

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/182650>

Copyright and reuse:

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

The Skeleton of Plasma Turbulence

Applications of Nonlinear Dynamical Systems Theory to a
Subcritical Plasma Turbulence Model

Oliver Joel Smith

A thesis submitted in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy in Physics

Department of Physics
University of Warwick January, 2023

Contents

List of Figures	iv
List of Tables	vi
Acknowledgements	vii
Declaration	viii
Abstract	ix
List of Symbols and Abbreviations	x
1 Introduction	1
1.1 Nuclear Fusion and Plasma Physics	2
1.2 Fusion Device Design	6
1.2.1 Coordinate Systems	9
1.3 Plasma Frameworks	10
1.4 Instabilities in Plasma	15
1.4.1 Ballooning and Interchange Instabilities	15
1.4.2 Drift Waves	18
1.5 Anomalous Transport and Drift Wave Turbulence	19
1.5.1 Coherent Structure Dominated Turbulence	20
1.5.2 Shear Flows and Turbulence	21
1.6 Review of Simple Models Containing Propagating Structures	23
1.7 Review of Literature Concerning Plasma Turbulence and Avalanche-like Structures	28
1.8 Relative Periodic Orbits	33
1.8.1 Dynamical Systems Background	33
1.8.2 Chaotic Dynamical Systems and Turbulence	38
1.8.3 Significance of Relative Periodic Orbits in Studying Turbulence	38

1.8.4	Previous Use of RPO Based Methods in Neutral Fluids	40
1.8.5	Subcriticality and the Edge of Chaos	41
1.8.6	Transient Growth	42
1.8.7	Bifurcations in Dynamical Systems	43
1.9	Summary Of Key Ideas and Thesis Outline	47
2	The PI Model Initial Value Solver	50
2.1	Plasma Interchange Model	50
2.2	Symmetries of the PI Model	55
2.3	Spectral PI Model Code	57
2.3.1	Basis Functions	58
2.3.2	Hybrid Timestepping	60
2.3.3	Nonlinear Terms	61
2.3.4	Timestepping Scheme	64
2.3.5	Remapping Wavenumbers	64
2.3.6	Convergence of PI Model Code	65
3	Turbulence and the Edge of Chaos in the PI Model	66
3.1	Nature of Turbulence in the PI Model	66
3.1.1	Saturation and Turbulence in Large ($L = 50$) Domain	67
3.1.2	Saturation and Turbulence in Small ($L = 5$) Domain	74
3.2	Locating the Edge of Chaos with the Bisection Method .	74
3.2.1	Edge State in Large ($L = 50$) Domain	76
3.2.2	Edge States in a Small Domain	78
4	Locating RPOs in The PI Model System	83
4.1	Introduction to the Newton-Krylov-Hookstep Method	83
4.1.1	Krylov Method	85
4.1.2	GMRES	87
4.1.3	Hookstep	87
4.1.4	Application to PI Model	88
4.1.5	Symmetries of RPOs	89
4.1.6	Tracking RPOs Through Parameter Space and Determining RPO Stability	89
4.2	RPOs In Large Boxes	90
4.2.1	EOC RPOs	91

4.2.2	Saddle Node Bifurcation at $S=1.6234$, and Stable RPO	91
4.2.3	First Hopf Bifurcation Pair	95
4.2.4	Secondary Hopf Bifurcation Pair	98
4.2.5	Comparisons to Turbulent Simulations With Chang- ing Shear	103
4.3	Early Results in Small $L = 5.0$ Domain	104
5	Conclusions	108
	Bibliography	109
A	Newton-GMRES-Hookstep Algorithm	117

List of Figures

1.1	A tokamak design.	8
1.2	Concentric magnetic surfaces in a tokamak geometry. . .	9
1.3	Coordinate systems for tokamak geometries.	10
1.4	The Rayleigh-Taylor instability.	16
1.5	The drift wave mechanism.	19
1.6	Stretching and splitting of an eddy by a shear flow. . . .	22
1.7	Pressure profiles against radius in L and H mode.	24
1.8	A representation of a dynamical system mapping.	33
1.9	Periodic orbits in the 2D and 3D Van der Pol systems. . .	35
1.10	2D representation of a subcritical dynamical system. . .	40
1.11	Subcriticality in the time reversed Van der Pol system . .	42
1.12	Demonstration of transient growth.	43
1.13	Logistic map and pitchfork bifurcation	45
1.14	Canonical Hopf bifurcation	46
1.15	Saddle node bifurcation	47
2.1	Demonstration of aliasing	63
2.2	Convergence of PI model simulations with N_t and N_x . .	65
3.1	Flux against position and time for typical turbulence sim- ulations	68
3.2	Decay rate, λ , of turbulence in $L = 50$ simulations at various values of S	69
3.3	Fourier transform and autocorrelation of turbulent flux. .	70
3.4	Collision between travelling waves and propagating pair. .	71
3.5	Spectral density of the angular frequency of the domain averaged flux in turbulent simulations against S , for high values of S	72
3.6	Spectral density of the angular frequency of the domain averaged flux in turbulent simulations against S , for low values of S	73

3.7	Small domain turbulence simulations.	75
3.8	Bisection method applied to Van der Pol oscillator.	77
3.9	Example of the bisection method.	78
3.10	Large domain edge state amplitudes against S	79
3.11	Amplitudes of closest-to-edge solutions in the small domain found using the bisection method against time for a range of S	81
3.12	A collection of near edge state solutions identified in the $L = 5.0$ box.	82
4.1	RPO amplitudes and stabilities in large domain.	92
4.2	RPO velocities and stabilities in large domain.	93
4.3	Profiles of travelling waves in the edge.	94
4.4	Profiles of travelling waves at saddle node bifurcation and a stable RPO.	95
4.5	Amplitude and stability of RPOs close to first Hopf bifurcation.	96
4.6	Eigenvalue properties of RPOs close to the first Hopf bifurcation.	97
4.7	Hopf branch RPOs	99
4.8	Main branch RPO allowed to go unstable and begins to oscillate in amplitude and leaves a zonal density behind it.	100
4.9	Eigenvalue properties of RPOs close to second Hopf bifurcation.	101
4.10	Flux against position and time for RPOs (i), (j), and (k)	102
4.11	Comparison of turbulent frequency power spectra against S and RPO positions in S	105
4.12	Small domain RPO amplitudes against S and stability.	107

List of Tables

1.1	Particle drifts experienced by charged particles in electromagnetic fields.	6
2.1	Linear and nonlinear parts of the time evolution equations for the Fourier components of the PI variables	60
4.1	Hopf eigenvalue angles associated with the three Hopf bifurcations on the primary Hopf branch, and associated valid frequencies.	100

Acknowledgements

Firstly, I want to thank my supervisor, Dr Ben McMillan for the generous support, technical guidance and patient explanations he provided throughout the completion of this PhD, as well as my collaborator Dr Chris Pringle for many useful conversations and providing useful learning material and critical code. Secondly, I thank the Centre for Postgraduate Training in Plasma Physics and High Energy Density Science for providing a thorough education in plasma physics through their Autumn and Winter schools. Thirdly, I want to thank the TDoTP collaboration for their all-round support, providing space to share my work and giving critical feedback and ideas to improve it. Fourthly, I thank all the friends and family, too numerous to list here, who have been so motivational and encouraging throughout.

Finally, and most of all, I want to thank my wife, Emily, who has been patiently and graciously supporting me through every stage of the PhD process, and had the misfortune of marrying a PhD student three months into his study. She has taught me more about navigating chaos than all the dynamical systems textbooks combined and I could not have got through it without her. This thesis is dedicated to her.

Declaration

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The work presented (including data generated and data analysis) was carried out by the author except in the cases outlined below:

List of data provided and/or analysis carried out by collaborators.

- GMRES-Hookstep code provided by C. T. Pringle¹ and written by A. P. Willis²

Parts of this thesis have been published by the author:

- Figures related to relative periodic orbits in a small domain were published on January 2nd 2023 ³

List of publications including submitted papers.

- Coauthor on ‘An overview of dynamical methods for studying transitions between states in sheared plasma flows’³

¹Centre for Fluid and Complex Systems, University of Coventry, Coventry, CV1 5FB

²School of Mathematics and Statistics, The University of Sheffield, Sheffield, S3 7RH, United Kingdom

³Teaca B, Pringle CCT, Smith OJ, McMillan BF. 2022 An overview of dynamical methods for studying transitions between states in sheared plasma flows. *Phil. Trans. R. Soc. A* **381**: 20210238

Abstract

The formation of localised travelling structures in plasma turbulence is investigated through the application of dynamical systems methods novel to plasma physics contexts. These methods involve studying a plasma model by identifying model solutions known as relative periodic orbits (RPOs). RPOs are a generalisation of periodic solutions which allow for symmetries, and collectively form a ‘skeleton’ in phase space representing the set of solutions the system can approach and move between. Locating and categorising these RPOs and their overall structure in parameter space gives a new framework through which plasma turbulence can be analysed, allowing complex fully nonlinear structures to be identified and isolated, and the structure of turbulence to be better understood. This RPO-based approach has been applied in neutral fluid systems and has aided the understanding of similar bursty behaviour in contexts such as pipe flow [Pringle *et al.*, *Phil. Trans. R. Soc. A* **367**, 457 (2009)] and Couette flow [Viswanath, D., *J. Fluid Mech.* **580**, 339 (2007)].

A pseudospectral code was developed for a simple plasma model known as the plasma interchange model. The ‘edge of chaos’ (EOC) - a manifold separating the basins of attraction of the equilibrium and turbulence - is located through a bisection method, and is found to contain attractor RPOs with similar characteristics to structures seen in turbulence. An implementation of the Newton-Krylov-Hookstep algorithm is used to extend the manifold of RPOs beyond the EOC, where travelling-wave-like RPOs are identified, some of which are stable, providing an intuitive explanation as to why such structures appear in plasma turbulence at high shear. Quasi-travelling waves with oscillatory structure are also identified, and found to generate spatially periodic zonal structures. It is demonstrated that RPO-based methods are a powerful tool for identifying coherent structures in plasmas and understanding their impact on turbulence.

List of Symbols and Abbreviations

Abbreviations

ASDEX	Axially Symmetric Divertor Experiment
EOC	Edge of chaos
GK	Gyrokinetics
ICF	Inertial Confinement Fusion
ITG	Ion temperature gradient (instability)
MCF	Magnetic Confinement Fusion
MCF	Magnetic Confinement Fusion
MHD	Magnetohydrodynamics
NKH	Newton-Krylov-Hookstep Algorithm
PI	Plasma Interchange (Model)
RPO	Relative Periodic Orbit
UKAEA	UK Atomic Energy Authority

Symbols

\bar{n}	Zonal density in the PI model
η	Plasma Resistivity
γ	Adiabatic Index
λ_D	Debye Length
\mathbf{B}	Magnetic field vector
\mathbf{E}	Electric field
\mathbf{F}	Force
\mathbf{g}	Gravitational vector, or effective gravity
\mathbf{J}	Plasma current
\mathbf{R}	Force on a species due to collisions with others
\mathbf{R}_c	Radius of Curvature of Magnetic Field
\mathbf{v}	Velocity of a particle
\mathbf{v}_\perp	Particle velocity component perpendicular to the magnetic field
∇	Vector gradient operator
ω	Frequency

ω_c	Cyclotron Frequency
ϕ	Electric Scalar Potential
ψ	Poloidal flux function
ρ	Charge density of plasma
τ	Period
τ_c	Confinement time
τ_{BC}	Period of the shear periodic boundary conditions
$\tilde{\phi}$	Wave potential in the PI model
\tilde{n}	Wave density in the PI model
B	Magnetic field strength
B_p	Poloidal magnetic field
B_t	Toroidal magnetic field
e	Electron, electron charge, or the mathematical constant e
f_s	Distribution function of particles of species s in phase space
m_s	Mass of particle of species s
n_s	Number density of a species s
p	Scalar pressure
P_α	Fusion heating rate
P_H	Plasma external heating rate
P_L	Energy loss rate
q	Safety factor
q_s	Charge of particles of species s
r_L	Larmor Radius
S	Flow shear rate
T_s	Temperature of species s
v_{\parallel}	Particle velocity parallel to the magnetic field
W	Total plasma energy
Z	Atomic Number of an Ion

Chapter 1

Introduction

Nuclear fusion has long been investigated as a potential effective source of sustainable and continuously available power. The leading approach being developed to achieve this is known as *magnetic confinement fusion* (MCF), in which a plasma is contained using the effects of a strong magnetic field, and heated enough that some particles will collide with a suitable amount of collision energy to fuse. The complication within this approach is that this magnetic confinement is not perfect; heat and particles can undergo transport across the magnetic field to the edge of the device where they are lost. This loss of energy from the plasma must be compensated for, either by fusion reactions themselves, or by external heating, and so minimising this transport is essential to a reactor design sustainable to operate. This transport is not fully avoidable - as long as collisions occur there will always be some kind of diffusion - however in fusion devices transport has been seen to be orders of magnitude larger than simple collisional effects alone would allow [1]. This excess ‘anomalous’ transport is primarily caused by the effects of turbulence driven by microinstabilities [2–4].

A critical development in the study of transport in fusion devices was the identification in the ASDEX reactor of a high confinement mode, known as ‘H-mode’, in which confinement is greatly enhanced [5] once a threshold heating power is reached. The mechanism for the suppression of transport was linked to the self-organisation of sheared flows at the edge of the plasma [6, 7] which acts as an energy sink for the turbulence, effectively suppressing it and allowing steep profile gradients to be generated. The interaction between turbulence and shear flows is thus of critical importance to containing plasma effectively.

One important feature of turbulence that is intricately linked to shear

flows is intermittent transport events, known as ‘avalanches’, which occur across a range of spatial scales and correspond to large scale transport events observed in experimental data [8, 9] and a range of plasma models [10–14]. These avalanches are radially localised structures that propagate through the plasma, predominantly outwards, carrying matter and energy with them.

This thesis examines the behaviour of these structures through a non-linear dynamical systems framework in which a generalisation of exactly periodic solutions, known as ‘relative periodic orbits’ (RPOs) populate the dynamical system space, and act as a sort of ‘skeleton’ about which turbulent solutions are structured. In this picture turbulent solutions will approach and move between these RPOs so that turbulence will approximate these RPOs intermittently. This analysis adapts techniques developed in neutral fluid theory [15, 16] and applies them to a simple 1D plasma model known as the plasma interchange model (PI) to demonstrate how avalanche-like structures can be better understood by using these methods.

This introduction will cover the relevant background to the thesis, including relevant fusion and plasma physics history and theory, dynamical systems theory, and a review of relevant literature on plasma turbulence and coherent structures.

1.1 Nuclear Fusion and Plasma Physics

The notion of harnessing nuclear fusion reactions to generate large amounts of energy can be traced back as far as 1920, when Arthur Eddington speculated that the fusion of small nuclei, and the associated mass deficit identified by F. W. Aston [17] provided the mechanism by which energy is generated in the Sun [18]. Even then, Eddington had the great foresight to envision the significance harnessing the power of nuclear fusion could have on human society, well encapsulated by this quote of his: ‘If, indeed, the sub-atomic energy in the stars is being freely used to maintain their great furnaces, it seems to bring a little nearer to fulfilment of our dream of controlling this latent power for the well-being of the human race - or for its suicide.’ Over the following century, Eddington’s intuitions of the nature of both the Sun and of humanity proved to be correct, with the development of plasma physics unlocking the Sun’s secrets and translating them into the realisation of fusion reactors and

thermonuclear weapons.

In today's world, climate change is widely understood to be the greatest threat to humanity as a whole, with 3.5 billion highly vulnerable to the effects of climate change [19]. This climate catastrophe is directly linked anthropogenic greenhouse gas emissions, of which approximately two-thirds is linked to the global energy system [19]. Rapidly reforming the energy sector, moving away from fossil fuels, will play an essential role in mitigating the damage of climate change and building a sustainable future for humanity. Nuclear fusion could form a crucial part of this transition as a powerful, efficient, continuously operable energy source with minimal hazardous waste products.

Nuclear fusion is a nuclear reaction in which two atomic nuclei collide with enough energy to overcome the electromagnetic repulsive force between them, getting close enough to each other that the short-ranged strong nuclear force can act to combine them into a single, larger nucleus, with smaller byproducts such as individual neutrons or protons. This process leaves some latent mass which is released as kinetic energy of the products. The most energetic fusion reaction, a deuterium-tritium fusion (D-T fusion), releases 17.6 MeV of energy in the kinetic energy of the products (a neutron and alpha particle). This gives fusion fuel an extremely high energy density; just one kilogram of fuel is enough to generate energy equivalent to 10 million kilograms of fossil fuel, and 1 tonne is enough to power a one gigawatt fusion reactor for a whole year [20]. The scale of fuel efficiency, scalability and sustainability offered by fusion has naturally maintained both academic and commercial interest since its first conception as an idea.

The earliest research focused on understanding the form of matter inside stars, so that we could hope to build machines within which its conditions could be mimicked. Our Sun, and indeed all stars, were found to consist of fully ionised gas - a sea of nuclei and dissociated electrons experiencing the influence of collective electric and magnetic fields. This ionised gas was dubbed 'plasma' by Irving Langmuir (anecdotally due to the similarity with the white and red blood cells in blood plasma) [21] and understanding its behaviour is essential to designing a reactor which can achieve sustainable fusion, while also having other applications in space physics, accelerator physics and engineering. The first successful fusion reaction observation came from research by Oliphant and Rutherford [22] who accelerated a stream of deuterons into stationary deuterons, generating tritium and ^3He .

The primary challenge of nuclear fusion is that two particles of like sign charge must get within close enough range of each other ($\sim 10^{-15}\text{m}$) that the strong nuclear force can act between them and fuse them together. The significant barrier to this is the Coulomb repulsion force between the two particles, and so to overcome this the ions must collide with sufficient energy to overcome this barrier. Were it not for quantum mechanics, the required collision energy would be of the order of 400 keV (for D-T fusion), however, it is possible that the colliding particles will tunnel through the Coulomb barrier, and this lowers the energy requirements. For deuterium-tritium (D-T) fusion, the cross section, σ , for this interaction peaks at a centre of mass energy of around 100 keV, and the reaction rate, $\langle \sigma v \rangle$ in a Maxwellian distribution peaks at a temperature of 70 keV, or around 800 million Kelvin [23]. In reality, these high temperatures are too difficult to achieve with current technology; modern fusion reactors are designed to operate around 10 keV, however the fusion rate can still be substantial due to the long tail in the Maxwellian distribution at high velocities.

There are several key properties of plasma which distinguish it from neutral fluids, and will be relevant to the discussion in this thesis. The most obvious is that being composed of free moving charged particles, a plasma is electrically conductive, and its bulk behaviour is governed by collective electromagnetic effects, allowing particles to interact at long ranges. The first collective behaviour is that a plasma is *quasineutral*. This means charged particles will move quickly to counter any electric potential so that no significant excess of charge will be allowed to build up, and we can typically assume that the number densities, denoted n_j for species j , satisfy the quasineutrality relation,

$$n_e \approx Z n_i, \tag{1.1}$$

for ions with atomic number Z . It's important to note that electric fields can still exist within the plasma, but the charge excess generating them can be very small, and yet still generate significant potentials.

The second collective behaviour of plasmas is the screening of excess charge by the bulk plasma. Over short distances, plasma temperature allows quasineutrality to be locally violated. The internal energy of the plasma competes with the restoring force of any charge imbalance, and

the particles of species j follow a Boltzmann distribution

$$n_j(x) = n_{j\infty} \exp(-q_j\phi/T_s), \quad (1.2)$$

where $n_{j\infty}$ is the density of species j far from any potential, q_j is the charge of a particle of species j , and T_j is the temperature of species j . This distribution acts to screen any potential at long distances. If a test ion is added to an otherwise neutral plasma, electrons will be attracted to it, and other ions repelled, so that the potential from the test ion is screened outside a short distance. Working through this fully by solving Poisson's equation and the Boltzmann distribution self-consistently gives the result that the potential of a single test charge is $\phi(r) = e^{-(r/\lambda_D)}\phi_v$, where ϕ_v is the potential of the test charge in a vacuum and

$$\lambda_D = \sqrt{\epsilon_0 T/n_e e^2} \quad (1.3)$$

is the **Debye length** [23]. This screening depends on there being a large number of particles in a sphere of radius Debye length, i.e $n\lambda_D^3 \gg 1$. This parameter is known as the **plasma parameter** and a plasma that satisfies this limit is called an **ideal plasma**.

As well as the collective behaviours described above, individual particles will also be subject to magnetic fields. A charged particle in a uniform magnetic field will experience a force $\mathbf{F}_B = \mathbf{v} \times \mathbf{B}$ which is perpendicular to the velocity, causing the particle to orbit the magnetic field line, with a **cyclotron frequency** $\omega_c = qB/m$ and **cyclotron or Larmor radius** $\rho_L = mv_\perp/|q|B$, where v_\perp is the velocity perpendicular to the field. This effectively 'traps' the particle on the magnetic field line. It is this mechanism that is used to trap particles with a magnetic field in fusion devices. When other forces are present, or there is curvature or variation in the magnetic field, the **guiding center** of the orbit will move relative to the field line due to values of either \mathbf{v} or \mathbf{B} changing within an orbit leading to a net motion. These drift effects are called **guiding center drifts**, and give rise to collective plasma velocities and currents. A summary of the key drifts is given in Table 1.1.

These drifts act on individual particles, but collectively drifts give rise to bulk plasma flows and currents, and hence are important to linking individual particle behaviour and fluid behaviour. The $\mathbf{E} \times \mathbf{B}$ drift is charge independent, meaning both ions and electrons drift in the same direction, at the same velocity, leading to bulk plasma flows. The other

Name	Velocity	Notes
$\mathbf{E} \times \mathbf{B}$ Drift	$\frac{\mathbf{E} \times \mathbf{B}}{B^2}$	
Gravitational Drift	$\frac{m\mathbf{g} \times \mathbf{B}}{qB^2}$	\mathbf{g} is the gravitational vector
Curvature Drift	$\frac{mv_{\parallel}^2}{qB} \frac{\mathbf{R}_c \times \mathbf{B}}{R_c^2 B}$	\mathbf{R}_c is the radius of curvature
∇B Drift	$\frac{mv_{\perp}^2}{2qB} \frac{\mathbf{B} \times \nabla B}{B^2}$	\mathbf{R}_c is the radius of curvature
Polarisation Drift	$\frac{m}{qB^2} \frac{d\mathbf{E}_{\perp}}{dt}$	\mathbf{E}_{\perp} is the electric field perpendicular to \mathbf{B}

Table 1.1: Particle drifts experienced by charged particles in electromagnetic fields.

drifts are charge dependent, and so generate bulk currents. These bulk currents give rise to the bulk $\mathbf{j} \times \mathbf{B}$ force that acts against the pressure force in order to contain the plasma in fusion devices. The polarisation drift plays a role whenever charge separation occurs due to other drifts, and acts to give the plasma an effective inertia; any perturbation which changes the electric field will generate a polarisation drift which acts to reduce it, until the perturbation and polarisation drift balance each other exactly.

1.2 Fusion Device Design

The objective of a thermonuclear fusion device is to create a plasma, hot and dense enough to support a substantial fusion reaction rate. In the absence of any confinement mechanisms, a plasma will rapidly expand through its pressure and the temperature will drop quickly, and so confinement plays a key role in creating and sustaining thermonuclear plasmas. There are many approaches to device design, which fall into two main categories: inertial confinement fusion (ICF); and magnetic confinement fusion (MCF). The ICF process heats and compresses spherical thermonuclear fuel pellets so rapidly that fusion can occur at the center on timescales shorter than the diffusion timescales. This approach is not the focus of this thesis and will not be discussed. The MCF approach instead contains a plasma by applying strong magnetic fields so that particles become effectively bound to field lines by the Lorentz force, greatly limiting transport in the direction perpendicular to the field.

Using the MCF approach, maintaining a plasma requires reaching a steady state in which the energy being added to the plasma - a combination of heating from external sources and energy generated by fusion reactions - is balanced against the energy escaping the plasma. Taking a

naive approach, we assume the rate of energy loss can be approximated by the expression $P_L = W/\tau_c$ where W is the total energy in the plasma and τ_c is the ‘confinement time’. In order to maintain the energy in the plasma, it is necessary that the total heating power combining external sources and the heating from energetic α -particles generated in fusion reactions exceeds, or at least matches, the losses from escaping particles and minimising P_L reduces the power required to maintain the plasma. **Ignition** is achieved when P_α , the α -particle heating *alone* is enough to balance the confinement losses, and the plasma can self-sustain. This requires a temperature at which the cross section for fusion reactions is suitably large, so $T \sim 10 - 50\text{keV}$. Performing the full calculation, one finds that for these temperatures the condition on ignition, known as the *triple product* is $nT\tau_c > 3 \times 10^{21} m^3 \text{keV s}$ [23].

Throughout the 20th Century, research focused on the understanding of plasma, and how the confinement time, τ_c , could be improved enough to achieve ignition. In the late 1930s and 1940s reactor concepts and designs were rapidly suggested, iterated and improved, and the first patent on a fusion reactor was registered in 1946 by the UK Atomic Energy Authority (UKAEA), designed by G. P. Thomson and M. Blackman based on the Z-pinch concept [24]. The Z-pinch existed in both cylindrical and toroidal geometries. In the cylindrical case, a plasma was generated by applying a high voltage across a gas which would then ionise and sustain a current along its axis (in the ‘z-direction’, giving the concept its name). In the toroidal case this current could be induced via an external electromagnet in a transformer-like configuration. In both configurations the current would induce a magnetic field encircling the primary axis of the pinch. The Lorentz forces associated with this current and field act collectively on the plasma via the $\mathbf{J} \times \mathbf{B}$ force to compress the plasma so that the plasma is forced towards the central axis of the device, creating a hot, dense core where fusion could occur. Physicists exploring these device designs quickly identified that there are inherent instabilities in the plasma, in particular the kink and flute (or interchange) instabilities. These instabilities caused rapid perturbation growth in these kinds of setup, greatly reducing the confinement time, and so an alternative design was needed to mitigate them as much as possible.

In the early 1950s, Soviet researchers worked on an alternative machine for plasma confinement, the **tokamak** (toroidal’naya kamera s magnitnymi katushkami - toroidal chamber with magnetic coils). A tokamak takes the design of the toroidal pinch and introduces a strong magnetic




Figure 1.1: A tokamak design. Poloidal coils induce a strong magnetic field in the toroidal direction, while a current through the plasma itself induces a poloidal field. These combine to give field lines which spiral around the torus. Image by S. Li, H. Jiang, Z. Ren, C. Xu [25]. License: CC BY 4.0¹ via Wikimedia Commons.

field in the toroidal direction by passing current through coils in the poloidal direction around the device. This strong toroidal field combines with the weaker poloidal field generated by toroidal current, and the net magnetic field winds around the torus, see Figure 1.1. This has the effect of stabilising many of the instabilities present in the previous pinch designs. A measure of the winding of the magnetic field, known as the **safety factor** for its association with the stability of the device, is defined as [23]

$$q = \frac{aB_t}{RB_p} \quad (1.4)$$

for a major radius R and minor radius a , where B_p and B_t are the magnetic field strength in the poloidal and toroidal directions respectively. The safety factor effectively counts the number of times a field line wraps the long way around the torus for each time it wraps around the short way. Crucially, kink modes are stabilised for values of $q > 1$ (i.e field lines wrap faster in the long direction than short direction)[23].

Like in the Z-pinch, the currents in the device interact with the magnetic fields to generate a $\mathbf{J} \times \mathbf{B}$ force, which must be equal and opposite to the pressure force in equilibrium. This gives the pressure balance force:

$$\mathbf{J} \times \mathbf{B} = \nabla p. \quad (1.5)$$

¹<https://creativecommons.org/licenses/by/4.0>




Figure 1.2: Nested magnetic surfaces in a tokamak geometry. Each surface has a constant pressure and contains a set of magnetic field (red) and current (blue) lines which lie within it.

A consequence of this is that $\mathbf{B} \cdot \nabla p = 0$ and $\mathbf{J} \cdot \nabla p = 0$ and so the pressure is constant along field lines and field lines and currents lie in concentric **magnetic surfaces**, as seen in Figure 1.2. Many properties of the plasma can be shown to be constant on a magnetic surface in equilibrium, including the safety factor q , and the poloidal flux per toroidal radian [23], ψ , and the latter is often used as a label for a magnetic surface, giving it the alternate name, ‘flux surface’. Magnetic surfaces do not have circular cross-sections due to the field closer to the central magnetic axis being stronger than the field further away. This makes representing tokamak geometries nontrivial, and different coordinate systems are used in different contexts.

1.2.1 Coordinate Systems

Different coordinate systems are used to represent tokamak geometries depending on the context. When considering large scale physics, the cylindrical coordinates (R, ϕ, Z) , with axes shown in Figure 1.2 are used. When examining physics across a toroidal cross section, either toroidal coordinates (r, θ, ϕ) , defined relative to the magnetic axis as in Figure 1.3a, or magnetic surface coordinates which use a magnetic surface label such as ψ in place of the radial coordinate r as in Figure 1.3b. Finally, when working with local physics, a flux-tube geometry is often used. A

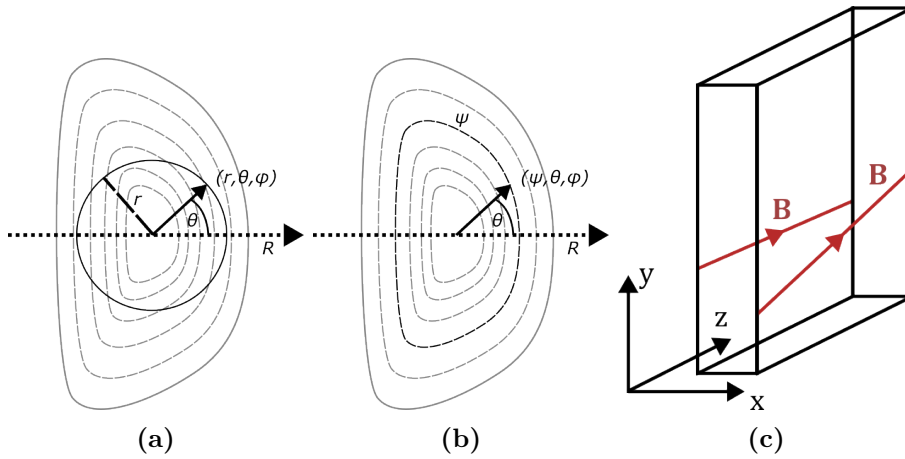


Figure 1.3: Coordinate systems used to describe magnetic geometries. ϕ is a toroidal angle coordinate defined as in Figure 1.2. (a): Toroidal coordinates. (b): Magnetic surface coordinates. (c): Flux tube coordinates.

flux-tube is a tube whose walls are parallel to the field everywhere, so the tube spirals with the field lines around the tokamak. The coordinates in the flux tube are (x, y, z) , where the z -axis is aligned with the magnetic field, the x -axis is directed out of the magnetic surface, and the y -axis is directed perpendicular to both. A flux tube geometry is shown in Figure 1.3c.

