M., Gavriil F. P., Woods P. M., Jensen J. B., Roberts M. S. E., Chakrabarty D., 2003, *ApJ*, 588, L93
- Kaspi V. M., Kramer M., 2016, arXiv e-prints, arXiv:1602.07738
- Kerr M., Hobbs G., Johnston S., Shannon R. M., 2016, *MNRAS*, 455, 1845
- Khalatnikov I. M., Hohenberg P. C., 1965, *An introduction to the theory of superfluidity*. WA Benjamin New York
- Khomenko V., Antonelli M., Haskell B., 2019, *Phys. Rev. D*, 100, 123002
- Khomenko V., Haskell B., 2018, *PASA*, 35, e020
- Konar S., Arjunwadkar M., 2014, ArXiv e-prints
- Kondaurova L., L’vov V., Pomyalov A., Procaccia I., 2014, *Phys. Rev. B*, 90, 094501
- Kondaurova L., Nemirovskii S. K., 2012, *Phys. Rev. B*, 86, 134506
- Kramer M., Stappers B., 2010, in *ISKAF2010 Science Meeting*
- Kwon W. J., Moon G., Choi J.-y., Seo S. W., Shin Y.-i., 2014, *Phys. Rev. A*, 90, 063627
- Landau L., 1941, *Physical Review*, 60, 356
- Landau L. D., 1932, *Phys. Zs. Sowjet*, 1, 285
- Landau L. D., Lifshitz E. M., 1959, *Fluid mechanics*
- Large M. I., Vaughan A. E., Mills B. Y., 1968, *Nature*, 220, 340
- Lattimer J. M., Prakash M., 2004, *Science*, 304, 536
- Levin L. et al., 2018, in *IAU Symposium, Vol. 337, Pulsar Astrophysics the Next Fifty Years*, Weltevrede P., Perera B. B. P., Preston L. L., Sanidas S., eds., pp. 171–174
- Lórinicz K. A., Wijngaarden R. J., 2007, *Phys. Rev. E*, 76, 040301
- Lin C. C., 1941, *Proceedings of the National Academy of Science*, 27, 570
- Link B., 2009, *Physical Review Letters*, 102, 131101

- Link B., 2012a, MNRAS, 421, 2682
- Link B., 2012b, MNRAS, 422, 1640
- Link B., Epstein R. I., Lattimer J. M., 1999, Phys. Rev. Lett., 83, 3362
- Link B. K., Epstein R. I., 1991, ApJ, 373, 592
- Liu K., Verbiest J. P. W., Kramer M., Stappers B. W., van Straten W., Cordes J. M., 2011, MNRAS, 417, 2916
- Lönnborn J. R., Melatos A., Haskell B., 2019, MNRAS, 487, 702
- Lorimer D. R., Kramer M., 2004, Handbook of Pulsar Astronomy
- Lower M. E. et al., 2020, MNRAS, 494, 228
- Lyne A., Graham-Smith, 2006, Pulsar Astronomy, Cambridge Astrophysics. Cambridge University Press
- Lyne A. G., Jordan C. A., Graham-Smith F., Espinoza C. M., Stappers B. W., Weltevrede P., 2015, MNRAS, 446, 857
- Lyne A. G., Pritchard R. S., Graham-Smith F., 1993, MNRAS, 265, 1003
- Lyne A. G., Shemar S. L., Smith F. G., 2000, MNRAS, 315, 534
- Manchester R. N., Hobbs G., 2011, ApJ, 736, L31
- Manchester R. N. et al., 2013, PASA, 30, e017
- Manchester R. N., Hobbs G. B., Teoh A., Hobbs M., 2005, AJ, 129, 1993
- Manchester R. N. et al., 2001, MNRAS, 328, 17
- Manchester R. N. et al., 1996, MNRAS, 279, 1235
- Manchester R. N. et al., 1978, MNRAS, 184, 159
- Mannarelli M., Rajagopal K., Sharma R., 2007, Phys. Rev. D, 76, 074026
- Marron J., Chung S., 2001, Computational Statistics, 16, 195
- Marron J., Nolan D., 1988, Stat. Probab. Lett., 7, 195
- Mastrano A., Melatos A., 2005, MNRAS, 361, 927
- Mastrano A., Suvorov A. G., Melatos A., 2015, MNRAS, 453, 522
- McKee J. W. et al., 2016, MNRAS, 461, 2809

