

Bistable turbulent transport in fusion plasmas with rotational shear



Nicolas Christen
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Theoretical Physics

Trinity 2021

Bistable turbulent transport in fusion plasmas with rotational shear

Nicolas Christen

Lincoln College, University of Oxford

Submitted for the degree of Doctor of Philosophy in Theoretical Physics

Trinity 2021

Turbulent transport is a major limiting factor for the energy yield of magnetic-confinement fusion experiments. Turbulence is driven by gradients in the profiles of equilibrium plasma parameters, and it can be regulated by the presence of sheared flows. Direct numerical simulations of turbulent transport are challenging because of the broad range of temporal and spatial scales characterising fusion plasmas. In this thesis, a novel algorithm is developed to include an externally imposed, mean rotational flow shear in linearly implicit, local δf gyrokinetic codes. Unlike the discrete-in-time approach currently followed by a number of codes, it treats flow shear continuously in time, allowing for sparser numerical grids to be used. In some cases, the new algorithm leads to a difference in the level of turbulent transport compared to the discrete-in-time approach. In the prevailing paradigm for plasma turbulence, a unique, stationary turbulent state is associated to given equilibrium plasma parameters. This thesis reports the first instance of bistable turbulence in a fusion plasma with rotational shear. Two distinct states, obtained with identical equilibrium parameters in first-principle gyrokinetic simulations, have turbulent fluxes of particles, momentum and energy that differ by an order of magnitude – with the low-transport state agreeing with experimental observations. The occurrence of the two states is regulated by the competition between the externally imposed, mean rotational flow shear and “zonal” flows generated by the plasma. With small turbulent amplitudes, zonal flows have little impact on transport, and the mean shear causes turbulence to saturate in a low-transport state. With larger amplitudes, the zonal shear can substantially oppose the effect of the mean shear, allowing the system to sustain a high-transport state. This poses a new challenge for research that, so far, has assumed a uniquely defined turbulent state.

Acknowledgements

As the completion of this DPhil draws nearer, I realise how grateful I am to be here, writing these words. Above all, I wish to thank everyone who made this journey possible.

First, I would like to thank my supervisor Michael Barnes for his truly inspiring work, his modesty, and for his unwavering support throughout my DPhil. I also thank Alex Schekochihin and Felix Parra for their insightful feedback, and for giving me the opportunity to join the plasma theory group in Oxford. Many thanks to Henri Weisen and Paula Sirén for providing the experimental data used in this thesis. Thanks to Christian Theiler, Paolo Ricci and others at the Swiss Plasma Center for their trust and their highly enjoyable lectures. And a special thank you to my high-school teacher Éric Nidegger for his entertaining physics classes.

I wish to thank my friends for the moments that we shared. Virgile, Colin, Nathan, Nico, Val, Vincent and Lily, for those amazing years at EPFL. Kevin, did you see the dolphins? Matthias, for those glimpses of Oxonian extravagance and for our expeditions to Wimbledon. Michael, for helping me in my attempts to understand gyrokinetics and the British way of life. Adwiteey, Ollie, Mantas, Alessandro and all the plasmaniaks, for our ping-pong sessions, our constant supply of Ben's cookies, and for all the moments that made this DPhil such a nice experience.

My thoughts go to my entire family. To Papa, Maman, Alain and Isabelle, with whom I always fostered a longing for new adventures. To Grand-papa and Grand-maman for these wonderful years spent together in Lausanne as I was starting my studies. To Tonton, for sharing his enthusiasm for

reading, cycling, and amazing breakfasts. To everyone who came to visit me in Oxford.

And finally, I wish to thank my partner Andrea for her unconditional love and support – I am already looking forward to our next adventures together.

The work in this thesis was supported by a scholarship from the Berrow Foundation, a scholarship from the Steppes Fund for Change, and a grant from the Fondation Hélène et Victor Barbour. The costs of travel and attendance to conferences were funded by Lincoln College and the Department of Physics of the University of Oxford. The author acknowledges use of the high performance computing cluster Hydra at the Theoretical Physics subdepartment of the University of Oxford, use of the ARCHER High Performance Computer through the Plasma HEC Consortium, EPSRC grant EP/L000237/1 under the projects e281-gs2 and e607, the use of the EUROfusion High Performance Computer (Marconi-Fusion) under the projects FUA32_MULTEI, FUA33_MULTEI, FUA34_MULTEI and FUA35_OXGK, and the use of the JFRS-1 supercomputer system at the International Fusion Energy Research Centre's Computational Simulation Centre (IFERC-CSC) at the Rokkasho Fusion Institute of QST (Aomori, Japan) under the project MULTEI.

Contents

1	Introduction	17
1.1	The need for sustainable energy production	18
1.2	Magnetic-confinement fusion and the tokamak design	18
1.3	Flow shear imposed by neutral beam injection	21
1.4	Transport with flow shear	21
1.5	Thesis motivation and outline	24
2	Simulating turbulent transport with a mean flow shear	26
2.1	Introduction	26
2.2	Model	27
2.2.1	The gyrokinetic approximation	27
2.2.2	Normalisation	30
2.2.3	Simulation domain, boundary conditions and system of coordinates	31
2.2.4	Spectral representation in the presence of flow shear	34
2.2.5	Fourier-transformed system of equations	35
2.2.6	Linear instabilities and Floquet oscillations	37
2.2.7	Mode-rational magnetic-flux surfaces	40
2.3	GS2 algorithm in the presence of flow shear	43
2.3.1	Discretised system of equations	44
2.3.2	Overview of the GS2 algorithm	44
2.3.3	Challenges in the presence of flow shear	46
2.3.4	Discrete-in-time approach to flow shear	46
2.3.5	Continuous-in-time approach to flow shear	49
2.4	Comparison of the approaches to flow shear	51

2.4.1	Linear simulations	52
2.4.2	Test cases for Fourier transforms	53
2.4.3	Nonlinear simulations	58
2.5	Conclusions	63
3	Bistable turbulence in the presence of flow shear	67
3.1	Introduction	67
3.2	Running nonlinear simulations with a mean flow shear	68
3.3	Existence of two distinct turbulent states	71
3.3.1	Identifying two distinct turbulent states	72
3.3.2	Importance of zonal modes	77
3.3.3	Nonlinear decorrelation times	79
3.3.4	Identifying the region of bistability	82
3.3.5	On the role of mode-rational surfaces	85
3.4	Consequences of bistability	89
3.4.1	Bifurcations in turbulent transport	89
3.4.2	Mean-gradient relaxation cycles	91
3.4.3	Implications for the numerical modelling of turbulence	92
3.5	Conclusions	93
4	Summary and conclusions	95
A	Using JETPEAK equilibria with the GS2 code	100
A.1	Sign conventions	100
A.2	Normalisation for the magnetic field	102
B	JET discharge	104
C	Discretisation in the GS2 code	106
D	Source code and input parameters	109
D.1	Source code	109
D.2	GS2 input parameters	110

E Reconstruction method for toroidal rotation profiles	115
E.1 Transport equations	115
E.1.1 Particle transport	116
E.1.2 Momentum transport	116
E.1.3 Energy transport	117
E.2 Computing fluxes from equilibrium profiles	117
E.3 Reconstruction method	118
E.4 Predicting rotation profiles	120
References	122

List of Figures

1.1	Sketch of the toroidal magnetic geometry in the core of a tokamak. A magnetic field line is shown in red, and the trajectory of a charged particle is illustrated in blue. The toroidal angle is denoted by ϕ and the poloidal angle by θ . The concentric circles represent nested magnetic flux surfaces. The center of the flux surfaces, where the magnetic field becomes purely toroidal, is known as the magnetic axis. The major radius is the distance from the vertical axis of symmetry and is denoted by R . The minor radius measures the distance from the magnetic axis and is denoted by r	20
2.1	In 2.1a, the shaded grey area corresponds to the outboard mid-plane. Looking down onto this plane from above in 2.1b, the edges of the flux-tube are shown in solid blue lines as it intersects with the plane at $\theta = 0$ (light blue) and $\theta = 2\pi$ (dark blue). The dashed blue lines indicate exact copies of the interSection at $\theta = 0$	32
2.2	Illustration of the basic mechanism of toroidal ITG at the outboard midplane of the tokamak, with the curvature of magnetic field lines denoted by $\kappa = \hat{\mathbf{b}} \cdot \nabla \hat{\mathbf{b}}$. At the inboard midplane, ∇T points in the opposite direction and \mathbf{V}_E tends to stabilise the temperature perturbation.	37
2.3	Linear growth rate and associates real frequencies. In the top panels $a/L_{T_i} = 0.87$ and in the bottom panels $a/L_{T_i} = 0.1.913$. Here flow shear has been neglected, i.e. $\gamma_E = 0$. All other physical parameters are taken from the JET discharge presented in Appendix B.	38

2.4	Typical Floquet oscillations in the electrostatic potential averaged along the magnetic field. Here, $\rho_i k_y = 0.5$, $a/L_{Ti} = 2.09$ and $\gamma_E = -0.1$. All other parameters are set to their values presented in Appendix B. . . .	39
2.5	Sketch illustrating the twist-and-shift parallel boundary condition. The light-blue square and dark-blue trapezoid show the cross-Section of the flux tube at $\theta = 0$ and $\theta = 2\pi$, respectively. The dashed light-blue squares represent periodic duplicates of the $\theta = 0$ cross-Section. The position of a given magnetic-field line at $\theta = 0$ is shown by the light-blue dot, and its position at $\theta = 2\pi$ is indicated by the dark-blue dot. The light-blue cross indicates the field line of the $\theta = 0$ cross-Section whose position coincides with the dark-blue dot in one of the periodic duplicates.	41
2.6	Cartoon illustrating how the wave-number re-mapping amounts to moving the computational domain in θ . Initially (2.6a), the structure is twisted by magnetic shear. Over a period of time T_{map} , flow shear contributes to the twisting, enhancing it at one end and decreasing it at the other (2.6b). In order to always include radially extended structures in the simulation, the computational domain is shifted in θ via a wave-number re-map (2.6c).	53
2.7	Time evolution of the normalised electrostatic potential for a single twist-and-shift chain, summed over the parallel coordinate. Panel 2.7a shows the results obtained with the widely used discrete-in-time approach for flow shear. Panel 2.7b shows the results obtained with our new approach, which treats flow shear continuously over time. We point out that the inserts have a different scale for their y -axis. Here, $\rho_i k_y = 0.3$, $\gamma_E = 0.2$ and all other physical parameters are set to their Cyclone Base Case values.	54

2.8	Time evolution of Ψ in the lab frame, according to equation (2.56). The discrete-in-time approach (2.8a) does not include the phase factor appearing in Fourier transforms, whereas the continuous-in-time approach (2.8b) does. Both methods become equivalent at $t = T$, when the NGP approximation is exact. The dashed black line corresponds to $y^* = 0$ and shows the tilt that the mean flow shear should be producing. The numerical parameters for this first test case are given in Section 2.4.2.	56
2.9	Analytical solution to equations (2.60) and (2.61) for Ψ , and comparison to numerical solutions obtained with the discrete- and continuous-in-time approaches. At $t = T$, the NGP approximation is exact, making both approaches equivalent for that time step. The numerical parameters for this second test case are given in Section 2.4.2.	59
2.10	$ \Delta\Psi = \Psi_{\text{num}} - \Psi $ for the nominal Δk_x and for Δk_x halved, with both discrete- and continuous-in-time approaches.	60
2.11	Turbulent heat flux, plotted for different Δk_x in 2.11a, and for different $T_{\text{map}}(k_y)/T_F$ in 2.11b (here we choose $\rho_i k_y = 0.5$ which is the wave number with the largest amplitudes). All points in 2.11a and 2.11b are identical, except for the ones highlighted by square boxes, for which both Δk_x and Δk_y were modified. Here, $a/L_{T_i} = 1.756$, $\gamma_E = -0.063$, and all other parameters are set to their values presented in Appendix B.	64
2.12	Amplitude of the electrostatic potential for the box-scale zonal mode in solid black and for the dominant non-zonal mode in dashed red. Here, $a/L_{T_i} = 1.756$, $\gamma_E = -0.063$, and all other parameters are set to their values presented in Appendix B.	65
3.1	Two distinct saturated turbulent states are obtained depending on whether amplitudes are small (low-transport state in blue) or large (high-transport state in red) when flow shear is included in the simulation. Here, $a/L_{T_i} = 1.76$ and all other parameters are set to their values presented in Appendix B.	70