1.3 Plasma Frameworks

Plasma, like a neutral fluid, is made up of a large number of particles, and hence a full simulation tracking the positions, momenta and interactions of these particles is completely infeasible. To make plasma physics tractable, different kinds of approximations and limits can be applied to provide significantly simpler models. An overview of the most important model frameworks is presented here.

Kinetic Model

The first step that can be taken to reduce the problem is to use a kinetic model, where a distribution function $f_s(\mathbf{x}, \mathbf{v}, t)$ of the particles in position and velocity space (phase space) is defined so that the number of particles of a species s with position in the range $[x, x + \delta x]$ and velocity in the range $[v, v + \delta v]$ is given by $f_s(x, v)\delta x\delta v$. A kinetic model is expressed in terms of this distribution function, rather than tracking the individual particles. The equation governing the distribution function is derived

directly from the conservation of particles and Hamilton's equations, and is known as the **Fokker-Planck equation** [23]

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \frac{\partial f_s}{\partial \mathbf{x}} + (e_s/m_s)(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f}{\partial \mathbf{v}} = \left(\frac{\partial f}{\partial t} \right)_c, \quad (1.6)$$

where f_s is the distribution function for particles of species s , \mathbf{E} and \mathbf{B} are the coarse-grained (over lengths of order $\sim \lambda_D$) electric and magnetic fields respectively. The RHS term represents the rate of change in the distribution function due to collisions between particles. This term contains all information about collisions between particles, and is dealt with via one of several different mathematical approaches, if not neglected completely in a collisionless limit. When combined with Maxwell's equations and the conversions from the distribution function to the charge and current densities,

$$\rho = \sum_s q_s \int f_s(\mathbf{x}, \mathbf{v}) d^3v \quad (1.7)$$

$$\mathbf{J} = \sum_s q_s \int \mathbf{v} f_s(\mathbf{x}, \mathbf{v}) d^3v, \quad (1.8)$$

one ends up with a system of seven integro-differential equations which cannot be solved analytically and would have to be solved computationally. In practice the range of length and time scales at play would require vast computational resources and so separation of these scales is implemented and different analytical techniques applied to reduce the kinetic model to one more tractable.

The kinetic model separates the collisional scales from the scale of collective effects, and so is valid over scales larger than the Debye length λ_D . The kinetic model retains effects related to the particle drifts and finite Larmor radius as well as the effects of trapped particles and fast particles which have not thermalised (for example from neutral beam injection).

Gyrokinetic Models

One common method for reducing the complexity of the kinetic equations is to separate the scales of particle gyrobehaviour and the background profile scales and average over the gyro-orbits of the particles to create a popular model known as gyrokinetics (GK) [26]. If a plasma is strongly

magnetised ($\rho_i \ll L$, where L is the length scale of the plasma and ρ_i is the ion Larmor radius), the perpendicular length scales are on the order of the gyroradius ($k_\perp \rho_i \sim 1$), and the frequencies of interest are much lower than the ion gyrofrequency ($\omega \ll \omega_i$) the particle motion can be decomposed into the fast circular gyromotion, and the slow motions parallel to the magnetic field and drifts. The explicit dependence on the gyrophase of each particle (i.e its exact position in its orbit around the field line) is assumed to be insignificant in the overall evolution of the system, and can be dropped (via a complex perturbative method) to leave $f = f(\mathbf{R}, v_\perp, v_\parallel)$ where \mathbf{R} is the gyrocenter position, and v_\perp and v_\parallel are the perpendicular and parallel scalar velocities respectively. This equates to the gyrokinetic ordering:

$$\frac{\delta f}{f} \sim \frac{\omega}{\Omega} \sim \frac{k_\parallel}{k_\perp} \sim \frac{|\delta \mathbf{B}|}{|\mathbf{B}|} \sim \frac{|\delta \mathbf{E}|}{|\mathbf{E}|} \sim \frac{\rho_i}{a} = \epsilon \quad (1.9)$$

where $\rho_i = mv_\perp/q_i B$ is the ion Larmor radius. The distribution function is now a function of five variables rather than the six of full kinetic theory. The evolution of f is combined with Maxwell's equations and the quasineutrality relation, Equation 1.1 to provide the full gyrokinetic model.

Gyrokinetic models have been used to simulate a range of turbulence modes in a range of limits (for a full review see [27]), including but not limited to: ion-temperature-gradient mode turbulence [28–30]; trapped electron mode turbulence [31–33]; and microtearing modes [34]. In all, gyrokinetic theory has played a huge role in the understanding of core turbulence and continues to provide an essential tool for investigating core plasma through simulations. Several gyrokinetic codes exist, such as GENE [35], and GS2 [36].

Fluid Model

The next level of simplification of the kinetic model is to take moments of the kinetic equation in velocity space, giving fluid variables as a function of position, and not velocity. The fluid variables are the density n , fluid velocity \mathbf{u} , and pressure tensor \mathbf{P} and are calculated separately for each species. These variables are defined in terms of the distribution function

moments [23]:

$$n_s(\mathbf{x}, t) = \int f_s(\mathbf{x}, \mathbf{v}', t) d\mathbf{v}' \quad (1.10)$$

$$\mathbf{u}_s(\mathbf{x}, t) = \frac{1}{n} \int f_s(\mathbf{x}, \mathbf{v}', t) \mathbf{v}' d\mathbf{v}' \quad (1.11)$$

$$\mathbf{P}_s(\mathbf{x}, t) = m \int f_s(\mathbf{x}, \mathbf{v}', t) (\mathbf{v}' - \mathbf{u})(\mathbf{v}' - \mathbf{u}) d\mathbf{v}' \quad (1.12)$$

The pressure tensor \mathbf{P} contains the force per unit area acting on an element of fluid due to the thermal distribution of particles, and for an isotropic distribution function, is diagonal with entries

$$p = (m/3) \int f(\mathbf{x}, \mathbf{v}', t) (\mathbf{v}' - \mathbf{u})^2 d\mathbf{v}'. \quad (1.13)$$

Taking equivalent moments of the kinetic equation gives the continuity equation

$$\frac{\partial n_s}{\partial t} + \nabla \cdot (n_s \mathbf{u}_s) = 0, \quad (1.14)$$

which is stating that the number density of each particle is conserved. The moment of the kinetic equation with \mathbf{v}' gives the momentum equation:

$$mn \left(\frac{\partial \mathbf{u}_s}{\partial t} + \mathbf{u}_s \cdot \nabla \mathbf{u}_s \right) = -\nabla \cdot \mathbf{P}_s + nZe(\mathbf{E} + \mathbf{u}_s \times \mathbf{B}) + \mathbf{R} \quad (1.15)$$

where \mathbf{R} here is the force on the species due to collisions with other species. In a collisionless plasma this is zero. Next, an equation for \mathbf{P} is required. This can be obtained from the second moment of the kinetic equation, however this would in turn contain a term that requires the third moment of the kinetic equation, and so ad infinitum. To close the set of equations, a simplifying approximation is usually made in the equation for \mathbf{P} , such as assuming adiabatic motion, such that $pn^{-\gamma}$ is constant, where $\gamma = C_p/C_V$ is the ratio of specific heats at constant pressure and volume. Combining with Maxwell's equations gives a full set of equations for three fluid variables n_s , p_s and \mathbf{u}_s . The fluid model continues to demonstrate the particle drifts, through the fluid velocity and current. An additional fluid drift also arises if a pressure gradient is present, called the diamagnetic drift. This differs from the previously discussed particle drifts, as it is not associated with the drift of particle guiding centres. Instead it arises due to an imbalance between particles

moving in one direction and the opposite direction in a plane due to the imbalance in the number of particles either side of the plane. This drift for species s is [23]

$$\mathbf{v}_{ds} = \frac{\mathbf{B} \times \nabla p}{n_s B^2}. \quad (1.16)$$

Fluid models are formally valid provided behaviour is localised, which is true if the mean free path of particles is small compared to the macroscopic length scales of the system.

Magnetohydrodynamics (MHD)

Combining the fluid equations for the two species into a single fluid equation gives a description called magnetohydrodynamics (MHD), and the separate behaviour of the ions and electrons is not explicitly kept. Typically the plasma is taken to be in the adiabatic limit, so the pressure is isotropic and viscous terms become zero and the model is dissipation free. The equations become:

$$\frac{\partial \rho_m}{\partial t} = -\nabla \cdot (\rho \mathbf{u}) \quad (1.17)$$

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = \mathbf{J} \times \mathbf{B} - \nabla p + \mathbf{F} \quad (1.18)$$

$$\frac{\partial p}{\partial t} + \mathbf{u} \cdot \nabla p = -\gamma p \nabla \cdot \mathbf{u} \quad (1.19)$$

where ρ_m is the mass density, \mathbf{u} is the fluid velocity, \mathbf{J} is the plasma current and \mathbf{F} is a force term containing external forces on the plasma, such as gravity. These equations are combined with Maxwell's equations in the limit where the displacement current is neglected, and an Ohm's law equation

$$\mu_0 \mathbf{J} = \nabla \times \mathbf{B} \quad (1.20)$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \quad (1.21)$$

$$\mathbf{E} + \mathbf{u} \times \mathbf{B} = \eta \mathbf{J} \quad (1.22)$$

where η is the resistivity of the plasma. A common assumption taken at this point is that the plasma is a perfect conductor, so that $\eta = 0$ in Equation 1.22, in which case the magnetic field lines are 'frozen in' to the plasma, moving with it.

MHD is the simplest plasma framework to use when developing a model, and can be thought of as an extension to the Navier-Stokes equa-

tions to include Maxwell's equations. MHD accurately represents many plasma phenomena such as the kink instability, and Rayleigh-Taylor instability, but some physics, such as drift waves and trapped particles are lost.

1.4 Instabilities in Plasma

Like neutral fluids, plasma can demonstrate a plethora of different types of instabilities, across a range of configurations. These are typically divided into macroinstabilities, which are ascribable to MHD physics, and microinstabilities, which require physics associated with the finite Larmor radius and kinetic dissipation effects. These microinstabilities have wavelengths on the order of the ion Larmor radius, and are responsible for the fine scale transport associated with anomalous transport (see Section 1.5). This thesis focuses on this fine scale microturbulence and a review of relevant modes is presented here.

1.4.1 Ballooning and Interchange Instabilities

The most significant instability relevant to the global profile of tokamak plasmas is the ballooning instability, a pressure driven instability analagous to the Rayleigh-Taylor instability in neutral fluids, but with a curvature force in place of the gravitational one.

Consider a plasma with a density (and hence pressure) gradient in the y -direction, being supported by a magnetic field in the z -direction. For a static equilibrium (in which the velocities are zero) the MHD momentum equation, Equation 1.19 gives a pressure balance equation:

$$-\nabla p + \mathbf{J} \times \mathbf{B} + \rho \mathbf{g} = 0 \quad (1.23)$$

$$-\nabla p + (1/\mu_0)(\nabla \times \mathbf{B}) \times \mathbf{B} + \rho \mathbf{g} = 0 \quad (1.24)$$

$$-\nabla \left(p + \frac{B^2}{2\mu_0} \right) + \rho \mathbf{g} = 0 \quad (1.25)$$

where $\mathbf{g} = -g\hat{\mathbf{y}}$ is the gravitational field. Note that $B^2/2\mu_0$ acts as an effective pressure, so is often given the name **magnetic pressure**. Since p increases with y , B must decrease with y such that the magnetic pressure gradient opposes the kinetic pressure gradient and gravitation terms.

This equilibrium is unstable due to the availability of perturbations




Figure 1.4: The Rayleigh-Taylor instability in a plasma. **Top:** Equilibrium configuration. A lower density region (red), lies beneath a denser region (blue) within a magnetic field \mathbf{B} pointing into the page. This is a simplified picture of a gradient in the pressure and density, pointed in the y -direction. The denser plasma is held in (unstable) equilibrium by a gradient in the magnetic field strength. A force, \mathbf{g} , representing either a gravitational or centrifugal force acts on the plasma. **Bottom:** A small perturbation, aligned perpendicular to both the gradients and the magnetic field. Charged particles experience gravitational (or curvature) drifts, ions to the right, and electrons to the left. The imbalanced x -distributions mean these charges are able to build up and establish an electric field. The $\mathbf{E} \times \mathbf{B}$ -drift created reinforces the perturbation.

which can reduce the potential energy of the system. If a perturbation occurs, with the wavevector directed in the x direction, the background density gradient will carry on top of it a small perturbation regions of excess and reduced density along the x axis. The gravitational drift of the particles will no longer balance each other, allowing charge to accumulate. This build up of charge creates an electric field, which in turn causes the particles, and hence the plasma as a whole, to experience an $\mathbf{E} \times \mathbf{B}$ drift in the direction that reinforces the perturbation, leading to growth. The growth of the electric field also causes a polarisation drift which acts to reduce the rate at which the \mathbf{E} -field is generated. The polarisation drift will cancel out the gravitational drift, and this observation can be used to identify the rate of change of the \mathbf{E} field. The instability is shown in Figure 1.4.

This mechanism also applies to the scenario where a plasma sits within a curved magnetic field, such as a tokamak. To zeroth order, the motion of the charged particles follows the curved magnetic field

lines, and so in their rotational rest frame they experience a centrifugal force in the outwards direction, parallel to the radius of curvature of the \mathbf{B} -field, which causes the **curvature drift** in Table 1.1. This is completely analogous to the gravitational case, with the centrifugal force replacing the gravitational term, however is significantly more relevant to tokamak plasmas since the gravitational drift is miniscule in comparison, so gravitational effects are typically completely neglected in tokamak physics. Note that when using the standard x , y , and z directions in flux tube coordinates, the perturbations have wavenumber in the y direction rather than the x direction, and the density gradient is in the x direction rather than the y direction - this equates to simply flipping the x and y directions in the above discussion and in Figure 1.4.

The curvature-driven Rayleigh-Taylor instability is also known as the **flute instability** since the perturbations in a cylindrical configuration extend out of the surface of the cylinder, much like flutes on architectural columns. The curvature only drives the plasma unstable when the radius of curvature is directed against the pressure gradient. In a tokamak, the toroidal field alone would cause the outboard side of the device to be unstable, and the inboard side to be stable, since the curvature is convex to the density. By twisting the field about the magnetic axis however, the (non-trapped) plasma particles will follow the field lines from the outboard side of the plasma, to stable inboard side, and since their parallel velocity is lower on the inboard side, they spend more time in the stable region, and thus the overall effect is that the flute mode is stable in a global sense. However, the local instability on the outboard side of the plasma still causes perturbations to grow locally, to an extent moderated by the parallel speed of the electrons, and this growth can still be significant enough to transport plasma to the edge of the device, if not carefully moderated. This mode is then called the **ballooning mode** since its behaviour is comparable to the elongations observed in a long balloon when squeezed.

In the limit where the flute-like perturbations do not perturb the magnetic field, the instability is known as the **interchange instability**, since the perturbations are equivalent to exchanging a bundle of plasma with a bundle of field lines.

1.4.2 Drift Waves

Drift waves are a type of mode which are responsible for fine scale plasma turbulence. The basic drift wave mechanism can be seen by examining a slab of plasma in a uniform magnetic field with a density gradient, such as the scenario set up in Figure 1.5, where a density gradient is directed in the x -direction and the magnetic field points in the z -direction. We consider a wave of the form $e^{ik_z z + ik_y y - i\omega t}$. Electrons have a low inertia, so assuming a low frequency, the electrons will move along the field lines to adiabatically respond to the potential perturbation, giving the Boltzmann relation for the electron density

$$\frac{\tilde{n}_e}{n_e} = \frac{e\phi}{T_e}, \quad (1.26)$$

which gives a parallel pressure which exactly opposes the parallel electric force. The wave perturbation gives a perturbed electric field in the y direction, which in turn causes an $\mathbf{E} \times \mathbf{B}$ drift in the x -direction. This perturbation is $\pi/2$ out of phase with respect to the ϕ (and hence n) and hence the wave propagates in the y -direction. Since there is a pressure gradient, one would expect the diamagnetic drift (Equation 1.16) to play a role, however since it is divergenceless, it cannot cause a density perturbation. The frequency of the drift wave is found by solving the system consisting of the ion fluid equations, Equation 1.26, and the quasineutrality relation eventually gives a frequency

$$\omega_{*e} = -\frac{k_y T_e}{e B n} \frac{dn}{dx} \quad (1.27)$$

which is real and so shows no growth or decay. This is because the simple fluid model being used does not include enough physics for the instability to be apparent. For the drift wave to be unstable, a mechanism is required to separate charges in the y -direction so that the $\mathbf{E} \times \mathbf{B}$ drift can enhance the perturbation, rather than only causing it to propagate. One such mechanism is the ∇B -drift given in Table 1.1. If there is a temperature gradient in the same direction as the density gradient, the ∇B -drift causes hotter particles to drift faster than colder particles allowing net charge separation. This is a very common and effective driver of drift wave instability, and the whole mechanism is called the **ion-temperature gradient (ITG) instability**. The fluid model cannot describe the effects of temperature on charge separation, and so to see this effect, a more advanced model, such as a gyrokinetic or full kinetic model,

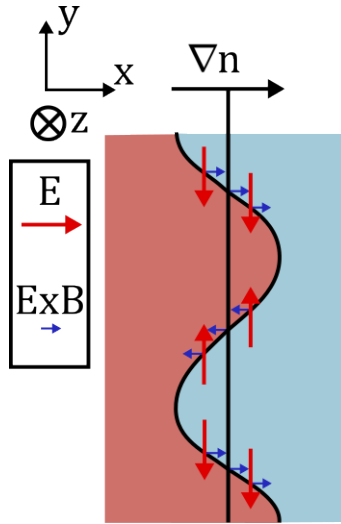


Figure 1.5: The drift wave mechanism. A density gradient in the x direction is perturbed by a wave with wavevector in the y and z directions. A magnetic field points into the page in the z direction. High density plasma is shown in light blue, and low density in light red. The potential is proportional to the density, through the Boltzmann electrons, so the perturbed E field (red arrows) points into areas of high density. The associated $\mathbf{E} \times \mathbf{B}$ drifts are shown in blue.

is required. A full calculation can be found in [23] or [37]. The ITG instability is thought to be a leading cause of plasma microturbulence and the associated turbulent transport[3, 11, 28, 37].

1.5 Anomalous Transport and Drift Wave Turbulence

To maximise the confinement time of a plasma, it is critical to understand and minimise the amount of heat and matter transport from the core of the plasma to the edge. For the most part, charged particles are effectively bound to magnetic field lines by the Lorentz force, however drift motion leads to particles having larger excursions than simple gyromotions, and so collisional effects lead to larger transport than would be expected based on gyration alone. An improved transport model, called **neoclassical transport**, accounts for these effects provide a more full understanding of these effects.

The neoclassical model provides an improved understanding of the transport effects in plasma, yet still falls short of accurately describing the transport rates in devices. Experimental measurements of transport in tokamak devices were found to be as much as an order of magni-

tude larger than the neoclassical theory predicted [2, 38]. Referred to as **anomalous transport**, this remained unexplained for some time, however the scaling laws of the transport in physical systems has led to a widely held consensus that anomalous transport is generated primarily by small-scale turbulence [1]. This turbulence is driven by microinstabilities with ITG driven drift wave turbulence understood to be a significant contribution to this transport. [2, 37].

Turbulence leads to the formation of complex structures, such as eddies, within the plasma, changing the mechanism, length scales and timescales over which transport occurs, drastically increasing the transport rates from the core to the edge of the plasma. The diffusion coefficient can be estimated through a simple back of the envelope calculation; a simple random walk approach to diffusion gives a diffusion coefficient of $D = \delta^2/\tau$, where δ is the mean free path, and τ is the mean collision time, or equivalently, $D = \delta v$, where v is the mean particle velocity. In a turbulent system, the velocity is given by the $\mathbf{E} \times \mathbf{B}$ -drift velocity, which has magnitude E/B . Next, E can be approximated as ϕ/δ as the eddy size and mean free path will be of the same order. Next one assumes that $e\phi \sim k_B T$ since the thermal energy available to the particles is what allows them to separate and form a potential gradient. Combining these gives the so-called **Bohm scaling** for a species s ,

$$D_s \sim k_B T_s / q_s B, \quad (1.28)$$

which gives, for a typical tokamak regime of $T_s = 10\text{keV}$, $B = 3\text{T}$, a diffusion coefficient of $D = 3 \times 10^7 \text{cm}^2 \text{s}^{-1}$.

The high diffusion rate of turbulent transport makes it essential to stabilise these instabilities as much as possible, or to at least reduce their saturation amplitudes to minimise the amounts of heat and particle losses from the core of the plasma.

1.5.1 Coherent Structure Dominated Turbulence

When physicists discuss turbulence, they are often referring to an idealised version of turbulence - **fully developed homogenous turbulence** - in which the fluid is characterised by a hierarchy of length scales which cascade energy from long length scales to short scales where energy is eventually dissipated through viscous terms. Under this framework, turbulence is treated by making assumptions about the statistics

of the spectrum of the turbulent fluctuations. Real turbulence however, whether experimental or simulated, is often instead built around the formation and evolution of unstable **coherent structures** [16, 39–41] - structures which are spatially localised, sometimes recurrent, and may propagate through the system, such as vortices and streamers. The propagation mechanisms behind these structures are outlined in Section 1.6. This turbulent phenomenology is particularly characteristic of subcritical systems [40, 42] such as the plasma interchange model used throughout this thesis. These are systems in which there is a fully linearly stable equilibrium, but a large enough perturbation pushes the system unstable, and will be explored more in Section 1.8.5.

Turbulence dominated by coherent structures is different to the turbulent cascade behaviour (though both types are chaotic) that a reader may be more familiar with, and so it is important to emphasise that the turbulence under consideration in this thesis is of the coherent structure dominated kind, rather than the idealised homogenous kind and henceforth will simply be called ‘turbulence’ (note that in some fields the term turbulence is reserved exclusively for homogenous fully developed turbulence - no such reservation is made here). Coherent structure dominated turbulence cannot be studied using all the same statistical methods of homogenous turbulence, and so alternative methods are required, particularly those of dynamical systems theory, which this thesis focuses on.

1.5.2 Shear Flows and Turbulence

It has been observed that sheared flows, where the fluid velocity has a strong gradient in the direction perpendicular to the flow, can minimise the amplitudes of turbulence and reduce turbulent transport [43–45]. The mechanism for this is quite straightforward [46]; when a plasma eddy placed in a stable shear flow, fluid elements within the eddy are advected in the flow direction at different velocities, distorting the structure of the eddy, stretching it in the direction of the flow. If the spatial extent of the eddy then exceeds the coherence length of the turbulence (i.e the structure is stretched across to a separate eddy) then it will become advected by other eddies, effectively breaking up the eddy, and the ‘original’ structure is no longer identifiable. This is demonstrated in Figure 1.6. The stretching and breaking up of the turbulent structures reduces the diffusive step size, the eddy lifetime is greatly reduced, and




Figure 1.6: Stretching and splitting of an eddy by a shear flow.

the turbulent intensity is reduced. All of these factors contribute to reducing the turbulent transport across the flow. Shear flows can arise in a plasma in a variety of ways:

- One method for heating the plasma injects a beam of neutral atoms close to the central axis. This causes the plasma to rotate in the toroidal direction, with the core moving faster than the edge. [47]
- A mean poloidal flow can be generated by turbulence through the Reynolds stress mechanism. [48]
- A biased probe can charge the plasma edge [49] to deliberately induce a shear flow.
- Sheared flows, called zonal flows can spontaneously be generated via nonlinear wave interaction processes in drift wave turbulence [7, 50].

Zonal flows play a particularly important role in the nonlinear saturation of turbulence, and have been found to be significant in simulations across an extensive range of scenarios and models [7, 51]. A zonal flow is a flow consisting of azimuthally and toroidally symmetric ($n = m = 0$ modes) bands of poloidal flow with radially varying velocity, with zero frequency. Due to their symmetry properties, they cannot access the free energy available in the pressure gradients, and are instead driven exclusively by nonlinear coupling with drift waves. Thus energy is taken directly from drift-wave turbulence and its intensity is reduced.

Zonal flow energy is dissipated primarily through collisional friction since they experience no Landau damping¹. Collisionality is typically

¹Landau damping is a kinetic effect in which waves can be damped due to there being a greater number of particles with thermal velocities below the wave velocity than above it - the slower particles are accelerated and the faster ones are slowed - so that the net flow of energy is from the wave to the particles. This is not relevant to the zonal flows since they do not propagate. [7]

small in tokamak conditions and so large zonal flows are able to build significant amplitudes and become an energy sink for the system. In addition, the zonal flow mode is able to grow faster than other $\mathbf{E} \times \mathbf{B}$ -flow modes as $k_{\parallel} = 0$, and therefore electrons do not flow along the field lines to screen the generating electric potential, giving them a lower effective inertia than modes with $k_{\parallel} \neq 0$.

Combined with the eddy-shearing effect discussed above, the drift-wave-zonal-flow interaction becomes a complex, but critical interaction to understand.

The effect of flow shear on turbulent transport has been used both intentionally and inadvertently to generate **transport barriers** within fusion devices, usually within the edge of the plasma, which have proved to be effective in reducing transport of heat out of the device. Of particular note, the transport barrier mechanism is believed to be responsible [6] for the so called ‘High Confinement Mode’ (H-mode) in tokamak devices, where energy confinement time is enhanced by a factor of two or more [52] when compared to the ‘Low Confinement Mode’ (L-Mode) in which turbulent transport is not suppressed and dominates transport. A characteristic feature of H-mode is the presence of a ‘pedestal’, which is a global, roughly uniform, increase in pressure compared to L-mode throughout the core of the device, with a steep gradient close to the edge, on the outside of the transport barrier, as shown in figure 1.7.

Transition from L-mode to H-mode is spontaneously triggered when a threshold heating power is exceeded, and though it is widely believed that the transport barrier is caused by shear flows, the details of the mechanism by which these shear flows are generated are not obvious and continue to be investigated [6, 53, 54].

In addition to reducing the magnitude of turbulence, shear flows can also fully linearly stabilise plasma equilibria in some configurations [39, 55–58].

1.6 Review of Simple Models Containing Propagating Structures

There exist a range of simple mechanisms by which a model can contain structures which propagate through space. These can be demonstrated using simple canonical fluid dynamics models. Here several of these models are reviewed.

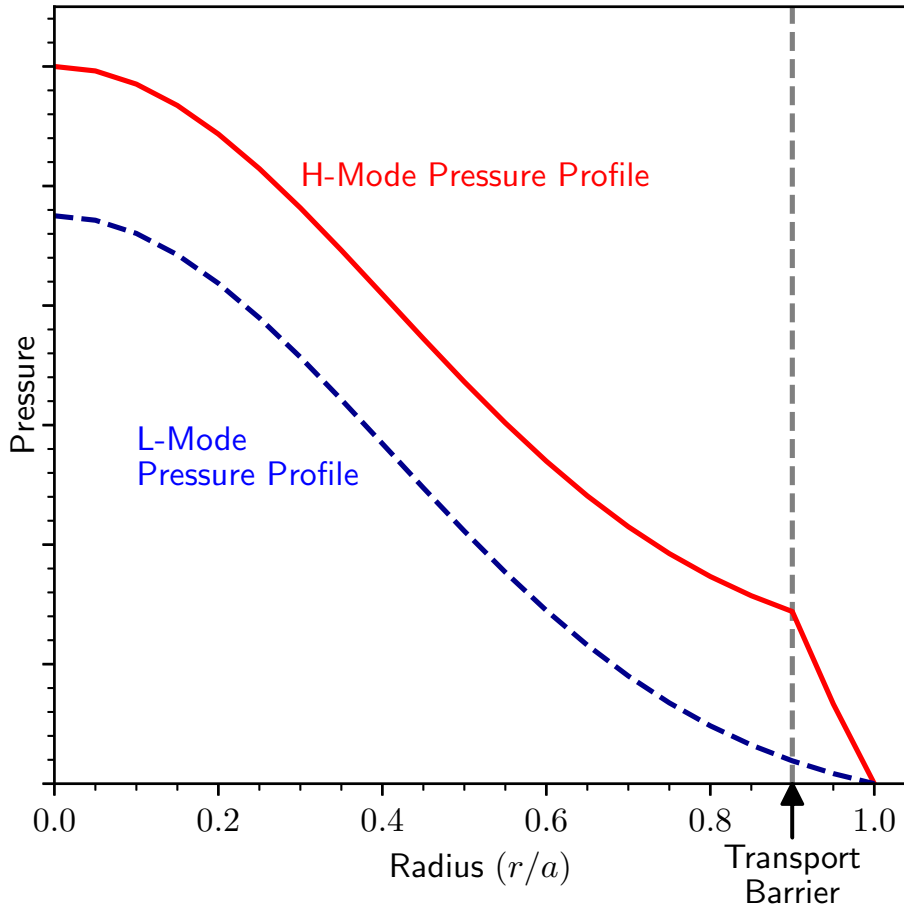


Figure 1.7: Pressure profiles against radius in L and H mode.

Linear Advection

The simplest mechanism by which propagation can occur is for a number density field $n(x)$ to be advected by a constant velocity c . This equates to a linear, single directional wave equation:

$$\partial_t n + c \partial_x n = 0 \quad (1.29)$$

In this scenario $n(x, t)$ will always be of the form $n(x, t) = n(x - ct)$ and hence the shape of n does not change but is simply translated with time.