- Melatos A., 1997, MNRAS, 288, 1049
- Melatos A., 2012, ApJ, 761, 32
- Melatos A., Douglass J. A., Simula T. P., 2015, ApJ, 807, 132
- Melatos A., Dunn L. M., Suvorova S., Moran W., Evans R. J., 2020, ApJ, 896, 78
- Melatos A., Howitt G., Fulgenzi W., 2018, ApJ, 863, 196
- Melatos A., Link B., 2014, MNRAS, 437, 21
- Melatos A., Peralta C., 2007, ApJ, 662, L99
- Melatos A., Peralta C., Wyithe J. S. B., 2008, ApJ, 672, 1103
- Melrose D. B., Rafat M. Z., 2017, in Journal of Physics Conference Series, Vol. 932, Journal of Physics Conference Series, p. 012011
- Mendell G., 1991, ApJ, 380, 515
- Mereghetti S., Pons J. A., Melatos A., 2015, Space Sci. Rev., 191, 315
- Middleditch J., Marshall F. E., Wang Q. D., Gotthelf E. V., Zhang W., 2006, ApJ, 652, 1531
- Migdal A. B., 1959, Nuclear Physics, 13, 655
- Monaghan J. J., 2005, Reports on Progress in Physics, 68, 1703
- Monard L. A. G., 2010, Central Bureau Electronic Telegrams, 2289, 1
- Mongiovì M. S., Jou D., Sciacca M., 2018, Phys. Rep., 726, 1
- Mongiovì M. S., Russo F. G., Sciacca M., 2017, MNRAS, 469, 2141
- Montoli A., Antonelli M., Pizzochero P. M., 2020, MNRAS, 492, 4837
- Morley P. D., García-Pelayo R., 1993, Europhysics Letters, 23, 185
- Morton D. C., 1964, ApJ, 140, 460
- Nakabayashi K., Zheng Z., Tsuchida Y., 2002, Physics of Fluids (1994-present), 14, 3973
- Namkham N., Jaroenjittichai P., Johnston S., 2019, MNRAS, 487, 5854
- Nelson J., Hills R., Cudaback D., Wampler J., 1970, ApJ, 161, L235
- Nemirovskii S. K., 2013, Phys. Rep., 524, 85

- Newton W. G., Berger S., Haskell B., 2015, MNRAS, 454, 4400
- Onsager L., 1949, *Il Nuovo Cimento*, 6, 279
- Onuchukwu C. C., Chukwude A. E., 2016, *Ap&SS*, 361, 300
- Osborne D. V., 1950, *Proceedings of the Physical Society A*, 63, 909
- Ostriker J. P., Gunn J. E., 1969, *ApJ*, 157, 1395
- Page D., Geppert U., Weber F., 2006, *Nuclear Physics A*, 777, 497
- Page D., Prakash M., Lattimer J. M., Steiner A. W., 2011, *Physical Review Letters*, 106, 081101
- Palfreyman J., Dickey J. M., Hotan A., Ellingsen S., van Straten W., 2018, *Nature*, 556, 219
- Parthasarathy A. et al., 2021, MNRAS, 502, 407
- Pavlov G. G., Stringfellow G. S., Cordova F. A., 1996, *ApJ*, 467, 370
- Peralta C., 2006, PhD thesis, University of Melbourne
- Peralta C., Melatos A., 2009, *ApJ*, 701, L75
- Peralta C., Melatos A., Giacobello M., Ooi A., 2005, *ApJ*, 635, 1224
- Peralta C., Melatos A., Giacobello M., Ooi A., 2006, *ApJ*, 651, 1079
- Peralta C., Melatos A., Giacobello M., Ooi A., 2008, *Journal of Fluid Mechanics*, 609, 221
- Perera B. B. P. et al., 2019, MNRAS, 490, 4666
- Pilkington J. D. H., Hewish A., Bell S. J., Cole T. W., 1968, *Nature*, 218, 126
- Pines D., Alpar M. A., 1985, *Nature*, 316, 27
- Pitaevskii L. P., 1961, *Sov. Phys. JETP*, 13, 451
- Pizzochero P. M., Antonelli M., Haskell B., Seveso S., 2017, *Nature Astronomy*, 1, 0134
- Posselt B., Pavlov G. G., 2018, *ApJ*, 864, 135
- Posselt B., Pavlov G. G., Suleimanov V., Kargaltsev O., 2013, *ApJ*, 779, 186
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical recipes in FORTRAN. The art of scientific computing*