3.2	Radial profiles of the experimentally observed fluxes, compared to the turbulent fluxes obtained from simulations performed by either including flow shear from the start (i.e., small initial amplitudes), or by restarting from a zero-shear saturated state (i.e., large initial amplitudes). Panel (a) shows the heat flux, and panel (b) the momentum-to-heat-flux ratio. Here, apart from $a/L_{T_i} = 1.76$ and the varying value for r_ψ/a , all parameters are set to their values presented in Appendix B.	70
3.3	Dependence of the turbulent ion heat flux on the inverse ion-temperature-gradient scale length. In the simulations labelled by green ‘+’ signs, the externally imposed mean flow shear was set to zero. For all other simulations, $\gamma_E = -0.079$ was used. Zonal modes are artificially zeroed out in simulations labelled by black crosses. The black circle denotes a simulation where amplitudes decay in time and no saturated turbulent state is observed. The dashed line marks the temperature gradient below which there is no effective linear instability with $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.	73
3.4	Occurrences of the two distinct turbulent states. Purple diamonds indicate parameters at which both a high- and a low-transport state can be obtained. Black circles denote simulations where amplitudes are decaying in time and no saturation of the turbulence is observed. Lines indicate contours of constant $\langle\gamma\rangle_t/\gamma_E$, where $\langle\gamma\rangle_t$ is the effective linear growth rate. The dashed line marks the temperature gradient below which $\langle\gamma\rangle_t < 0$. The solid line indicates $T_F\langle\gamma\rangle_t = 1$ (recalling that $T_F = 2\pi\hat{s}/\gamma_E$). Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. $\langle\gamma\rangle_t$ is obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$	74

3.5	Consecutive snapshots of the turbulence in real space for the low-transport state where $a/L_{T_i} = 1.76$ and $\gamma_E = -0.079$. In the top panels, the fluctuating electrostatic potential is plotted at three successive times at the outboard of the torus, in the plane perpendicular to \mathbf{B} . In the bottom panels, the zonal flow is plotted at the same times. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.	75
3.6	Consecutive snapshots of the turbulence in real space for the high-transport state where $a/L_{T_i} = 1.76$ and $\gamma_E = -0.079$. In the top panels, the fluctuating electrostatic potential is plotted at three successive times at the outboard of the torus, in the plane perpendicular to \mathbf{B} . In the bottom panels, the zonal flow is plotted at the same times. The zonal shear averaged between the two vertical dashed lines is compared to the externally imposed mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.	76
3.7	Energy spectra for a high-transport (a,c,e) and a low-transport (b,d,f) state obtained with identical equilibrium parameters. Panels (c) to (f) compare the spectra to predictions of the critical-balance argument, which does not take into account the effect from mean flow shear. Here, $a/L_{T_i} = 2.0$, $\gamma_E = -0.105$, and all other parameters are set to their values presented in Appendix B.	77
3.8	Dependence of the turbulent ion heat flux on the inverse ion-temperature-gradient scale length. Here, $\gamma_E = -0.079$. Zonal modes are artificially zeroed out in simulations labelled by black crosses. The dashed line marks the temperature gradient below which there is no linear instability in the presence of mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.	79

- 3.9 Two-point spatial correlation function for a low-transport state (a), a high-transport state (b) and a state with no mean flow shear (c). The correlation lengths ℓ_x and ℓ_y correspond to the e -folding of $\text{Cor}[\varphi_{\text{NZ}}]$ along the solid and dashed lines, respectively. In the three states shown here, $a/L_{T_i} = 1.76$. In (a) and (b), $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. . . . 80
- 3.10 Comparison between linear time scales and nonlinear decorrelation time scales. Here $\gamma_E = -0.079$, except for the cases indicated by green squares which have no mean flow shear. Simulations to the left of the dashed vertical line are linearly stable on average with $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. Linear growth rates are obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest instantaneous and average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$ 81
- 3.11 Dependence of the product of the Floquet period and the time-averaged linear growth rate on the mean flow shear rate and the ion-temperature-gradient-length scale. We indicate the region where the low-transport state is observed by the dotted line with downward-pointing triangles (nonlinear simulations were only performed for $|\gamma_E| \leq 0.105$). The necessary condition (3.4) for the existence of low-transport states is satisfied below the solid black line. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. Linear growth rates $\langle \gamma \rangle_t$ is obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$. 83

- 3.12 Estimates used for the scaling argument at the end of Section 3.3.4. Panel (a) shows the estimated zonal flow shear rate. Panel (b) shows the relative magnitudes of zonal and non-zonal modes. Panel (c) compares the non-zonal radial correlation length ℓ_x to the zonal radial correlation length $\ell_{x,Z}$. Panel (c) compares the non-zonal correlation lengths in x and y . Here, $\gamma_E = -0.079$, except for simulations indicated by green squares which are obtained without mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. . 84
- 3.13 Time traces of the turbulent ion heat flux, showing the fragility of the low-transport state close to the threshold above which low-transport states cease to exist. Panel (a) shows a scan in the maximum radial wave number K_x , and panel (b) shows how these results change when the spacing Δk_y for wave numbers in the y direction is halved. Here, $\rho_i K_x = 1.7$, $\rho_i \Delta k_y = 0.09$, $\rho_i K_y = 0.9$, $a/L_{T_i} = 2.087$, $\gamma_E = -0.105$ and all other parameters are set to their values presented in Appendix B. . 87
- 3.14 Effect of mode-rational surfaces in simulations with fully kinetic electrons, and in simulations with an imposed Boltzmann response for passing non-zonal electrons (labelled as “adiabatic passing electrons”). Panel (a) shows the time traces of the turbulent ion heat flux, where the horizontal bars indicate the windows used for time averages. Panel (b) shows the time-averaged radial-wave-number spectrum of the zonal electrostatic potential. The dots indicate $2\pi/\Delta_{x,m}$, where $\Delta_{x,m}$ is the distance between two mode-rational surfaces of order m . Panel (c) shows the corresponding time-averaged radial profiles of the zonal flow shear. Here, $\rho_i K_x = 1.7$, $\rho_i \Delta k_y = 0.09$, $\rho_i K_y = 0.9$, $a/L_{T_i} = 2.087$, $\gamma_E = -0.105$ and all other parameters are set to their values presented in Appendix B. . 88

3.15	Threshold in temperature gradient above which low-transport states cease to exist, for simulations with fully kinetic electrons, and simulations with an imposed Boltzmann response for passing non-zonal electrons (labelled as “adiabatic passing electrons”). Here, apart from $\gamma_E = -0.079$ and the varying a/L_{T_i} , all other parameters are set to their values presented in Appendix B.	89
3.16	Illustration of bifurcations and hysteresis in the turbulent heat flux as a/L_{T_i} is modified. The different stages, labelled as A-B-C-D on the sketch in panel (a), each correspond to a stage in the sequence of simulations shown on the right. Panel (b) indicates the sequence of temperature gradients that are given as inputs to the code. Each change in the gradient requires a new simulation, where the final state of the previous simulation is used as an initial condition. Panel (c) shows the corresponding turbulent heat flux that is obtained. Note that step D has been run for a longer time (not shown in this Figure) to ensure that it does indeed correspond to a saturated low-transport state. Here, apart from $\gamma_E = -0.079$ and the varying a/L_{T_i} , all other parameters are set to their values presented in Appendix B.	90
3.17	Dependence of the ion heat flux (a,b) and the momentum-to-heat-flux ratio (c,d) on the imposed flow shear and the inverse ion-temperature-gradient scale length. The top panels show results for the high-transport states, and the bottom panels show the low-transport states. We point out that there is a gap between the values of the heat flux obtained in (a) and those obtained in (b). The constant Π_i/Q_i path defined by points 1, 2, 3 and 4 gives an example of the successive stages of a gradient-relaxation cycle, when the heat injected into the plasma corresponds to a steady-state flux within the aforementioned gap. Here, apart from γ_E and a/L_{T_i} , all other parameters are set to their values presented in Appendix B.	91

B.1	Radial profiles of the JET discharge #68448. Panel (a) shows the temperature, panel (b) the particle density and panel (c) the mean rotation frequency.	104
B.2	Radial profiles of the energy deposition (a) and of the torque deposition (b) in the discharge #68448. $S_{\varepsilon,i}$ is the external source of energy, and $S_{\Delta,i}$ is the energy transfer with electrons (see equation (E.8)). The deposition is computed based on the mean plasma parameter profiles, using the ASCOT code. Section E.2 explains how the turbulent fluxes in the experiment can be computed from the profiles (a) and (b).	105
E.1	Comparison between the experimental rotation profile (black line), and the rotation profile reconstructed by following the method presented in Section E.3 (red circles). The blue ‘+’ symbols are obtained via the same reconstruction method, but by replacing the momentum flux obtained from local gyrokinetic simulations by the actual momentum-flux profile from the experiment. For cases where the simulations match the transport levels of the experiment, the red circles and the blue ‘+’ symbols would give the same result.	120

Chapter 1

Introduction

The aim of research on magnetic-confinement fusion plasmas is to develop a steady source of sustainable energy. A major challenge in increasing the energy yield of fusion experiments is the transport of particles, momentum and heat by turbulent processes. The turbulence is driven by gradients in the mean plasma parameters (e.g. temperature), and can be regulated by the presence of sheared flows. Until now, it was widely assumed that, for a given set of mean plasma parameters, a unique stationary turbulent state should be obtained. The focus of this thesis is a study of turbulent transport in the presence of an externally imposed, mean rotational flow shear, guided by numerical simulations. The main result is that two distinct turbulent states can be obtained from identical mean plasma parameters – a phenomenon known as bistability.

The structure of this introductory Chapter is the following. In Section 1.1 we present the motivation behind research on magnetic-confinement fusion. In Section 1.2 we give a brief overview of the fundamental principles of magnetic-confinement fusion, and explain the so-called tokamak design concept for fusion experiments. In Section 1.3 we briefly explain how the injection of particle beams is used to fuel and heat the plasma, and how the beams generate a mean rotational flow in the plasma. In Section 1.4 we discuss the wide range of spatio-temporal scales characterising turbulence in fusion plasmas. We mention the challenges this poses to numerical simulations and briefly summarise some of the progress that has been made over the past decades in understanding the impact of sheared flows on turbulent transport. Finally, in Section 1.5, we present the main motivations for this work, as well as the outline of the thesis.

1.1 The need for sustainable energy production

The Intergovernmental Panel on Climate Change (IPCC) has gathered evidence from a broad body of scientific work indicating that the rise of the global mean surface temperature since pre-industrial levels is almost certainly caused by human activity [1]. The work shows that this rise in temperature has increased the frequency and magnitude of negative impacts from weather events on human societies.

Since the Paris Agreement was signed in 2015, the United Nations Framework Convention on Climate Change pledged to limit global warming to less than 2°C above pre-industrial levels, and to work towards limiting it to 1.5°C [2]. If the current rate of emissions is maintained, the IPCC estimates that the 1.5°C threshold is likely to be exceeded before the year 2052 [3].

In 2018, about three quarters of the global greenhouse gas emissions stemmed from energy production [4]. Thus, in order to achieve the goal set by the Paris Agreement, the world's energy consumption needs to shift towards electricity that is produced with lower greenhouse gas emissions. The energy yield from renewable sources such as wind turbines or solar panels is however subject to large variations throughout days and seasons. For these renewable sources to meet the global demand for energy, storage systems would need to be greatly expanded and new storage technologies would need to be developed [5]. Therefore, it is likely that another, more steady source of sustainable energy will be required to complement the varying contribution from renewables.

Research on magnetic-confinement fusion aims to develop a new, steady source of sustainable energy, which could become a solution on the long term together with renewable sources, without many of the drawbacks of nuclear fission.

1.2 Magnetic-confinement fusion and the tokamak design

The fundamental principle in magnetic-confinement fusion is to capture the energy of fast neutrons produced by fusing hydrogenic nuclei. In order for the positively charged nuclei to overcome their mutual repulsion from Coulomb forces and for them to fuse,

the density of nuclei and their energy must become large enough. At realistic energies, the most favourable reaction for fusion is given by



where a deuterium nucleus and a tritium nucleus fuse to produce an energetic α -particle (3.5 MeV) and an energetic neutron (14.1 MeV) [6]. The aim for energy production is to generate as many fusion reactions as possible in a system at steady-state.

Currently, the favoured approach towards a fusion power plant relies on the reactor design called a tokamak. In tokamaks, a hydrogenic plasma – i.e., an ionised gas where electrons are stripped from their nucleus – is confined by a toroidally shaped magnetic cage. Magnetic field lines form this cage by winding around the torus, tracing out nested toroidal surfaces known as magnetic flux surfaces (see Figure 1.1). Near the edge of the cage, flux surfaces cease to exist as magnetic field lines intersect with the inner wall of the tokamak. Throughout this work, we refer to the central region of the plasma, where magnetic field lines do not intersect with the wall, as “the core”. The rapid gyromotion of charged particles about magnetic field lines, due to the Lorentz force, mostly constrains the particles to evolve in the vicinity of a given flux surface. This confinement reduces outward transport of particles, heat and momentum from the core to the edge of the torus, allowing temperatures and densities to rise in the core. This in turn increases the probability for the ions to fuse and, ultimately, improves the energy yield of the device [6].