Nonlinear Advection

The most simple nonlinear advection equation is seen when the field being advected is the velocity itself. The equation governing this system is

$$\partial_t u + u \partial_x u = 0 \tag{1.30}$$

where $u(x, t)$ is the velocity field. In this system, the faster sections of the velocity catch up to the slower ones, eventually leading to the former overtaking the latter, and u either becoming multivalued (as in the case of an ocean wave breaking) or a discontinuity can be inserted, modelling a shock. Shocks appear in many physical systems, including ICF plasmas, explosions, and sonic booms. While the advection equation above does not include dissipation, in physical systems the steep gradients lead to rapid dissipation of energy through viscous effects, causing heating. This heating is harnessed to initiate fusion reactions in ICF devices.

Dissipation and Burgers' Equation

There are many additional effects that can be added to the simple advection equation that reflect effects common to models across all fluid dynamics. One such term is viscous dissipation, giving Burgers' equation, first introduced by Bateman [59] and later studied in detail by Burgers [60]. The equation is

$$\partial_t u + u \partial_x u = \nu \partial_{xx} u \tag{1.31}$$

where ν is a control parameter analogous to the viscosity in fluid models, and controls the diffusion of the velocity field and acts dissipatively. The presence of the dissipation term prevents gradients becoming too steep, mitigating the discontinuity of the shock and causing structures to ultimately decay with time.

Solitons and Korteweg-de Vries Equation

A second physical effect that can be added to the simple advection model is dispersion, where waves of different wavelengths travel at different phase velocities. Dispersion will naturally cause a wave packet to spread out and reduce in amplitude as it travels, however this spreading effect can be countered by the bunching effect of advection described above. The simplest model incorporating these two competing effects is the Korteweg-de Vries (KdV) system [61] which relates to the physical case of long wavelength solitary waves on the surface of shallow water. The canonical non-dimensionalised form of the equation is

$$\partial_t u - 6u\partial_x u + \partial_{xxx} u = 0. \quad (1.32)$$

The nonlinear and dispersion terms compete in such a way that for a travelling wave of the correct form and amplitude, their effects cancel (aside from the translation of the wave) and the wave maintains its form for a long time - this solution is usually known as a **soliton**. Though a strict definition of the term does not exist. Solutions labelled solitons will usually be localised, of constant form, and will interact in such a way with other solitons that their form remains unchanged [62] as they pass each other, though they may be displaced relative to their original trajectory. In the KdV system, solitons have the form

$$u(x, t) = -\frac{c}{2} \operatorname{sech}^2[c^{1/2}(x - ct - x_0)] \quad (1.33)$$

where c and x_0 are parameters of the solution corresponding to the speed and initial peak of the soliton. The amplitude of a soliton is linearly related to its speed and quadratically to the inverse of its length scale. The existence of these solutions can be explained intuitively as a balance between the nonlinear term, which acts to focus the function locally, and the dispersive term, which acts to defocus the function locally.

Nonlinear Schrödinger Equation

The nonlinear Schrödinger equation (NLSE) is another standard PDE that demonstrates propagating soliton-like solutions. The NLSE has been linked to the propagation of waves in optical fibres and the envelope function for modulated surface waves in deep water. The equation is

$$i\partial_t \psi = -\frac{1}{2}\partial_{xx}\psi + \kappa|\psi|\psi^2 \quad (1.34)$$

where $\psi(x, t)$ is a wavefunction and κ is a parameter of the model. For $\kappa < 0$, the nonlinear term is a focussing nonlinearity and the system allows solitons. In the case $\kappa > 0$ the system is defocussing and instead allows ‘dark soliton’ solutions which have a local dip in amplitude and a constant amplitude at infinity.

Self Organised Criticality: Sandpiles and Avalanches

In 1992 a new paradigm was proposed by Bak, Tang, and Wiesenfeld [63] to explain the $1/f$ -spectrum of noise that appears across dynamical systems in several physics disciplines, including plasma physics. They demonstrate that dynamical systems with spatial degrees of freedom will self-organise into a critical configuration - this is where energy accumulates locally, steepening gradients until the local gradient exceeds a critical value and becomes locally unstable. Energy is then transferred to neighbouring areas, causing them to become unstable and so on, so that energy is ultimately transferred through space in a domino-like effect. This means small local perturbations can trigger large events, called **avalanches**. The authors demonstrate this phenomenon with a sandpile model. The model is a simple cellular automata that models a discrete variable z on a discrete grid with coordinates x, y . If z exceeds a model parameter K the grid point ‘collapses’ and donates a unit of z to each of its neighbours, i.e

$$z(x, y) \rightarrow z(x, y) - 4 \quad (1.35)$$

$$z(x \pm 1, y) \rightarrow z(x, y) + 1 \quad (1.36)$$

$$z(x, y \pm 1) \rightarrow z(x, y) + 1 \quad (1.37)$$

This collapse can in turn trigger further collapses, and so a front of collapse events can propagate from a single event. The authors identified that in models of this type across infinite domains, if small local perturbations are repeatedly added to the system, the system gets attracted into a critical state where the correlation length of the system goes to infinity, meaning there is no single identifiable length scale in the system, in the same way as the Ising model at critical temperature. Similarly, the timescales of the fallout from a single perturbation also have a correlation time of infinity. It is this phenomena that causes the $1/f$ frequency distribution of noise.

Hwa and Kardar [64] described a method to create a continuous rather than discrete sandpile-like model which has hydrodynamic properties. They show through symmetry arguments that Burgers’ equation is the general form of continuous, coarse-grained SOC models.

1.7 Review of Literature Concerning Plasma Turbulence and Avalanche-like Structures

Hasegawa-Mima Model

The Hasegawa Mima model [65] was developed to describe the pseudo-three-dimensional dynamics of nonuniform magnetized plasmas in the cold ion limit $T_i \ll T_e$. The model is a reduced fluid model which accounts for three-dimensional effects on the electron, but reduces the ion-dynamics to two dimensions (those perpendicular to the magnetic field). The authors argue that the parallel electron dynamics are important since if there is a slow variation in potential along the parallel direction, the electrons are free to flow along the field lines to screen this potential, obeying the Boltzmann distribution $\tilde{n}/n_0 = e\phi/T_e$. The high inertia of the ions makes them too slow to respond in the parallel direction and their dynamics becomes fully two-dimensional. The model assumes phenomena have low frequency compared to the ion gyrofrequency, high B_0 so that the model can be treated as electrostatic (i.e. $\mathbf{E} = -\nabla\phi$), and the plasma remains quasineutral. Finally, for simplicity, the pressure tensor is assumed to be isotropic, and so there is no viscosity term in the equations.

The electric and magnetic fields give an $\mathbf{E} \times \mathbf{B}$ drift as in Table 1.1 and the time-varying electric field gives rise to a polarisation drift. The magnitude of the polarisation drift is small compared to that of the $\mathbf{E} \times \mathbf{B}$ drift, however the divergence of the latter is zero and so the former is retained in divergence terms to allow compression of the ions. The sum of these two drifts becomes the fluid velocity of the ions and the ion continuity equation becomes

$$\begin{aligned} \frac{d \ln n}{dt} + \nabla \cdot (\mathbf{u}_E + \mathbf{u}_P) \\ = \frac{d \ln n}{dt} + \frac{1}{\omega_{ci} B_0} \nabla \cdot \left[-\frac{\partial}{\partial t} \nabla_{\perp} \phi - (\mathbf{u}_E \cdot \nabla_{\perp}) \nabla_{\perp} \phi \right] = 0 \end{aligned} \quad (1.38)$$

and writing the density as $n = n_0(1 + e\phi/T_e)$ and assuming a small potential, so that $e\phi \ll T_e$ leads to

$$\frac{d}{dt} \left[\frac{e\phi}{T_e} \right] + \mathbf{u}_E \cdot \nabla \ln n_0 + \nabla \cdot \mathbf{u}_P = 0, \quad (1.39)$$

and this can be normalised and rearranged to give the Hasegawa-Mima equation,

$$\frac{\partial}{\partial t} (\nabla^2 \phi - \phi) - [(\nabla \phi \times \hat{b}) \cdot \nabla] \left[\nabla^2 \phi - \ln \left(\frac{n_0}{\omega_{ci}} \right) \right] = 0. \quad (1.40)$$

The HM equation is similar to Euler's equation for an incompressible 2D fluid,

$$\left(\frac{\partial}{\partial t} - \nabla \psi \times \hat{z} \cdot \nabla \right) \nabla^2 \psi = 0, \quad (1.41)$$

where $\psi(x, y)$ is the stream function satisfying $\mathbf{u} = \hat{z} \times \nabla \psi$. The differences between the HM model and a 2D Euler fluid are the anisotropy that appears in the $\nabla \ln(n_0/\omega_{ci})$ term and the lack of scale invariance in the HM model, as both ϕ and $\nabla^2 \phi$ appear in the time derivative on the LHS. The HM model has become a widely used model that features drift waves, though being a fluid model with adiabatic electrons, no drift wave instability is present and so one must add a forcing term by hand to see instability occur.

The Hasegawa-Mima equation has been found to be equivalent to the Charney equation which governs the geophysical Rossby waves [66], a type of atmospheric and oceanic waves, driven unstable by planetary rotation and surface curvature. The equivalence arises due to the mathematical equivalence of the $\mathbf{J} \times \mathbf{B}$ forces of the plasma and the Coriolis forces of the Earth's atmosphere. Like drift waves, Rossby waves also interact with zonal flows in an analogous way to drift waves.

Drift waves exhibit spectra close to the Kolmogorov-Kraichnan spectrum [67] of 2D incompressible fluid turbulence [66]. Analysis of the HM system and Charney system (often combined into the Charney-Hasegawa-Mima equation) showed that there are two primary ways zonal flows can be generated by turbulence. The first is via an inverse cascade of energy from small scales to large scales [66, 68], which is only possible due to the turbulence being approximately 2D. The second way zonal flows are generated is via a modulational instability [69–71]. The modulational instability is an instability which occurs when a monochromatic wave has an amplitude modulated by a wave with wavelength much longer than its

own. This modulational instability acts as a secondary instability to the primary drift instability. This mechanism has been used to demonstrate how the transition from L-mode to H-mode could occur via exchange of energy between drift waves and zonal flows [72].

The Hasegawa-Mima system is also found to contain propagating dipole vortex-like coherent structures [73]. These solutions are a pair of vortices close to one another rotating in opposite directions so that the flow between them causes them to propagate. The appearance of coherent structures in this model indicates that the RPO methods being developed in this thesis could also be applied to the Hasegawa-Mima system to locate, isolate and characterise these structures.

Hasegawa Wakatani Model

The HM equation proved to be useful in developing theories of strong drift wave turbulence [74, 75] in the core of tokamaks and explaining its broad wavenumber spectrum. Observations from tokamaks [76, 77] identified that turbulent fluctuations in tokamak plasmas increased in amplitude at the edge of the plasma, where the plasma is cooler and collisional effects matter more. To better describe this important plasma regime, Hasegawa and Wakatani (HW) built on the HM equations to include dissipative and resistive effects to derive a model for edge plasma [78].

$$\left(\frac{\partial}{\partial t} - \nabla\phi \times \hat{z} \cdot \nabla\right)\nabla^2\phi = c_1(\phi - n) + c_2\nabla^4\phi \quad (1.42)$$

$$\left(\frac{\partial}{\partial t} - \nabla\phi \times \hat{z} \cdot \nabla\right)(n + \ln n_0) = c_1(\phi - n) \quad (1.43)$$

where $c_1 = -\frac{T_e}{e^2 n_0 \eta \omega_{ci}} \nabla_{\parallel}^2$, $c_2 = \frac{\mu}{\rho_s^2 \omega_{ci}}$. In the collisionless limit, $c_1 \gg 1$, and $c_2 \ll 1$ and the model simplifies to the HM equations. In the incompressible limit, $c_1 \ll 1$ and $c_2 \ll 1$ and the model reduces to the Euler equations for an incompressible fluid. The HW model is a more physically accurate model than the HM equation, and unlike the HM model, features an instability that drives drift waves.

Plasma as an SOC Model

Bursty transport in plasma has been approached from several different angles in attempts to understand the mechanisms causing bursts, and their interaction with zonal flows and turbulence. The earliest models

adapted the idea of SOC avalanches to plasma at marginal stability. References [79, 80] speculated that the dynamics of turbulent transport near marginal stability (where an increase in local gradient can greatly increase the flux, thereby reducing it again) could be described by a SOC paradigm, inspired by the BTW sandpile model and the work by Hwa and Kardar [64]. The authors later identified a model that exhibits SOC behaviour [10] derived from a pressure gradient driven model [81] and has the form of a nonlinear Burgers' equation. The authors showed that with certain dissipative parameters, and in the absence of a source term, the pressure profile self-organises to one which is critical. When the noise source term is reintroduced, local instabilities are triggered which trigger avalanches that propagate through the system as in the sandpile model, and transport heat and particles with them.

Interchange Turbulence Model

The HM and HW models outlined above describe turbulence consisting of drift waves. While drift wave turbulence is accepted to be the most dominant form of turbulence in tokamak plasmas, other waves and instabilities can be excited and can play a significant role in transport in some configurations. Beyer and Spatschek introduced a model in 1996 [82] which described the interchange instability in a reduced MHD model. The non-dimensionalised model equations are

$$\partial_t \nabla_{\perp}^2 \phi + \{\phi, \nabla_{\perp}^2 \phi\} = -\nabla_{\parallel}^2 \phi - \partial_y p + \nu \nabla_{\perp}^4 \phi \quad (1.44)$$

$$\partial_t p + \{\phi, p\} = -\partial_y \phi + \chi \nabla_{\perp}^2 p \quad (1.45)$$

where the Poisson brackets $\{\cdot, \cdot\}$ represent advection terms, $\{\phi, A\} = (\vec{u}_{E \times B} \cdot \nabla A) = \partial_x \phi \partial_y A$. This model is similar to a host of other models available at the time, however the model described here has formed the basis for many further reduced models that demonstrate propagating structures and hence remains the most relevant model to this work. This model was later used to simulate ballooning turbulence in tokamak plasmas with stochastic field lines [83], to explore the interplay of zonal flows with avalanches and streamers [51]. Later a low-dimensional reduced form of the model was derived using proper orthogonal decomposition [84] that is used to study shear flow generation, and the impact of shear flows on turbulent amplitudes.

Intermittent Transport in Gyrokinetics and Experimental Data

Avalanchelike bursts of turbulence were identified in gyrokinetic simulations in 2003 [4] and comparisons to experimental data from the DIII-D tokamak showed consistent energy diffusivities.

Scrape Off Layer Interchange Turbulence

In Reference [13] a reduced 2D model for interchange turbulence in the scrape-off layer was derived by allowing the radial profiles to vary in time, while keeping a constant flux source term. The model exhibits poloidal modulations which have linear growth rates maximised at $k_y = k_{y0}$. This model was compared in [12] with 2D and 3D global simulations of the core region and yielded a consistent description of transport processes across multiple scales, demonstrating that the flux-driven model could be appropriate for these kinds of analyses, while forcing a fixed profile does not allow the profile to self-organise into a critical state. They go on to demonstrate the statistical properties of the associated propagating fronts, noting that their timescale is longer than the lifetime of vortices in the fixed gradient models, and that the time averaged profiles are supercritical everywhere, differing from the simpler sandpile models but the frequency spectrum for intermediate timescales approximately follows a $1/f$ distribution characteristic of SOC. The scrape off layer model was then reduced in [85] by separating the poloidally fluctuating density and electric potential fields from the poloidally symmetric modes, and selecting only the modes with poloidal wavenumber k_{y0} , which corresponds to the highest growth rate. The reduced model reveals a transition from dynamics dominated by diffusion to dynamics dominated by avalanches as the relative weight of nonlinear mode coupling is reduced.

Local Models of Structures

McMillan et al in 2009 [39] discussed that SOC models, while demonstrating avalanche-like behaviour, do not accurately describe some qualitative features of ITG turbulence observed in full GK simulations of ITG. The primary differences are that the frontlike solutions in the works described above have amplitude dependent velocities and are linear in nature while it is the nonlinear behaviour of bursts that is of interest since the bursts appear only in turbulent regimes. The authors determine that the reason for this is that the models rely on global physics, being driven either by a constant flux or global pressure gradient, but the bursts observed in




Figure 1.8: A representation of a dynamical system mapping. A system initially in state x_0 at $t = 0$ will be in state $f^t(x_0)$ at time t . The path parameterised by t is called the orbit of x_0 .

ITG turbulence appear to be associated with turbulence scales on short spatial and temporal scales, and as such the coarse-graining procedures that prior approaches took cannot be fully justified. They propose an alternative formulation which uses local, rather than global physics to explain the propagation of bursts, building a model that is closer to an autosoliton-like model in which soliton-like structures exist for a specific velocity and amplitude, and have zero value at $\pm\infty$.

1.8 Relative Periodic Orbits

The fundamental idea behind the work in this thesis is rooted in the theory of nonlinear dynamical systems. In particular, nonlinear coherent structures, such as travelling waves, in plasma and their stability are examined by describing them as relative periodic orbits, a feature of dynamical systems theory. This approach has not had previous use within the plasma theory community, and hence it will likely be of benefit to the reader to include a brief introduction to the relevant terminology and physics of dynamical systems theory here.

1.8.1 Dynamical Systems Background

A **dynamical system** is a system whose state space can be fully represented by a manifold, \mathcal{M} , such that for any x representing a possible state of the system, $x \in \mathcal{M}$. As the system's state evolves dynamically with time, the position $x(t)$ on the manifold evolves analogously to the evolution of the state. This time evolution is represented by an operator f^t which maps a point x_0 on \mathcal{M} to another point $x(x_0, t) \in \mathcal{M}$ representing the system after t time units. Mathematically,

$$x(x_0, t) = f^t(x_0) \tag{1.46}$$


where x_0 on the LHS is a label that can be dropped if the context allows, then $x(t) = f^t(x_0)$. This definition allows for both discrete (where $t \in \mathbb{Z}$) and continuous (where $t \in \mathbb{R}$) dynamical systems. The discrete case is sometimes called a map. If f^t is continuous and well-behaved in time, it is also possible to define the local **velocity field**, $v(x)$, by the limit

$$\dot{x} = v(x) = \lim_{dt \rightarrow 0} \frac{f^{dt}(x) - x}{dt}. \quad (1.47)$$

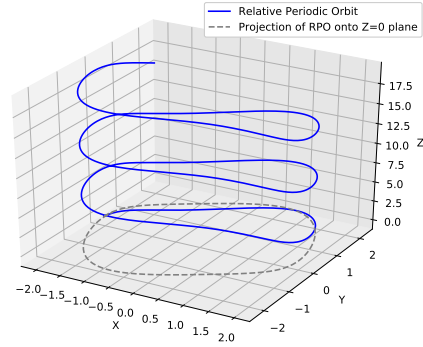
In other words, v represents the ‘velocity’ through the system space that a particle at that point and time experiences (to reiterate, this is the system space velocity, and not necessarily a ‘physical’ one). Note that it is possible for f^t to be dependent on time, as well as on state space position, however it is straightforward to simply treat t as another coordinate in system space, and take $v(x)$ in that direction to always be equal to one. This allows the form given in equation to be a general form without losing the possibility of time dependence.

This framework allows one to access the full future and, if time-reversible, full past of a system in a state x_0 . The path sketched out by $x(t)$ forms a continuous set of points, called the **orbit** through the point x_0 , and can be represented by the manifold $\mathcal{M}_{x_0} \subset \mathcal{M}$. Every point on \mathcal{M} will have an associated orbit. We can now define orbits that have special properties:

- **Stationary point:** A point is stationary if $x(t) = x(0)$ for all t . Expressed differently, the orbit of $x(0)$ contains only the point $x(0)$.
- **Periodic orbit:** The orbit through a point x_0 is periodic if there exists a non-zero real T such that $x(x_0, t + T) = x(x_0, t)$ for all t . The period of the orbit is the smallest non-zero value of T for which this is true for a given x_0 .
- **Aperiodic orbit:** An orbit is aperiodic if there exists no pair of $t_1 \neq t_2$ for which $x(t_1) = x(t_2)$. This means the orbit never passes through the same point twice.
- **Relative periodic orbit (RPO):** A more generalised periodic orbit which allows for discrete or continuous symmetry operations. An orbit is an RPO if there exists a well defined combination (T, σ) such that $x(t+T) = G(\sigma)x(t)$, for all t , where the operator $G(\sigma)$ is a



(a) The periodic orbit and Stationary Point in the 2D Van der Pol oscillator governed by Equations 1.48 and 1.49



(b) A relative periodic orbit (blue) in the 3D form of Van der Pol oscillator governed by Equations 1.48-1.50 and its projection onto the plane $Z = 0$ (dashed grey).

Figure 1.9

symmetry operator (translation, rotation etc.) and σ parameterises the symmetry operator (translation length, rotation angle, etc.)

These different types of orbit can be exemplified using the Van der Pol oscillator system [86], which takes a state (X, Y) which lies on the manifold \mathbb{R}^2 , and evolves according to the Van der Pol equations:

$$\dot{X} = Y \quad (1.48)$$

$$\dot{Y} = \mu(1 - X^2)Y - X \quad (1.49)$$

Here μ is a parameter of the model, and is set to $\mu = 0.5$ for this example. The model contains a stationary point at $X = Y = 0$, since here $\dot{X} = \dot{Y} = 0$. It also contains a periodic orbit that encloses the stationary point. These points are shown in Figure 1.9a. If a third direction, Z , is introduced to the Van der Pol system, with the evolution in this direction given by

$$\dot{Z} = 1 + X \quad (1.50)$$

then the periodic orbit becomes instead a **relative periodic orbit**; the orbit never passes through the same point twice, and hence is no longer considered a periodic orbit, yet it maintains the same periodicity as in the 2D case when allowing for translations in the Z -direction. In this case translation by approximately 6.38 in the positive Z direction maps the orbit onto itself. This RPO is shown in Figure 1.9b.

These invariant features of a dynamical system can also be categorised according to their stability. It is straightforward to define a Jacobian for

a given trajectory (periodic or not) starting at x_0 and evolving for t :

$$J^t(x_0)_{ij} = \frac{\partial x_i(t)}{\partial x_{0j}}. \quad (1.51)$$

This Jacobian describes how the local neighbourhood of x_0 expands and contracts along different directions through its eigenvalues and eigenvectors in the time period t . For stationary points, a very small time is chosen $t = \delta t$ and this will give eigenvalues $\Lambda_i = e^{\lambda_i \delta t}$ from which the growth rates λ_i can be extracted. In periodic orbits, we are usually concerned with the expansion *per orbit*, rather than per unit time. For an orbit \mathcal{S} , which has period τ , we choose $x_0 \in \mathcal{S}$ and $t = \tau$ in Equation 1.51 to determine the Jacobian of interest.

The eigenvalues and eigenvectors, $\Lambda^{(i)}$ and $\mathbf{e}^{(i)}$, of the Jacobian determine the stability of the feature. For a given eigenvalue-eigenvector pair, one of three outcomes is possible:

- $|\Lambda^{(i)}| > 1$: perturbations in the direction $\mathbf{u}^{(i)}$ grow. This direction is *unstable*
- $|\Lambda^{(i)}| < 1$: perturbations in the direction $\mathbf{u}^{(i)}$ contract. This direction is *stable*
- $|\Lambda^{(i)}| = 1$: perturbations in the direction $\mathbf{u}^{(i)}$ stay a constant size. This direction is *marginal or a center* and is of particular relevance to symmetries of the system. Stepping in the the direction of the orbit causes one to stay on the same RPO, and after a single orbit return to the initial displaced position

and the overall stability of a feature is stable if all eigendirections are stable, or unstable if at least one is not. All eigenvalues will either be real, or part of conjugate pairs, with the latter case corresponding to a decay spiral into the structure. The stability of the feature is important as it gives a qualitative indication of how likely the system is to wander into a region close to it, and how long it might remain there.

This definition of stability eigenvalues differs from another commonly used definition, which instead uses the Jacobian

$$J_{ij} = \partial f_i / \partial x_j, \quad (1.52)$$

where $\partial_t \mathbf{x} = \mathbf{f}(\mathbf{x})$ is model's evolution equation. The eigenvalues, λ_i of this Jacobian have slightly different properties, with $\Re(\lambda_i) > 0$ indicating

a direction is unstable, and $\Re(\lambda_i) < 0$ indicating a direction is stable. The reason for using the former definition in this work is that the structures of interest are periodic orbits specifically, and so we are more concerned with what happens to a perturbation to an orbit over a full period than what happens at short time.

Another important property of dynamical systems is their symmetry. A transformation g acting on $x \in \mathcal{M}$ is a symmetry of the system if the behaviour of the system is *equivariant* with respect to g , meaning the transformation has no effect on the physics of the system. Mathematically this is expressed as

$$gf^t(x) = f^t(gx), \quad (1.53)$$

so that g commutes with the physics of the system. g can be either a continuous or discrete operator. The 2D Van der Pol system has a discrete rotational symmetry; rotation by π radians about the Z -axis: $g_R = C^{1/2} : (X, Y) \mapsto C^{1/2}(X, Y) = (-X, -Y)$. The 3D version loses this discrete rotational symmetry, however gains a continuous symmetry, translation in the Z -direction. The invariant solutions themselves also have a symmetry group which will be a subset of the system's symmetry group. The 2D VDP system's RPO also has the discrete rotational symmetry, however each of the individual RPOs in the 3D VDP system do not carry the full translational symmetry of the model, only a translation by the precise amount needed to map it back onto itself. The 'lost' symmetry direction then translates RPOs onto separate but related RPOs.

The principles of dynamical systems apply to systems of any dimension. Systems governed by partial differential equations, such as the Navier-Stokes equations for example, can be thought of as having infinite dimension [41] in the sense there are infinitely many degrees of freedom (the model quantities defined at every point in space) and so having a state space with infinite dimensions. In practice, this dimensionality is reduced by translating the PDE equation into an ODE one by only considering the model quantities at set grid points, or by considering amplitudes of a fixed set of modes.

1.8.2 Chaotic Dynamical Systems and Turbulence

The term ‘chaos’ has a variety of definitions depending on the field of application. It is a fundamental concept underlying this thesis, and it is therefore important to outline the definition being used here:

- **Chaotic System:** A deterministic dynamical system containing regions of state space in which all trajectories are unstable (so that small perturbations increase in size exponentially) and exhibit mixing - nearby trajectories locally separate, but a small volume moving with the flow will eventually intersect its initial volume.

The implication of this definition is that trajectories are *highly sensitive to initial perturbations*. Chaotic systems, though deterministic, are in practice highly unpredictable as any measurement of a state will always have some uncertainty which translates to a large uncertainty at late time. It is possible, and in fact exhibited by the plasma interchange model used throughout this thesis, for a system to contain subspaces which are chaotic, and some which are not.

Systems governed by partial differential equations, particularly viscous systems, can exhibit low-dimensional dynamical behaviours where motion is restricted to an inertial manifold of lower dimension embedded within the higher dimensional state space. Provided this manifold has dimensionality greater than three, it is possible that the dynamical behaviour within this manifold is chaotic. Turbulence is one example of this principle. For example, at late time, homogenous turbulence will settle onto a manifold of states which have a spectrum following the inertial range. Coherent structure dominated turbulence will settle onto a manifold of states containing the coherent structures. Note that though not all chaotic systems are turbulent (the Lorenz attractor, for instance), for a behaviour to be considered turbulent, it must be chaotic.

1.8.3 Significance of Relative Periodic Orbits in Studying Turbulence

Relative periodic orbits have been found to be especially useful in studying turbulence in neutral fluids [15, 16, 87, 88]. They have long been known to collectively form a kind of ‘backbone’ or ‘skeleton’ [89] in chaotic systems, about which the more complex trajectories evolve. A representation of this understanding is shown in Figure 1.10. In this picture, a subcritical system (this will be defined in Section 1.8.5) has a

stable equilibrium represented by the origin, an unstable ‘edge of chaos’ (this will be defined in Section 1.8.5) between the basins of attraction of the equilibrium and turbulence, and a set of RPOs embedded in the turbulence about which the turbulence evolves. As discussed above, turbulence can be understood as a chaotic dynamical system with infinite dimensionality - this dimensionality is reduced by projecting the model onto a grid, or onto a set of modes. This relationship allows aspects of dynamical systems theory to be applied to studies of turbulence. In chaotic systems, the number of RPOs is infinite, but for most purposes, the low period orbits carry more weight in calculation than longer orbits and calculating these simpler orbits is sufficient to calculate many properties of turbulence to high order. This can be seen most clearly via the dynamical zeta function formalism, which states that for an observable A one can define a dynamical zeta function $\zeta_A(z, \beta)$ as

$$\zeta_A(z, \beta) = \prod_p \left(1 - \frac{z^{\tau_p}}{\Lambda_p} e^{-\beta A_p} \right)^{-1} \quad (1.54)$$

where the product is over all *prime periodic orbits*, p_i , τ_p is the period of the orbit, and Λ_p is the leading unstable eigenvalue of its Jacobian, and A_p is the average of A over the orbit [90]. Prime periodic orbits are the periodic orbits which form a basis for all other periodic orbits. For example, a periodic orbit following a trajectory $x(t)$ with period τ can be simply repeated n times to give a PO with period $n\tau$, but only the lowest period orbit is a prime periodic orbit. The arguments of the function, z and β , are a complex and real variable respectively and don’t directly represent anything, however the function of the variables encapsulates information about the observable A in the location of its poles. For a given β The DZF will have a pole at some value of z in the complex plane, $z = s_A(\beta)$. The position of this pole changes with β and the rate of change gives the dynamical average (over all chaotic paths) of A for general ergodic orbits,

$$\langle A \rangle = \frac{\partial}{\partial \beta} \ln s_A(\beta). \quad (1.55)$$

Hence, measurable averages in chaotic (turbulent) orbits can be recovered fully from the DZF, in much the same way statistical properties of a thermal system can be recovered from the partition function. Crucially, the DZF can be calculated from the knowledge of all of the prime pe-

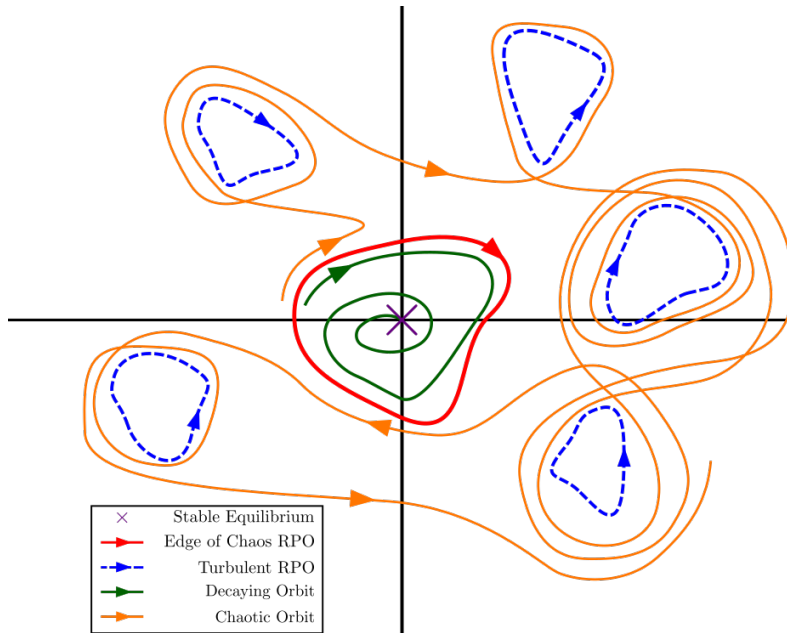


Figure 1.10: 2D representation of a subcritical dynamical system populated by RPOs. A stable equilibrium at the origin, and its basin of attraction, is separated from the turbulent manifold by the edge of chaos, which itself contains an RPO. In the turbulent region of state space there exist several unstable RPOs (blue dashed) and the chaotic orbits (example in orange) can be thought of as moving between the local areas of these orbits. Note that a real dynamical system cannot have orbits which intersect themselves and this diagram is purely to build intuition.

riodic orbits, and furthermore, the full expansion contains terms which are highly convergent with increasing period of the prime periodic orbits. This means that even calculating a handful of the lowest period orbits can give a reasonably accurate picture of the overall behaviour of a chaotic system. This framework can be seen in more detail in [41] and is applied to the Lorenz attractor in [90].