Price S., Link B., Shore S. N., Nice D. J., 2012, MNRAS, 426, 2507

Prix R., 2004, Physical Review D, 69, 043001

R Core Team, 2014, R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria

Radhakrishnan V., Cooke D. J., Komesaroff M. M., Morris D., 1969, Nature, 221, 443

Radhakrishnan V., Manchester R. N., 1969, Nature, 222, 228

Ray P. S. et al., 2019, ApJ, 879, 130

Reichley P. E., Downs G. S., 1969, Nature, 222, 229

Reisenegger A., 1993, Journal of Low Temperature Physics, 92, 77

Ricca R. L., 1996, Fluid Dynamics Research, 18, 245

Rosado P. A., Sesana A., Gair J., 2015, MNRAS, 451, 2417

Rosendahl J., Vekić M., Kelley J., 1993, Phys. Rev. E, 47, 1401

Rots A. H., Jahoda K., Lyne A. G., 2004, ApJ, 605, L129

Ruderman M., 1969, Nature, 223, 597

Ruderman M., 1976, ApJ, 203, 213

Ruderman M., Zhu T., Chen K., 1998, ApJ, 492, 267

Saffman P., 1995, Vortex Dynamics, Cambridge Monographs on Mechanics. Cambridge University Press

Sanidas S. et al., 2019, A&A, 626, A104

Sauls J. A., Stein D. L., Serene J. W., 1982, Phys. Rev. D, 25, 967

Schwarz K. W., 1985, Phys. Rev. B, 31, 5782

Schwarz K. W., 1988, Phys. Rev. B, 38, 2398

Scott D. W., 1992, Multivariate Density Estimation

Sedrakian A. D., 1995, MNRAS, 277, 225

Seveso S., Pizzochero P. M., Grill F., Haskell B., 2016, MNRAS, 455, 3952

Seveso S., Pizzochero P. M., Haskell B., 2012, MNRAS, 427, 1089

- Shabanova T. V., 2005, MNRAS, 356, 1435
- Shabanova T. V., 2010, ApJ, 721, 251
- Shannon R. M., Cordes J. M., 2010, ApJ, 725, 1607
- Shannon R. M., Lentati L. T., Kerr M., Johnston S., Hobbs G., Manchester R. N., 2016, MNRAS, 459, 3104
- Shapiro S., Teukolsky S., 2008, Black Holes, White Dwarfs and Neutron Stars: The Physics of Compact Objects. Wiley
- Shaw B. et al., 2018, MNRAS
- Sheather S. J., Jones M. C., 1991, J. R. Stat. Soc. Ser. B, 53, 683
- Shemar S. L., Lyne A. G., 1996, MNRAS, 282, 677
- Shternin P. S., Yakovlev D. G., Heinke C. O., Ho W. C. G., Patnaude D. J., 2011, MNRAS, 412, L108
- Sidery T., Passamonti A., Andersson N., 2010, MNRAS, 405, 1061
- Silverman B. W., 1986, Density estimation for statistics and data analysis. London: Chapman and Hall
- Simula T., Davis M. J., Helmersen K., 2014, Phys. Rev. Lett., 113, 165302
- Simula T. P., Mizushima T., Machida K., 2008, Phys. Rev. Lett., 101, 020402
- Sornette A., Sornette D., 1989, EPL (Europhysics Letters), 9, 197
- Sourie A., Chamel N., 2020, MNRAS, 493, L98
- Springel V., 2005, MNRAS, 364, 1105
- Stagg G., Parker N., 2017, Physical Review Letters, 118, 135301
- Strong S. A., Carr L. D., 2012, Journal of Mathematical Physics, 53, 033102
- Thomson J. J., 1883, A Treatise on the Motion of Vortex Rings. Macmillan
- Tilley D., Tilley J., 1990, Superfluidity and Superconductivity, Graduate Student Series in Physics. Taylor & Francis
- Trümper J., Becker W., 1998, Advances in Space Research, 21, 203
- Tsakadze J. S., Tsakadze S. J., 1980, Journal of Low Temperature Physics, 39, 649
- Tsubota M., Kobayashi M., Takeuchi H., 2013, Phys. Rep., 522, 191