However, the magnetic confinement of the plasma is not perfect, and some particles still diffuse across magnetic flux surfaces. These particles eventually escape from the core and stop contributing to the fusion reactions. In order to compensate for these losses, the plasma must be continuously fuelled and heated. The heating can come either from external sources, or from the energetic α -particles produced by the fusion reactions. The typical time scale associated with the confinement of energy in a tokamak can then be expressed as

$$\tau_E = \frac{\int_V d^3r \frac{3}{2} \sum_s n_s T_s}{P_{\text{heat}}}, \quad (1.2)$$




Figure 1.1: Sketch of the toroidal magnetic geometry in the core of a tokamak. A magnetic field line is shown in red, and the trajectory of a charged particle is illustrated in blue. The toroidal angle is denoted by ϕ and the poloidal angle by θ . The concentric circles represent nested magnetic flux surfaces. The center of the flux surfaces, where the magnetic field becomes purely toroidal, is known as the magnetic axis. The major radius is the distance from the vertical axis of symmetry and is denoted by R . The minor radius measures the distance from the magnetic axis and is denoted by r .

where n_s and T_s respectively denote the density and temperature of particle species s , V is the volume of the confined plasma and P_{heat} is the total heating power [7]. Note that, throughout this work, we follow the standard notation by absorbing the Boltzmann constant k_B in the definition of the temperatures T_s . The energy confinement time (1.2) must be maximised to reduce energy losses.

Present-day tokamaks were built with the intention to focus on research rather than to produce energy, and so far, all experiments require a higher energy input than they can produce from fusion reactions. At the time of writing, the closest attempt towards breakeven was a discharge at the Joint European Torus experiment (JET) in 1999, where the fusion power reached 67% of the input power [8]. Several projects are aiming to surpass this record, and to show that magnetic-confinement fusion can become viable for energy production. At the time of writing, the most advanced of these projects are the large, experimental reactor ITER, with deuterium-tritium experiments due to start in 2036 [9], and the compact, high-magnetic-field tokamak SPARC whose construction is due to start in 2021 [10].

1.3 Flow shear imposed by neutral beam injection

One of the main methods to heat and fuel the plasma in tokamaks consists of injecting beams of energetic neutral atoms of hydrogen isotopes into the confined plasma [11]. The neutral atoms can penetrate the plasma unaffected by the magnetic field, until they are ionised through collisions and become part of the confined plasma. However, accelerating neutral atoms is a complicated process: first, the atoms need to be ionised; the ions are then accelerated by an electric field; and finally they need to recombine into neutral atoms before penetrating into the tokamak. The devices required to achieve this can be very large and costly. For example, each of the three neutral beam injectors for ITER will measure 26m in length [12] (in comparison, the major radius of the tokamak to the magnetic axis will be about 6m), and the construction of the first full-scale prototype is still under way at the time of writing [13, 14].

While neutral beam injection may be challenging from an engineering perspective, one particularity that is interesting about this heating technique is its ability to make the confined plasma rotate. As the neutral atoms collide with the plasma, ionise, and eventually slow down, they transfer some of their momentum to the plasma, making it spin in the toroidal direction. Typically, most of the energy and momentum of the beam is deposited around the centre of the confined plasma, which results in the plasma spinning faster in the core of the confined region than near the edge. As we will discuss in the following Sections, this shear in the toroidal plasma flow can be used to increase the energy confinement time (1.2) in fusion experiments.

1.4 Transport with flow shear

A major limitation to the energy yield of tokamaks is the transport of particles, momentum and energy across magnetic flux surfaces. This transport, which regulates the energy confinement time, can be due to a variety of processes.

Classical transport describes how particles diffuse across flux surfaces via Coulomb collisions, resulting in outward fluxes of particles, momentum and heat. A typical diffusion coefficient associated with classical transport can be estimated as the product of the collision frequency times the radius of a particle's gyromotion. However, classical

diffusion coefficients are known to be orders of magnitude too small to explain the transport levels observed in tokamaks.

Neoclassical transport [15, 16] takes into account the geometrical effects on the diffusion of particles. A key point is that particles can become trapped on the outboard side of a tokamak due to magnetic mirroring. Indeed, the magnetic field strength in tokamaks is larger on the inboard side of the torus, which prevents particles with low velocities parallel to the magnetic field from sampling the whole torus poloidally. Trapped particle orbits effectively lead to larger diffusivities than classical transport, but those remain orders of magnitude below diffusivities typically observed in tokamaks.

Transport in tokamaks is dominated by another phenomenon: turbulence [17–19]. Hence, this thesis focuses specifically on turbulent transport. In the context of magnetic-confinement fusion, “turbulence” denotes the chaotic and collective dynamics of charged particles and electromagnetic fields. As the core of the plasma is continuously fuelled and heated, the ensuing gradients in the radial profiles of mean plasma parameters (e.g., particle density, toroidal flow and temperature) represent sources of free energy that can drive microinstabilities. Turbulent transport in tokamaks is sustained by these microinstabilities, establishing a feedback loop with the mean plasma parameter profiles: the mean profiles drive radial turbulent transport; the turbulent transport modifies the mean profiles; and the process repeats itself until a stationary state is reached. In order to make quantitative estimates of the energy confinement time, and to guide the design of future fusion devices, it is therefore essential to develop our understanding of turbulent transport. This work focuses specifically on turbulent transport within the core region of the plasma.

Direct numerical simulations of turbulent transport are challenging because of the wide range of temporal and spatial scales that characterise tokamak plasmas. The mean profiles typically evolve on the energy confinement time scale (estimate for ITER: $\tau_E = 3.7\text{s}$ [20]), and they have a characteristic length scale of the order of the minor radius of the torus (e.g., $\sim 2\text{m}$ in ITER [20]). Meanwhile, the periods of gyromotion of particles around magnetic field lines are much shorter than τ_E (estimate for electrons in ITER: $\sim 10^{-12}\text{s}$ [20]), and the gyroradii of particles are typically much shorter than the minor radius (estimate for the electron gyroradius in ITER: $\sim 10^{-4}\text{m}$, assuming a

core temperature of 20keV). In addition, collisions are weak in the hot core of fusion plasmas. Hence, fluid models such as Braginskii’s equations [21] and their successive corrections [22–25] can only be used to simulate the plasma at the edge of the tokamak, where the plasma is cooler and more collisional. In the hotter core that is the focus of this thesis, kinetic effects must be retained, which increases the computational cost of simulations. Present-day computing capabilities are by far insufficient to run a simulation with a kinetic approach that includes all the wide ranges of spatio-temporal scales discussed here.

Significant progress in direct numerical simulation of plasma turbulence has been achieved through the development of scale-separated methods, such as the widely used δf gyrokinetic approach [26–29]. As we explain in Chapter 2, this approach takes advantage of the separation between scales associated with the mean plasma parameters and those associated with the turbulence. In particular, the approach enables direct numerical simulations to explore the impact of flow shear on turbulence. It is important to note, however, that simulations of plasma turbulence remain computationally expensive. As an example, a single simulation such as we describe in Chapter 2 requires $\sim 1.5 \times 10^4$ CPU-hours to compute the level of turbulent transport for some given mean plasma parameters.

Over the past decades, it was established that sheared flows play an important role in regulating turbulent transport. In the absence of an externally imposed mean flow shear, the plasma is known to spontaneously generate sheared “zonal” flows, which contribute to the saturation of turbulence [30–36]. So far, it has therefore been assumed that the effect of zonal flows is always to suppress turbulence. When a mean flow shear is imposed externally, it provides an additional mechanism for saturation. Specifically, the shear in the flow perpendicular to the mean magnetic field was found to reduce turbulent fluctuations. This effect was initially studied with linearised models [37–39], and was later observed in a vast number of nonlinear gyrokinetic simulations, see e.g. [40–46]. The suppression of turbulence by an externally imposed flow shear has also been confirmed by experiments in tokamaks with neutral beam injection [47, 48].

In this context, the contributions of the present thesis is twofold: to improve how mean flow shear is included in simulations of turbulent transport, and to shed new light on turbulence suppression by flow shear.

1.5 Thesis motivation and outline

In order to determine the level of turbulent transport in the presence of a mean flow shear, the vast majority of δf gyrokinetic simulations rely on an algorithm developed by Hammett *et al* [49]. As we describe in Chapter 2, a limitation of this algorithm is that the perpendicular flow shear is treated discretely in time: simulations feature periods of time where flow shear has no effect, interspersed with discrete jumps to account for the shearing. This is especially concerning for simulations of experiments with moderate levels of flow shear, as the time periods between discrete jumps become longer with weaker flow shear. A first attempt to address this concern was made in [50], but the proposed method is computationally expensive and raises new issues near the boundaries of the computational domain. A more efficient solution was proposed in [51]. While this solution focuses on correcting nonlinear interactions in gyrokinetic simulations with flow shear, it can be extended to also cover linear physics. However, this extension is not straightforward for codes whose algorithm is designed to speed up linear calculations – known as linearly implicit codes. The first contribution of this thesis is to develop a new algorithm to treat flow shear continuously over time in linearly-implicit gyrokinetic codes.

This work follows the widely used δf gyrokinetic approach, which we describe in Chapter 2. In this approach, the evolution of turbulence is simulated on top of an effectively fixed mean background. Given some mean profile of plasma parameters, the simulation determines the associated turbulent fluxes of particles, momentum and heat. In experiments, however, it is the inverse process that occurs: the fluxes are fixed by external sources, and the plasma then determines the associated mean profiles of density, flow and temperature. Until now, the paradigm in magnetic-confinement fusion has been to assume a one-to-one correspondence between the mean plasma parameters and the stationary turbulent state. Assuming a one-to-one correspondence,

it becomes essentially equivalent to solve the gradient-to-flux problem or its inverse, i.e., for some given mean plasma parameters, the simulated fluxes must correspond to the fluxes observed in the experiment. However, studies of turbulence in other contexts show that distinct turbulent states can sometimes be obtained with identical mean parameters – a phenomenon known as multistability. Examples of multistability were for instance reported in neutral fluid turbulence [52–56] and in magnetohydrodynamics [57, 58]. The second contribution of this thesis is to challenge the assumption of a one-to-one correspondence between mean gradients and fluxes in the context of magnetic-confinement fusion, and to show that two distinct turbulent states can be obtained with identical mean parameters in the presence of mean flow shear.

The thesis is organised as follows. In Chapter 2, a novel algorithm is developed to simulate flow shear in linearly implicit gyrokinetic codes. The implementation of the new algorithm is validated, and its advantages compared to the widely used method by Hammett *et al* are highlighted. In Chapter 3, it is reported that two distinct, stationary turbulent states can be sustained with identical mean plasma parameters in the presence of flow shear. The two states and their respective mechanisms for turbulence saturation are characterised. The potential consequences for turbulent transport in tokamaks are then discussed. Finally, the conclusions of this work are drawn in Chapter 4.

The Appendices to this thesis are structured in the following manner. Appendix A gives some useful relations needed for the comparison of experimental data and gyrokinetic simulations. Appendix B presents the experimental data used for this thesis. Appendix C explains the discretisation employed in the simulations. Appendix D lists the numerical resolutions used in this work, and gives a typical input file for the simulations shown in this thesis. Finally, Appendix E briefly presents a method to reconstruct profiles of the mean rotational flow using gyrokinetic simulations.

The content of Chapter 2 is based on the publication [59] by N. Christen et al. The content of Chapter 3 is based on a paper in preparation by N. Christen et al.

Chapter 2

Simulating turbulent transport with a mean flow shear

2.1 Introduction

In Section 1.4, we gave an overview of the long-standing challenge posed by turbulent transport to the energy yield of nuclear fusion devices. We underlined the importance of sheared mean flows in regulating the turbulence, as well as the vast range of temporal and spatial scales that make direct numerical simulations exceedingly difficult to design. In this Chapter, we describe a scale-separated approach known as the gyrokinetic approximation, which eases the computational requirements for simulating turbulence in fusion plasmas. We then propose a new and improved algorithm to account for mean flow shear in a gyrokinetic code.

The first implementation of mean flow shear for local, δf gyrokinetic simulations was developed by Hammett *et al* [49] in the linearly implicit-in-time code GS2 [60], and was later applied to others including GENE [61], GWK [62] and CGYRO [50]. In this approach, the system of gyrokinetic and Maxwell equations is expressed in the frame that rotates and shears along with the flow, allowing the equations to be Fourier analysed across the mean magnetic field. The effective radial wave number in the lab frame – a constant without flow shear – becomes time-dependent in the presence of flow shear. To retain the benefits of using a code that is implicit in time, this approach to flow shear approximates such time-dependences by combining a nearest point approximation on a fixed grid, together with a wave number re-mapping method. This ensures that the computationally expensive part of the implicit scheme only needs to be executed at

initialisation, rather than at every time step. However, many experiments only sustain modest levels of flow shear, for which the radial wave number evolves very slowly: simulations then produce long periods of time where the perpendicular flow shear has no effect, separated by discrete jumps in wave number. We wish to avoid this unphysical behaviour, and consider here an improved, continuous-in-time approach to flow shear.

The Chapter is organised as follows. First, we give a brief overview in Section 2.2 of the gyrokinetic approximation, as well as the equations that are solved with the GS2 code. In the following Section 2.3, we summarise the challenges posed by flow shear to the main algorithm of GS2, we explain the discrete-in-time approach to flow shear developed by Hammett *et al*, and we present our new continuous-in-time algorithm for flow shear. Finally, in Section 2.4, we show numerical results obtained with GS2 that test the implementation of the continuous-in-time approach, and we compare it to the discrete-in-time method.