1.8.4 Previous Use of RPO Based Methods in Neutral Fluids

The work in this thesis was inspired by the application of RPO based methods to neutral fluid problems. The earliest application of the methods used in this thesis was seen in Viswanath’s work in 2007 on bursts of turbulence in Couette flow [15], in which the author was able to identify recurrent structures which were able to explain bursting phenomena. Similar work was performed to identify RPOs in a pipe flow scenario by Duguet, Pringle and Kerswell [16, 87, 88]. The Newton-Krylov-Hookstep algorithm developed for these works is the same as the one used in this

paper.

1.8.5 Subcriticality and the Edge of Chaos

The term **subcritical** is often used (see, for example [42]) to describe configurations which are linearly stable, but a suitably large perturbation will be allowed to grow due to nonlinear terms. This perturbation could grow indefinitely in some models, such as the 2D toy model below, or in the case of a fluid model grow until nonlinear saturation kicks in and the system becomes turbulent. An example toy model that demonstrates this behaviour can be seen if one takes the 2-dimensional Van der Pol system described above, Equations 1.48 and 1.49 [86] but flip the sign of the state space velocities, so that the fixed point at the origin is now stable rather than unstable, and the periodic orbit becomes unstable, rather than stable. The model is hence described by the equations

$$\dot{x} = -y \tag{1.56}$$

$$\dot{y} = -\mu(1 - x^2)y - x \tag{1.57}$$

where for this case we set $\mu = 2$. A plot of this system is shown in Figure 1.11 showing in green the domain of initial perturbations which lie in the basin of attraction of the origin and in red, those in the basin of attraction of infinity (or in higher dimensional chaotic models, the basin of attraction of chaos). There is a manifold of codimension one that lies on the boundary between these two basins, known as the ‘edge of chaos’ (EOC) [42] shown in the toy example as a red dotted line. Note that in this simple model the outer basin is not actually chaotic, and perturbations go to infinity rather than saturating as would be the case in a turbulent system - this toy model is deliberately simple to make visualisation more intuitive.

In some higher dimensional subcritical models, including fluid models, the edge of chaos manifold (which has codimension one) has been found to contain RPOs, and in the plasma interchange model (see Chapter 2) these RPOs were found to have a curious qualitative similarity to the coherent propagating structures seen in full turbulence [42]. It was this observation that motivated the work of Pringle et al [42] to investigate the edge of chaos in this system further. The work in this thesis builds upon this work and demonstrates that there is indeed a strong connection

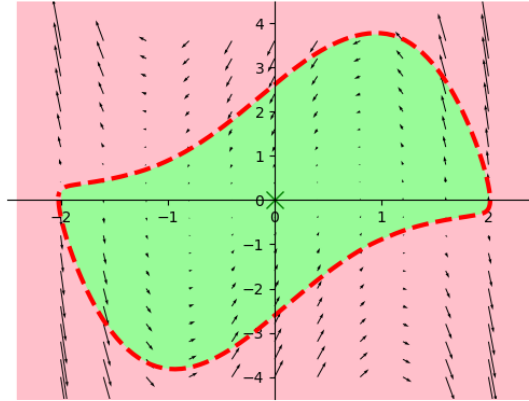


Figure 1.11: Subcriticality in the time reversed Van der Pol oscillator system. The equilibrium point at the origin is stable, with a basin of attraction shown in green. The basin of attraction of infinity is shown in light red, and the boundary between the two basins is a relative periodic orbit, shown in red dashed lines.

between the edge of chaos manifold and the structures seen in turbulence.

1.8.6 Transient Growth

Subcritical configurations are often not considered in descriptions of transition to turbulence, since their linear stability implies that small perturbations will decay, and therefore not be able to grow to a level at which nonlinear terms become important. The reality however is that perturbations can grow in amplitude, even in fully linearly stable systems, via a process called transient growth. This may appear counter intuitive; in a linear system all perturbations are a linear combination of eigenvectors, the coefficients of which decay monotonically and so at first glance it seems the total amplitude must also decay. However, this is not necessarily true. Consider for instance the system governed by the simple linear equation

$$\frac{d}{dt} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{bmatrix} -1 & -18 \\ 0 & -10 \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \quad (1.58)$$




Figure 1.12: Left: A path (blue) taken by a two-dimensional linear system governed by equation 1.58, demonstrating that it undergoes a period of transient growth before decaying. The eigenvectors are shown in red along with a grid aligned to their directions. **Right:** The distance from the origin of the orbit shown to the left against time.

This is straightforwardly solved to get solutions of the form

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = a_1(t)\vec{v}_1 + a_2(t)\vec{v}_2 \quad (1.59)$$

$$= a_1(0)e^{-10t} \begin{pmatrix} 2/\sqrt{5} \\ 1/\sqrt{5} \end{pmatrix} + a_2(0)e^{-t} \begin{pmatrix} 1 \\ 0 \end{pmatrix} . \quad (1.60)$$

Clearly each term individually decays exponentially, however it is still possible for the total amplitude to grow before decaying, as demonstrated with the example in Figure 1.12 with the initial conditions $x(0) = 0, y(0) = 1$. This is possible only when the eigenvectors are not orthogonal. It is this transient growth process that allows small perturbations to grow enough that linear stability analysis can fail in some scenarios, such as subcritical fluid configurations.

1.8.7 Bifurcations in Dynamical Systems

When a dynamical system is governed by equations including a model parameter, it is possible for equilibrium solutions and periodic orbits to be created, destroyed, split or merged as the model parameter changes, changing the nature of the system. The points at which these changes in behaviour occur are called **bifurcations** [91], and they play a critical role in dynamical systems theory, particularly with regards to the transition to chaos. There are many types of bifurcation that can occur, but

bifurcations of almost all types are associated with a change in stability of an equilibrium or periodic orbit. For example, the time-reversed Van der Pol system described in Equations 1.48 and 1.49 is controlled by a parameter μ . For $\mu < 0$ there is a single solution which is non-aperiodic, $X(t) = Y(t) = 0$, and this solution is stable. For $\mu > 0$, the stability of the origin solution changes, and is now unstable. In addition, a new stable periodic solution is present. The change in behaviour at $\mu = 0$ is associated with a bifurcation.

There are four bifurcation types that show up throughout this thesis, which will be described below. Note that in these descriptions, the stability eigenvalues are those defined in Section 1.8.1 corresponding to the Jacobian defined in Equation 1.51 and differ from the other frequently used definition which uses the eigenvalues of the Jacobian defined in Equation 1.52 and the specific stability conditions on the eigenvalues may differ from those the reader is familiar with, though are ultimately equivalent.

Period Doubling Bifurcation

A stationary point or periodic orbit undergoes a period doubling bifurcation when one of its stability eigenvalues passes through -1 as the model parameter is changed. Taking a perturbation from the original orbit in the corresponding eigendirection will give an orbit with double the original period, since the perturbation will become the negative of itself in a single period, and return to its original value after two periods. The typical example model of the period doubling is the logistic map [92] which is a mapping

$$x_{n+1} = rx_n(1 - x_n) \tag{1.61}$$

that repeatedly demonstrates period doubling bifurcations as r is increased. For $r < 3$ there is a single, stable, stationary solution. For $3 < r \lesssim 3.4495$ the previously stable solution is unstable and a new solution with period 2 exists. For $3.4995 \leq r \lesssim 3.5541$ a period four solution appears, and this trend continues with the period of the orbit doubling at successively closer values of r . Successive period doubling bifurcations such as this are known as period doubling cascades and lead to increasingly chaotic behaviour in a system. The points which are part of the periodic orbits in the logistic map are plotted against r in Figure 1.13a.

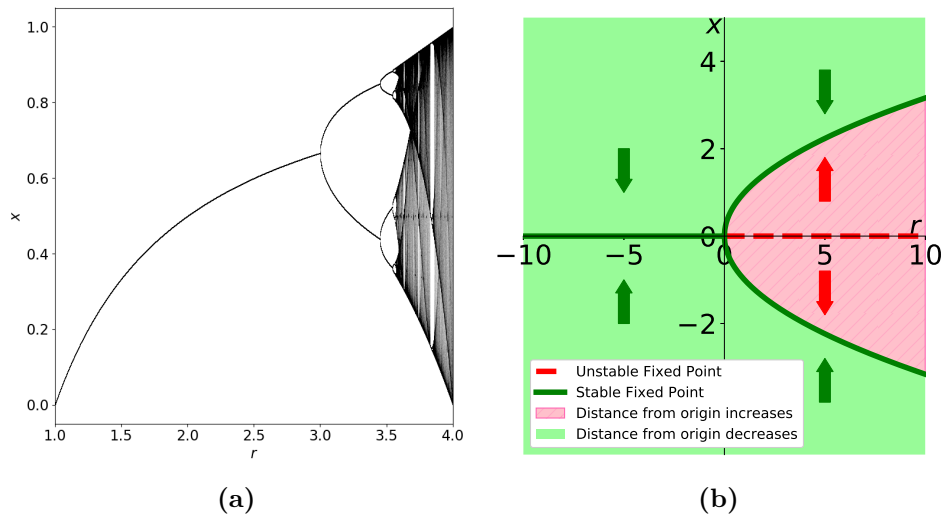


Figure 1.13: (a) Logistic map: x value of Points which belong to periodic orbits in the logistic map plotted against the model parameter r . Period doubling occurs where the branch splits in two. **(b) Pitchfork bifurcation:** Representation of the system described by Equation 1.62. The position x is plotted against the model parameter r , with fixed points shown by lines and the distinct behavioural regions of system space by colour. The pitchfork bifurcation occurs at $r = 0$.

Pitchfork Bifurcation

A pitchfork bifurcation occurs when a stationary or periodic orbit has a stability eigenvalue which passes through 1, either from above or below, and leads to two more solutions being born. The canonical example of a pitchfork bifurcation is the system governed by the equation

$$\frac{dx}{dt} = rx - x^3 \quad (1.62)$$

where r is the model parameter. For $r < 0$ a single stationary solution exists at $x = 0$, with stability eigenvalue $\Lambda = 1 + r < 1$ and hence stable. As r passes through zero from below, this eigenvalue passes through one from below, so that the $x = 0$ solution becomes unstable. In addition, the model equation admits two more stationary solutions at $x = \pm\sqrt{r}$, which are stable. This example is represented in Figure 1.13b - the horizontal axis represents the model parameter r , and the vertical axis represents x . The fixed points are shown as lines with colour representing the stability (green is stable, dashed red is unstable) of the fixed point, and the non-fixed points are coloured based on whether the system evolves towards (green) or away from (red) $x = 0$.

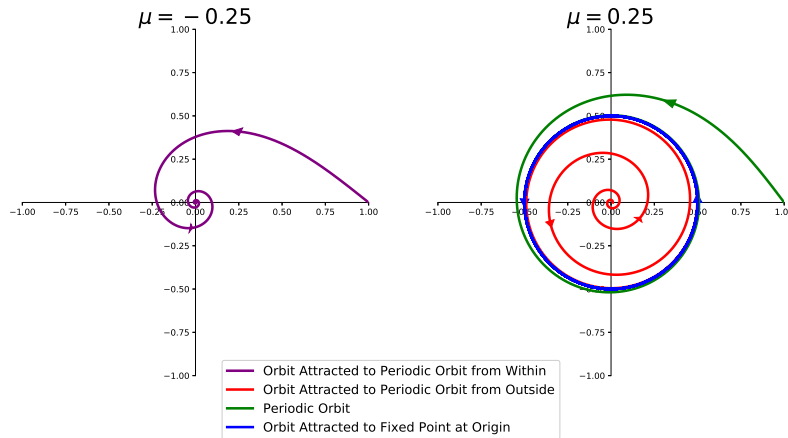


Figure 1.14: Example orbits in the canonical Hopf bifurcation system, Equations 1.63 and 1.64, at $\mu = -0.25$ (left) and $\mu = 0.25$ (right). In the system with $\mu > 0$ a periodic orbit attractor is present and the origin is an unstable fixed point.

Hopf Bifurcation

A Hopf bifurcation occurs when a periodic or stationary solution has a complex eigenvalue pair whose magnitude passes through one. This leads to the appearance of a periodic solution which oscillates around the original solution at a frequency dependent on the angle of the complex eigenvalue pair. The canonical form of a Hopf bifurcation is a system in polar coordinates with evolution equations

$$\frac{d\omega}{dt} = \omega \tag{1.63}$$

$$\frac{dr}{dt} = r(\mu - r^2). \tag{1.64}$$

For $\mu < 0$, the origin is a stable fixed point, with two stable directions (e.g, the x and y directions). For $\mu > 0$ the origin becomes unstable and a stable periodic solution at constant $r = \sqrt{\mu}$ exists. This system is shown in Figure 1.14 for $\mu < 0$ and $\mu > 0$.

Saddle Node Bifurcation

A saddle node bifurcation differs from the previously described bifurcations in that there is no direct change in stability as the model parameter is changed, and on one side of the bifurcation no solution exists at all. The saddle node bifurcation occurs when two fixed points/periodic orbits with different stability approach each other in system space as the parameter changes, ultimately colliding with and annihilating each other.

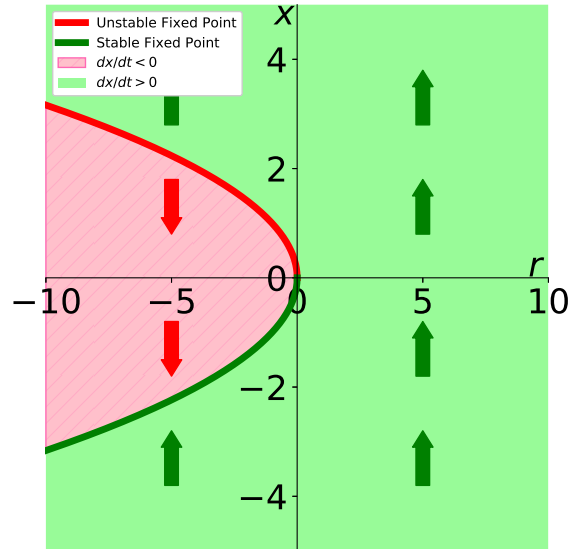


Figure 1.15: Representation of the system described by Equation 1.65. The position x is plotted against the model parameter r , with fixed points shown by lines and the distinct behavioural regions of system space by colour. The saddle node bifurcation occurs at $r = 0$.

As the solutions approach each other, one of their stability eigenvalues approaches one from below, the other from above, with the corresponding eigendirections being exactly aligned at the bifurcation point as they become the same solution. The canonical model for a saddle node bifurcation is

$$\frac{dx}{dt} = r + x^2 \quad (1.65)$$

which has two stationary solutions, $x_{\pm} = \pm\sqrt{-r}$, for $r < 0$, and no stationary solutions for $r > 0$. The stability eigenvalues for the two solutions are $\Lambda_{\pm} = 1 \pm 2\sqrt{-r}$, with the eigenvalue for each approaching unity as $r \rightarrow 0^-$. This system is shown in Figure 1.15.

1.9 Summary Of Key Ideas and Thesis Outline

This introduction has introduced several concepts which may be new to the reader and have been introduced as disparate concepts, so here these concepts will be brought together in a succinct summary.

This work aims to convince the reader that **nonlinear dynamical systems theory** has the potential to be applied in plasma physics con-

texts to unlock new understanding of **plasma turbulence**, while also demonstrating some interesting results the approach has generated in the context of turbulence containing **coherent structures**. One such coherent structure is a radially propagating ‘burst’ of turbulence which has approximately constant velocity and amplitude and are localised, and has been seen in simulations of ITG turbulence. These structures have previously been described as **avalanches in self organised criticality**, however it has been argued [39] that these models do not accurately describe the phenomenology of coherent structures seen in ITG turbulence in gyrokinetic simulations. A simpler 1D model known as the **plasma interchange model** has been proposed as a minimal model for describing this turbulence. Chapter 2 describes the plasma interchange model in detail, and describes the spectral methods used to implement it computationally. The beginning of Chapter 3 explores the form turbulence takes in the PI model.

The plasma interchange model incorporates a background shear flow, which acts to shear perturbations and relaminarise them, so that the system is subcritical. A **subcritical** system is one which is linearly stable, but has some nonlinear instability so that perturbations which are large enough will be unstable and grow, either indefinitely or until saturating. Subcritical systems feature an **edge of chaos** which separates the perturbations in state space which are small enough to stabilise from those which grow due to nonlinearities. States lying precisely on the edge of chaos will neither stabilise nor grow nonlinearly - instead they will lie in the edge forever. A system governed by partial differential equations (such as the PI model) can be thought of as an infinite-dimensional dynamical system, or a finite- but high-dimensional system if examined on a discrete grid. In the second half of Chapter 3, the **bisection method** will be applied to locate the edge of chaos in the PI model.

A **relative periodic orbit** (RPO) is a generalisation of a periodic orbit which allows for symmetry operations such as translation. In a PDE-described system, an RPO could take many forms, including many coherent structures such as travelling waves, oscillations, vortices, and breathers. Studying these structures by thinking of them as RPOs will allow us to isolate nonlinear structures, identify new ones, and explore the transition to turbulence through their **bifurcations**.

In the PI model, the EOC has previously been found to contain stable (aside from the instability in the direction out of the edge) RPOs which exhibit similar phenomena (travelling waves with constant ampli-

tude and velocity) to the coherent structures in turbulence, and so it has previously been hypothesised that the edge of chaos RPOs and the coherent structures seen in plasma turbulence could somehow be related. In Chapter 4, RPOs found in the EOC will provide a starting point to search for other RPOs not in the EOC, and it will be seen that the EOC RPOs and RPOs embedded in turbulence are in fact connected. In addition, a series of bifurcations builds increasing complexity in the transition to turbulence, from simple travelling structures, to turbulence dominated by coherent structures.

Chapter 2

The PI Model Initial Value Solver

In this chapter, we introduce the plasma interchange model, discuss its properties, and the pseudospectral methods and modifications used to implement a background shear flow into a code implementation of the model.

2.1 Plasma Interchange Model

As discussed in Chapter 1, a broad collection of models have been used in attempts to describe the propagating structures in plasma turbulence, each selecting different physics to highlight, and indeed physics to neglect in order to simplify the model. The judgement on which physics to include or exclude depends on the phenomena being investigated, and what kind of investigation is being performed. Gyrokinetic models are used to provide the most true-to-life simulations, with great complexity and computational cost. The complexity of gyrokinetic models can make it difficult to understand the physics behind the phenomena which are observed and so simpler models are constructed in an attempt to identify which physics is enough to construct a ‘minimal model’ of the phenomena at hand.

Gyrokinetic simulations in the collisionless, electrostatic, adiabatic electron limit in a flux tube geometry [39] showed that travelling wave-like solutions continue to exist in these limits, and that the travelling waves observed are all of very similar amplitude and velocity in a soliton-like manner. These observations differed from the characteristic avalanche structures of SOC models, such as that of Reference [13], which carry a

range of velocities based on their amplitudes. In addition, the travelling waves in the gyrokinetic simulations had the same values at positive and negative infinity, indicating they are localised, while SOC models do not. The authors conclude that the soliton-like travelling waves in ITG turbulence cannot be explained as SOC avalanches and argue that this is due to the dependence of these structures on local physics, while the SOC based models such as that in reference [13] use coarse-graining procedures to model transport on scales larger than the scales of local turbulent fluctuations. They propose that a slightly more complex model that incorporates local physics is required to capture accurate phenomenology, and suggest that a reduced 1-dimensional form of the plasma interchange (PI) model derived in [51] is a suitable candidate.

2D Plasma Interchange Model

The 2D PI model [42] was originally derived from the resistive MHD equations by applying the flute approximation [82], which assumes wavenumbers in the direction perpendicular to the field are much larger than in the parallel direction, as is the case in the interchange instability. The 2D PI model was originally used to describe interchange driven turbulence in the scrape off layer of a tokamak [13], as well as studying propagating structures in the presence of a shear flow [51]. The coordinates for the PI model is similar to the flux tube setup discussed in Section 1.2.1, but there is no shear in the magnetic field and the model is fully 2D so there is no dependence on z . The 2D PI model equations are expressed in terms of the fluid density n and the vorticity $\nabla \times \mathbf{u} = \nabla \times (\mathbf{E} \times \hat{b}) = \nabla_{\perp}^2 \phi$:

$$\frac{dn}{dt} = D\nabla^2 n \quad (2.1)$$

$$\frac{d(\nabla^2 \phi)}{dt} = -\frac{g}{n} \frac{\partial n}{\partial y} + \nu \nabla^4 \phi. \quad (2.2)$$

where $\frac{d}{dt}$ is the convective derivative ($\frac{\partial}{\partial t} - \nabla \phi \times \hat{z} \cdot \nabla$). The model is effectively a 2D fluid model with an additional buoyancy term, and though derived from resistive MHD, is very generic in representing curvature driven instabilities, and is also similar to equations for drift wave instabilities such as the Hasegawa Wakatani equation, (Equation 1.42), and is able to capture some of the crucial mechanisms of drift wave-zonal flow interactions: the instability of drift waves; Reynolds stress generation of zonal flows; back-reaction of zonal flows on zonal perturbations; and saturation of zonal flows. The ability of the model to capture these generic be-

haviours makes it a good candidate for a minimal model for investigating coherent structures in drift wave turbulence, as well as interchange turbulence. A shear flow can be implemented in the 2D PI model by imposing shear-periodic boundary conditions: $\phi(x, y, t) = \phi(x + L, y + SLt, t)$ and an equivalent relation for n .

1D Plasma Interchange Model

In microinstability driven turbulence, the drift wave k_y spectrum is very narrow and can be approximated by keeping just a single wavenumber [39, 51] with $k_y = \pm k_{y0}$, the most unstable mode, along with the $k = 0$ zonal mode. The 1D PI model is derived by projecting the 2D PI model equations onto the reduced space spanned by the basis consisting of these three modes and solving the equations in this reduced space. This process is known as Galerkin projection [93]. The inclusion of both the positive and negative k_{y0} modes retains zonal flow generation by the coupling between them. The density and electric potential can thus be written as a sum of a complex wave field, represented by a tilde, and a zonal quantity, represented by a bar

$$\begin{aligned} n(x, y, t) &= \bar{n}(x, t) + \tilde{n}(x, t)e^{ik_{y0}y} + \tilde{n}^*(x, t)e^{-ik_{y0}y} \\ \phi(x, y, t) &= \bar{\phi}(x, t) + \tilde{\phi}(x, t)e^{ik_{y0}y} + \tilde{\phi}^*(x, t)e^{-ik_{y0}y} \end{aligned}$$

and this reduces the 2D PI model equations, 2.1 and 2.2, to a 1D model of the four quantities:

$$\partial_t \tilde{n} = -ik_{y0} \tilde{n} \partial_x \bar{\phi} + ik_{y0} \tilde{\phi} \partial_x \bar{n} + D \partial_{xx} \tilde{n} \quad (2.3)$$

$$\partial_t \tilde{\phi} = -ik_{y0} \tilde{\phi} \partial_x \bar{\phi} + \frac{ig\tilde{n}}{k_{y0}n_0} + \nu \partial_{xx} \tilde{\phi} \quad (2.4)$$

$$\partial_t \bar{n} = ik_{y0} \partial_x (\tilde{n}^* \tilde{\phi} - \tilde{n} \tilde{\phi}^*) + D \partial_{xx} \bar{n} \quad (2.5)$$

$$\partial_t E = ik_{y0} \partial_x (\tilde{\phi} \partial_x \tilde{\phi}^* - \tilde{\phi}^* \partial_x \tilde{\phi}) + \nu \partial_{xx} E \quad (2.6)$$

where the substitution $E \equiv \partial_x \bar{\phi}$ has been applied for conciseness. Beware that unlike some conventions, here the bar quantities refer to the zonal $k_y = 0$ mode amplitudes and not to an average over the domain. In the 2D model, n and ϕ are real quantities, so zonal components $\bar{n}(x)$ and $E(x)$ must be real, while wave terms \tilde{n} and $\tilde{\phi}$ can be complex. Following McMillan et al [39] we here set the diffusion of the zonal density \bar{n} to zero on the basis that tokamak transport is in the collisionless limit. Diffusion is retained for $\bar{\phi}$ to act as a simple saturation mechanism for the

zonal flows which otherwise build up excessively. The drift mode diffusion terms are kept as a simple implementation of k_x -dependent growth rates, as the $\partial_{xx}\tilde{n}$ term becomes $-k_x^2\tilde{n}_k$ in Fourier space. The removal of the \bar{n} diffusion has the effect of making turbulent flux dominant over diffusive flux, as is seen in tokamaks and gyrokinetic simulations. Assuming the equilibrium density n_0 has a constant negative gradient (this is valid since the scales of the turbulence are much shorter than the global profiles, and a negative gradient is chosen as this is necessary for the interchange instability to even occur) the parameters of the model can be reduced by rescaling the x coordinate, time and the quantities. Next, we nondimensionalise the equations through the transformations

$$\begin{aligned} t &\rightarrow \tau t = \sqrt{\frac{n_0}{g\partial_x n_0}} t, & x &\rightarrow x_0 x = \sqrt{D\tau} x, \\ E &\rightarrow \frac{1}{\tau k_{y0}} E, & \tilde{\phi} &\rightarrow \tilde{\phi}_0 \tilde{\phi} = \sqrt{\frac{x_0^2}{k_{y0}^2 \tau^2}} \tilde{\phi}, \\ \tilde{n} &\rightarrow \frac{k_{y0} n_0}{\tau g} \tilde{\phi}_0 \tilde{n}, & \bar{n} &\rightarrow x_0 \partial_x n_{eq} \bar{n}, \end{aligned}$$

where n_{eq} is the equilibrium density. The model timescale is the Brunt-Väisälä time; the time associated with convection oscillations. The length scales are the distance of diffusion in a time unit for x and the drift wave wavelength, $1/k_{y0}$ in y . In dimensionless form, the model equations are

$$\partial_t \tilde{n} = -i\tilde{n}E + i\tilde{\phi}(\partial_x \bar{n} - 1) + \partial_{xx} \tilde{n} \quad (2.7)$$

$$\partial_t \tilde{\phi} = -i\tilde{\phi}E + i\tilde{n} + \partial_{xx} \tilde{\phi} \quad (2.8)$$

$$\partial_t \bar{n} = i\partial_x (\tilde{n}^* \tilde{\phi} - \tilde{\phi}^* \tilde{n}) \quad (2.9)$$

$$\partial_t E = i\partial_x (\tilde{\phi} \partial_x \tilde{\phi}^* - \tilde{\phi}^* \partial_x \tilde{\phi}) + \nu' \partial_{xx} E \quad (2.10)$$

where $\nu' = \nu/D$ remains the only model parameter. ν' controls the size of zonal flows compared to the diffusion length. In Reference [39] it was discussed that for zonal flows to be generated at all $\nu' \lesssim 8$, and in the limit $\nu' \ll 1$ zonal flows become narrower than the wave diffusion length and bursts do not show up at all. In the limit $\nu' \gg 1$ the model still exhibits turbulence consisting of coherent structures however there is no generation of zonal flows limiting the applicability of the model in this limit. The authors decide to select $\nu' = 1$ as an intermediate ordering that was found post hoc to reflect characteristic behaviour of ITG turbulence, and the same selection is made here.