- Turcotte D. L., 1999, Reports on Progress in Physics, 62, 1377
- van Eysden C. A., 2015, Journal of Fluid Mechanics, 783, 251
- van Eysden C. A., Melatos A., 2010, MNRAS, 409, 1253
- van Eysden C. A., Melatos A., 2011, Journal of Low Temperature Physics, 165, 1
- van Eysden C. A., Melatos A., 2013, Journal of Fluid Mechanics, 729, 180
- van Eysden C. A., Melatos A., 2014, Journal of Fluid Mechanics, 744, 89
- Vinen W. F., 1957, Proceedings of the Royal Society of London Series A, 240, 114
- Vinen W. F., 1961, Proceedings of the Royal Society of London Series A, 260, 218
- Vinen W. F., Niemela J. J., 2002, Journal of Low Temperature Physics, 128, 167
- Wand M., 2014, KernSmooth: Functions for kernel smoothing for Wand & Jones (1995). R package version 2.23-12
- Wand M. P., Jones M. C., 1995, Kernel smoothing. London: Chapman and Hall
- Wang N., Manchester R. N., Pace R. T., Bailes M., Kaspi V. M., Stappers B. W., Lyne A. G., 2000, MNRAS, 317, 843
- Warner B., Nather R. E., Macfarlane M., 1969, Nature, 222, 233
- Warszawski L., Melatos A., 2011, MNRAS, 415, 1611
- Warszawski L., Melatos A., 2012, MNRAS, 423, 2058
- Warszawski L., Melatos A., 2013, MNRAS, 428, 1911
- Warszawski L., Melatos A., Berloff N. G., 2012, Phys. Rev. B, 85, 104503
- Watkins N. W., Pruessner G., Chapman S. C., Crosby N. B., Jensen H. J., 2015, Space Sci. Rev.
- Weber F., Negreiros R., Rosenfield P., Stejner M., 2007, Progress in Particle and Nuclear Physics, 59, 94
- Wheatland M. S., 2000, ApJ, 536, L109
- Wijngaarden M. J. P., Ho W. C. G., Chang P., Heinke C. O., Page D., Beznogov M., Patnaude D. J., 2019, MNRAS, 484, 974
- Wong T., Backer D. C., Lyne A. G., 2001, ApJ, 548, 447
- Yarmchuk E. J., Gordon M. J. V., Packard R. E., 1979, Phys. Rev. Lett., 43, 214

- Yarmchuk E. J., Packard R. E., 1982, *Journal of Low Temperature Physics*, 46, 479
- Yavorskaya I., Belyaev Y. N., Monakhov A., Astafeva N., Scherbakov S., Vvden-skaya N., 1980, NASA STI/Recon Technical Report N, 81, 25327
- Younes G., Ray P. S., Baring M. G., Kouveliotou C., Fletcher C., Wadiasingh Z., Harding A. K., Goldstein A., 2020, *ApJ*, 896, L42
- Yu M., Liu Q. J., 2017, *MNRAS*, 468, 3031
- Yu M. et al., 2013, *MNRAS*, 429, 688
- Yui S., Kobayashi H., Tsubota M., Guo W., 2020, *Phys. Rev. Lett.*, 124, 155301
- Yui S., Tsubota M., Kobayashi H., 2018, *Phys. Rev. Lett.*, 120, 155301
- Zampieri L. et al., 2014, *MNRAS*, 439, 2813
- Zhang S.-N., Xie Y., 2012, *ApJ*, 761, 102