2.2 Model

We now summarise the gyrokinetic approximation, and we present the key aspects of the GS2 code. In Section 2.2.1, we show the standard orderings assumed in the δf gyrokinetic approach and the equations that we wish to solve. Normalisations are presented in Section 2.2.2. In Section 2.2.3, we describe the simulation domain as well as two key boundary conditions used by the code. We also define the system of coordinates that is employed in the absence of mean flow shear. In Section 2.2.6 we present a physical picture of the dominant linear instability and how it is affected by the mean flow shear. In Section 2.2.4 we show how the coordinates that are used in the code get adapted in the presence of flow shear, and how this modifies the spectral representation that is employed in the plane perpendicular to the magnetic field. Finally, the Fourier-analysed equations solved by the code are presented in Section 2.2.5.

2.2.1 The gyrokinetic approximation

Following our discussion in Section 1.4 of the wide range of time and length scales appearing in fusion devices, the δf gyrokinetic approach [26–29] uses the ratio of gy-

roradius to machine size $\rho_{*s} = \rho_s/a \ll 1$ for species s as an asymptotic-expansion parameter. Here, $\rho_s = |\boldsymbol{\rho}_s| = |\hat{\mathbf{b}} \times \mathbf{v}/\Omega_s|$ is the particle's gyroradius and a is the equilibrium length scale (chosen to correspond to the minor radius of the tokamak). Furthermore, $\hat{\mathbf{b}} = \mathbf{B}/B$ denotes the unit vector in the direction of the magnetic field \mathbf{B} , B is the norm of \mathbf{B} , \mathbf{v} is the velocity of the particle and $\Omega_s = eZ_s B/m_s c$ is the gyrofrequency with Z_s the charge number of the particle, e the elementary charge, m_s the particle mass and c the speed of light. In the frame moving with the background plasma flow \mathbf{u} , the turbulent fluctuations are assumed to have a typical frequency $\omega_{\text{turb}} \sim \rho_{*s} \Omega_s$ [28]. The length scale of fluctuations is assumed to be $O(a)$ along magnetic field lines but $O(\rho_s)$ across field lines. On the other hand, the equilibrium evolves on the energy-confinement-time scale τ_E which is ordered to be $O(1/\rho_{*s}^2 \omega_{\text{turb}})$. The equilibrium length scale is ordered to be $O(a)$, both along and across field lines. The distribution function f_s can be described as the sum of an equilibrium piece $\langle f_s \rangle$ and a fluctuating piece $\delta f_s = f_s - \langle f_s \rangle$, where $\langle \cdot \rangle$ denotes an average over the turbulent time scale and over the turbulent length scales across the magnetic field. The fluctuating part is ordered to be much smaller than the equilibrium part:

$$\frac{\delta f_s}{\langle f_s \rangle} \sim \rho_{*s}. \quad (2.1)$$

Finally, the orderings in time and space can be summarised as

$$\frac{d}{dt} \ln(\delta f_s) \sim \rho_{*s}^{-2} \frac{d}{dt} \ln \langle f_s \rangle \sim O(\rho_{*s} \Omega_s) \quad (2.2)$$

$$\hat{\mathbf{b}} \cdot \nabla \ln(\delta f_s) \sim \rho_{*s} |\nabla \ln(\delta f_s)| \sim |\nabla \ln \langle f_s \rangle| \sim O(1/a). \quad (2.3)$$

Here, $d/dt = \partial/\partial t + \mathbf{u} \cdot \nabla$ is the convective time derivative with respect to the mean flow \mathbf{u} .

Within this asymptotic expansion in ρ_{*s} , we restrict ourselves to the high-flow regime, in which the background flow u is first ordered to be comparable to the ion thermal speed $v_{\text{th},i} = \sqrt{2T_i/m_i}$, with T_i the ion temperature. We then consider a subsidiary expansion in the Mach number $M \equiv u/v_{\text{th},i}$ such that $\rho_{*s} \ll M \ll 1$, and we neglect second order effects in M . In the high-flow regime, the background

plasma flow is toroidal to lowest order in ρ_{*s} , and the associated angular frequency Ω_ϕ is species independent and constant on a magnetic flux surface [29, 63]: $\mathbf{u} = \Omega_\phi(\psi)R^2\nabla\phi + O(\rho_*v_{\text{th},i})$ with ψ the flux surface label, R the major radius and ϕ the toroidal angle. We further restrict ourselves to cases with low thermal-to-magnetic-pressure ratio (plasma β) and hence only retain electrostatic fluctuations.

Making use of the assumptions above, one can express the Fokker-Planck equation order by order in $\rho_{*s}^n\Omega_s f_s$ for $n = 0, 1, 2$ and average over the particles' rapid gyromotion to obtain the gyrokinetic equation governing the evolution of fluctuations of the distribution function:

$$\begin{aligned} \frac{\partial g_s}{\partial t} + \left(\mathbf{u} + w_{\parallel}\hat{\mathbf{b}} + \mathbf{V}_{d,s} + \langle \mathbf{V}_E \rangle_{\mathbf{R}_s} \right) \cdot \nabla g_s = \langle C[\delta f_s] \rangle_{\mathbf{R}_s} \\ - \frac{eZ_s F_{0,s}}{T_s} \left(w_{\parallel}\hat{\mathbf{b}} + \mathbf{V}_{d,s} \right) \cdot \nabla \langle \varphi \rangle_{\mathbf{R}_s} - \langle \mathbf{V}_E \rangle_{\mathbf{R}_s} \cdot \left(\frac{Im_s w_{\parallel}}{BT_s} F_{0,s} \nabla \Omega_\phi + \nabla F_{0,s} \right) \end{aligned} \quad (2.4)$$

in $(\mathbf{R}_s, \varepsilon_s, \mu_s, \vartheta)$ coordinates, where $\mathbf{R}_s = \mathbf{r} - \boldsymbol{\rho}_s$ is the particle's gyrocenter, \mathbf{r} is its position, $\varepsilon_s = m_s w^2/2$ is its kinetic energy, $\mu_s = m_s w_{\perp}^2/2B$ its magnetic moment and ϑ its gyrophase. Here, $\langle \cdot \rangle_{\mathbf{R}_s} = (2\pi)^{-1} \int_0^{2\pi} (\cdot) d\vartheta$ denotes a ‘‘gyro-average’’ over the gyrophase ϑ at fixed \mathbf{R}_s . We denote by $g_s(\mathbf{R}_s, \varepsilon_s, \mu_s) = \langle \delta f_s \rangle_{\mathbf{R}_s}$ the gyro-averaged fluctuating distribution function, and by φ the fluctuating electrostatic potential. The particle's velocity relative to the background flow is $\mathbf{w} = \mathbf{v} - \mathbf{u}$, of which $w_{\parallel} = \mathbf{w} \cdot \hat{\mathbf{b}}$ is the component along the magnetic field. Furthermore, C is the collision operator, $I(\psi) = R^2 \mathbf{B} \cdot \nabla \phi$, and $F_{0,s}$ is a Maxwellian distribution of velocities. The species-dependent guiding centre drift velocity $\mathbf{V}_{d,s} = \mathbf{V}_{B,s} + \mathbf{V}_{C,s}$ includes the ∇B and curvature drifts $\mathbf{V}_{B,s} = \hat{\mathbf{b}}/\Omega_s \times \left[w_{\perp}^2 \nabla \ln(B)/2 + w_{\parallel}^2 \hat{\mathbf{b}} \cdot \nabla \hat{\mathbf{b}} \right]$, as well as the Coriolis drift $\mathbf{V}_{C,s} = (2w_{\parallel}\Omega_\phi/\Omega_s)\hat{\mathbf{b}} \times (\hat{\mathbf{Z}} \times \hat{\mathbf{b}})$, where $\mathbf{w}_{\perp} = \mathbf{w} - w_{\parallel}\hat{\mathbf{b}}$ is the component of \mathbf{w} perpendicular to the magnetic field, and where $\hat{\mathbf{Z}}$ is the unit vector in the vertical direction. The fluctuating $\mathbf{E} \times \mathbf{B}$ drift is given by $\mathbf{V}_E = c\hat{\mathbf{b}}/B \times \nabla \varphi$. The nonlinearity in equation (2.4) stems from the gyro-averaged fluctuating $\mathbf{E} \times \mathbf{B}$ drift advecting g_s on the left-hand side. Perpendicular flow shear enters through \mathbf{u} , while the drives from the shear in the parallel flow and from the temperature gradient respectively enter via the $\nabla \Omega_\phi$ and $\nabla F_{0,s}$ terms on the right-hand side. Throughout this work, we specify the temperature gradient via the normalised inverse gradient length $a/L_{T_s} = -a d(\ln T_s)/dr_\psi$. We denote by r_ψ the half-diameter of a flux surface at the height of the magnetic axis.

The set of equations is closed by the quasineutrality condition:

$$\sum_s Z_s \int d^3w \langle g_s \rangle_{\mathbf{r}} = \sum_s \frac{Z_s^2 e}{T_s} \left(n_s \varphi - \int d^3w \langle \langle \varphi \rangle_{\mathbf{R}_s} \rangle_{\mathbf{r}} F_{0,s} \right) \quad (2.5)$$

where n_s is the density of particles, and the brackets $\langle \cdot \rangle_{\mathbf{r}}$ denote an average over the gyrophase at fixed particle position. Lastly, we neglect all effects associated with impurities in the plasma, and only consider two kinetic species – the electrons and the main hydrogenic ion species.

The turbulent contribution to the heat and momentum fluxes are given by

$$Q_s = \left\langle \int d^3\mathbf{v} \frac{m_s v^2}{2} \delta f_s \mathbf{V}_E \cdot \nabla \psi \right\rangle_{\psi}, \quad (2.6)$$

$$\Pi_s = \left\langle m_s R^2 \int d^3\mathbf{v} (\mathbf{v} \cdot \nabla \phi) \delta f_s \mathbf{V}_E \cdot \nabla \psi \right\rangle_{\psi}, \quad (2.7)$$

respectively, with $\langle \cdot \rangle_{\psi}$ denoting a volume average over the flux tube.

2.2.2 Normalisation

All quantities appearing in codes such as **GS2** are normalised in order to avoid errors associated with machine precision. In this work, the reference species chosen for normalisations are the main ions, i.e. quantities such as masses, temperatures, densities or electric charge are normalised to their value for the main ion species. Following the gyrokinetic ordering, frequencies of the order of the turbulent frequency ω_{turb} are normalised by $v_{\text{th},i}/a$. Lengths of the order of the machine size are normalised by the minor radius a of the tokamak. Estimating the turbulent diffusivities through a random-walk argument with a typical length scale ρ_i and a typical frequency ω_{turb} yields the so-called gyro-Bohm normalisation factors. For the turbulent fluxes of heat and momentum, they are respectively:

$$Q_{\text{gB}} = \langle |\nabla \psi| \rangle_{\psi} n_i T_i v_{\text{th},i} \rho_{*i}^2 \quad (2.8)$$

and

$$\Pi_{\text{gB}} = \langle |\nabla \psi| \rangle_{\psi} n_i m_i a v_{\text{th},i}^2 \rho_{*i}^2. \quad (2.9)$$

2.2.3 Simulation domain, boundary conditions and system of coordinates

Numerically solving the gyrokinetic and quasineutrality equations for a substantial portion of a fusion device’s volume is computationally very costly, and it poses a challenge for choosing appropriate radial boundary conditions. An alternative approach is followed by so-called “local” codes, which take advantage of the scale separations appearing in fusion devices by focusing on a thin, filament-like simulation domain known as a flux tube [64]. A flux tube extends along a magnetic-field line as it wraps around the torus, which we illustrate in Figure 2.1a. In directions perpendicular to \mathbf{B} , the flux tube is assumed to be wide enough (several ρ_s) for the turbulent fluctuations on one side of the domain to be decorrelated from the ones on the opposite side. The flux tube is still kept much narrower than the equilibrium length scale a , so that equilibrium quantities and associated drive terms can be considered as constant across \mathbf{B} in the simulation. Formally, this corresponds to the limit $\rho_{*s} \rightarrow 0$. This local approximation ensures that fluctuations at either end of the perpendicular domain are not only decorrelated, but also statistically identical (since identical drives yield the same turbulence on average over time). It is further assumed that two such instances of decorrelated, statistically identical turbulence can be set exactly equal to one another without affecting the statistical properties of the simulation domain. Hence, periodic boundary conditions are enforced perpendicular to \mathbf{B} [64].