These equations are solved in a periodic box $x \in [0, L]$. To investigate

the relationship between shear flows and turbulence, a background shear flow is applied by imposing a linear gradient in E . This gradient in E corresponds to a gradient in $\mathbf{E} \times \mathbf{B}$ flow, and hence a large scale linear shear flow. This is implemented by rewriting $E \rightarrow Sx + E(x)$ where S measures the strength of the shear and becomes a parameter of the model. The value of S is taken to be constant; there is no feedback into S from the model equations, so the model cannot be used to study the generation of large scale shear flows and phenomena such as the L- to H-mode transition, but rather is intended to study the impact shear flows have on the generation and behaviour of coherent structures. Shorter scale zonal flows can however still be represented through $E(x)$. This establishes a background flow shear parameterised by S , and changing S has a large impact on the behaviour of the system. In order to maintain periodicity, the boundary conditions are modified to maintain a continuous gradient in E at the boundaries. The adjusted boundary conditions are $\tilde{n}(L) = \tilde{n}(0)e^{-iStL}$, and $\tilde{\phi}(L) = \tilde{\phi}(0)e^{-iStL}$. The boundary conditions are therefore time-dependent, and repeat with a period

$$\tau_{BC} = \frac{2\pi}{LS}. \quad (2.11)$$

The final PI model equations are finally written:

$$\partial_t \tilde{n} = -i\tilde{n}(Sx + E) + i\tilde{\phi}(\partial_x \bar{n} - 1) + \partial_{xx} \tilde{n} \quad (2.12)$$

$$\partial_t \tilde{\phi} = -i\tilde{\phi}(Sx + E) + i\tilde{n} + \partial_{xx} \tilde{\phi} \quad (2.13)$$

$$\partial_t \bar{n} = i\partial_x (\tilde{n}^* \tilde{\phi} - \tilde{\phi}^* \tilde{n}) \quad (2.14)$$

$$\partial_t E = i\partial_x (\tilde{\phi} \partial_x \tilde{\phi}^* - \tilde{\phi}^* \partial_x \tilde{\phi}) + \partial_{xx} E \quad (2.15)$$

From a dynamical system perspective, the state space of the system is the infinite dimensional direct sum of all possible functions each quantity can take:

$$\mathcal{M}_{\mathcal{I}, \mathcal{L}} = \{[\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \mid \tilde{n}, \tilde{\phi} \in \mathcal{F}(X, \mathbb{C}) \text{ and } E, \bar{n} \in \mathcal{F}(X, \mathbb{R})\}$$

where X is the set of positions in the box.

$$X = \{x \in \mathbb{R} \mid 0 \leq x \leq L\}$$

For conciseness, at times solutions will be represented by notation of the form $\mathbf{U}(x, t) = [\tilde{n}(x, t), \tilde{\phi}(x, t), \bar{n}(x, t), E(x, t)]$ throughout this thesis.

2.2 Symmetries of the PI Model

The symmetry properties of a model provide crucial context for understanding its physics and solutions, particularly in the field of nonlinear dynamical systems which is used throughout this thesis. An operation is a symmetry of the system if the equations of the system are *equivariant* under the operation. The PI model has four key symmetries which are outlined below.

Discrete \mathbb{Z}_2 Symmetry

The first symmetry of the system is a discrete symmetry operator belonging to the group \mathbb{Z}_2 . The operator reflects the system so $x \rightarrow L - x$. On its own, this reflection is not sufficient to be a symmetry of the system, however and the full operator acts on a state as:

$$Q_Z : [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \rightarrow [-\tilde{n}^*(L - x), -\tilde{\phi}^*(L - x), -\bar{n}(L - x), -E(L - x)] \quad (2.16)$$

This can easily be verified to be a symmetry by plugging this transformation into Equations 2.12 to 2.15 and noting that the equations are consistent with the original ones.

Continuous Phase Symmetry

The wave quantities \tilde{n} and $\tilde{\phi}$ only enter the model equations for the averaged quantities through products of conjugated and non conjugated wave quantities. This means there is a degree of freedom in the choice of phase of these quantities and the operator

$$Q_\theta(\delta\theta) : [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \rightarrow [\tilde{n}(x)e^{i\delta\theta}, \tilde{\phi}(x)e^{i\delta\theta}, \bar{n}(x), E(x)] \quad (2.17)$$

leaves the PI equations invariant.

Translation Pseudosymmetry

Translation symmetry is broken by the presence of the shear flow. The term $xS\tilde{\phi}$ for example has explicit dependence on position, causing the equation to no longer be equivariant under translation. However, the effects of the shear on translation are only seen in the term that rotates

the wave quantities in the complex plane, and so translation corresponds to changing the rate at which the phase of the waves changes, but leaves the amplitudes and, crucially, the relative phase of the two wave quantities unchanged. This leaves the dynamics of the system unaffected through the same argument as in the discussion about phase symmetry above. Expressed differently, if $u(x, t)$ is a solution, the translated solution $u'(x, t) = u(x - \delta x, t)$ is not a solution since the phase of the wave quantities, \tilde{n}' and $\tilde{\phi}'$ do not behave correctly. These phases can be corrected by operating on the translated system with a phase rotation operator, and it transpires that $Q_\theta(-S \delta x t)u(x - \delta x)$ is a solution. This modified transformation cannot itself be called a ‘symmetry’ due to the time dependence, but it does allow a new solution to be obtained from a known one by translating it. We call this type of operation a ‘pseudosymmetry’. We define the translation operator

$$Q_{dx}(\delta x) : [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \rightarrow [\tilde{n}(x - \delta x), \tilde{\phi}(x - \delta x), \bar{n}(x - \delta x), E(x - \delta x)] \quad (2.18)$$

Global E Increase

As in the case of translation, increasing $E(x)$ by an additive constant δE also only changes the phase rotation speed of a solution by a constant. Thus, another pseudosymmetry operator is defined by:

$$Q_E(\delta E) : [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \rightarrow [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x) + \delta E]. \quad (2.19)$$

As with translation, the global E increase takes a known solution $u(x, t)$ and generates another solution $u'(x, t) = Q_\theta(-\delta E t)Q_E(\delta E)u(x, t)$.

For certain parameters, the combination of translation and E increase does leave the equations invariant creating an exact symmetry. Under operators $Q_{dx,E}(\delta x) = Q_{dx}(\delta x)Q_E(-S \delta x)$ the PI model is exactly equivariant and so $Q_{dx,E}$ is a symmetry of the system.

Global Increase in \bar{n}

Since only the gradients of the zonal density \bar{n} are present in the PI model equations, a global uniform increase in $\bar{n}(x)$ has no effect on the

physics of the system. There is therefore a symmetry operator

$$Q_n(\delta\bar{n}) : [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x), E(x)] \rightarrow [\tilde{n}(x), \tilde{\phi}(x), \bar{n}(x) + \delta\bar{n}, E(x)]. \quad (2.20)$$

In total, the PI model has a single class of discrete symmetry, Q_Z , three classes of continuous symmetry Q_θ , $Q_{dx,E}$, and Q_n , and two pseudosymmetric classes Q_{dx} , and Q_E .

2.3 Spectral PI Model Code

The implementation of the PI model computationally is complicated by the presence of the background shear flow, which has the physical effect of shearing perturbations so that their wavenumbers change in time, and raises the computational issue of ‘shear-periodic’ boundary conditions; the boundary conditions that evolve with time according to the background shear. In previous work [39, 42] a finite difference code was used to simulate the PI model, however for this work a pseudospectral method was used as it exhibits exponential convergence [94] and hence allows increased accuracy compared to finite difference methods for equivalent computation time, and this improved accuracy is essential for the analyses performed in this work, due to the chaotic nature of the simulations being performed. The pseudospectral method is however more complex to implement than a straightforward finite difference code, particularly in this scenario due to the shear flow and nonlinearities and this is the first time the specific approaches used here have been applied to a shear-periodic model. While pseudospectral methods are very standard and well documented, an overview of the methods used is outlined here to demonstrate how the complications of the model were resolved.

The key idea behind the pseudospectral method is to solve a differential equation so that it is satisfied at a series of grid points, called the collocation grid. If there are N grid points, a solution $\mathbf{U}_N(x, t)$ is found, which will be an approximation to the true solution. As N increases, this approximation will improve and $\mathbf{U}_N(x, t)$ will converge on the true solution $\mathbf{U}^*(x, t)$. This differs from a regular spectral method only in how the nonlinear terms are treated, a process that will be described in detail below, and is chosen due to the relative efficiency that can be achieved through the use of the Fast Fourier Transform (FFT) algorithm.

2.3.1 Basis Functions

First, some Fourier-like basis functions are defined. These are used due to the periodicity of the boundary conditions so that the real quantities are assumed to be of the form:

$$\bar{n}(x) = \sum_{j=-\infty}^{\infty} \bar{n}_j e^{ik_j x} \quad (2.21)$$

$$E(x) = \sum_{j=-\infty}^{\infty} E_j e^{ik_j x} \quad (2.22)$$

where $k_j = 2\pi j/L$. The complex quantities have a modified form which satisfies the shear periodic boundary conditions

$$\tilde{n}(x) = \sum_{j=-\infty}^{\infty} \tilde{n}_j e^{i(k_j - St)x} \quad (2.23)$$

$$\tilde{\phi}(x) = \sum_{j=-\infty}^{\infty} \tilde{\phi}_j e^{i(k_j - St)x} \quad (2.24)$$

In this form, it is clear that one interpretation of the background shear is that it has the effect of continuously evolving the wavenumbers of the Fourier modes of the complex quantities. For example, in a shear of $S = 0.5$, at $t = 0$ the basis function wavenumbers will be $k = \dots, -4\pi/L, -2\pi/L, 0, 2\pi/L, 4\pi/L, \dots$ while at $t = 1$, the basis function wavenumbers will be $k = \dots, -(4\pi/L) - 0.5, -(2\pi/L) - 0.5, -0.5, (2\pi/L) - 0.5, (4\pi/L) - 0.5, \dots$. We can introduce an effective wavenumber $k_{j,eff}(t) = k_j - St$ which is used as a shorthand, but it is important to note that this effective wavenumber is time dependent.

It is straightforward to calculate the coefficients of these series given the quantities as functions of the real space coordinate x . For the zonal averaged quantities one can use a Fourier transform, and for the complex quantities one takes a Fourier transform after premultiplying by e^{iStx} , i.e

$$\tilde{n}_j = \mathcal{F}[e^{iStx} \tilde{n}(x)] \quad (2.25)$$

The time dependence of the basis functions is unconventional and not typical of pseudospectral methods, however as shown below, this time dependence is ultimately eliminated from the equations for the coefficients, and so this approach remains suitable.

Computationally, we want to simulate these quantities at discrete

points on a quadrature grid rather than across the whole domain. In Fourier space this equates to only keeping a finite number of modes. If, in real space one has N grid points, then N spectral modes should be kept to maintain the same level of information when moving to Fourier space. There are different conventions that can be used to select these k -modes, but we use the range $-\pi N/L \leq k < \pi N/L$. This choice is made as the diffusion terms will rapidly reduce high wave modes and so by keeping this band centred on zero we maximise the amount of useful information retained as time is evolved. Since we are working with a discrete grid we can work with the very efficient ‘Fast Fourier Transform’ algorithm [95] (FFT) to convert between the real components at each point and the Fourier coefficients. The FFT was implemented using the subroutine library FFTW [96].

By inserting the series form of each quantity into the first PI model Equation 2.12, we get

$$\partial_t \sum_j \bar{n}_j e^{ik_j x} = i \partial_x \left[\sum_{l,m} \tilde{n}_l^* \tilde{\phi}_m e^{i(k_m - k_l)x} - \tilde{n}_l \tilde{\phi}_m^* e^{i(k_l - k_m)x} \right] \quad (2.26)$$

$$= - \sum_{j,l} (k_j \tilde{n}_l^* \tilde{\phi}_{l+j} - k_j \tilde{n}_l \tilde{\phi}_m^*) e^{ik_j x} \quad (2.27)$$

$$\partial_t \bar{n}_j = - \sum_{l=-\infty}^{\infty} k_j (\tilde{n}_l^* \tilde{\phi}_{l+j} - \tilde{n}_l \tilde{\phi}_{l-j}^*). \quad (2.28)$$

Following a similar process for Equation 2.13 gives

$$\partial_t E_j = - \sum_{l=-\infty}^{\infty} k_j (-i(k_{l-j} - St) \tilde{\phi}_l \tilde{\phi}_{l-j}^* - i(k_{j+l} - St) \tilde{\phi}_l^* \tilde{\phi}_{j+l}) - k_j^2 E_j \quad (2.29)$$

$$= i \sum_{l=-\infty}^{\infty} (k_j(k_{l-j} - St) \tilde{\phi}_l \tilde{\phi}_{l-j}^* + (k_l - St) \tilde{\phi}_{l-j}^* \tilde{\phi}_l) - k_j^2 E_j \quad (2.30)$$

$$= i \sum_{l=-\infty}^{\infty} k_j (k_{2l-j} - 2St) \tilde{\phi}_{l-j}^* \tilde{\phi}_l - k_j^2 E_j. \quad (2.31)$$

For the wave potential and wave density quantities, one must remember to perform the full derivative on the LHS, and this cancels with the shear term on the RHS.

LHS	Linear part (L)	Nonlinear part (N)
$\partial_t \tilde{n}_j$	0	$-\sum_l (k_j)(\tilde{n}_l^* \tilde{\phi}_{l+j} - \tilde{n}_l \tilde{\phi}_{l+j}^*)$
$\partial_t E_j$	$-k_j^2 E_j$	$i \sum_{l=-\infty}^{\infty} k_j (k_{2l-j} - 2St) \tilde{\phi}_{l-j}^* \tilde{\phi}_l$
$\partial_t \tilde{\phi}_j$	$i\tilde{n}_j - k_j^2 \tilde{\phi}_j$	$-i \sum_l (\tilde{\phi}_l E_{j-l})$
$\partial_t \tilde{n}_j$	$-i\tilde{\phi}_j - k_j^2 \tilde{n}_j$	$i \sum_l (-\tilde{n}_l E_{j-l} + \tilde{\phi}_l k_{j-l} \tilde{n}_{j-l})$

Table 2.1: Linear and nonlinear parts of the time evolution equations for the Fourier components of the PI variables

$$\begin{aligned} \sum_j (\partial_t \tilde{\phi}_j - iSx \tilde{\phi}_j) e^{i(k_j - St)x} &= \sum_j (-i\tilde{\phi}_j Sx) e^{i(k_j - St)x} \\ &- i \sum_{j,l} (\tilde{\phi}_l E_{j-l}) e^{i(k_j - St)x} + \sum_j (i\tilde{n}_j - k_j^2 \tilde{\phi}_j) e^{i(k_j - St)x} \end{aligned} \quad (2.32)$$

$$\partial_t \tilde{\phi}_j = -i \sum_l (\tilde{\phi}_l E_{j-l}) + i\tilde{n}_j - k_j^2 \tilde{\phi}_j \quad (2.33)$$

and similarly for the \tilde{n} equation

$$\partial_t \tilde{n}_j = i \sum_l (-\tilde{n}_l E_{j-l} + \tilde{\phi}_l k_{j-l} \tilde{n}_{j-l}) - i\tilde{\phi}_j - k_j^2 \tilde{n}_j. \quad (2.34)$$

We now split the equations for each quantity, into a linear part, $L(x, t)$ and nonlinear part, $N(x, t)$ which are shown in Table 2.1. The treatment of linear and nonlinear terms is very different, and a mixture of explicit and implicit methods are applied to maintain code stability.

2.3.2 Hybrid Timestepping

The timestepping code uses a combination of an implicit Crank-Nicholson scheme for linear terms and a half time-step method for nonlinear terms. The implicit procedure removes the CFL restriction on time-stepping and maintains numerical stability. The Crank-Nicholson scheme was chosen because of the second-order convergence in time. The Crank-Nicholson scheme is given by the equation

$$\frac{da}{dt} \approx \frac{a^{j+1} - a^j}{dt} = \frac{1}{2} \left[F^{j+1}(a, \partial_x a, \dots) + F^j(a, \partial_x a, \dots) \right] \quad (2.35)$$

where superscripts indicate the timestep numbers and $F(a, \partial_x a, \dots)$ represents the function on the RHS of the time evolution equation. The

Crank-Nicholson scheme is only beneficial in stabilising linear terms, so the nonlinear terms are estimated at half-timesteps and inserted separately. For example, the $\tilde{\phi}$ implicit equation becomes

$$\frac{\tilde{\phi}_k^{j+1} - \tilde{\phi}_k^j}{dt} = \frac{1}{2}[L_{\tilde{\phi},k}^{j+1}(\tilde{\mathbf{U}}^{j+1}) + L_{\tilde{\phi},k}^j(\tilde{\mathbf{U}}^j)] + N_{\tilde{\phi},k}^{j+1/2}(\tilde{\mathbf{U}}^{j+1/2}), \quad (2.36)$$

where $\tilde{\mathbf{U}}^r$ represents the vector containing all Fourier coefficients of the four quantities. To determine $N^{r+1/2}$ the same scheme is used for a half timestep but the nonlinear term is calculated at the original time r to give $\tilde{\mathbf{U}}^{r+1/2}$. The implicit Equation 2.35 is solved for $\tilde{\mathbf{U}}^{j+1}$ in terms of $\tilde{\mathbf{U}}^j$ and $N^{j+1}(\tilde{\mathbf{U}})$, to give

$$\bar{n}_k^{j+1} = \bar{n}_k^j + dt N_{\bar{n},k}^{j+1/2} \quad (2.37)$$

$$E_k^{j+1} = \frac{(1 - dtk^2/2)E_k^j + dt N_{E,k}^{j+1/2}}{1 + dtk^2/2} \quad (2.38)$$

$$\begin{aligned} \tilde{\phi}_k^{j+1} = \frac{1}{1 - (dt^2/4a_+^2)} & \left[(a_-/a_+ + (dt^2/4a_+^2))\tilde{\phi}_k^j \right. \\ & + i(dt/2a_+)[(1 + a_-/a_+)\tilde{n}_k^j] \\ & \left. + i(dt^2/2a_+^2)N_{\bar{n},k}^{j+1/2} + (dt/a_+)N_{\tilde{\phi},k}^{j+1/2} \right] \quad (2.39) \end{aligned}$$

$$\tilde{n}_k^{j+1} = (a_-/a_+)\tilde{n}_k^j - i(dt/2a_+)[\tilde{\phi}_k^j + \tilde{\phi}_k^{j+1}] + (dt/a_+)N_{\bar{n},k}^{j+1/2} \quad (2.40)$$

where $a_+ = 1 + \frac{dt}{2}(k - S(t + dt))^2$ and $a_- = 1 - \frac{dt}{2}(k - St)^2$. These equations are used each timestep to calculate the coefficients at the next. First, the nonlinear terms need to be calculated.

2.3.3 Nonlinear Terms

The nonlinear terms in the PI model evolution are all products of two of the four PI model quantities, or their derivatives. These products become convolutions in the spectral space and if we calculate these convolutions explicitly - a pure spectral method - the computational cost scales as $\mathcal{O}(n^2)$. Using the pseudospectral method instead, we inverse Fourier transform the nonlinear terms back to real space, where the explicit product is calculated only at each grid point, before being Fourier

transformed once more back into the Fourier space. Using this method reduces the scaling of the computation cost to that of the FFT algorithm, that is, $\mathcal{O}(n \log n)$. For instance, to calculate the nonlinear term $\sum_l (\tilde{\phi}_l E_{j-l})$ one first applies an inverse Fast Fourier Transform to both $\tilde{\phi}_l$ and E_l and multiplies them, before applying an FFT to the product. This can be summarised by

$$\sum_l (\tilde{\phi}_l E_{j-l}) = \mathcal{F}[\mathcal{F}^{-1}[\tilde{\phi}_l] * \mathcal{F}^{-1}[E_l]] \quad (2.41)$$

Here it is important to note that when performing the inverse FFT on a complex quantity's coefficients, one doesn't end up with the 'true' real space values of $\tilde{\phi}$ at the grid points, because the FFT is assuming typical Fourier modes rather than the time dependent ones described above. This is of no consequence for the calculation, since the convolution is independent of the basis used for the calculation, but for the purposes of understanding what the calculation is doing, the multiplication is not being performed in the 'physical' spatial domain.

Dealing with Aliasing

When using this method, one has to be aware of the effects of *aliasing* on the calculation. The discretisation of the domain sets an upper limit on the magnitude of the values of k , so that $|k| \leq K = N\pi/L$ but in the products of functions these frequencies combine and generate modes at higher frequencies than those available. These higher frequencies instead wrap around and alias to the valid frequency $k_A = k \pm 2mK$ so that k_A is in the valid range.

For example, consider a periodic system of $N = 5$ grid points, $x_j = 0, 1, 2, 3, 4$ spanning a domain of length 5. The allowed values of k in a periodic domain of this size are $2n\pi/5$, for integer n , however at the grid points, k -values separated by multiples of 2π appear equivalent, so that the only distinct Fourier modes are given by $n = -2, -1, 0, 1, 2$, with all other modes being equivalent to one of these. If two functions are defined in the domain, $f(x) = \exp(i2\pi x/5)$, and $g(x) = \exp(i4\pi x/5)$ the product of these two functions, fg , would be $f(x)g(x) = e^{i6\pi x/5}$ however this lies outside the valid range of k values for the chosen grid, and so the values of fg at the grid points alias to the $k = -4\pi/5$ mode as seen in Figure 2.1. This will cause errors in calculation if not treated properly.

To counter the effects of aliasing, Orszag's 'two-thirds rule' is applied [97]. This rule states that for a quadratically nonlinear equation,




Figure 2.1: Demonstration of the aliasing of the $k = 6\pi/5$ mode onto the $k = -4\pi/5$ mode on a grid of five points spread evenly across a domain of size $L = 5$. **Red:** The function $fg = \exp(i6\pi x/5)$. **Blue:** The function fg aliases to, $\exp(-i4\pi x/5)$. **Black:** The values of fg at the grid points $x = 0, 1, \dots, 4$ which also lie on the aliased mode.

one should filter the high wavenumbers so that one is left with $(2/3)N$ wavenumbers in the range $k \in [-(2/3)K, (2/3)K]$ where as above, $K = N\pi/L$. To demonstrate, on a grid spanning a domain of length $L = 6$ with $N = 60$ uniform grid points, the unique mode wavenumbers are $k = 2n\pi/6$ for $n = -30, -29, \dots, 29, 30$. The two-thirds rule says in quadratic terms we must remove all modes with $|k| < 40\pi/L$. Doing this leaves the valid modes corresponding to $n = -19, -18, \dots, 18, 19$. Multiplying any two of these modes gives a mode either inside this range, or, if outside of it, one which does not alias back inside the range. The $n = 17$ ($k = 34\pi/6$) and $n = 19$ ($k = 38\pi/6$) modes multiply to give the $k = 72\pi/6$ mode, which would incorrectly alias to the $k = -48\pi/6$ mode, but since this lies outside the two-thirds range, can be ignored. The two-thirds rule does not solve the problem of information being lost (this will be dealt with below however), but does ensure the modes in the two-thirds range are treated correctly.

Zero Padding

As discussed above, dealing with aliasing through the two-thirds rule means information about the system is lost, firstly because we only con-

sider two-thirds of the wavenumbers in the multiplicands, and secondly because after multiplying we only keep two-thirds of the wavenumbers in the product. The first of these can be solved by first padding the set of Fourier coefficients, a_n , with $N/4$ zeros before and $N/4$ zeros after:

$$\begin{array}{c}
 \underbrace{[a_{-N/2}, a_{-N/2+1}, \dots, a_0, \dots, a_{N/2-1}]}_N \\
 \downarrow \text{Zero padding} \downarrow \\
 \underbrace{[0, \dots, 0]}_{N/4}, \underbrace{[a_{-N/2}, a_{-N/2+1}, \dots, a_0, \dots, a_{N/2-1}]}_N, \underbrace{[0, \dots, 0]}_{N/4}
 \end{array}$$

This is done so that when applying the two-thirds rule, the only terms ignored would be zero anyway. The second issue is not directly solvable; the information about high frequency modes generated by nonlinear terms either has to incorrectly map into aliased coefficients or be ignored, the latter being the only sensible option. The impact of ignoring these modes can be minimised by using a suitable number of grid points (and equivalently Fourier modes) such that the highest frequency modes have small enough amplitude that ignoring them has no significant impact. This is made easier in the PI model by diffusion, which acts to reduce high wavenumber mode amplitudes, with the rate of decay proportional to k^2 .

2.3.4 Timestepping Scheme

The nonlinear terms are calculated at half timesteps by taking a half timestep based on the current system state and calculating the half-timestep nonlinear terms using these new values. The nonlinear terms are therefore calculated twice per timestep.

2.3.5 Remapping Wavenumbers

An important complication of the background shear remains to be discussed. The allowed wavenumbers in the discrete Fourier expansion should be restricted to be in the range $-K < k < K$ where $K = N\pi/L$, however the complex quantities have evolving wavenumbers due to the background shear, giving the effective $k_{eff,j} = k_j - St$. This then raises the question of what ought to happen when a $k_{eff}(t)$ is no longer in this




Figure 2.2: Convergence of PI model simulations with N_t and N_x

range. If simply ignored, the retained wavenumbers will simply continue to shift until they are outside this desired range. To combat this a remapping system is used so that when a wavenumber has evolved to the point it is now closer to a different k_j than it was at $t = 0$, the labels for all coefficients shift along one mode, and the value of t used in transforming between real and Fourier space is adjusted accordingly.

This means at $t = K/SN$, we can slide the elements of the vectors containing $\tilde{\phi}_j$ and \tilde{n}_j along, and set $t = -K/SN$. For example at $t = 0$ the coefficient $\tilde{\phi}_4$ corresponds to the $e^{(4i\pi/L)x}$ term in the expansion, but at $t = \pi/SNL$ this will have shifted to be the $e^{i(4\pi/L - S(\pi/SNL))x}$ coefficient.

2.3.6 Convergence of PI Model Code

The pseudospectral method used, and the implicit timestepping procedure, lead to the expectation of second order convergence in δt and exponential convergence in δx . This was tested by initialising the system with a Gaussian at the center of a domain with size $L = 50$ and evolving for 10 boundary condition periods in a background shear of $S = 0.5$ and comparing the results across simulations performed with different numbers of gridpoints, N_x , and timesteps N_t . This showed the code converged to second order in time and to exponential order in position, as seen in Figure 2.2.

Chapter 3

Turbulence and the Edge of Chaos in the PI Model

3.1 Nature of Turbulence in the PI Model

When $S \neq 0$ the shear in the PI model acts to increase the wavenumbers of small perturbations until dispersive terms cause them to decay, as discussed in Section 2.3.1. This stabilising effect acts in competition with the nonlinear effects - for example, $\mathbf{E} \times \mathbf{B}$ advection can act to cause density to accumulate rather than disperse, however these effects increase nonlinearly with perturbation size so that they overbalance dispersion only if a critical amplitude has been reached, analogously to the 2D toy model in Section 1.8.5. The PI model is hence subcritical for $S \neq 0$. As discussed in Section 1.8.5, this means a stable equilibrium exists, but given a large enough perturbation, nonlinear instabilities take over the linear stability and perturbations will grow until saturating nonlinearly. This also means the PI model features an ‘edge of chaos’ which lies between basins of attraction of the equilibrium basin and the high amplitude saturated state. Note that the term ‘edge of chaos’ is used for this manifold regardless of whether the behaviour on the high amplitude side is actually chaotic or not.

In this chapter, simulations in the PI model will be examined and be seen to demonstrate a range of complex behaviours in the saturated state, with the phenomenology of the saturated state strongly dependent on both box size, L , and shear, S . In general we will see that at low S values the saturated state will be turbulence dominated by propagating coherent structures, while at higher S values, the saturated state consists of a small number of simpler isolated travelling coherent structures which are

stable, do not interact with each other, and behave nonchaotically. The objective of this thesis is to demonstrate that the turbulent behaviour seen at lower values of S is connected to both the edge of chaos and the simpler behaviour at higher S using the framework of RPOs.

The behaviour of turbulence in the PI model is studied using simulations performed using the spectral-PI code in a typical small (relative to the size of coherent structures) domain ($L = 5$) and a typical large domain ($L = 50$). Some studies were repeated in $L = 100$ and $L = 200$ sized domains and found to give very similar results to those in the $L = 50$ box, indicating that the $L = 50$ box should be a suitable representative for the general behaviour of the PI model in large domains. The $L = 50$ domain has a size that is large enough to contain multiple structures, while not being so large that computational time is compromised. Large domains give a representation of how realistic system scale turbulence behaves, while smaller domains reveal a different type of structure, that doesn't show up organically in larger simulations, but could still theoretically play a role in the dynamics in larger domains since any solution in the small domain also appears in larger ones by simply repeating the solution spatially. The work on turbulence in the smaller domain is very preliminary and the results are in an early stage of analysis so it is not yet clear whether they present any useful insights into the full domain turbulence. Nevertheless, the results show some interesting phenomena and are thus included.

3.1.1 Saturation and Turbulence in Large ($L = 50$) Domain

In the large domain, turbulent simulations were performed in the range $S = 0.1$ to $S = 1.5$ with a step size of 0.1. Simulations at $S = 0$ were not performed at the time of writing as the code developed is linked closely to the existence of shear - the periodic boundary conditions for instance have infinite period at $S = 0$ and the code would require modification to run - and the primary point of interest at this stage is in the edge of chaos, which does not exist in the PI model for $S = 0$ as here the system is not subcritical, it is simply unstable. The behaviour at $S = 0$ and its transition into $S \neq 0$ is relevant to the study of turbulence and would be a natural next step for this work. The behaviour of the saturated state in the large domain is found to be in agreement with References [39] and [42]. It is characterised by localised avalanche-




Figure 3.1: Flux against position and time for typical turbulence simulations in the PI model with values of S 0.5 (top), 1.1 (left), and 1.5 (right).

like structures which propagate with equal propensity in directions of both increasing and decreasing x , as would be expected from the Q_Z symmetry (Section 2.2). These avalanches become increasingly sparse and isolated with increasing S . At lower values of S , avalanches are numerous and avalanche-based events (collisions, spontaneous splitting, spontaneous decay) are common, and the system is turbulent. At higher shears the avalanche becomes stable, maintaining a constant amplitude and not experiencing spontaneous splitting or decay - in this region of parameter space the saturated state is no longer chaotic and the behaviour is very predictable. Typical examples of these behaviours can be seen in Figure 3.1 for $S = 0.5, 1.1$, and 1.5 . Notice that at the higher shear, the system evolves into a stable travelling wave which lasts indefinitely. At the intermediate shear, a similar travelling wave can be seen, however it leaves a wake behind it which can itself grow into a travelling wave and propagate parallel to the original wave. At the lower shear, the existence of travelling structures is still evident, and they are more numerous and transient than in the higher shear cases. A multitude of interactions between waves can also be seen at low shear, with collisions, splitting, and decay of waves contributing to the turbulent behaviour.

Turbulence at shear below $S \approx 0.6$ was found to be able to sustain itself indefinitely, and above $S \approx 1.1$ the system can exist in a state containing a single structure indefinitely. In the interim region, simulations showed that the system is able to return from turbulence to equilibrium, through annihilation of structures in collisions or decay of structures. Long turbulent simulations were performed at each shear and the domain-averaged rightward density flux against time was calculated. This is equivalent to averaging $Q(x, t) = i(\tilde{n}^* \tilde{\phi} - \tilde{n} \tilde{\phi}^*)$ across the domain. Examining the ensemble-averaged flux over many simulations gives an average flux decay rate λ approximately satisfying

$$\langle \bar{Q} \rangle(t) \approx \bar{Q}(0) e^{-\lambda t} \quad (3.1)$$

where angled brackets here indicate an ensemble average and a bar indicates a spatial average. The values of λ for each shear are shown in Figure 3.2. Above a critical shear of $S_{crit} \approx 1.623$ no high amplitude state can be sustained and all perturbations decay to zero.




Figure 3.2: Decay rate, λ , of turbulence in $L = 50$ simulations at various values of S .

A potential explanation for this behaviour is that at low shear the travelling structures split at a higher rate than they decay, so there is no net decrease in the number of structures, and hence in the amplitude of the state. In the intermediate region, the rate of decay and annihilation becomes significant enough that all structures are eventually destroyed. At higher shear, the travelling waves appear to be stable and hence long-lived - this will be confirmed in the RPO analysis in Chapter 4. The lack




Figure 3.3: Left: 2-dimensional Fourier transform of the flux of turbulence in a $L = 200$ box against wavenumber and frequency. Right: 2-dimensional autocorrelation of turbulence in a $L = 200$ box against displacement and time delay.

of any saturated state for $S > 1.623$ would be explained by no RPOs existing for these values and again this will be confirmed in the RPO analysis in Chapter 4.

The travelling waves in PI model turbulence demonstrate several important differences from the avalanches seen in other models such as those in [10, 85]. These differences are: a) the bursts in the PI model exist in the absence of a spatially varying source term, b) there is no preferred propagation direction, and c) bursts have roughly equal velocity (at a constant value of S). Property a) is obvious from the lack of source term in the PI model equations. Properties b) and c) can be seen by looking at the space and time Fourier transform or the two-dimensional autocorrelation of turbulence. Examples of these can be seen in Figure 3.3 for $S = 0.5$. The symmetry of both the autocorrelation and the Fourier spectra indicate the balance between forwards and backwards propagating structures. The linear relationship between ω and k in the Fourier spectrum indicates the propagation mechanism is non-dispersive and the linear relationship between the δt and δx values of the peaks in the autocorrelation indicate the constant velocity of the bursts.

As discussed in Section 1.7 and [39], the reason for this difference in behaviour is that SOC models depend on coarse-graining and large scale gradients while the PI model retains local physics, and the structures seen in ITG turbulence are more qualitatively similar to the PI model waves. The existence of travelling wave solutions in the absence of dispersion, indicating that the propagation mechanism is entirely nonlinear, is one of the main motivations for using the model to attempt to capture the




Figure 3.4: Flux (colour) against position, x (horizontal), and time t (vertical) for simulations in a $L = 50$ box showing **(a)**: A collision between two travelling waves resulting in annihilation of both. This occurred when $S = 1.092$. **(b)**: Two travelling waves travelling at the same velocity at $S = 1.098$

minimal physics of the structures.

Finally, another set of simulations in $L = 50$ were performed where turbulence was generated at $S = 1.5$ and the shear was then incrementally both increased and decreased, with turbulence allowed to settle after each step in shear. A Fourier transform of the domain-averaged flux time series was taken at each shear step. At $S = 1.096$, the single travelling wave split into two waves travelling in opposite directions, before colliding and annihilating. This event is shown in Figure 3.4a. The power spectrum of the average flux time series is shown in 3.5 and shows several key frequencies in the oscillations of the travelling waves. Understanding the origin of these frequencies is a part of the motivation of this work, and we will see in Section 4.2 that the RPO framework indeed explains these frequencies as arising due to Hopf bifurcations in the RPO skeleton.

The incremental shear method was repeated starting with turbulence at $S = 0.5$, again running two simulations with slowly changing shear in the increasing and decreasing directions. This gives the frequency spectrum for the domain averaged average flux seen in Figure 3.6. There are no obvious frequencies in this region of S , supporting the idea that this regime is dominated by many interacting travelling waves, rather than just one or two. In the simulation with increasing S , the turbulence quickly decayed once S reached 0.584. The cutoffs at which turbulence is long-lived is consistent with the zero decay rate ranges in Figure 3.2.

In summary, there is a range of values of S for which there is no long-lived saturated solution, yet above and below this range long-lived solutions exist. Above the decaying range the saturated solution becomes very localised, with behaviour a single travelling wave, or a pair




Figure 3.5: Spectral density of the angular frequency of the domain averaged flux in turbulent simulations against S , for high values of S .

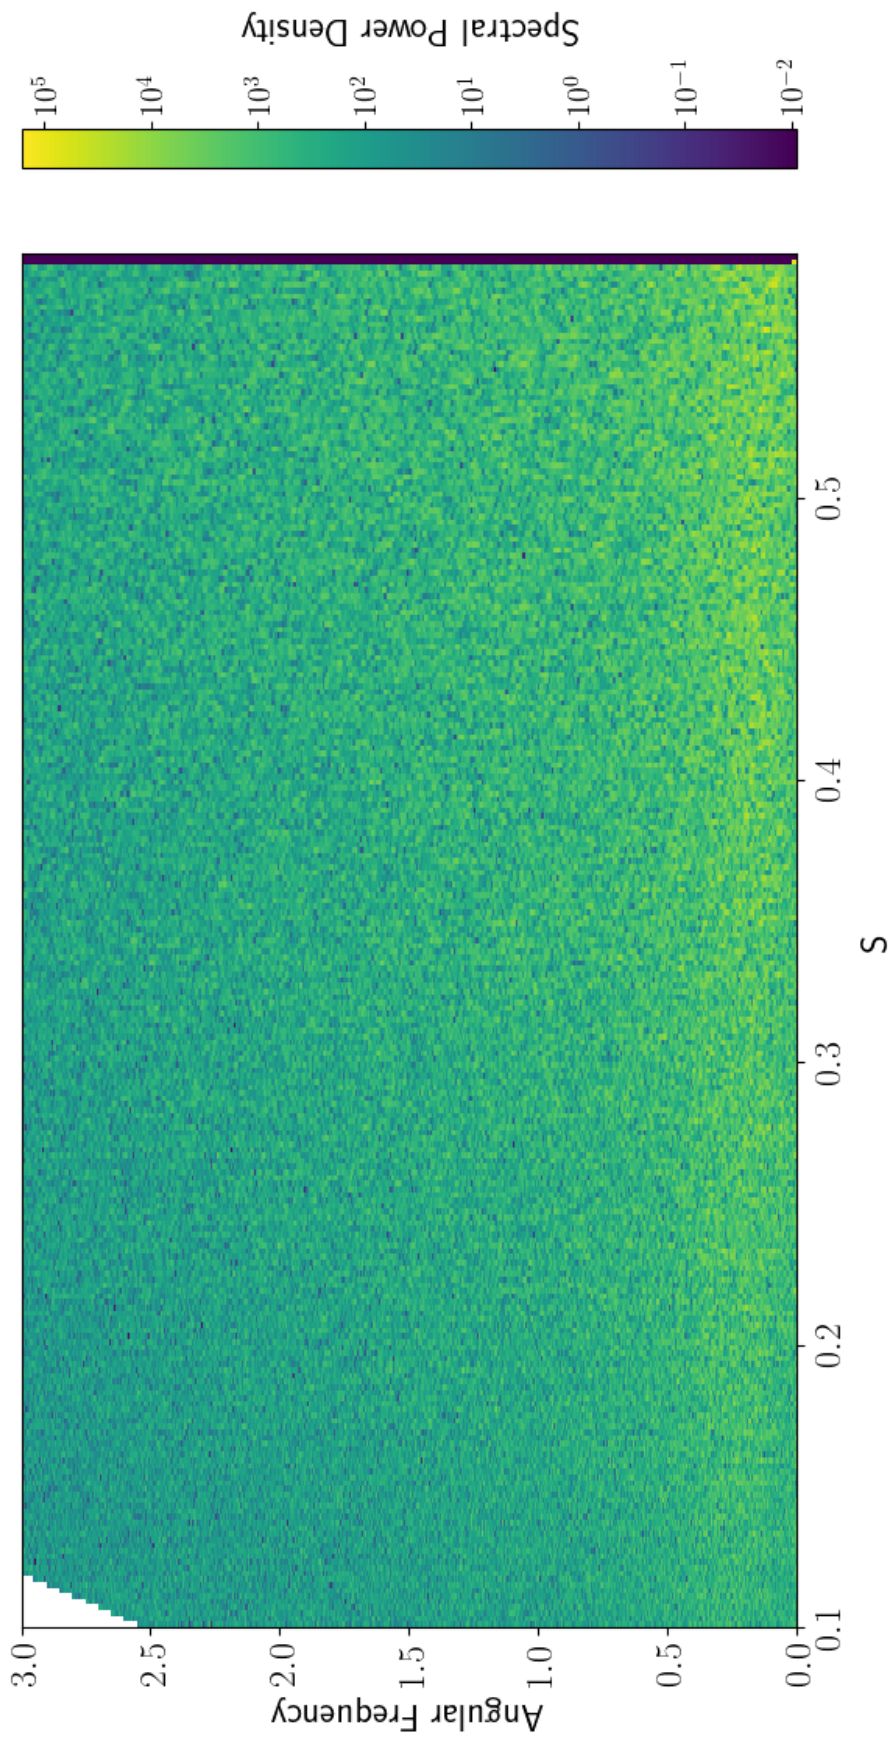


Figure 3.6: Spectral density of the angular frequency of the domain averaged flux in turbulent simulations against S , for low values of S . The top left corner has no data as the range of k is limited by the size of the time-step, which was decreased proportionally to the boundary condition period $\tau_{BC} = 2\pi/LS$ (Equation 2.11) to keep a constant number of timesteps in each simulation, so at low S , the range of k is smaller than the range on the figure.

of propagating waves propagating side by side (see Figure 3.4b). Below the intermediate range, turbulence seems to be sustained by having many travelling waves, with low decay rates, and high survival chance in collisions. It was hoped that the behaviour at lower shears might be explained in part by the structure of RPOs in small domains, however this was not achieved in time for this thesis, and it is not yet known if these would be useful.

3.1.2 Saturation and Turbulence in Small ($L = 5$) Domain

Turbulence in small domains has only been briefly explored due to time limitations. The results obtained so far have been included, but there has been limited analysis. The turbulence in small domain takes on significantly different characteristics to the large box. In small boxes the turbulence is characterised by stationary oscillations across all shear, see for example Figure 3.7. The reason for this change in behaviour is that in the smaller domain, the travelling waves seen in the large domain cannot fit inside the box and wrap around, interacting with itself. It was hoped, due to the fact the small domain is able to demonstrate this ‘interaction’, albeit in a contrived and unrealistic way, that the small domain might prove to be a computationally accessible way to study the interactions between two structures when they collide, or to study single structures splitting into two. Unfortunately this hypothesis has not had the time to be explored fully yet. The boundary conditions also play more of a direct role in determining the dynamics in the $L = 5$ box. Another observation regarding the turbulence in the $L = 5$ domain are that it has finite lifetime for $S > 0.3$, with no apparent high S range at which structures can self sustain perpetually, unlike the results in the $L = 50$ domain.

3.2 Locating the Edge of Chaos with the Bisection Method

In subcritical systems, such as the PI system, there exists a stable equilibrium and a high amplitude (possibly turbulent) saturated state attractor in system space. The basins of attraction of these two states are separated by a manifold of codimension unity, referred to as the ‘edge of




Figure 3.7: Flux against position and time for typical turbulence simulations in the PI model. Domain size $L = 5.0$, S takes values 0.3 (left), 0.6 (center), and 1.0 (right).

chaos' (EOC). Within the EOC manifold, RPOs can exist. These RPOs must lie fully within the EOC, since any orbit which leaves the EOC cannot return to it. It has been found [42] that in a range of fluid systems, the EOC contains RPOs with just a single unstable eigendirection - the direction out of the EOC. These RPOs are therefore stable attractors in the EOC subspace, and hence at long time, orbits starting in the EOC will be attracted into one of these RPOs, and in this case the RPO is called an **edge state**.

It has been observed that the qualitative behaviour of edge state RPOs in the EOC in the PI system is similar to the structure of RPOs in full turbulence. [42]. The edge states are soliton-like and propagate at a constant velocity dependent on S , similar to the travelling waves observed in turbulence or the single high amplitude travelling wave in the saturated state at high S . These similarities between the edge states and the structures in full turbulence led to the suggestion that studying the edge of chaos and the edge states within may aid in illuminating the nature of the turbulence itself.

Following Pringle *et al.* [42], a straightforward algorithm known as the **bisection method** was applied to the PI model code to locate the

edge of chaos in the PI model across a range of box sizes and shear values. The bisection method takes a shape of initial state, such as a Gaussian blob of density and potential (the precise form of this does not matter) and gives it an amplitude A . If an edge exists, there will exist an amplitude $A = A_{Edge}$ for which the initial state lies within the edge and remains in it forever. The bisection method iteratively trials different values of A and examines late time behaviour to determine whether A is above or below A_{Edge} and adjusting accordingly.

To demonstrate the bisection method, an instance of it applied to the inverse Van der Pol oscillator (Equations 1.48 and 1.49) can be seen in Figure 3.8. First, an arbitrary direction \hat{u} is chosen, in this case $\hat{u} = \left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right)$. The direction of \hat{u} is not important in two-dimensional systems such as this one as the edge is itself an RPO, so the vector direction only changes the point on the RPO the system begins in. In higher dimensional systems, while the edge of chaos will be identifiable using any \hat{u} , any attractive RPO lying within the edge of chaos will need to be close enough to the point at which \hat{u} pierces the edge of chaos that the system has time to be attracted to it before the instability of the edge takes over, so several trial vectors can be used to increase the chance of locating any edge states.

Once a direction of \hat{u} has been chosen, initial conditions of the form $u_i(0) = A_i\hat{u}$ are trialled, giving orbits $u_i(t)$. These orbits are classified as either ‘laminar’ or ‘turbulent’ (though in this two-dimensional toy example turbulence doesn’t actually occur) based on the long time behaviour, with orbits going to zero belonging to the former class, and those going to infinity (in the PI model, it will instead be those going to turbulence or other high amplitude saturated states) to the latter. The next trial value of A is then the bisection of the smallest A_i that gives a turbulent orbit, and the largest A_i that gives a laminar orbit.

The previous work of [42] examining the edge of chaos and its attractor is here recreated using the spectral code, as well as expanded to smaller box sizes which exhibit domain scale standing oscillations rather than propagating solutions in the edge.

3.2.1 Edge State in Large ($L = 50$) Domain

An example instance of the bisection method applied to the PI model is represented in Figure 3.9a for a large box of size $L = 50$ with shear $S = 0.5$. The rightward flux against position and time of the longest




Figure 3.8: Example of the bisection method applied to the time-reversed Van der Pol system for 30 iterations. **Left:** Trial orbits for bisection method applied to the time reversed Van der Pol oscillator. Trial initial conditions lie along the direction of the vector $(1, 1)$ (marked in grey) with the distance from the origin A acting as the trial variable. Trial orbits which end at the origin are marked in orange, and those that end at infinity in blue. The longest-lived trial is shown in green. **Right:** Values of A_i trialed and their iteration numbers.

surviving trial - the closest to the true edge state - is shown in Figure 3.9b. After an initial transient period where the system is in the edge (but yet to be attracted to the edge attractor), the system takes the form of a propagating wavepacket before eventually the system leaves the edge along the unstable direction and relaminarises. Note that the amplitude of the edge is not completely constant - this is because the wave propagates around the edge of the box and interacts with the \bar{n} field which affects the amplitude of the wave as it passes by. While there is nonuniformity in \bar{n} across x the wave will not have constant amplitude, and the system is still not in the edge state. The \bar{n} field remains mostly unaffected by the passing wave but is shifted by the flux as it passes through. The flux increases with the gradient of $\partial_x \bar{n}$ and so the passing wave acts to spread out the \bar{n} distribution. The exact edge state RPO will have a uniform background \bar{n} however this takes a long time to homogenise and so even the best solution found with the bisection method will diverge from the edge before being fully attracted to the edge state.

The bisection method was applied to large box simulations for a range of values S , and the amplitudes against time of the closest solution to the edge for each value is shown in Figure 3.10a. For all values of S examined, except $S = 0.1$, the system was attracted close to an edge state taking the form of a travelling wave of form similar to that in Figure

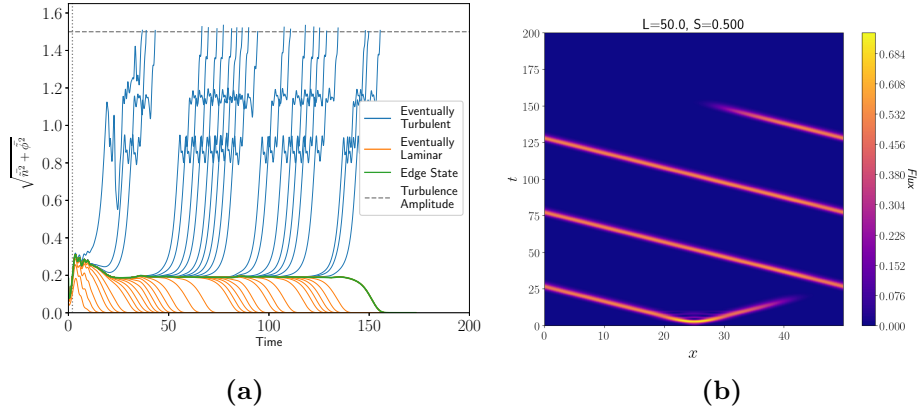


Figure 3.9: **a)** Amplitudes against time of trial solutions calculated during a completion of the bisection method at $S = 0.5$. Blue lines show solutions that were deemed to become turbulent by the algorithm, and orange lines those that relaminarise. **b)** The flux, $(\tilde{n}^* \tilde{\phi} - \tilde{\phi}^* \tilde{n})$ against position and time for the longest closest-to-edge solution found by the bisection method at $S = 0.5$.

3.9b, though with different amplitudes and velocities. As in Figure 3.9a, the amplitude of the orbit in the edge is not constant due to the back effects of the \tilde{n} distribution on the amplitude of the travelling wave, and hence a direct precise measurement of the amplitude of the edge state is not directly calculable from these simulations. However, these simulations will provide a useful starting point in the RPO solving code (Section 4) where the initial conditions will be further refined to converge onto the exact edge state. Despite the limitations on accurate amplitude analyses, the general trend of increasing amplitude with increasing S is clear from Figure 3.10a, and the velocity of the travelling waves can easily be deduced by tracking the position of the peak of the flux in the edge state with time. This calculation gives the velocity of the edge state travelling wave as a function of S as shown in Figure 3.10b.

3.2.2 Edge States in a Small Domain

As mentioned in Section 3.1.2, the results in the small box are still in an early stage, and have not been analysed in detail. Since all solutions existing in the small box will exist in large boxes (by simply tiling the smaller domain across the larger one) it is hoped that the smaller box will provide a more computationally accessible route to analyse some aspects of the state space structure in larger boxes, perhaps the collisions between waves or their splitting.

The bisection method was applied to the smaller $L = 5.0$ box, and like in the large domain, the edge attractors look similar to the turbulence,

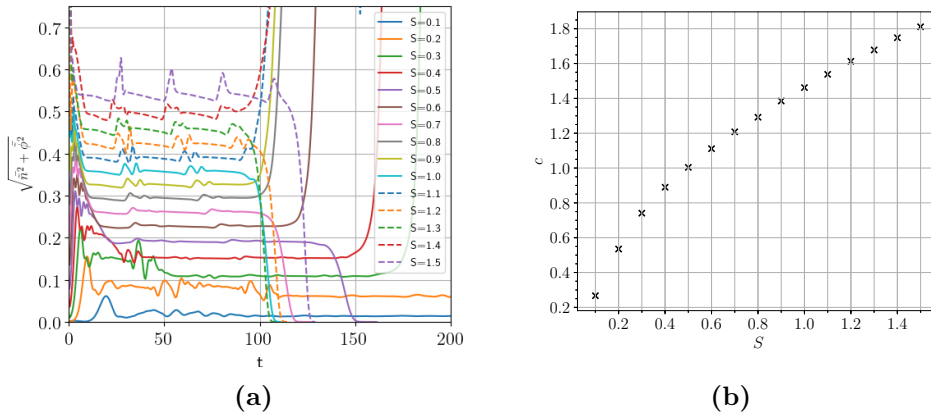


Figure 3.10: a) Amplitudes against time for the closest-to-edge solution found using the bisection method for a range of values of S from 0.1 to 1.5. b) Plot of the velocities of the travelling wave attractor the edge against S .

taking the form of domain-scale oscillations similar to those in Figure 3.7. The amplitude of some the closest-to-edge solutions against time is shown in Figure 3.11, and a selection of edge states are shown in Figure 3.12. Figure 3.11 demonstrates that many of the solutions do not fully settle on an edge RPO before decaying or blowing up, and as S transitions from above 0.5 to below it shows evidence of a period doubling bifurcation; many of the solutions have a period of $2\tau_{BC}$, where $\tau_{BC} = 2\pi/LS$ is the period of the changing boundary conditions. Above $S = 0.5$ this period 2 orbit slowly converges towards a period 1 orbit, while for $S < 0.5$ the period 2 solutions appear to be attractive.

To attempt to converge on an edge RPO before the unstable direction out of the edge becomes significant, the bisection method was used again, however the initial state chosen was the state of the best edge solution immediately before it decays. This was then be repeated until the state converges on an edge RPO. At a test value of $S = 0.47$, a period $2\tau_{BC}$ RPO within the edge was converged on. This indicates a period doubling bifurcation is likely present, potentially indicating a period doubling cascade is present, which could provide a path to turbulence in a manner similar to period doubling cascade in the logistic map system [91].

There is also evidence of a bifurcation between $S = 1.0$ and $S = 1.1$, where the form of the edge RPO transitions from the stationary Q_Z -symmetric oscillation to a travelling wave-like solution, as seen in Figure 3.12. The significance of this transition is not clear, but it seems that the origin of the transition is that as S increases, the width of the travelling wave narrows, and at some point becomes narrow enough to fit inside

the smaller domain. This bifurcation will also be studied in more detail in Section 4.3.




Figure 3.11: Amplitudes of closest-to-edge solutions in the small domain found using the bisection method against time for a range of S .




Figure 3.12: A collection of near edge state solutions identified in the $L = 5.0$ box using the bisection method at S values of 0.2 (top left), 0.5 (top right), 0.7 (bottom left), and 1.1 (bottom right). For each RPO, the flux is plotted against position and time.

Chapter 4

Locating RPOs in The PI Model System

This section of the thesis describes the general methods used to locate RPOs in dynamical systems, adaptations made to suit the PI model, and presents the interesting structures identified using these methods. The RPO finding code is built on the *Newton-Krylov-Hookstep* (NKH) algorithm [15], which combines the Newton-Krylov optimisation algorithm [98, 99] with the locally constrained hookstep procedure [100].

4.1 Introduction to the Newton-Krylov-Hookstep Method

In a generic dynamical system with a state space manifold \mathcal{M} and time evolution operator f^t , a point $\mathbf{u} \in \mathcal{M}$ is a member of an RPO if there exists a period τ and symmetry operator $g(\sigma)$ - parameterised by symmetry variables σ_i condensed into a vector ω - for which the residual between initial state \mathbf{u} and the symmetry-operated final state is zero. Mathematically, this condition on RPOs can be written

$$R(\mathbf{u}, \tau, \boldsymbol{\sigma}) \equiv g(\boldsymbol{\sigma})f^t(\mathbf{u}) - \mathbf{u} = 0. \quad (4.1)$$

The algorithm is indifferent to the specific details of how each argument of R affects its value, so the arguments can be combined into a single vector argument

$$\mathbf{X} = [u_1, \dots, u_m, \tau, \sigma_1, \dots, \sigma_s], \quad (4.2)$$

where m is the dimension of \mathcal{M} and s is the number of symmetry arguments, and \mathbf{X} hence contains $m+s+1$ elements. The space containing all possible vectors \mathbf{X} will be called the solution space \mathcal{X} , and is importantly distinct from the system space \mathcal{M} . Equation 4.1 can now be rewritten in the familiar form of a Newton-style optimisation problem,

$$R(\mathbf{X}) = 0 \tag{4.3}$$

which can be solved computationally for \mathbf{X} , provided one has an initial guess which is suitably close in the solution space, using the Newton-Krylov-Hookstep (NKH) algorithm [15]. Assuming an exact solution \mathbf{X}^* exists, one can attempt to converge iteratively on \mathbf{X}^* with a series of guesses \mathbf{X}_i which, if successful, approach \mathbf{X}^* as $i \rightarrow \infty$. Note that we distinguish between the i -th iterate and the i -th component of \mathbf{X} by keeping the bold vector notation for the former, and the non-bold notation for the latter. Expanding Equation 4.3 in terms of $d\mathbf{X}_i$ gives

$$0 = R(\mathbf{X}^*) = R(\mathbf{X}_i + d\mathbf{X}_i) = R(\mathbf{X}_i) + DR(\mathbf{X}_i)d\mathbf{X}_i + \mathcal{O}(d\mathbf{X}_i^2), \tag{4.4}$$

where DR is a matrix containing the $m \times (m+s+1)$ Jacobian of the current orbit $DR_{jk}(\mathbf{X}) = \frac{\partial R_j}{\partial X_k}(\mathbf{X})$. This gives m equations which must be supplemented by $s+1$ additional rows to give a fully unique solution, and these are constraints on the step $d\mathbf{u}$ that prevent it from being in the direction of time evolution or the directions corresponding to symmetries. Steps in these directions are useless since if $u(t)$ is an RPO, then by definition $u(t+t')$ and $g(\sigma)u(t)$ will also be RPOs. These constraints will be of the form

$$d\mathbf{u} \cdot \frac{d\mathbf{u}}{dt} = 0 \tag{4.5}$$

$$d\mathbf{u} \cdot \frac{d\mathbf{u}}{dx} = 0 \tag{4.6}$$

for the time evolution and spatial translation directions respectively; other equations of this form will be chosen based on the symmetry of the system. These additional constraints can be embedded in the matrix DR , bringing it to size $(m+s+1) \times (m+s+1)$. Assuming \mathbf{X}_i is suitably close to \mathbf{X}^* , we can neglect second order terms, and the problem becomes

$$DR(\mathbf{X}_i)d\mathbf{X}_N = -R(\mathbf{X}) \tag{4.7}$$

which must be solved to obtain an approximation of $d\mathbf{X}_i$ which can then be used as a step towards \mathbf{X}^* ; the next guess is taken to be $\mathbf{X}_{i+1} = \mathbf{X}_i + d\mathbf{X}_i$. This is repeated to converge on \mathbf{X}^* until a desired residual accuracy is achieved, or until improvement is no longer seen with each step.

The task is to solve Equation 4.7 efficiently enough to be practical. Since 4.7 is in the form of a standard matrix problem $Ax = b$ there are a range of standard techniques and methods available, however if the dimensionality of the system is huge - as in the PI model - it is not practical to explicitly calculate the full matrix DR and invert it as inversion scales as $\mathcal{O}(N_x^3)$ which is not practical for computation. Instead, Equation 4.7 is projected onto successive subspaces $\mathcal{K}_n \subset \mathcal{X}$, known as Krylov subspaces, which contain the most unstable eigendirections. The calculation of the Krylov subspaces and projection operation is done by the Krylov algorithm [98], which is outlined below in Section 4.1.1, and the solution in this space is calculated using the GMRES algorithm [101], described in Section 4.1.2. The NKH algorithm is further enhanced by incorporating a hookstep scheme [100], which identifies the optimal step in a restricted volume close to the current guess, with a size dependent on the accuracy of the Newton step, which effectively limits the size of a step to ensure the reduction in residual is optimised. The hookstep algorithm is discussed in Section 4.1.3. The combination of these algorithms is known as the Newton-Krylov-Hookstep (NKH) method and was previously used to great effect in fluid models, vastly improving the identification of RPOs in Couette flow [15], and pipe flow [87] compared to the damped Newton methods used prior.

4.1.1 Krylov Method

In general, when iteratively solving linear algebra equations of the form $Ax = b$, it is most efficient to project the equation onto a subspace of the full system, and minimise the residual in this space. A particularly convenient subspace to use is known as a **Krylov subspace**, which is a subspace spanned by the vectors generated by successive applications of A to b . The n -th Krylov subspace, \mathcal{K}_n is given by

$$\mathcal{K}_n = \text{span}(\{b, Ab, A^2b, A^3b, \dots, A^nb\}) \quad (4.8)$$

and is convenient for problems of this type since applying A successively will emphasise the eigendirections corresponding to the largest eigenvalues of A and it is these directions that can provide the most leverage in solving these equations. Iterative processes project problems onto successive Krylov subspaces, with higher n giving better convergence. An important algorithm common to Krylov space methods is **Arnoldi Iteration**, a method for constructing an orthonormal basis of \mathcal{K}_n .

The algorithm is very similar to the widely known Gram-Schmidt procedure for creating an orthonormal basis for a generic set of n arbitrary vectors. The Arnoldi iteration process starts by assigning the first basis vector q_1 to be equal to \hat{b} , the unit vector in the direction of b . Next one takes the linear product Aq_1 and splits it via projection into its components in the direction q_1 and the perpendicular direction to q_1 , which is then defined to be the direction q_2 . Explicitly,

$$Aq_1 = h_{1,1}q_1 + h_{2,1}q_2 \quad (4.9)$$

from which q_2 can be extracted. In general, q_n is found by solving

$$Aq_n = \sum_{i=1}^{n+1} h_{i,n}q_i \quad (4.10)$$

for q_{n+1} . While performing this algorithm, one needs to calculate and store the components $h_{i,j}$ of an $(n+1) \times n$ upper Hessenberg matrix

$$\tilde{H}_n = \begin{pmatrix} h_{1,1} & h_{1,2} & \cdots & h_{1,n-1} & h_{1,n} \\ h_{2,1} & h_{2,2} & \cdots & h_{2,n-1} & h_{2,n} \\ 0 & h_{3,2} & \cdots & h_{3,n-1} & h_{3,n} \\ \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & \cdots & 0 & h_{n,n-1} & h_{n,n} \\ 0 & \cdots & 0 & 0 & h_{n+1,n} \end{pmatrix} \quad (4.11)$$

which has components $h_{i,j} = 0$ for $i > j + 1$. The columns of the matrix give the coordinates in the basis q_1, \dots, q_n of successive projections onto successive Krylov subspaces. From the derivation of \tilde{H}_n it is clear that it satisfies

$$AQ_n = Q_{n+1}\tilde{H}_n \quad (4.12)$$

where Q_n is the $m \times n$ matrix whose columns are q_1, \dots, q_n . This matrix form makes calculations significantly faster, and is a flexible efficient tool for solving numerical linear algebra problems computationally.