Within the local approximation, the toroidal angular frequency is given by $\Omega_\phi(\psi) \simeq \Omega_{\phi,0} + \Omega'_{\phi,0}(\psi - \psi_0)$, where the subscript 0 indicates that a quantity is evaluated along the central field line in the simulation domain, and the prime denotes differentiation with respect to the flux label ψ . In the absence of flow shear ($\Omega'_{\phi,0} = 0$), one can work in the co-rotating frame (θ, x, y) , where the poloidal angle θ serves as the coordinate parallel to the field, x is the radial coordinate perpendicular to flux surfaces, and y is the coordinate perpendicular to the field line within a flux surface:

$$x = \frac{q_0}{r_{\psi,0} B_r} (\psi - \psi_0), \quad (2.10)$$




Figure 2.1: In 2.1a, the shaded grey area corresponds to the outboard mid-plane. Looking down onto this plane from above in 2.1b, the edges of the flux-tube are shown in solid blue lines as it intersects with the plane at $\theta = 0$ (light blue) and $\theta = 2\pi$ (dark blue). The dashed blue lines indicate exact copies of the intersection at $\theta = 0$.

$$y = \frac{1}{B_r} \frac{\partial \psi}{\partial r_\psi} \Big|_{r_{\psi,0}} (\alpha - \alpha_0). \quad (2.11)$$

B_r is a reference magnetic-field strength defined in Appendix A. The poloidal magnetic flux that is often used to label flux surfaces is defined as

$$\psi = \int_0^r dr' r' R \mathbf{B} \cdot \nabla \theta, \quad (2.12)$$

where r is the minor radius of the torus. The safety factor is given by

$$q(\psi) = \frac{1}{2\pi} \int_0^{2\pi} d\theta \frac{\mathbf{B} \cdot \nabla \phi}{\mathbf{B} \cdot \nabla \theta} \Big|_\psi, \quad (2.13)$$

and represents the number of toroidal turns required by a field line to travel once around the torus poloidally. Finally,

$$\alpha = \phi - \int_0^\theta d\theta' \frac{\mathbf{B} \cdot \nabla \phi}{\mathbf{B} \cdot \nabla \theta} \Big|_\psi \quad (2.14)$$

labels field lines on a given flux surface. Note that the x and y coordinates are not in general orthogonal.

In [64], Beer *et al* also introduce a boundary condition along the magnetic field lines, often referred to as the “twist-and-shift” boundary condition. It allows flux-tube simulations with minimal computational domains to be representative of an entire

flux surface. To understand this boundary condition, we consider a (ϕ, x) surface cutting through the torus, illustrated by the grey plane in Figure 2.1a. While the actual choice for this surface is usually the inboard mid-plane, in this illustration we pick the outboard mid-plane for simplicity. As the flux-tube wraps around the torus, it periodically intersects with this surface after every 2π turn in the poloidal angle. The periodicity in the y direction perpendicular to the magnetic field implies that fluctuations at such an intersection with the flux-tube can be equivalently described by exact copies of any other intersection. Here, we consider two consecutive intersections, for example at $\theta = 0$ and $\theta = 2\pi$ (illustrated in Figure 2.1b). For any fluctuating quantity Φ , one must then have

$$\Phi(t, \theta = 0, x, y(\theta = 0, x, \phi)) = \Phi(t, \theta = 2\pi, x, y(\theta = 2\pi, x, \phi)) \quad (2.15)$$

Notice that, to account for magnetic shear,

$$\hat{s} = (r_{\psi,0}/q_0) dq/dr_{\psi}|_{r_{\psi,0}}, \quad (2.16)$$

that deforms the cross section of the flux tube, the equality holds for fixed toroidal angle ϕ , not for fixed field-line label y .

As we mentioned above, this parallel boundary condition for flux tubes hinges on perpendicular periodicity, and hence on the assumption that the box is wider than the perpendicular correlation length of the turbulence. This in turn means that, for a flux tube to represent a whole flux surface, the parallel correlation length of fluctuations cannot be longer than the parallel size of the computational domain (in our example, we have chosen this size to be 2π when projected onto θ). In the present work, we always assume that the parallel correlation length is shorter than 2π , in agreement with [65]. We should however mention that this is not always necessarily the case, as was for instance found in the work of [66–72] on the topic of mode-rational flux surfaces, which we discuss in Sections 2.2.7 and 3.3.5.

2.2.4 Spectral representation in the presence of flow shear

In the presence of flow shear ($\Omega'_{\phi,0} \neq 0$), the time-dependent terms appearing in the gyrokinetic equation (2.4) through $\mathbf{u} \cdot \nabla$ are not periodic in x . Thus, the periodic solution obtained by Fourier transforming in (x, y) would not be consistent with the equations that we wish to solve. The solution that is adopted by GS2 and other codes is to restore periodicity by working in the shearing frame (θ, x, y^*) [49], with

$$y^* = y - x\gamma_E t. \quad (2.17)$$

Here, we have defined the normalised shearing rate of the mean flow as

$$\gamma_E = \frac{a}{v_{\text{th},i}} \frac{r_{\psi,0}}{q_0} \left. \frac{\partial \Omega_\phi}{\partial r_\psi} \right|_{r_{\psi,0}}. \quad (2.18)$$

From here on, the star superscript will be used to indicate a quantity that becomes time dependent in the presence of flow shear (e.g., y^*). It is not to be confused with the star subscript of the expansion parameter ρ_{*s} , or with the usual notation for complex conjugates. On the few occasions where we wish to denote a complex conjugate, we explicitly clarify our notations for the reader.

Given the periodic boundary conditions in the perpendicular directions discussed in the previous Section, any fluctuating quantity Φ can hence be expressed as a Fourier series in the shearing frame:

$$\Phi(t, x, y^*, \theta) = \sum_{k_x, k_y} \hat{\Phi}_{\mathbf{k}}(t, \theta) e^{ik_x x + ik_y y^*}, \quad (2.19)$$

where we define the wave vector $\mathbf{k} = k_x \nabla x + k_y \nabla y^*$, with wave numbers given by $k_x = j_x \Delta k_x$ and $k_y = j_y \Delta k_y$. The grid-spacings are $\Delta k_x = 2\pi/L_x$ and $\Delta k_y = 2\pi/L_y$, where L_x and L_y are the perpendicular sizes of the computational domain. The integer indices j_x and j_y span $-(N_x - 1)/2 \leq j_x \leq (N_x - 1)/2$ and $0 \leq j_y < N_y$, respectively. We recall that our model describes real-valued quantities ($\Phi \in \mathbb{R}$), so that $\hat{\Phi}_{k_x, -k_y} = \hat{\Phi}_{-k_x, k_y}^*$, where the star superscript is exceptionally used to denote a complex conjugate: this allows us to only consider $k_y > 0$. Note that the exponent in (2.19) can also be written in terms of (x, y) as $ik_x^*(t)x + ik_y y$ where $k_x^*(t) \equiv k_x - k_y \gamma_E t$ is the effective, time-dependent radial wave number in the lab frame. With this spectral representation in the

shearing frame, time dependences due to flow shear in the gyrokinetic-quasineutrality system of equations (2.4)-(2.5) appear when either a gradient or a gyro-average is applied to a fluctuating quantity:

$$\nabla_{\perp} \Phi = \sum_{k_x, k_y} i\mathbf{k}^*(t) \hat{\Phi}_{\mathbf{k}} e^{ik_x x + ik_y y^*}, \quad (2.20)$$

with $\mathbf{k}^*(t) = k_x^*(t) \nabla x + k_y \nabla y$, and

$$\langle \Phi(\mathbf{r}) \rangle_{\mathbf{R}_s} = \sum_{k_x, k_y} J_{0s}^*(t) \hat{\Phi}_{\mathbf{k}} e^{ik_x x + ik_y y^*}, \quad (2.21)$$

where $J_{0s}^*(t) \equiv J_0(|\mathbf{k}^*| \rho_s)$ is the zeroth Bessel function of the first kind evaluated using the time-dependent, lab-frame wave vector.

In Fourier space, the parallel boundary condition (2.15) becomes

$$\hat{\Phi}_{k_x, k_y}(t, \theta = 0) = \hat{\Phi}_{k_x + 2\pi \hat{s} k_y, k_y}(t, \theta = 2\pi) e^{-i2\pi j_y q_0 \frac{2\pi}{L_{\phi}}} \quad (2.22)$$

with $L_{\phi} = L_y B_r / (d\psi/dr_{\psi})|_{r_{\psi,0}}$. The extra phase factor can be set to unity by making an $O(\rho_{*s})$ change in q_0 to lie on a flux surface where $q_0 2\pi / L_{\phi} \in \mathbb{Z}$ (remembering that j_y is an integer labelling the different k_y 's). In simulations, (2.22) effectively creates twist-and-shift chains of linked radial wave numbers. In order to apply the twist-and-shift boundary condition (2.22), we point out that the quantity

$$j_{\text{twist}} = 2\pi \hat{s} L_x / L_y \quad (2.23)$$

always needs to be an integer.

2.2.5 Fourier-transformed system of equations

Having restored periodicity across the magnetic field by working in the shearing frame, the gyrokinetic equation (2.4) can be expressed in Fourier space:

$$\begin{aligned} & \frac{\partial \hat{g}_{\mathbf{k},s}}{\partial t} + w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{\partial \hat{g}_{\mathbf{k},s}}{\partial \theta} + i\mathbf{V}_{d,s} \cdot \mathbf{k}^* \hat{g}_{\mathbf{k},s} + \hat{\mathcal{N}}_{\mathbf{k}} = \langle C[\delta f_s] \rangle_{\mathbf{R}_s} \\ & - \frac{eZ_s F_{0,s}}{T_s} \left(w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{\partial}{\partial \theta} + i\mathbf{V}_{d,s} \cdot \mathbf{k}^* \right) (J_{0s}^* \hat{\varphi}_{\mathbf{k}}) - ik_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_{\psi}} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) J_{0s}^* \hat{\varphi}_{\mathbf{k}}. \end{aligned} \quad (2.24)$$

In GS2, the nonlinear term $\hat{\mathcal{N}}_{\mathbf{k}}$ is treated pseudo-spectrally: spatial derivatives are computed in Fourier space; they are then transformed individually to real space, where the nonlinear term is computed; and the result is transformed back to Fourier space,

$$\hat{\mathcal{N}}_{\mathbf{k}} = \frac{cq_0}{r_{\psi,0}B_r^2} \frac{\partial\psi}{\partial r_{\psi}} \Big|_{r_{\psi,0}} \mathcal{F}_{\mathbf{k}} \left[\{g_s, \langle\varphi\rangle_{\mathbf{R}_s}\}_{x,y} \right], \quad (2.25)$$

where the Poisson bracket $\{\cdot, \cdot\}_{x,y}$ is defined as

$$\{g_s, \langle\varphi\rangle_{\mathbf{R}_s}\}_{x,y} = \frac{\partial g_s}{\partial x} \Big|_y \frac{\partial \langle\varphi\rangle_{\mathbf{R}_s}}{\partial y} \Big|_x - \frac{\partial g_s}{\partial y} \Big|_x \frac{\partial \langle\varphi\rangle_{\mathbf{R}_s}}{\partial x} \Big|_y, \quad (2.26)$$

and the derivatives are computed as:

$$\frac{\partial \langle\varphi\rangle_{\mathbf{R}_s}}{\partial x} \Big|_y = \mathcal{F}_{x,y}^{-1} [ik_x^* J_{0s}^* \hat{\varphi}_{\mathbf{k}}], \quad (2.27)$$

$$\frac{\partial \langle\varphi\rangle_{\mathbf{R}_s}}{\partial y} \Big|_x = \mathcal{F}_{x,y}^{-1} [ik_y J_{0s}^* \hat{\varphi}_{\mathbf{k}}], \quad (2.28)$$

$$\frac{\partial g_s}{\partial x} \Big|_y = \mathcal{F}_{x,y}^{-1} [ik_x^* \hat{g}_{\mathbf{k},s}], \quad (2.29)$$

$$\frac{\partial g_s}{\partial y} \Big|_x = \mathcal{F}_{x,y}^{-1} [ik_y \hat{g}_{\mathbf{k},s}], \quad (2.30)$$

with $\mathcal{F}_{x,y}^{-1}$ the inverse discrete Fourier transform to the shearing frame defined by (2.19), and with $\mathcal{F}_{\mathbf{k}}$ the associated forward transform. At this stage, it is worth pointing out that the Poisson bracket (2.26) is invariant under the transformation from the lab frame (x, y) to the shearing frame (x, y^*) . We choose to express the bracket using (x, y) . In Fourier space, the quasineutrality equation (2.5) becomes:

$$\sum_s Z_s \int d^3w J_{0s}^* \hat{g}_{\mathbf{k},s} = \sum_s \frac{Z_s^2 e n_s}{T_s} (1 - \Gamma_s^*) \hat{\varphi}_{\mathbf{k}} \quad (2.31)$$

with $\Gamma_s^* = (1/n_s) \int d^3w F_{0,s} (J_{0s}^*)^2$.

2.2.6 Linear instabilities and Floquet oscillations

One of the dominant instabilities sustaining turbulence in the core is the toroidal ion temperature gradient instability (ITG) which occurs at the scale of the ion gyroradius [73, 74]. Ions in hotter regions of the tokamak can drift faster than in colder regions (since $\mathbf{V}_{B,i} \sim \rho_{*i} v_{th,i} \propto T_i$). It follows that fluctuations in temperature also generate fluctuations in the particle density. This sets up an electric field with an associated $\mathbf{E} \times \mathbf{B}$ drift. On the outboard side of the torus ($\theta = 0$), the geometry is such that this drift reinforces the temperature fluctuations, leading to an instability. This is illustrated in Figure 2.2. On the inboard side ($\theta = \pi$), the geometry is such that temperature fluctuations are stabilised.