4.1.2 GMRES

Krylov methods can be used to solve $Ax = b$ equations through the ‘Generalised Minimal Residual’, or GMRES method. Put simply, at each iteration, representing the iteration number by n , x^* is approximated by finding the $x_n \in \mathcal{K}_n$ that minimises the residual $r = \|Ax_n - b\|$. Now, x_n is in \mathcal{K}_n so must be able to be written in terms of the basis given by columns of Q_n ,

$$x_n = Q_n y_n. \quad (4.13)$$

And so the task now is to find the vector y_n that minimises $\|AQ_n y_n - b\|$, or, by inserting 4.12, $\|Q_{n+1} \tilde{H}_n y_n - b\|$. The norm remains unchanged if its contents are premultiplied by Q_{n+1}^* , since in general $\|Q_{n+1}^* M\| = \sqrt{(M^* Q_{n+1})(Q_{n+1}^* M)} = \sqrt{M^* M} = \|M\|$ and so the residual function can be further simplified to

$$r_n = \|\tilde{H}_n y_n - |b\rangle \hat{e}_1\|. \quad (4.14)$$

This is then minimised using standard QR-factorisation methods, the details of which are not of specific relevance to this work so will not be expanded on here, [101] contains a full explanation of the method. The QR-factorisation methods were implemented using the standard LAPACK library [102]. Importantly, the upper Hessenberg form of \tilde{H}_n , and the simplicity of b in this Krylov basis greatly improves the time it takes to solve this step.

4.1.3 Hookstep

Finally, when the initial guess at a solution is not sufficiently close to an exact solution for the linear approximation to apply, the step taken could very easily ‘overshoot’ the solution, or be in a direction such that the residual actually increases with the step, amongst other undesirable outcomes. For this reason, a ‘hookstep’ algorithm is also used, which limits the size of the step to a value δ for which the linear approximation remains valid.

The hookstep algorithm still aims to minimise the residual $r(y)$ however the magnitude of y (and hence the magnitude of dx since they are related by orthogonal operators) is restricted by $|y| \leq \delta$. This minimisation cannot be solved by the straight QR-decomposition method of the regular Newton-GMRES algorithm, and instead a singular value decomposition (SVD) is required. The SVD splits the matrix \tilde{H} into a form $\tilde{H}_k = UDV^\dagger$, reducing the residual to

$$r(y) = \|UDV^\dagger y_n - |b|\hat{\mathbf{e}}_1\| = \|DV^\dagger y_n - |b|U^\dagger \hat{\mathbf{e}}_1\| = \|D\hat{y} - \hat{b}\|, \quad (4.15)$$

and since D is diagonal this is easily solved using Lagrange multipliers with solution $\hat{y}_i = (\hat{b}_i D_{i,i}) / (D_{i,i}^2 + \mu)$ where μ is chosen so that $\|\hat{y}^2\| = \delta^2$.

The optimum size of δ is found by trialling different values and comparing the reduction in residual to that predicted by the linear theory. If the reduction is consistent with the linear theory, δ can be increased to take a larger step, and if inconsistent it can be decreased. For a detailed discussion of the algorithm used to do this, see [100].

A summary of the full method in pseudocode is shown in Appendix A.

Code written in FORTRAN by A. P. Willis and provided by C. T. Pringle implementing the GMRES and Hookstep sections of the algorithm were modified by myself to work with the PI model and the Newton section of the code was written entirely by myself with significant adaptations made for the PI model.

4.1.4 Application to PI Model

In the PI model, the shear periodic boundary conditions restrict the possible values the period, τ , can take. This follows from the observation that the only way two states, $u_1(x, t_1)$ and $u_2(x, t_2)$, can be equal is if they each have the same boundary conditions, restricting their time arguments via $t_2 - t_1 = n\tau_{BC} = 2n\pi/LS$, from Equation 2.11. This naturally leads to the conclusion that the period, τ of an RPO must satisfy $\tau = 2n\pi/LS$. The NKH scheme depends strongly on having continuous rather than discrete variables, and so solving directly for the period of an RPO is not possible using the standard methods. This leaves two options: a) Don't solve for the period at all. Select a value for n that makes sense for the RPO being looked for, and only search for RPOs with the corresponding

period; or b) solve for the period indirectly by allowing S itself to be solved for by NKH. This is implemented by making S an element of \mathbf{X} in place of the period τ . Since S is continuous, τ_{BC} is now allowed to take any value. One still needs to choose a sensible value for n_{BC} to converge on a solution, but the flexibility in S makes convergence more likely. Both of these options were written into the RPO finding code as part of this work. The first option is used when the properties of RPOs at a specific shear are being examined, while the latter is used when exploring the solution space as a whole, such as tracking manifolds of solutions through \mathcal{X} with changing S (see Section 4.1.6 below). The possible symmetry operators that could be involved in RPOs reflect the model symmetries outlined in Section 2.2. Both the full symmetries and pseudosymmetries will be permitted, and thus the symmetry operator in the residual function is $g(\sigma) = g_{PI}(\delta\theta, \delta x, \delta E, \delta\bar{n})$. It is straightforward to show that the $\delta\bar{n}$ and δE arguments must be zero for an RPO, since the average zonal density and flow cannot evolve in time. This leaves just the $\delta\theta$ and δx parameters to be solved for as part of \mathbf{X} .

4.1.5 Symmetries of RPOs

Recall the discussion in Section 2.2 regarding the symmetries of the PI model. An RPO orbit \mathcal{M}_p will be invariant under some subset of $G_p \subset G_{PI}$ and this allows the splitting G_{PI} into G_p and the coset group G_{PI}/G_p of symmetries of the model which are not symmetries of the orbit. The orbits $c\mathcal{M}_p$ for $c \in G_{PI}/G_p$ are also RPOs which are distinct from \mathcal{M}_p . Each RPO located is therefore a single representative from a set of distinct RPOs. Only a single member of each RPO set needs to be calculated as the rest are dynamically equivalent and easily calculable by applying elements of G_{PI}/G_p .

4.1.6 Tracking RPOs Through Parameter Space and Determining RPO Stability

A key objective of the methods used in this thesis is to be able to locate new RPOs from those that have already been found, so that the evolution and bifurcation of RPOs with changing S can be investigated. This can be done by taking two known solutions, close to each other in system space and in S , extrapolating their initial conditions and symmetry parameters to another nearby value of S . The solver then uses this as its

initial guess at a solution and attempts to converge on a new RPO at a different value of S . If convergence is not successful, then smaller steps are attempted, and if convergence is consistently good, larger ones can be used. If an RPO cannot be found, the code stops. This can happen if an RPO is so unstable that numerical precision limits become significant, or if solutions cease to exist suddenly due to discontinuity in the manifold of solutions.

To measure the stability of an RPO, the eigenvectors of the RPO must be identified. Since the valid periods of the RPO are discretised, it is convenient to treat time evolution as a map, rather than a continuous operator. The Jacobian of the RPO, defined in Equation 1.51, is calculated using the second order form:

$$J_{ij}^* = \frac{[g(\sigma^*)f^\tau(u^* + \epsilon\hat{u}_i) - G(\sigma^*)f^\tau(u^* - \epsilon\hat{u}_i)]_j}{2\epsilon} \quad (4.16)$$

from which the eigenvalues are extracted using standard methods from the LAPACK linear algebra package [102]. One expects the Jacobian to have at least four eigenvalues equal to unity, one for each of the four continuous model symmetries. Numerically these values will be slightly above or below unity, and must therefore be ignored when counting the number of unstable eigenmodes. One way to do this is to ignore eigenvectors in the symmetry subspace, however a simpler solution was used in which the four closest eigenvalues to unity were ignored.

4.2 RPOs In Large Boxes

In Chapter 3, the edge of chaos was located in the large box for the parameter range $S = 0.1$ to $S = 1.5$ and the edge was found to contain an RPO which had a single unstable direction, pointing out of the edge. The edge RPO takes the form of a travelling wave with constant amplitude and velocity. The NKH algorithm described above was used to take the approximate edge RPOs identified in Chapter 3 as initial guesses and converge on precise RPOs. These precise RPOs were then tracked through parameter space as described in Section 4.1.6, extending the set of RPOs beyond the end of the edge of chaos, where a second, higher amplitude, more complex branch of solutions was identified. An overview of the RPOs identified, including their average flux, velocity, value of S for which the RPO exists, and number of unstable eigendirections is presented in Figures 4.1 and 4.2. The period of the RPOs is $1\tau_{BC}$ for

most RPOs, and for those for which this is not the case, the period is shown (as a multiple of τ_{BC}) in the inset axes of Figure 4.1. The features of this set of RPOs, and their associated bifurcations are examined in more detail in the discussion that follows.

4.2.1 EOC RPOs

As predicted by the results presented in Section 3.2.1, the RPO solver was able to identify RPOs with a single unstable eigendirection within the edge of chaos. In the large system, these RPOs could be identified for the full range of values of S , up to $S \approx 1.6234$ where the manifold of solutions turns back on itself, creating the higher amplitude stable branch of RPOs which are stable. The point at which these two branches meet is a saddle node bifurcation (see Section 4.2.2 below). The edge state RPOs are period $\tau = \tau_{BC}$ travelling waves with a constant envelope function; the winding of the complex wave fields changes through the period but the amplitudes remain constant.

The profiles of a handful of these RPOs, labelled (a), (b) and (c) in Figure 4.1 are shown in Figure 4.3. In all cases the profile represents a travelling wave propagating to the right. At low S the form of the envelope resembles that of a soliton, and as S increases the width of the wave decreases, its amplitude increases and the form of the solution becomes less symmetric and more complex.

4.2.2 Saddle Node Bifurcation at $S=1.6234$, and Stable RPO

At the high shear end of the EOC at $S = 1.6234$ a saddle node bifurcation was identified. At this point, labelled (d) on Figure 4.1, the RPO curve turns back on itself and one sees the existence of multiple RPOs at each value of S . Associated with this turning point is a saddle node bifurcation. A saddle node bifurcation occurs in parameter space when two different solutions are simultaneously created as a control parameter changes. In this case, the upper and lower (EOC) RPO branches are created as S is decreased and the number of RPOs goes from 0 for $S > S_{(d)}$ to 2 for $S < S_{(d)}$. The upper branch is higher in both amplitude and velocity. For $S > S_{(d)}$, neither branch exists and there is no high amplitude state at all, with all initial perturbations quickly decaying to zero. This is in agreement with the results presented in Section 3.1.1, where we saw

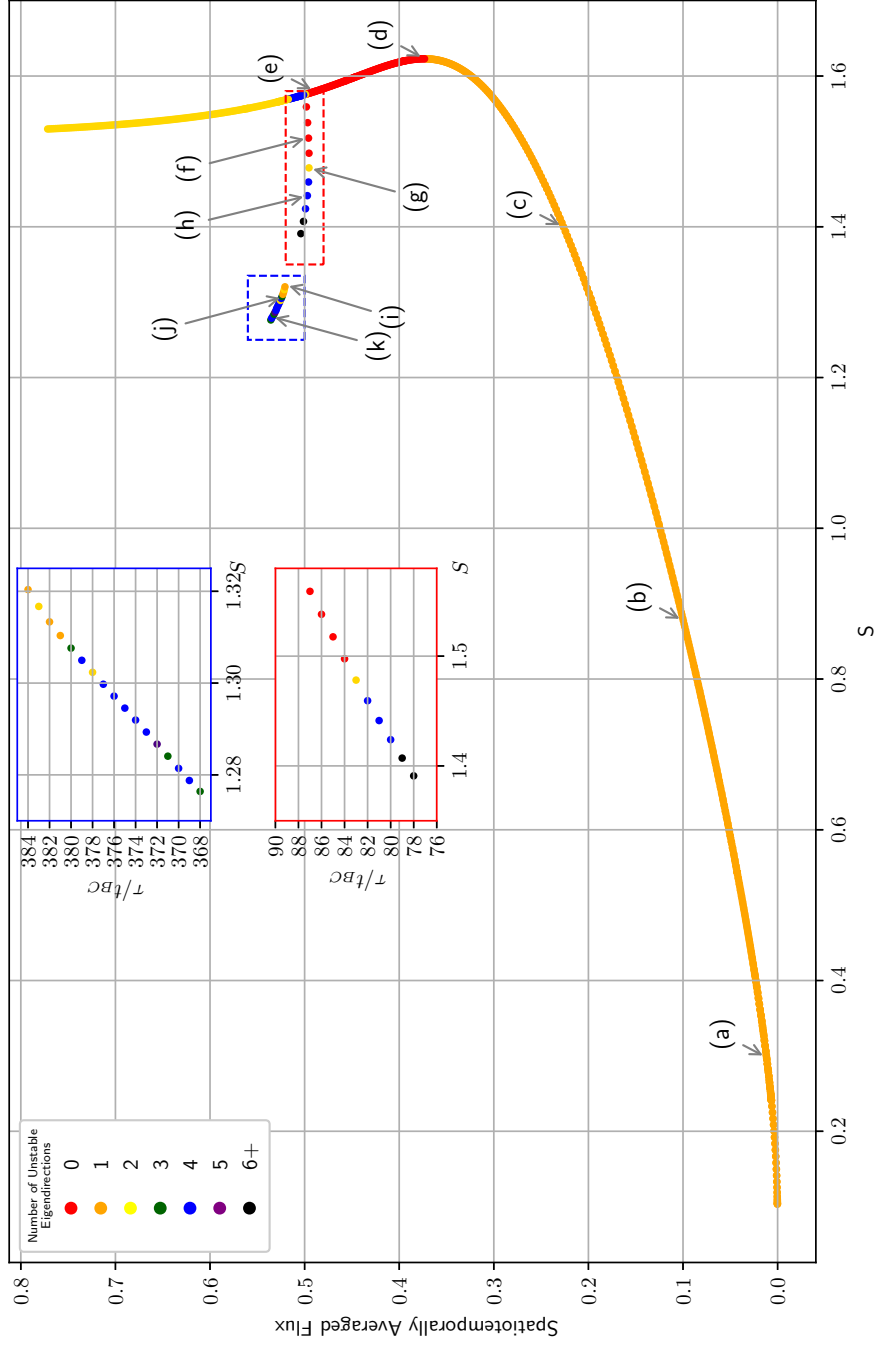


Figure 4.1: The set of RPOs identified in the large $L = 50$ box, plotted as points against their mean flux over the whole domain and period, and shear S at which they were located. The number of unstable eigendirections of each point is represented by its colour. The period of each orbit on the main branch is a single boundary period, $\tau = \tau_{BC} = \frac{2\pi}{L^5}$. The period of RPOs off of the main branch in the regions bounded by blue/red boxes is given in inset plots demarked by blue/red borders.




Figure 4.2: The set of RPOs identified in the large $L = 50$ box, plotted as points against the velocity of the travelling wave, and shear S at which they were located. The number of unstable eigendirections of each point is represented by its colour. The period of each orbit on the main branch is a single boundary period, $\tau = \tau_{BC} = \frac{2\pi}{L\dot{S}}$. The period of RPOs off of the main branch in the regions bounded by blue/red boxes is given in inset plots demarcated by blue/red borders.




Figure 4.3: Profiles of \bar{n} , E , $|\bar{n}|$, $|\partial_x \tilde{\phi}|$ for RPOs (a), (b), and (c). These profiles keep are constant in time and propagate to the right at velocities shown in Figure 4.2

that above this critical value of S , turbulence quickly decayed.

Associated with the saddle node is a change in the stability of the associated RPO, with the single unstable eigenvalue crossing one from above and becoming stable. Consequently, the RPOs on the high-amplitude branch of the bifurcation are completely stable. The existence of these stable attractors in the PI model is significant and interesting in itself, as it begins to provide an explanation as to why structures of this form appear so frequently in turbulence in some regimes: simply, they are attractors and so a system finding itself close to these solutions will remain there indefinitely. In noisy initial conditions, the system may locally look like such a structure, and be attracted to the travelling wave RPO, while elsewhere the same thing is happening, so the domain is filled with these structures which propagate and interact with each other.

The profile of the travelling wave RPO at the bifurcation point (d) and an RPO in the stable set of RPOs (e) are shown in Figure 4.4. Like the edge state, these RPOs maintain their amplitudes with time, and the only time evolution is in translation, and the changing phase of the wave quantities.

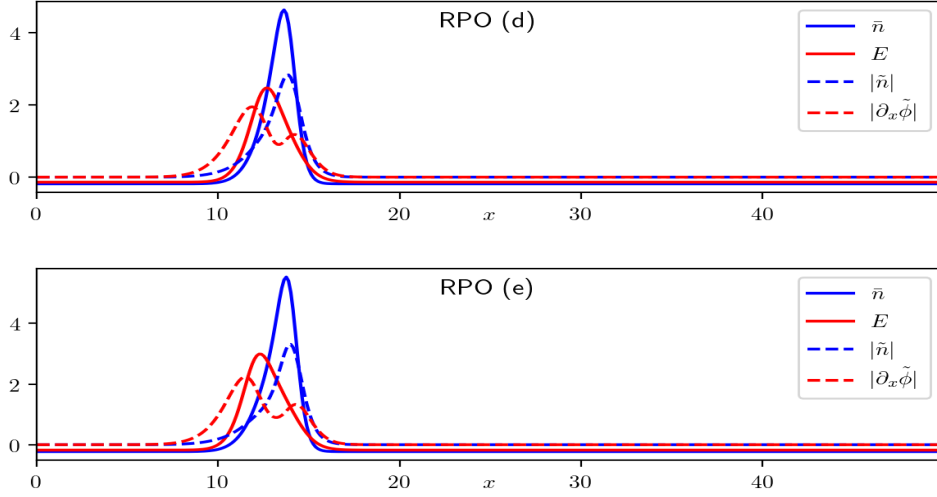


Figure 4.4: Profiles of \bar{n} , E , $|\tilde{n}|$, $|\partial_x \tilde{\phi}|$ for RPOs (d), the located RPO closest to the saddle node, and (e), an example of a stable RPO. These profiles keep are constant in time and propagate to the right at velocities shown in Figure 4.2.

4.2.3 First Hopf Bifurcation Pair

Around $S = 1.576$, two Hopf bifurcations occur close to one another in parameter space. The relevant region in parameter space is shown in more detail in Figure 4.5, and occurs when the stability of the ‘main’ branch (to the right of the figure) has a change in stability in two directions, corresponding to an eigenvalue pair becoming unstable. The two bifurcations are seen by a transition from red to yellow (from 0 to 2 unstable directions) and from yellow to blue (from 2 to 4 unstable directions). The RPO eigenvalue magnitudes are shown in Figure 4.6a for RPOs located near the bifurcation points in parameter space. The eigenvalue corresponding to the first Hopf bifurcation increases rapidly in magnitude as S decreases, never returning below magnitude one. The second Hopf bifurcation relates to a separate eigenvalue going unstable, however this second pair quickly restabilises and drops below amplitude of one.

Due to the stability change, the RPOs on this branch beyond the bifurcation are not stable from this point on, and as the unstable eigenvalue pair increases in value it becomes increasingly difficult to converge on a solution with the RPO solver. These states were not investigated further, however since the amplitude of the states quickly exceeds the typical turbulence amplitudes, it is not thought that these states are relevant to turbulence.


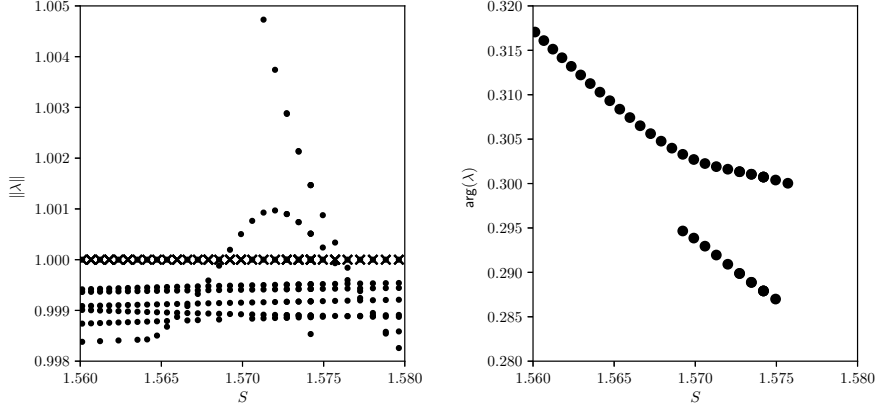


Figure 4.5: The average flux and number of unstable eigenvalues against S for each RPO close to the first Hopf bifurcation against S

When a system experiences a Hopf bifurcation, the pre-bifurcation invariant orbit continues to exist, with a change in stability, and a new invariant orbit is born, separating off from the original orbit. The new orbit carries an additional frequency associated with the angle of the eigenvalue pair that becomes unstable. Orbits on the original branch will be referred to as the ‘original branch’ RPOs and those on the new branch will be called ‘primary Hopf branch’ RPOs.

The values of the arguments of an eigenvalue pair can be used to predict the period of the branch RPOs close to the bifurcation point. The angle of all unstable pairs for each RPO near the bifurcation is shown in Figure 4.6b. These angles represent the phase difference between an initial perturbation lying in the plane associated with a complex eigenvector pair, and the final perturbation after the period of the RPO. Dividing this angle by the period will give the expected frequency of the Hopf RPOs associated with the bifurcation.

The bifurcation angles were approximately $\phi_1 = 0.300$ for the first Hopf bifurcation and $\phi_2 = 0.287$ for the second. These correspond to angular frequencies $\omega_{H1} = 3.760$ and $\omega_{H2} = 3.597$ radians per time unit, or periods of $\tau_{H1} = 1.67$ and $\tau_{H2} = 1.747$. These periods were then used to inform initial guesses in the RPO solver with the objective of locating branch RPOs. RPOs are only detectable when a Hopf period lines up with an integer multiple of boundary periods, so $\tau_H = n\tau_{BC}$. Inserting the boundary period for $S = 1.576$ gives $n \approx 20.9$ for the first



(a) Magnitudes of 10 largest eigenvalues for each RPO. Crosses denote real eigenvalue pairs against S . In the region where there are no unstable directions, no points are shown. (b) Complex angles of the unstable eigenvalues. Discs denote complex eigenvalue pairs against S . In the region where there are no unstable directions, no points are shown.

Figure 4.6

bifurcation and $n \approx 21.9$ and so we expected to find branch RPOs by trialling periods in the range $20\tau_{BC} - 22\tau_{BC}$.

Trialling these values of n , RPO solutions were found on a branch coming off of the main branch near the bifurcation points. This branch has solutions with a roughly constant frequency (with changing S), and since τ_{BC} is not constant, the RPO branch could be picked up at different S values by trialling different values of n . This method identified an RPO at each of $n = 20$ and $n = 21$, however they were far apart in shear making it not possible to track the solutions further. To improve the resolution, the trialled values of n were multiplied by 4 and the range of integers from 74 to 90 were trialled. Since an RPO with period τ will also be an RPO with period 4τ (or indeed any integer multiple) this does not change the details of the RPOs but does allow more detail in the parameter space structure of the RPOs to be explored. The larger period allowed several RPOs on the Hopf branch to be identified which can be seen in Figure 4.1 in the red box.

RPOs on Hopf Branch

The branch coming off the bifurcations identified above was found to contain RPOs with a frequency of $\omega = 3.584$, close to the value of $\omega_{H2} = 3.597$ obtained from the eigenvalue analysis, indicating that this branch comes from the second bifurcation, for which the unstable eigendirection quickly restabilises. It is not yet clear if such a branch exists for the other Hopf bifurcation, but no RPOs on such a branch could be found.

The origin of the low frequency in the Hopf RPOs is the oscillation in amplitude of the travelling wave. A significant consequence of this oscillation is that it leads to the wave leaving behind zonal flows. The zonal density \bar{n} , and the zonal perturbation to E (which in the dimensionless form of the PI model equates to the zonal flow) are shown plotted against time for RPOs (f), (g), and (h) in Figure 4.7. The density perturbations left behind remain unaffected until the wave wraps around the box to interact with them again, since there is no dissipation term in the \bar{n} equation (Equation 2.12). In contrast, the zonal flows left behind quickly decay due to the dissipative term in Equation 2.13. In other models where the saturation of zonal flows is handled differently, it may be possible for the zonal flow to be sustained long enough for it to interact with the wave as it passes by again. The formation of the zonal flows in the travelling wave indicates that they arise due to local physics, and while a detailed study of the zonal flow generation mechanism has not been completed here due to time limitations, the very fact such a state can be isolated using these methods is indicative of the power RPO-based methods can wield in approaching these kinds of problems in future.

A natural question to ask about these solutions is whether the travelling wave is in fact ‘leaving behind’ zonal flows as a result of its oscillation, or whether the zonal density perturbation is required for this kind of interaction to occur. This question can be answered by taking an unstable RPO on the ‘main’ branch and allowing it to evolve for some time so that the Hopf instability can set in. The results of this test can be seen in the plot of \bar{n} against position and time in Figure 4.8. To begin with there is no zonal density perturbation (aside from that within the travelling wave), however as the wave propagates a zonal density is generated, proving that the density wave is generated by the wave, and does not need to exist already for the oscillation in the travelling wave amplitude to occur.

As the bifurcated branch moves further from the original branch, the amplitude of the zonal density wave increases, and the magnitude of the oscillation in flux increases.

4.2.4 Secondary Hopf Bifurcation Pair

Examining the eigenvalues of the Hopf branch RPOs suggests the existence of a secondary sequence of Hopf bifurcations in the region $S \in [1.4, 1.5]$. The magnitudes and angles of these eigenvalues are seen in




Figure 4.7: Primary Hopf branch RPOs (f), (g), and (h) zonal densities (colour, left column) and zonal flows (colour, right column) against position (horizontal) and time (vertical). Vertical lines in the density field indicate the wave is leaving behind a zonal density perturbation.





Figure 4.8: Main branch RPO allowed to go unstable and begins to oscillate in amplitude and leaves a zonal density behind it.

Figure 4.9a. Here we follow the same procedure as described above in Section 4.2.3 to find the corresponding Hopf period and frequencies by dividing the Hopf angle by the RPO period, $\tau = 7.067$. Care must be taken to account for the possibility that the angles associated with the Hopf eigenvalues are aliased from other values. This can happen if the oscillation caused by the Hopf bifurcation is fast enough that multiple Hopf periods occur within a single RPO period. To account for this, frequencies are calculated as if the angles were any of $\phi_H + 2k\pi$ for $k = 0, \pm 1, \pm 2, \dots$. These frequencies are shown in Table 4.1.

Bifurcation	S	Eigenvalue Angle	Possible ω_H
H3	1.478	0.272	$0.0385+0.890k$
H4	1.460	1.691	$0.239+0.890k$
H5	1.407	2.669	$0.378+0.890k$

Table 4.1: Hopf eigenvalue angles associated with the three Hopf bifurcations on the primary Hopf branch, and associated valid frequencies.

The periods corresponding to these frequencies were trialled as initial guesses in the RPO solving code to search for RPOs close to the bifurcation that might lie on a ‘secondary Hopf’ branch. No RPOs could be identified close to the bifurcation, though this was not surprising for two reasons. Firstly, the periods of the Hopf branch could be very long, which makes finding an RPO difficult since longer simulations are required, and there is more time for instabilities to take over. Secondly, any RPO must



(a) Magnitudes of 10 largest eigenvalues for each RPO. Crosses denote real eigenvalues and discs denote complex eigenvalue pairs against S .
 (b) Complex angles of the unstable eigenvalue pairs against S .

Figure 4.9

have a period that is an integer multiple of all of the boundary condition period from Equation 2.11 ($2\pi/LS \approx 0.085$), the period of the first Hopf branch ($\tau_{H2} \approx 1.747$), and the period of a secondary Hopf bifurcation. It is rare for all of these periodicities to line up and hence unlikely for an RPO to be found.

An unexpected result was obtained however, when the RPO solver converged on a cluster of RPOs far from the bifurcation point with an approximate frequency of $\omega \approx 0.172$, shown in Figure 4.1 in the blue box. The flux against position and time for RPOs (i), (j), and (k) in this cluster are shown in Figure 4.10. These RPOs are more complex than those on the primary Hopf branch, showing more complex behaviour in the wake of the travelling wave, with some close to, but not exactly, periodic behaviour at a frequency $\omega \approx 0.86$, consistent with the secondary Hopf frequency H3 in Table 4.1 and an alias compensation of $k = -1$. This suggests that there is a secondary Hopf branch that comes off the primary Hopf branch, which cannot be found easily using the RPO solver, due to the misalignment of the primary and secondary Hopf bifurcation frequencies, but eventually this branch is findable with the RPO solving code, possibly because of another hidden bifurcation. With more time, this branch would be investigated further, and the reasons for the RPO solver being unable to detect it pinned down more concretely.