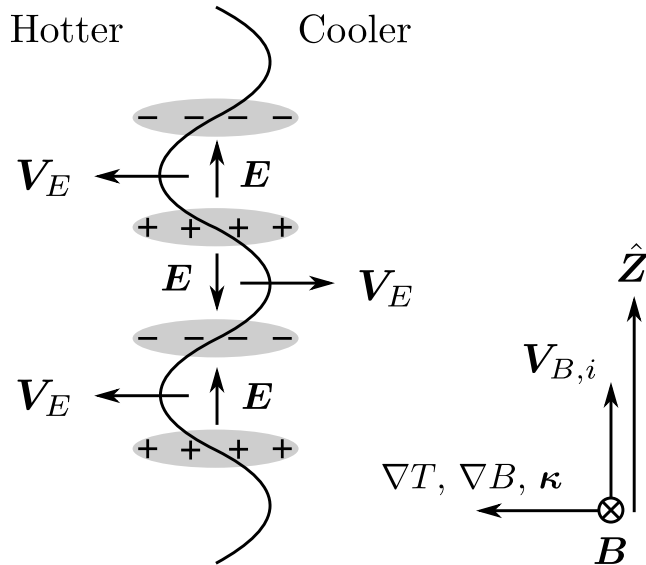


Figure 2.2: Illustration of the basic mechanism of toroidal ITG at the outboard mid-plane of the tokamak, with the curvature of magnetic field lines denoted by $\kappa = \hat{\mathbf{b}} \cdot \nabla \hat{\mathbf{b}}$. At the inboard midplane, ∇T points in the opposite direction and \mathbf{V}_E tends to stabilise the temperature perturbation.

Other instabilities are commonly observed to contribute to turbulent transport in tokamaks. Another instability occurring at length scales of the ion gyroradius is the trapped electron mode (TEM) [75]. Physically, the TEM is due to a resonance between electrostatic fluctuations and the toroidal precession drift of trapped electrons. It is driven by gradients in the electron pressure and its associated real frequency carries the opposite sign compared to the frequency associated with ITG. At length scales

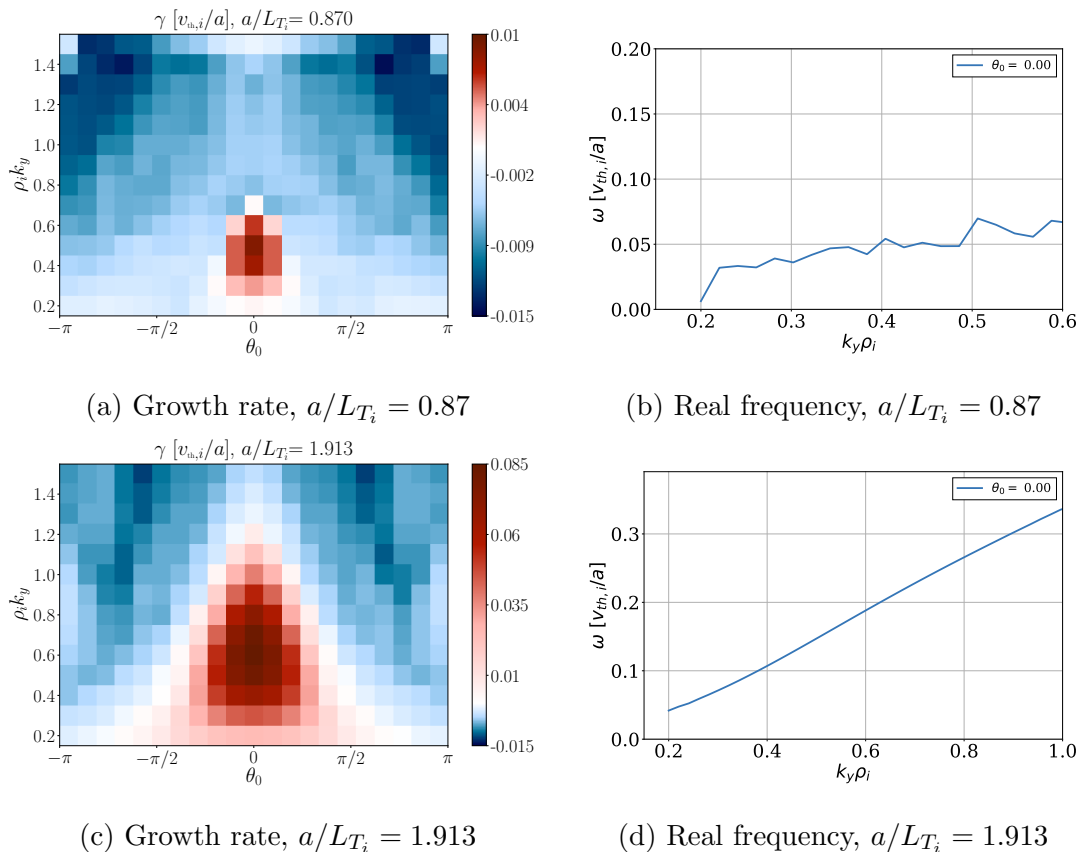


Figure 2.3: Linear growth rate and associates real frequencies. In the top panels $a/L_{T_i} = 0.87$ and in the bottom panels $a/L_{T_i} = 0.1.913$. Here flow shear has been neglected, i.e. $\gamma_E = 0$. All other physical parameters are taken from the JET discharge presented in Appendix B.

of the electron gyroradius, a common instability is the electron temperature gradient instability (ETG) [75–78] whose physical mechanism is essentially identical to ITG.

In this work, we model the plasma discharge #68448 conducted in the JET tokamak [79], which we describe in Appendix B. For this discharge, ITG is the dominant instability sustaining core turbulent transport. In Figure 2.3, we show the linear growth rate γ and the associated real frequency ω of modes when flow shear is neglected ($\gamma_E = 0$). Here, a positive value of ω indicates an instability related to ions rather than electrons. The ballooning angle $\theta_0 = k_x/(\hat{s}k_y)$ is the poloidal angle at which modes have no radial structure (for a fixed ϕ) [64]. As we briefly mentioned above, ITG modes typically grow at the outboard midplane ($\theta_0 \simeq 0$) and decay at the inboard midplane ($\theta_0 \simeq \pm\pi$).

Crucially, $\gamma_E \neq 0$ in this discharge, as it features a sheared flow that is induced by

the injection of neutral beams. In tokamaks with a sheared mean flow, fluctuations are advected along magnetic field lines. This is because the pitch of magnetic field lines varies radially in tokamaks. The variation in the pitch is measured by the magnetic shear \hat{s} . The perpendicular flow shear γ_E essentially shears fluctuations, and when $\hat{s} \neq 0$, the shearing from the mean flow can be undone by fluctuations moving along \mathbf{B} . As a result, fluctuations are advected along the field lines, passing through alternating regions with strong (outboard side) and weak (inboard side) ITG drive. The resulting modes, known as Floquet modes [37, 39], are analogous to Bloch states from condensed matter physics. A visual illustration of the advection of modes along \mathbf{B} under the combined effects of γ_E and \hat{s} can be found in Figure 2.6. The temporal evolution of Floquet modes can be described by an instantaneous growth rate γ (whose maximum we denote by γ_{\max}) with a Floquet period $T_F = |2\pi\hat{s}/\gamma_E|$, or by an effective average growth rate $\langle\gamma\rangle_t$, as is illustrated by Figure 2.4. We point out that γ_{\max} is usually much closer to the maximum ITG growth rate observed in the absence of mean flow shear than to $\langle\gamma\rangle_t$.




Figure 2.4: Typical Floquet oscillations in the electrostatic potential averaged along the magnetic field. Here, $\rho_i k_y = 0.5$, $a/L_{T_i} = 2.09$ and $\gamma_E = -0.1$. All other parameters are set to their values presented in Appendix B.

Figure 2.3 can be insightful to visualise how the time-dependent growth rate with $\gamma_E \neq 0$ is related to the growth rate in the absence of flow shear. As we discussed in Section 2.2.4, the radial wave number becomes time dependent when $\gamma_E \neq 0$. It follows

that the ballooning angle also becomes time dependent, $\theta_0^* = k_x^*/(\hat{s}k_y)$, which reflects the advection along field lines triggered by the flow shear. It can easily be checked that the Floquet period T_F is the time required for a mode to travel once poloidally around the tokamak (i.e., for θ_0^* to change by 2π). As a mode with a given k_y travels along the magnetic field, its instantaneous growth rate is roughly given by the corresponding $\gamma(\theta_0, k_y)$ obtained with $\gamma_E = 0$ (if, as it is the case here, the drive from the parallel velocity gradient is negligible). In other words, the time-dependent growth rate with $\gamma_E \neq 0$ can be obtained by moving with a period T_F along a row of a plot like ?? or 2.3c obtained with $\gamma_E = 0$. An estimate for the effective average growth rate $\langle \gamma \rangle_t$ can therefore be obtained by averaging along θ_0 the growth rate obtained in the absence of flow shear. At $\rho_i k_y = 0.5$ for the example in Figure 2.3a, this yields $\langle \gamma \rangle_t \simeq -0.003$ (i.e., flow shear suppresses ITG on average), whereas for Figure 2.3c $\langle \gamma \rangle_t \simeq 0.017$ (i.e., ITG still grows on average).

2.2.7 Mode-rational magnetic-flux surfaces

We now discuss an important consequence of the twist-and-shift parallel boundary condition used in flux-tube simulations, which we presented in Section 2.2.3. In essence, the cross-Section at one end of the flux tube ($\theta = 2\pi$) is mapped onto the periodically duplicated cross-Section at the other end of the flux tube ($\theta = 0$). In the presence of magnetic shear $\hat{s} \propto \partial q / \partial x \neq 0$, the pitch of magnetic field lines wrapping around the torus depends on which magnetic-flux surface (i.e., which x) they lie on. A detailed illustration is shown in the sketch of Figure 2.5. This implies that the cross-Section of the flux tube becomes twisted when moving along θ , and any given field line at $\theta = 2\pi$ (indicated by a dark-blue dot in Figure 2.5) is misaligned by a distance $d_y(x) = |y(x, \theta = 2\pi, \phi) - y(x, \theta = 0, \phi)|$ compared to its initial position at $\theta = 0$ (light-blue dot). Since periodicity is assumed in the y direction, the end of the field line at $\theta = 2\pi$ describes the same spatial location as the end of another field line at $\theta = 0$ (light-blue cross in Figure 2.5). The parallel boundary condition (2.15) sets turbulence at those two locations to be exactly identical.

An important consequence is that the magnetic field lines are artificially forced to close on themselves in a simulation. To understand this, we consider a magnetic-flux




Figure 2.5: Sketch illustrating the twist-and-shift parallel boundary condition. The light-blue square and dark-blue trapezoid show the cross-Section of the flux tube at $\theta = 0$ and $\theta = 2\pi$, respectively. The dashed light-blue squares represent periodic duplicates of the $\theta = 0$ cross-Section. The position of a given magnetic-field line at $\theta = 0$ is shown by the light-blue dot, and its position at $\theta = 2\pi$ is indicated by the dark-blue dot. The light-blue cross indicates the field line of the $\theta = 0$ cross-Section whose position coincides with the dark-blue dot in one of the periodic duplicates.

surface in the box where the misalignment d_y of the field lines is equal to a rational fraction times the size in y of the computational domain, i.e., $d_y/L_y = n/m$, with $n, m \in \mathbb{N}$. On this surface, the field lines in simulations are mapped back exactly onto themselves after m poloidal turns around the torus. Such magnetic-flux surfaces are known as mode-rational surfaces, or pseudo-rational surfaces, and have been studied extensively in [66–72]. We stress that mode-rational surfaces are a numerical artefact due to the parallel boundary condition. They need to be distinguished from the actual rational surfaces that exist in fusion devices, which are associated with a rational safety factor $q \in \mathbb{Q}$. Mode-rational surfaces with $m = 1$ and $m = 2$ are also known as mode-integer and mode-half-integer surfaces, respectively. In Figure 2.5, a mode-integer surface would correspond to a radial location at which the light-blue dot and the cross lie on top of each other; for a mode-half-integer surface, the dot and the cross would be separated by $L_y/2$.

In the local approximation, $q(x) = q_0 + x(\partial q/\partial x)|_{x=0}$, and the misalignment of field lines can be expressed as

$$d_y(x) = 2\pi \hat{s}x, \quad (2.32)$$

where we have used the definitions (2.10), (2.11), (2.13), (2.14) and (2.16) from Section 2.2.3. At the end of Section 2.2.4, we explained that the quantity $j_{\text{twist}} = 2\pi\hat{s}L_x/L_y$ must be an integer in order to apply the twist-and-shift boundary condition in a pseudo-spectral code such as GS2. We also recall that, on a mode-rational surface of order m , the ratio d_y/L_y is given by a multiple of $1/m$. Thus, equation (2.32) implies that the radial distance separating two mode-rational surfaces of order m can be expressed as

$$\Delta_{x,m} = \frac{L_x}{mj_{\text{twist}}}, \quad (2.33)$$

and the number of mode-rational surfaces of order m in a simulation is given by

$$N_m = \frac{L_x}{\Delta_{x,m}} = mj_{\text{twist}}. \quad (2.34)$$

In other words, $j_{\text{twist}} \in \mathbb{N}^+$ is the number of mode-integer surfaces that are present in a given simulation. To convince ourselves that this is the case, we can consider drawing analogous sketches to the one in Figure 2.5, but with $d_y(L_x/2) = L_y/2, L_y, 3L_y/2, \dots$ In doing so, we must remember to take into account that $x = 0$ is always a mode-integer surface, and that the box is periodic in x (i.e., $x = \pm L_x/2$ represent the same surface).