Figure 4.10: Flux against position and time for RPOs (i), (j), and (k)

4.2.5 Comparisons to Turbulent Simulations With Changing Shear

Summarising the results so far:

- The edge of chaos was found to consist of RPOs with a single unstable direction, as anticipated by the results of the bisection method.
- Tracking the RPOs beyond the EOC revealed a saddle node bifurcation at $S = 1.6234$ where the manifold of RPOs turns back on itself in parameter space and the stability changes from a single unstable direction to no unstable directions. The associated RPOs are attractors.
- Further tracking of the main branch of RPOs as the amplitude continues to increase with decreasing shear showed a pair of Hopf bifurcations around $S = 1.576$ where two pairs of complex eigenvalue pairs become unstable. These bifurcations are named $H1$ and $H2$, and have bifurcation frequencies $\omega_{H1} = 3.760$, $\omega_{H2} = 3.597$. The complex pair associated with $H2$ restabilises again on the main branch after a further decrease in S .
- A branch of RPOs was identified in which a travelling wave generates zonal flows behind it as it travels. The wave has an oscillation in amplitude at frequency $\omega \approx 3.584$ indicating that the branch originates from the $H2$ bifurcation. The branch remains stable for a range of S before the branch undergoes its own series of (at least) three Hopf bifurcations ($H3$, $H4$, and $H5$) between $S = 1.495$ and $S = 1.460$. The frequencies can be calculated from the angles of the Hopf eigenvalues, however due to the possibility of aliasing, the frequencies associated with each bifurcation could be one of several values, given by the values in Table 4.1.
- Another group of RPOs was inadvertently located at a shear between $S = 1.27$ and $S = 1.31$. These RPOs have a frequency of $\omega \approx 0.172$, but there is an almost-periodicity within these orbits at $\omega \approx 0.86$, which is consistent with the possibility that these RPOs ultimately connect to the branch breaking off from the primary Hopf branch at the $H3$ bifurcation.

The bifurcations and associated frequencies found in the above analysis is consistent with the results of the turbulent simulations in Section

3.1.1 in which turbulence at $S = 1.5$ was allowed to evolve while the shear was slowly increased and decreased. The time-frequency Fourier spectrum of the average flux across the domain as S changes, shown in Figure 3.5, is shown again in Figure 4.11 here alongside the plot of RPO amplitude and stability against S . The range of domain-averaged flux for the turbulence is shown in blue on the RPO section of the plot and for the RPOs is shown by error bars on the points representing RPOs.

There are several features of this figure that demonstrate that the turbulence is closely following the RPOs as S is changed. Firstly, the amplitude and range of values the domain averaged flux takes during the simulation in the high amplitude state and RPOs closely match, apart from a sudden increase in the range of the flux average as S is decreased at $S \approx 1.47$. This coincides with the H3 bifurcation, so is consistent with the theory that the high amplitude state follows a secondary Hopf branch coming off of the primary Hopf branch. Secondly, the frequency spectrum picks up many key frequencies associated with Hopf bifurcations. Firstly, from $S = 1.47$ to $S = 1.575$ the only dominant frequency is at $\omega = \omega_{H2} = 3.597$, then as S increases further, a mode with frequency $\omega = \omega_{H1} = 3.760$ briefly appears before disappearing. The average flux is then constant with time, but non-zero, following the amplitude of the stable RPOs exactly until suddenly dropping to zero when S reaches a value of 1.6234.

Moving in the other direction with decreasing S , at $S = 1.478$, the shear associated with the H3 bifurcation, a frequency at $\omega = 0.86$ appears, and shortly after, at $S \approx$ the spectrum becomes more complex, but is dominated by the frequencies at $\omega = 0.86$ and $\omega = 1.12$ which are respectively consistent with the H3 and H4 angles observed in 4.1, as well as harmonics of these two frequencies. The ω_{H2} mode also begins to decrease in amplitude as S is decreased.

Altogether the turbulent simulations are consistent with the RPO picture represented in Figure 4.12, where turbulence naturally forms around RPOs.

4.3 Early Results in Small $L = 5.0$ Domain

The process outlined above was repeated in the small, $L = 5$ domain. The set of RPOs located was significantly more complex than that found in the $L = 50$ domain. The located RPOs are shown as before, plotting




Figure 4.11: Top: The set of RPOs represented as in 4.1 with the range of spatially-averaged flux and spatiotemporally-averaged flux of the S -varying turbulent simulations shown in blue. Bottom: The spectral power density of the spatially-averaged flux against S of the S -varying turbulent simulations.

their mean flux against S and with the number of unstable eigendirections represented by the colour of the points. The bifurcation visible at $S \approx 0.5$ is a period doubling bifurcation, as predicted from the results of applying the bisection method in the small domain in Section 3.2.2. The bifurcation at $S \approx 1.0$ is a pitchfork bifurcation with the original branch continuing as a stationary oscillation and losing a stability direction. The branch coming off at the bifurcation contains RPOs that are travelling waves, with the pitchfork pair corresponding to a wave propagating in each direction. Due to time pressures, the detailed analysis and plots of these results could not be completed and written up before submission.




Figure 4.12: The set of RPOs identified in the large $L = 5$ box, plotted as points against their mean flux over the whole domain and period, and shear S at which they were located. The number of unstable eigendirections of each point is represented by its colour. The period of all orbits is $17BC$ except on the branch that bifurcates off the main branch at the period doubling bifurcation where RPOs have period $27BC$.

Chapter 5

Conclusions

The primary objective of this work was to explore the applicability of relative periodic orbit based methods, used previously in neutral fluid physics [15, 87], to plasma physics models, as well as providing a resource for those in the plasma physics community to aid in becoming familiar with these techniques. The secondary aim was to develop the understanding of the interactions between coherent structures and zonal flows. Regarding the first aim, the successful identification of a branch of RPOs that closely follow the behaviour of the PI model at long time is hopefully a persuasive enough result to demonstrate the power that these methods can wield in identifying nonlinear behaviour of plasma systems. In addition, it is hoped with further research, the RPOs identified in smaller domains could also prove to be useful in explaining other behaviours such as stationary oscillations and the formation mechanism of coherent structures. Regarding the second objective, the isolation of an RPO which is a quasi-travelling wave which generates zonal flows gives an opportunity to study and understand the mechanisms of such structures in future research.

The work in this thesis can be broken up into three connected sections. First, the plasma interchange model - a 1D model proposed as a minimal model for propagating coherent structures including interactions with zonal flows - was implemented in FORTRAN using a modified pseudospectral method. The adaptations made to the method, namely shifting wavenumbers with time and the remapping of these wavenumbers, allowed the system to be implemented with a background shear flow S , and the associated shear-periodic boundary conditions. This method was implemented from scratch, and was tested and confirmed to have second order convergence in timestep, and exponential order in grid spacing

(Chapter 2).

Second, the qualitative properties of turbulence in the PI model was studied using the aforementioned code in both large and small boxes. In large boxes, turbulence was found to be dominated by propagating coherent structures, with properties similar to those seen in gyrokinetic simulations in the collisionless, electrostatic limit with adiabatic electrons [39]. A 2D Fourier transform and 2D autocorrelation demonstrated the uniform velocity of these structures, and their non-dispersive nature indicating that their propagation mechanism is due to local nonlinear effects rather than diffusion. The analysis of the turbulence extended to identifying the edge of chaos in the PI model, recreating the work in Reference [42] and building upon it to identify and isolate an attractive (within the edge) edge RPO that takes the form of a travelling wave in a large domain, and a stationary oscillation in the small domain.

Finally, an RPO solving algorithm, the Newton-Krylov-Hookstep algorithm was adapted to work with the PI model code. This allowed precise RPOs to be located, and interpolated to reveal a manifold of RPO solutions. This also revealed bifurcations in this manifold, of Hopf type in the large domain, and of pitchfork and period doubling type in the small domain. The Hopf bifurcations corresponded with frequencies present in turbulent simulations, and the amplitude of the turbulence closely follows Hopf branches, indicating that turbulence is closely following the manifolds of RPO solutions. The RPO solver was also able to locate RPOs which were travelling waves with an oscillation in amplitude, which appear to generate zonal flows and density perturbations as they propagate. In the small box, early results show a period doubling bifurcation, which could be the beginning of a period doubling cascade, and a supercritical pitchfork bifurcation which transfers the stationary oscillation to a travelling wave RPO.

Bibliography

- ¹A. J. Wootton, B. A. Carreras, H. Matsumoto, K. McGuire, W. A. Peebles, C. P. Ritz, P. W. Terry, and S. J. Zweben, “Fluctuations and anomalous transport in tokamaks”, *Physics of Fluids B: Plasma Physics* **2**, 2879–2903 (1990).
- ²J. W. Conner and H. R. Wilson, “Survey of theories of anomalous transport”, *Plasma Physics and Controlled Fusion* **36**, 719–795 (1994).
- ³A. M. Dimits, T. J. Williams, J. A. Byers, and B. I. Cohen, “Scalings of Ion-Temperature-Gradient-Driven Anomalous Transport in Tokamaks”, *Physical Review Letters* **77**, 71–74 (1996).
- ⁴J. Candy and R. E. Waltz, “Anomalous Transport Scaling in the DIII-D Tokamak Matched by Supercomputer Simulation”, *Physical Review Letters* **91**, 10.1103/PhysRevLett.91.045001 (2003).
- ⁵F. Wagner, G. Fussmann, T. Grave, M. Keilhacker, M. Kornherr, K. Lackner, K. McCormick, E. R. Müller, A. Stäbler, G. Becker, K. Bernhardt, U. Ditte, A. Eberhagen, O. Gehre, J. Gernhardt, G. v. Gierke, E. Glock, O. Gruber, G. Haas, M. Hesse, G. Janeschitz, F. Karger, S. Kissel, O. Klüber, G. Lisitano, H. M. Mayer, D. Meisel, V. Mertens, H. Murmann, W. Poschenrieder, H. Rapp, H. Röhr, F. Ryter, F. Schneider, G. Siller, P. Smeulders, F. Söldner, E. Speth, K. Steuer, Z. Szymanski, and O. Vollmer, “Development of an Edge Transport Barrier at the H-Mode Transition of ASDEX”, *Physical Review Letters* **53**, 1453–1456 (1984).
- ⁶F. Wagner, “A quarter-century of H-mode studies”, *Plasma Physics and Controlled Fusion* **49**, 10.1088/0741-3335/49/12B/S01 (2007).
- ⁷P. H. Diamond, S.-I. Itoh, K. Itoh, and T. S. Hahm, “Zonal flows in plasma—a review”, *Plasma Physics and Controlled Fusion* **47**, R35–R161 (2005).
- ⁸P. A. Politzer, “Observation of Avalanchelike Phenomena in a Magnetically Confined Plasma”, *Physical Review Letters* **84**, 1192–1195 (2000).
- ⁹T. Rhodes, R. Moyer, R. Groebner, E. Doyle, R. Lehmer, W. Peebles, and C. Rettig, “Experimental evidence for self-organized criticality in tokamak plasma turbulence”, *Physics Letters A* **253**, 181–186 (1999).
- ¹⁰B. A. Carreras, D. Newman, V. E. Lynch, and P. H. Diamond, “A model realization of self-organized criticality for plasma confinement”, *Physics of Plasmas* **3**, 2903–2911 (1996).

- ¹¹X. Garbet and R. E. Waltz, “Heat flux driven ion turbulence”, *Physics of Plasmas* **5**, 2836–2845 (1998).
- ¹²X. Garbet, Y. Sarazin, P. Beyer, P. Ghendrih, R. Waltz, M. Ottaviani, and S. Benkadda, “Flux driven turbulence in tokamaks”, *Nuclear Fusion* **39**, 2063–2068 (1999).
- ¹³Y. Sarazin and P. Ghendrih, “Intermittent particle transport in two-dimensional edge turbulence”, *Physics of Plasmas* **5**, 4214–4228 (1998).
- ¹⁴P. Beyer, Y. Sarazin, X. Garbet, P. Ghendrih, and S. Benkadda, “2D and 3D boundary turbulence studies”, *Plasma Physics and Controlled Fusion* **41**, A757–A769 (1999).
- ¹⁵D. Viswanath, “Recurrent motions within plane Couette turbulence”, *Journal of Fluid Mechanics* **580**, 339–358 (2007).
- ¹⁶Y. Duguet, C. C. Pringle, and R. R. Kerswell, “Relative periodic orbits in transitional pipe flow”, *Physics of Fluids* **20**, 10.1063/1.3009874 (2008).
- ¹⁷F. W. Aston, “Isotopes and Atomic Weights”, *Nature* **105**, 617–619 (1920).
- ¹⁸A. J. Eddington, “The internal constitution of the stars”, *The Observatory* **43**, 341–358 (1920).
- ¹⁹IPCC, *IPCC Sixth Assessment Report*, tech. rep. (2022).
- ²⁰Culham Centre For Fusion Energy, *Bringing the power source of the stars down to Earth could give us low-carbon electricity for millennia to come.* (2022) ccfe.ukaea.uk/fusion-energy/fusion-in-brief/.
- ²¹L. Tonks, “The Birth of “Plasma””, *American Journal of Physics* **35**, 857–858 (1967).
- ²²M. L. E. Oliphant, E. Rutherford, and P. Harteck, “Transmutation effects observed with heavy hydrogen”, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* **144**, 692–703 (1934).
- ²³J. Wesson, *Tokamaks*, 3rd (Oxford University Press, 2004), pp. 377–377.
- ²⁴M. G. Haines, *Fifty years of controlled fusion research*, tech. rep. (1996), pp. 643–656.
- ²⁵S. Li, H. Jiang, Z. Ren, and C. Xu, “Optimal Tracking for a Divergent-Type Parabolic PDE System in Current Profile Control”, *Abstract and Applied Analysis* **2014**, 1–8 (2014).
- ²⁶A. J. Brizard and T. S. Hahm, “Foundations of nonlinear gyrokinetic theory”, *Reviews of Modern Physics* **79**, 421–468 (2007).
- ²⁷X. Garbet, Y. Idomura, L. Villard, and T. Watanabe, “Gyrokinetic simulations of turbulent transport”, *Nuclear Fusion* **50**, 043002 (2010).
- ²⁸R. D. Sydora, T. S. Hahm, W. W. Lee, and J. M. Dawson, “Fluctuations and transport due to ion-temperature-gradient-driven instabilities”, *Physical Review Letters* **64**, 2015–2018 (1990).

- ²⁹R. E. Waltz, G. D. Kerbel, and J. Milovich, “Toroidal gyro-Landau fluid model turbulence simulations in a nonlinear ballooning mode representation with radial modes”, *Physics of Plasmas* **1**, 2229–2244 (1994).
- ³⁰G. W. Hammett, M. A. Beer, W. Dorland, S. C. Cowley, and S. A. Smith, “Developments in the gyrofluid approach to Tokamak turbulence simulations”, *Plasma Physics and Controlled Fusion* **35**, 973–985 (1993).
- ³¹D. R. Ernst, P. T. Bonoli, P. J. Catto, W. Dorland, C. L. Fiore, R. S. Granetz, M. Greenwald, A. E. Hubbard, M. Porkolab, M. H. Redi, J. E. Rice, K. Zhurovich, and Alcator C-Mod Group, “Role of trapped electron mode turbulence in internal transport barrier control in the Alcator C-Mod Tokamak”, *Physics of Plasmas* **11**, 2637–2648 (2004).
- ³²A. G. Peeters, C. Angioni, M. Apostoliceanu, F. Jenko, F. Ryter, and the ASDEX Upgrade team, “Linear gyrokinetic stability calculations of electron heat dominated plasmas in ASDEX Upgrade”, *Physics of Plasmas* **12**, 022505 (2005).
- ³³T. Dannert and F. Jenko, “Gyrokinetic simulation of collisionless trapped-electron mode turbulence”, *Physics of Plasmas* **12**, 072309 (2005).
- ³⁴S. E. Parker, Y. Chen, W. Wan, B. I. Cohen, and W. M. Nevins, “Electromagnetic gyrokinetic simulations”, *Physics of Plasmas* **11**, 2594–2599 (2004).
- ³⁵GENE Development Team, *GENE Code Website*, <https://genecode.org>.
- ³⁶The GS2 Team, *GS2 Website*, <https://gyrokinetics.gitlab.io/gs2/>.
- ³⁷W. Horton, “Drift waves and transport”, *Reviews of Modern Physics* **71**, 735–778 (1999).
- ³⁸J. Hugill, “Transport in tokamaks – a review of experiment”, *Nuclear Fusion* **23**, 331–373 (1983).
- ³⁹B. F. McMillan, S. Joliet, T. M. Tran, L. Villard, A. Bottino, and P. Angelino, “Avalanchelike bursts in global gyrokinetic simulations”, *Physics of Plasmas* **16**, 10.1063/1.3079076 (2009).
- ⁴⁰F. Van Wyk, E. G. Highcock, A. A. Schekochihin, C. M. Roach, A. R. Field, and W. Dorland, “Transition to subcritical turbulence in a tokamak plasma”, *Journal of Plasma Physics* **82**, 10.1017/S0022377816001148 (2016).
- ⁴¹P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay, *Chaos: Classical and Quantum*, 17.6.2 (2022).
- ⁴²C. C. Pringle, B. F. McMillan, and B. Teaca, “A nonlinear approach to transition in subcritical plasmas with sheared flow”, *Physics of Plasmas* **24**, 10.1063/1.4999848 (2017).
- ⁴³P. W. Terry, “Suppression of turbulence and transport by sheared flow”, *Reviews of Modern Physics* **72**, 109–165 (2000).

- ⁴⁴G. Dangelmayr, J. Hettel, and E. Knobloch, “Parity-breaking bifurcation in inhomogeneous systems”, *Nonlinearity* **10**, 1093–1114 (1997).
- ⁴⁵R. C. Wolf, “Internal transport barriers in tokamak plasmas”, *Plasma Physics and Controlled Fusion* **45**, R1–R91 (2003).
- ⁴⁶H. Biglari, P. H. Diamond, and P. W. Terry, “Influence of sheared poloidal rotation on edge turbulence”, *Physics of Fluids B: Plasma Physics* **2**, 1–4 (1990).
- ⁴⁷S. Suckewer, H. Eubank, R. Goldston, J. McEnerney, N. Sauthoff, and H. Towner, “Toroidal plasma rotation in the PLT tokamak with neutral-beam injection”, *Nuclear Fusion* **21**, 1301–1309 (1981).
- ⁴⁸P. H. Diamond and Y. B. Kim, “Theory of mean poloidal flow generation by turbulence”, *Physics of Fluids B* **3**, 1626–1633 (1991).
- ⁴⁹R. Taylor, M. Brown, B. Fried, H. Grote, J. Liberati, G. Morales, P. Pribyl, D. Darrow, and M. Ono, “H-mode behavior induced by cross-field currents in a tokamak”, *Physical Review Letters* **63**, 2365–2368 (1989).
- ⁵⁰B. N. Rogers, W. Dorland, and M. Kotschenreuther, “Generation and Stability of Zonal Flows in Ion-Temperature-Gradient Mode Turbulence”, *Physical Review Letters* **85** (2000).
- ⁵¹S. Benkadda, P. Beyer, N. Bian, C. Figarella, O. Garcia, X. Garbet, P. Ghendrih, Y. Sarazin, and P. Diamond, “Bursty transport in tokamak turbulence: Role of zonal flows and internal transport barriers”, *Nuclear Fusion* **41**, 995–1001 (2001).
- ⁵²M. Keilhacker, “H-mode confinement in tokamaks”, *Plasma Physics and Controlled Fusion* **29**, 320 (1987).
- ⁵³R. Hong, G. Tynan, P. Diamond, L. Nie, D. Guo, T. Long, R. Ke, Y. Wu, B. Yuan, and M. Xu, “Edge shear flows and particle transport near the density limit of the HL-2A tokamak”, *Nuclear Fusion* **58**, 016041 (2018).
- ⁵⁴L. Schmitz, “Turbulence and $E \times B$ flow correlations across the L-H transition in DIII-D deuterium and hydrogen plasmas”, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **381**, 10.1098/rsta.2021.0237 (2023).
- ⁵⁵B. F. McMillan, C. C. Pringle, and B. Teaca, “Simple advecting structures and the edge of chaos in subcritical tokamak plasmas”, *Journal of Plasma Physics* **84**, 10.1017/S0022377818001216 (2018).
- ⁵⁶M. Barnes, F. I. Parra, E. G. Highcock, A. A. Schekochihin, S. C. Cowley, and C. M. Roach, “Turbulent transport in tokamak plasmas with rotational shear”, *Physical Review Letters* **106**, 10.1103/PhysRevLett.106.175004 (2011).
- ⁵⁷A. A. Schekochihin, E. G. Highcock, and S. C. Cowley, “Subcritical fluctuations and suppression of turbulence in differentially rotating gyrokinetic plasmas”, *Plasma Physics and Controlled Fusion* **54**, 055011 (2012).

- ⁵⁸E. G. Highcock, M. Barnes, A. A. Schekochihin, F. I. Parra, C. M. Roach, and S. C. Cowley, “Transport Bifurcation in a Rotating Tokamak Plasma”, *Physical Review Letters* **105**, 215003 (2010).
- ⁵⁹H. Bateman, “Some Recent Researches on the Motion of Fluids”, *Monthly Weather Review* **43**, 163–170 (1915).
- ⁶⁰J. Burgers, “A Mathematical Model Illustrating the Theory of Turbulence”, *Advances in Applied Mechanics* **1**, 171–199 (1948).
- ⁶¹D. J. Korteweg and G. de Vries, “On the change of form of long waves advancing in a rectangular canal, and on a new type of long stationary waves”, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **39**, 422–443 (1895).
- ⁶²P. G. Drazin and R. S. Johnson, “The scattering and inverse scattering problems”, in *Solitons* (Cambridge University Press, Feb. 1989), pp. 39–63.
- ⁶³P. Bak, C. Tang, and K. Wiesenfeld, “Self-organized criticality: An explanation of the $1/f$ noise”, *Physical Review Letters* **59**, 381–384 (1987).
- ⁶⁴T. Hwa and M. Kardar, “Avalanches, hydrodynamics, and discharge events in models of sandpiles”, *Physical Review A* **45**, 7002–7023 (1992).
- ⁶⁵A. Hasegawa and K. Mima, “Pseudo-three-dimensional turbulence in magnetized nonuniform plasma”, *Physics of Fluids* **21**, 87 (1978).
- ⁶⁶C. Connaughton, S. Nazarenko, and B. Quinn, “Rossby and drift wave turbulence and zonal flows: The Charney–Hasegawa–Mima model and its extensions”, *Physics Reports* **604**, 1–71 (2015).
- ⁶⁷R. H. Kraichnan, “Inertial Ranges in Two-Dimensional Turbulence”, *Physics of Fluids* **10**, 1417 (1967).
- ⁶⁸G. Boffetta, F. D. Lillo, and S. Musacchio, “Inverse cascade in Charney–Hasegawa–Mima turbulence”, *Europhysics Letters (EPL)* **59**, 687–693 (2002).
- ⁶⁹S. Gallagher, B. Hnat, C. Connaughton, S. Nazarenko, and G. Rowlands, “The modulational instability in the extended Hasegawa–Mima equation with a finite Larmor radius”, *Physics of Plasmas* **19**, 122115 (2012).
- ⁷⁰A. I. Smolyakov, P. H. Diamond, and V. I. Shevchenko, “Zonal flow generation by parametric instability in magnetized plasmas and geostrophic fluids”, *Physics of Plasmas* **7**, 1349–1351 (2000).
- ⁷¹D. Y. Manin and S. V. Nazarenko, “Nonlinear interaction of small-scale Rossby waves with an intense large-scale zonal flow”, *Physics of Fluids* **6**, 1158–1167 (1994).
- ⁷²C. Connaughton, S. Nazarenko, and B. Quinn, “Feedback of zonal flows on wave turbulence driven by small-scale instability in the Charney–Hasegawa–Mima model”, *EPL (Europhysics Letters)* **96**, 25001 (2011).
- ⁷³J. S. Hesthaven, J. P. Lynov, and J. Nycander, “Dynamics of nonstationary dipole vortices”, *Physics of Fluids A: Fluid Dynamics* **5**, 622–629 (1993).

- ⁷⁴A. Hasegawa, C. G. MacLennan, and Y. Kodama, “Nonlinear behavior and turbulence spectra of drift waves and Rossby waves”, *Physics of Fluids* **22**, 2122 (1979).
- ⁷⁵D. Fyfe and D. Montgomery, “Possible inverse cascade behavior for drift-wave turbulence”, *Physics of Fluids* **22**, 246 (1979).
- ⁷⁶S. Zweben and R. Taylor, “Phenomenological comparison of magnetic and electrostatic fluctuations in the Macrotor tokamak”, *Nuclear Fusion* **21**, 193–200 (1981).
- ⁷⁷R. E. Slusher and C. M. Surko, “Study of Density Fluctuations in the Absorption of Oxygen on Silicon”, *Physical Review Letters* **40**, 400–403 (1978).
- ⁷⁸A. Hasegawa and M. Wakatani, “Plasma Edge Turbulence”, *Physical Review Letters* **50**, 1493 (1983).
- ⁷⁹D. E. Newman, B. A. Carreras, P. H. Diamond, and T. S. Hahm, “The dynamics of marginality and self-organized criticality as a paradigm for turbulent transport”, *Physics of Plasmas* **3**, 1858–1866 (1996).
- ⁸⁰P. H. Diamond and T. S. Hahm, “On the dynamics of turbulent transport near marginal stability”, *Physics of Plasmas* **2**, 3640–3649 (1995).
- ⁸¹B. A. Carreras, L. Garcia, and P. H. Diamond, “Theory of resistive pressure-gradient-driven turbulence”, *Physics of Fluids* **30**, 1388 (1987).
- ⁸²P. Beyer and K. H. Spatschek, “Center manifold theory for the dynamics of the L-H-transition”, *Physics of Plasmas* **3**, 995–1004 (1996).
- ⁸³P. Beyer, X. Garbet, and P. Ghendrih, “Tokamak turbulence with stochastic field lines”, *Physics of Plasmas* **5**, 4271–4279 (1998).
- ⁸⁴P. Beyer, S. Benkadda, and X. Garbet, “Proper orthogonal decomposition and Galerkin projection for a three-dimensional plasma dynamical system”, *Physical Review E* **61**, 813–823 (2000).
- ⁸⁵Y. Sarazin, X. Garbet, P. Ghendrih, and S. Benkadda, “Transport due to front propagation in tokamaks”, *Physics of Plasmas* **7**, 1085–1088 (2000).
- ⁸⁶B. Van der Pol, “On “relaxation-oscillations””, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**, 978–992 (1926).
- ⁸⁷C. C. T. Pringle and R. R. Kerswell, “Asymmetric, Helical, and Mirror-Symmetric Traveling Waves in Pipe Flow”, *Physical Review Letters* **99**, 074502 (2007).
- ⁸⁸Y. Duguet, A. P. Willis, and R. R. Kerswell, “Transition in pipe flow: The saddle structure on the boundary of turbulence”, *Journal of Fluid Mechanics* **613**, 255–274 (2008).
- ⁸⁹P. Cvitanović, “Periodic orbits as the skeleton of classical and quantum chaos”, *Physica D* **51**, 138–151 (1991).

- ⁹⁰B. M. Boghosian, A. Brown, J. Lätt, H. Tang, L. M. Fazendeiro, and P. V. Coveney, “Unstable periodic orbits in the Lorenz attractor”, in *Philosophical transactions of the royal society a: mathematical, physical and engineering sciences*, Vol. 369, 1944 (June 2011), pp. 2345–2353.
- ⁹¹K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos* (Springer New York, New York, NY, 1996).
- ⁹²R. M. May, “Simple mathematical models with very complicated dynamics”, *Nature* **261**, 459–467 (1976).
- ⁹³S. Repin, “One Hundred Years of the Galerkin Method”, *Computational Methods in Applied Mathematics* **17**, 351–357 (2017).
- ⁹⁴J. Shen, T. Tang, and L.-L. Wang, *Spectral Methods*, Vol. 41 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011).
- ⁹⁵T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 4th (MIT Press, 2022).
- ⁹⁶*FFTW Homepage*, <https://www.fftw.org>.
- ⁹⁷S. A. Orszag, “On the Elimination of Aliasing in Finite-Difference Schemes by Filtering High-Wavenumber Components”, *J. Atmos. Sci.* **28**, 1074 (1971).
- ⁹⁸C. T. Kelley, *Solving Nonlinear Equations with Newton’s Method* (Society for Industrial and Applied Mathematics, Jan. 2003).
- ⁹⁹B. García-Archilla, J. Sánchez, and C. Simó, “Krylov methods and determinants for detecting bifurcations in one parameter dependent partial differential equations”, *BIT Numerical Mathematics* **46**, 731–757 (2006).
- ¹⁰⁰J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations* (Society for Industrial and Applied Mathematics, Jan. 1996).
- ¹⁰¹L. Trefethen and D. Bau, *Numerical Linear Algebra* (Society for Industrial and Applied Mathematics, 1997).
- ¹⁰²*LAPACK Linear Algebra Package*, netlib.org/lapack.

Appendix A

Newton-GMRES-Hookstep Algorithm

This Appendix contains a detailed outline of the NKH algorithm used in the thesis, described in Section 4.1. The algorithm is presented in pseudocode below.

```

/*Iterative search for RPOs. A is a m-by-m matrix.
  f(x) uses the parameter terms in x to evolve the
  state space terms. */
Input:  $f, initX$ 
Output:  $xRPO$ 
 $x \leftarrow initX;$ 
while True do
   $r \leftarrow f(x) - x;$ 
   $q_1 \leftarrow r/\|r\|;$ 
  foreach  $n = 1, 2, \dots, m$  do
     $v \leftarrow (f(x + q_i) - f(x - q_i)/2) + q_i;$ 
    foreach  $j = 1$  to  $n$  do
       $h_{j,n} \leftarrow q_j^*v;$ 
       $v \leftarrow v - h_{j,n}q_j$ 
    end
     $e1 \leftarrow \{1, 0, 0, \dots, 0\};$ 
     $h_{n+1,n} \leftarrow \|v\|;$ 
     $q_{n+1} \leftarrow v/h_{n+1,n};$ 
     $y \leftarrow \min_y(\|Hy - \|r\|e1\|);$ 
     $gres \leftarrow \|Hy - \|r\|e1\|;$ 
    if ( $gres < gtol$ ) or ( $n \geq gits$ ) then
       $U, D, VT = SVD(H);$ 
       $p \leftarrow \|r\|U_{1,:};$ 
       $\mu \leftarrow \min_{\mu}(\sum_i(p_i D_{i,i}/(\mu + D_{i,i}^2))^2 - \delta^2);$ 
      foreach  $j = 1$  to  $n$  do
         $q_j = p_j d_j / (\mu + D_{i,i}^2)$ 
      end
       $dx \leftarrow HVT^\dagger q$  break;
    end
  end
   $x \leftarrow x + dx;$ 
   $res \leftarrow |f(x) - x|;$ 
end

```

Algorithm 1: Newton-Krylov-Hookstep Algorithm