On a mode-rational surface, linear modes that are extended along the magnetic field are allowed to interact with themselves via the nonlinear term in the gyrokinetic equation. This fictitious self-interaction is strongest for mode-rational surfaces with low m , for which the field lines are forced to close on themselves after fewer poloidal turns. In simulations, the self-interaction has been found to generate steady (i.e., constant in time) layers of zonal shear in the vicinity of mode-integer surfaces [66–72]. As has been shown in [71], the impact of these numerical artefacts on simulations can be minimised by increasing L_y while keeping L_x fixed. Effectively, this means reducing j_{twist} , and thus decreasing the number of mode-integer surfaces in the simulation. The impact of these layers will be further discussed in the context of this work in Section 3.3.5.

2.3 GS2 algorithm in the presence of flow shear

We showed in Section 2.2 that the presence of flow shear introduces new time dependences in the Fourier-transformed gyrokinetic and quasineutrality equations, which we denoted by star superscripts. These time dependences would in principle be straightforward to include in codes that solve for $\hat{g}_{\mathbf{k},s}$ with an explicit time-advance algorithm: the only modification required would be to update the coefficients $\mathbf{V}_{d,s} \cdot \mathbf{k}^*$ and J_{0s}^* at every time step. Tests conducted with the GS2 code indicate that this increases runtime by $\sim 25\%$.

However, certain codes, including GS2, use an algorithm that is implicit in time. Unlike explicit algorithms, A-stable implicit algorithms have no stability criterion limiting the size of their time step (known as a Courant-Friedrichs-Lewy, or CFL condition). This can therefore lead to significant savings in computing time by taking larger time steps. Unfortunately, this also introduces additional challenges to the numerical treatment of flow shear, as we will now discuss.

In Section 2.3.1, we present the discretised equations solved by the code. We then give a summary of the implicit algorithm used in GS2 in Section 2.3.2. In the two following Sections 2.3.3 and 2.3.4, we discuss the challenges that the algorithm poses in the presence of flow shear, and we explain an elegant, widely used method that allows for the treatment of flow shear within this implicit algorithm. We also consider the limitations of this method that treats flow shear discretely in time. Finally, we propose in Section 2.3.5 a new, more accurate approach for treating flow shear continuously over time.

For simplicity, we omit collisions from our discussions in this Section. At the time of writing, collisions in the GS2 code are treated in the usual, discrete-in-time approach. An extension of our work could be to also apply the new, continuous-in-time approach to collisions in GS2. We do not expect this implementation to present any significant challenges, and the same steps as those presented in Section 2.3.5 could be followed. Furthermore, implementing this change in the code is likely to have a small impact on our work: all of our simulations are either collisionless or feature very low collision frequencies.

2.3.1 Discretised system of equations

At the end of Section 2.2.4, we stressed that the parallel boundary condition (2.22) creates for every k_y a set of linearly independent twist-and-shift chains of connected θ -segments. Here, we use bold-font vector notations to denote discrete arrays whose indices span every (θ, k_x) of a single twist-and-shift chain. Matrices with two such indices are denoted using straight, bold upper case letters. The nonlinear term is treated explicitly in time and is irrelevant for the present Section. Therefore, references to the gyrokinetic equation in this Section correspond to the collisionless gyrokinetic equation without the nonlinear term. Discretising in time and θ , the gyrokinetic equation (2.24) can be written schematically as

$$\mathbf{A}_{n+1}^* \hat{\mathbf{g}}_{n+1} + \mathbf{B}_n^* \hat{\mathbf{g}}_n = \mathbf{C}_{n+1}^* \hat{\varphi}_{n+1} + \mathbf{D}_n^* \hat{\varphi}_n, \quad (2.35)$$

with the index n labelling the time step, and where matrices \mathbf{A}^* , \mathbf{B}^* , \mathbf{C}^* and \mathbf{D}^* are defined in Appendix C. For simplicity of notation, we suppress indices labelling species, k_y , ε_s and μ_s , as well as the one distinguishing different twist-and-shift chains. Similarly, the quasineutrality condition becomes

$$\mathbf{E}_{n+1}^* \hat{\varphi}_{n+1} = \mathbf{W}_{n+1}^* \hat{\mathbf{g}}_{n+1} \quad (2.36)$$

with \mathbf{E}^* and \mathbf{W}^* given in Appendix C. Note that \mathbf{W}^* does not have the same dimensions as the other matrices since it represents an integration over velocity space. All quantities at time step n are known, and we wish to solve for $\hat{\mathbf{g}}_{n+1}$ and $\hat{\varphi}_{n+1}$.

2.3.2 Overview of the GS2 algorithm

An obvious way to solve the discretised system of equations would be to isolate $\hat{\varphi}_{n+1}$ in the quasineutrality equation (2.36), insert the resulting expression into the gyrokinetic equation (2.35) and find $\hat{\mathbf{g}}_{n+1}$. However, the velocity integration represented by \mathbf{W}_{n+1}^* leads to a dense matrix that would need to be inverted to obtain $\hat{\mathbf{g}}_{n+1}$. As this would be computationally costly, GS2 uses an alternative method [60] related to a Green's function approach. The general idea is to determine how the distribution function responds when the potential $\hat{\varphi}_n$ is set to zero and the potential $\hat{\varphi}_{n+1}$ is set to a

Kronecker delta in θ and k_x . By doing this for every θ and k_x , we will show that one can compute $\hat{\varphi}_{n+1}$ without knowing $\hat{\mathbf{g}}_{n+1}$, and then solve the gyrokinetic equation (2.35) for the distribution function.

Since (2.35) is a linear equation, $\hat{\mathbf{g}}_{n+1}$ can be written as a sum of two terms, where $\hat{\mathbf{g}}_{n+1}^{\text{old}}$ only depends on “old” fluctuating quantities from the current time step n , while $\hat{\mathbf{g}}_{n+1}^{\text{new}}$ also depends on the “new” potential from time step $n + 1$:

$$\hat{\mathbf{g}}_{n+1} = \hat{\mathbf{g}}_{n+1}^{\text{old}} + \hat{\mathbf{g}}_{n+1}^{\text{new}}. \quad (2.37)$$

The two terms are defined such that they satisfy

$$\mathbf{A}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{old}} + \mathbf{B}_n^* \hat{\mathbf{g}}_n = (\mathbf{C}_{n+1}^* + \mathbf{D}_n^*) \hat{\varphi}_n \quad (2.38)$$

and

$$\mathbf{A}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{new}} = \mathbf{C}_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}}, \quad (2.39)$$

respectively. Here, $\hat{\varphi}_{n+1}^{\text{dif}}$ is defined as the potential difference between the present and the next time steps

$$\hat{\varphi}_{n+1}^{\text{dif}} = \hat{\varphi}_{n+1} - \hat{\varphi}_n. \quad (2.40)$$

Equation (2.38) can be solved directly to find $\hat{\mathbf{g}}_{n+1}^{\text{old}}$ since all other quantities are evaluated at the previous time step and are therefore known. One can then introduce a Green’s function \mathbf{G}^* such that

$$\mathbf{A}_{n+1}^* \mathbf{G}_{n+1}^* = \mathbb{1}. \quad (2.41)$$

Multiplying from the right by $\mathbf{C}_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}}$, one identifies $\hat{\mathbf{g}}_{n+1}^{\text{new}}$ as

$$\hat{\mathbf{g}}_{n+1}^{\text{new}} = \mathbf{G}_{n+1}^* \mathbf{C}_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}} \equiv \left(\frac{\delta \hat{\mathbf{g}}_{n+1}^{\text{new}}}{\delta \hat{\varphi}_{n+1}^{\text{dif}}} \right)_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}} \quad (2.42)$$

where we defined the “response matrix” $\delta \hat{\mathbf{g}}_{n+1}^{\text{new}} / \delta \hat{\varphi}_{n+1}^{\text{dif}}$. The matrix can be computed by setting $\hat{\varphi}_{n+1}^{\text{dif}}$ to a Kronecker delta function in (θ, k_x) , solving (2.39) for $\hat{\mathbf{g}}_{n+1}^{\text{new}}$, filling the appropriate column of $\delta \hat{\mathbf{g}}_{n+1}^{\text{new}} / \delta \hat{\varphi}_{n+1}^{\text{dif}}$ with this result and iterating for every (θ, k_x) in the connected twist-and-shift chain. In the absence of flow shear, \mathbf{A} , \mathbf{B} , \mathbf{C} , \mathbf{D} and \mathbf{G} are independent of time, and one only needs to go through this process once per

simulation, unless the time step size changes to satisfy the CFL condition set by the computation of the nonlinear term.

Inserting expressions (2.37), (2.40) and (2.42) into the discrete quasineutrality equation (2.36) yields

$$\hat{\varphi}_{n+1}^{\text{dif}} = (\mathbf{M}^*)_{n+1}^{-1} (\mathbf{W}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{old}} - \mathbf{E}_{n+1}^* \hat{\varphi}_n), \quad (2.43)$$

with the matrix \mathbf{M}^* defined by

$$\mathbf{M}_{n+1}^* = \mathbf{E}_{n+1}^* - \mathbf{W}_{n+1}^* \left(\frac{\delta \hat{\mathbf{g}}^{\text{new}}}{\delta \hat{\varphi}^{\text{dif}}} \right)_{n+1}^*. \quad (2.44)$$

Once $(\mathbf{M}^*)_{n+1}^{-1}$ is known, $\hat{\varphi}_{n+1}$ can be obtained from (2.43) and $\hat{\mathbf{g}}_{n+1}$ can then be determined by solving the full gyrokinetic equation (2.35).

2.3.3 Challenges in the presence of flow shear

Based on the previous Sections, we can identify two important issues appearing in this implicit algorithm when flow shear is present. First, consider a simulation with a fixed, finite set of radial wave numbers k_x in the shearing frame. As the simulation advances in time, the associated wave numbers in the lab frame, $k_x^*(t, k_x) = k_x - k_y \gamma_E t$, would all tend to $\pm\infty$ depending on the sign of γ_E . In other words, the simulation would eventually contain only structures that have been highly sheared in the lab frame, and would hence be of little interest for transport studies where structures elongated in the radial direction ($k_x^* \simeq 0$) play an important role. The second issue with flow shear is related to the computational cost of a simulation. As discussed in Section 2.3.2, the response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\varphi}^{\text{dif}}$ required to compute $\hat{\varphi}_{n+1}$ is independent of time if $\gamma_E = 0$, in which case it has to be computed only once at the beginning of the simulation, and every time the size of the time step changes. But if $\gamma_E \neq 0$, this matrix becomes time dependent and would have to be re-computed at every time step, leading to prohibitively slow simulations.

2.3.4 Discrete-in-time approach to flow shear

We now describe how the two challenges associated with flow shear for the GS2 algorithm are addressed by a re-mapping approach developed by Hammett *et al* [49].

This approach has previously been implemented in **GS2** and other gyrokinetic codes including **GENE** [61], **GKW** [62] and **CGYRO** [50]. We then highlight the disadvantages associated with the discrete nature of this implementation.

To address the issue of modes getting more and more sheared over time in the lab frame, Hammett’s approach uses a wave-number re-mapping method [49]. In practice, the Fourier series (2.19) of a fluctuating quantity Φ is truncated to a finite sum. At $t = 0$, we have $-K_x \leq k_x \leq K_x$ and $0 \leq k_y \leq K_y$. The set of Fourier coefficients that are being simulated is updated over time, such that $-K_x \lesssim k_x^* \lesssim K_x$ at all times (i.e., such that modes radially elongated in the lab frame are always included in the simulation). For each k_y there is a corresponding time that we denote by $T_{\text{map}}(k_y) = \Delta k_x / |\gamma_E k_y|$, at which

$$k_x^*(t + T_{\text{map}}, k_x, k_y) = k_x^*(t, k_x, k_y) \mp \Delta k_x. \quad (2.45)$$

From here on, the upper sign will always correspond to $\gamma_E > 0$ and the lower to $\gamma_E < 0$. In the lab frame, $T_{\text{map}}(k_y)$ corresponds to the time required for modes with this particular k_y to regain radial periodicity on the edges of the box (for a visual illustration, see Figure 2.8b, which will be presented in Section 2.4.1). Over the interval of time T_{map} , the most sheared mode in the simulation, $\hat{\Phi}_{\mp K_x, k_y}$ is considered to shear into structures that are radially fine enough to get averaged out to a negligible amplitude by the gyromotion of particles; this mode is dropped from the simulation. Simultaneously, the mode $\hat{\Phi}_{\pm(K_x + \Delta k_x), k_y}$ has now been unsheared enough to become of interest and gets added into the simulation with zero amplitude. The Fourier sum at time t is then given by

$$\Phi(t, x, y^*, \theta) = \sum_{k_y = -K_y}^{K_y} \sum_{k_x = -K_x \pm S(t, k_y) \Delta k_x}^{K_x \pm S(t, k_y) \Delta k_x} \hat{\Phi}_{\mathbf{k}}(t, \theta) e^{ik_x x + ik_y y^*}, \quad (2.46)$$

where $S(t, k_y)$ is the number of times the simulation has dropped a k_y -mode and added a new one. In the code, the shift in radial wave number is triggered every time k_x^* crosses the mid-point between the two nearest-lying $n\Delta k_x$ values, with $n \in \mathbb{Z}$. This is implemented by incrementing S at every shift, and by using a “re-mapped” version of the Fourier transformed quantities, denoted by a tilde:

$$\tilde{\Phi}_{k_x, k_y} = \hat{\Phi}_{k_x \pm S \Delta k_x, k_y}, \quad (2.47)$$

where the new Fourier component entering the simulation $\tilde{\Phi}_{\pm K_x, k_y}$ is set to zero at each shift.

To address the computing time issue associated with the time dependence of the response matrix, Hammett's approach makes a nearest-grid-point (NGP) approximation [49]: $k_x^*(t, k_x, k_y) \simeq \bar{k}_x(t, k_x, k_y)$, where we define

$$\bar{k}_x(t, k_x, k_y) = k_x \mp S(t, k_y) \Delta k_x \quad (2.48)$$

as the $n\Delta k_x$ value nearest to the lab-frame wave number $k_x^*(t, k_x, k_y)$, with $n \in \mathbb{Z}$. From here on, the overbar will be used to denote quantities evaluated using the NGP approximation of the time-dependent lab-frame wave number.

Combined with the wave-number re-mapping, the NGP approximation gets rid of all time dependences due to flow shear in the code, and all star superscripts in Sections 2.3.1 and 2.3.2 can be ignored. Therefore, the response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\boldsymbol{\varphi}}^{\text{dif}}$ in this implementation only has to be computed at initialisation and when the time-step size changes, similarly to cases with $\gamma_E = 0$. As an illustration of the time dependences vanishing, the inverse Fourier transform of the radial derivative at fixed y of a quantity Φ is given by:

$$\left. \frac{\partial \Phi}{\partial x} \right|_y = \sum_{k_y=-K_y}^{K_y} \sum_{k_x=-K_x \pm S(t) \Delta k_x}^{K_x \pm S(t) \Delta k_x} k_x^*(t) \hat{\Phi}_{k_x, k_y}(t) e^{ik_x^*(t)x + ik_y y} \quad (2.49)$$

$$\stackrel{\text{remap}}{=} \sum_{k_y=-K_y}^{K_y} \sum_{\bar{k}_x=-K_x}^{K_x} k_x^*(t) \tilde{\Phi}_{\bar{k}_x, k_y}(t) e^{i\bar{k}_x x + ik_y y} e^{i(k_x^*(t) - \bar{k}_x)x} \quad (2.50)$$

$$\stackrel{\text{NGP}}{\simeq} \sum_{k_y=-K_y}^{K_y} \sum_{\bar{k}_x=-K_x}^{K_x} \bar{k}_x \tilde{\Phi}_{\bar{k}_x, k_y}(t) e^{i\bar{k}_x x + ik_y y}, \quad (2.51)$$

where we explicitly show all time dependences. In the first expression, k_x^* is time-dependent, and the set of k_x is changing over time. In the last expression, the set of \bar{k}_x values entering the sum is fixed, and all time dependences from flow shear vanish.

However, it is crucial to note that the NGP approximation will produce two different types of errors in simulations. First, unphysical, discrete jumps in time will occur in the linear evolution of fluctuations, with a period of $T_{\text{map}}(k_y)$. This effect vanishes in the limit of small spacing in the radial-wave-number grid, $\Delta k_x \rightarrow 0$. But it is made

worse for cases with weaker flow shear, where radial wave numbers in the lab frame k_x^* evolve very slowly in time and get approximated to the same nearest neighbour for a long period of time. This is of particular concern given that many experiments in fusion devices feature a slow shearing rate of the mean flow.

The second type of error associated with the NGP approximation arises when Fourier transforms are performed, e.g., in the computation of the nonlinear term. A phase factor $e^{i(k_x^* - \bar{k}_x)x}$ is missed by the discrete-in-time approach to flow shear. This can be seen in the example above, where the phase factor has been approximated to unity to go from (2.50) to (2.51). This type of error has been stated in [51] not to vanish in the limit $\Delta k_x \rightarrow 0$. Following this observation, the authors added the missing phase factor in the gyrokinetic code **GENE**. One could however argue that given that $|k_x^* - \bar{k}_x| \sim \Delta k_x$, the error at any fixed location in (x, y) should get smaller when Δk_x is decreased. Errors at the radial edges of the box would indeed be independent of Δk_x , since $L_x \sim 1/\Delta k_x$. But this should not affect the centre of the domain: two instances of turbulence separated by more than one correlation length should not influence one another, regardless of whether one instance suffers from the missing phase factor. This would suggest that the error associated with the missing $e^{i(k_x^* - \bar{k}_x)x}$ factor could be reduced by decreasing Δk_x , while only considering the same portion of the simulation domain and ignoring the erroneous radial edges. A simple test case presented in Section 2.4.2 supports this argument, but nonlinear gyrokinetic simulations shown in Section 2.4.3 suggest a more complicated picture.

In the next Section, we will present a new approach to include flow shear in a linearly implicit δf gyrokinetic code, which removes both types of errors that arise in the presence of such a shear.

2.3.5 Continuous-in-time approach to flow shear

We now present an alternative numerical approach for including flow shear continuously over time, which gets rid of the errors associated with the discrete approach detailed in Section 2.3.4.

Similarly to the discrete-in-time implementation, we use the wave number re-mapping method described in the previous Section to ensure that modes elongated radially in the lab frame are always included in the simulation.

The first novel idea is to remove time dependences from the implicit part of the algorithm by evaluating them explicitly in time. In the gyrokinetic equation (2.24), every term $\mathcal{L}[k_x^*] \hat{\Phi}_{\mathbf{k}}$ that depends on the wave number in the lab frame can be re-written as:

$$\mathcal{L}[k_x^*] \hat{\Phi}_{\mathbf{k}} = \mathcal{L}[\bar{k}_x] \hat{\Phi}_{\mathbf{k}} + (\mathcal{L}[k_x^*] - \mathcal{L}[\bar{k}_x]) \hat{\Phi}_{\mathbf{k}} \quad (2.52)$$

where we added and subtracted the same term, evaluated at the nearest $n\Delta k_x$ neighbour. For every term in the discretised gyrokinetic equation (2.35) that is evaluated at the next time step $n+1$, we perform (2.52). The first term $\mathcal{L}[\bar{k}_x] \hat{\Phi}_{\mathbf{k}}$ is discretised in the usual implicit way, and hence appears in the implicit part of the algorithm. This term is nothing but the NGP approximation that was used in the discrete-in-time implementation. In our new approach, we choose to treat the additional $(\mathcal{L}[k_x^*] - \mathcal{L}[\bar{k}_x]) \hat{\Phi}_{\mathbf{k}}$ term explicitly in time; i.e., to evaluate it at the current time step n . In the discrete-in-time approach, this term is neglected. Similarly to the explicit nonlinear term (2.25), there is an associated Courant-Friedrichs-Lewy condition that has to be fulfilled for the scheme to be stable. But the new explicit terms are always small compared to the corresponding implicit terms since $\rho_s |k_x^* - \bar{k}_x| \leq \rho_s (\Delta k_x)/2 \ll 1$. Hence, the new condition for numerical stability should be less stringent than the one already set by the nonlinear term. One should note, however, that if Δk_x is made large enough (much larger than any sensible choice for a nonlinear simulation), this scheme would eventually become unstable. Following the decomposition (2.52), the time dependences in \mathbf{A} and \mathbf{C} vanish from (2.38) and (2.39). Equations (2.41) and (2.42) then yield a time-independent response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\varphi}^{\text{dif}}$.

This decomposition alone does not remove all time dependences from the implicit part of the code: to compute $\hat{\varphi}_{n+1}$, one must evaluate the right-hand side of equation (2.43), where \mathbf{E}^* , \mathbf{W}^* and $(\mathbf{M}^*)^{-1}$ depend on time. We update \mathbf{E}^* and \mathbf{W}^* at every time step with little extra computational cost. The time dependence in the matrix \mathbf{M}^* could be computed exactly. However, inverting it at every time step would be very

costly. Instead, we use linear interpolation to approximate the time dependence of its inverse $(\mathbf{M}^*)^{-1}$. Note that while there is a simple dependence of \mathbf{M}^* on time, there is no such simple dependence for the inverse $(\mathbf{M}^*)^{-1}$. At the beginning of a simulation, we compute three matrices \mathbf{M}_ℓ^{-1} with $\ell = -1, 0, +1$, defined respectively by replacing all instances of k_x^* in (2.44) with $\bar{k}_x + \ell\Delta k_x$. Interpolating then gives

$$(\mathbf{M}^*)^{-1} \simeq \left(1 - \frac{|k_x^* - \bar{k}_x|}{\Delta k_x}\right) \mathbf{M}_0^{-1} + \frac{|k_x^* - \bar{k}_x|}{\Delta k_x} \mathbf{M}_{\text{sgn}(k_x^* - \bar{k}_x)}^{-1}. \quad (2.53)$$

Notice that we do not interpolate the response matrix $\delta\hat{\mathbf{g}}^{\text{new}}/\delta\hat{\varphi}^{\text{dif}}$: since it was made independent of time, each \mathbf{M}_ℓ^{-1} contains the same $\delta\hat{\mathbf{g}}^{\text{new}}/\delta\hat{\varphi}^{\text{dif}}$, and only depends on time through \mathbf{E}^* and \mathbf{W}^* .

As was argued in [51], the implementation of Fourier transforms also needs to be modified to take into account the continuous time evolution of k_x^* . This has direct implications for the computation of the nonlinear term (2.25). With the wave number re-mapping, the discrete inverse Fourier transform (2.19) can be written as:

$$\mathcal{F}_{x,y}^{-1} [\hat{\Phi}_{\mathbf{k}}] = \sum_{k_y = -K_y}^{K_y} e^{i(k_x^*(t,k_y) - \bar{k}_x)x_n} e^{ik_y y_m} \sum_{\bar{k}_x = -K_x}^{K_x} \tilde{\Phi}_{\bar{k}_x, k_y} e^{i\bar{k}_x x_n}, \quad (2.54)$$

with $x_n = nL_x/(2N_x+1)$ and $y_m = mL_y/(2N_y+1)$. Similarly for the forward transform:

$$\mathcal{F}_{\mathbf{k}} [\Phi] = \sum_{n=-N_x}^{N_x} e^{-i(k_x^*(t,k_y) - \bar{k}_x)x_n} e^{-i\bar{k}_x x_n} \sum_{m=-N_y}^{N_y} \Phi(x_n, y_m) e^{-ik_y y_m}. \quad (2.55)$$

With the NGP approximation made in the discrete-in-time approach, k_x^* is approximated to \bar{k}_x and the first exponential terms in (2.54) and (2.55) are set to unity. In the continuous-in-time approach, we update these phase factors at every time step and modify the Fourier transforms in the following way: for $\mathcal{F}_{x,y}^{-1}$, inverse transform in k_x , multiply by $e^{i(k_x^* - \bar{k}_x)x}$ and inverse transform in k_y ; for $\mathcal{F}_{\mathbf{k}}$, transform in y , multiply by $e^{-i(k_x^* - \bar{k}_x)x}$ and transform in x .

2.4 Comparison of the approaches to flow shear

In this Section, we compare numerical results obtained with the discrete-in-time and continuous-in-time approaches. In Section 2.4.1, we first focus on tests of the modifications made to the linear GS2 algorithm. We then present simplified cases in Section

2.4.2 that illustrate the changes in the Fourier transforms and how these can affect the computation of the nonlinear term. Finally, fully nonlinear gyrokinetic simulations of an experiment at the Joint European Torus are shown and discussed in Section 2.4.3, highlighting the impact of this work on simulating turbulent transport in the presence of flow shear.

2.4.1 Linear simulations

We first present linear simulations of the widely used, collisionless ‘‘Cyclone Base Case’’ [31], to which we add a mean flow shear with the rate $\gamma_E = 0.2$. Twist-and-shift chains of connected radial wave numbers, which we introduced at the end of Section 2.2.4, can be seen as a single ballooning structure, that can extend in θ by more than 2π . At any given time, this chain corresponds to a real-space structure that is twisted along θ by the presence of magnetic shear \hat{s} (illustrated by the sketch in Figure 2.6a). Now, if $\gamma_E \neq 0$, the background flow shear will contribute to the twisting, enhancing it at one end in θ and decreasing it at the other end (Figure 2.6b). After a period of time T_{map} , the wave number re-mapping is applied in order to retain the radially extended part of the structure in the simulation. As shown in Figure 2.6c, this amounts to moving the computational domain in real space along θ . In the absence of flow shear, the ballooning mode typically peaks at $k_x^* = 0$ and $\theta = 0$, where the mode is aligned with the ∇B and curvature drives for the ion temperature gradient instability, and where finite gyroradius stabilisation is weakest. When $\gamma_E \neq 0$, it modifies the twist of the structure over time, and the peak moves along the field line in θ towards the smallest k_x^* in the chain. As the peak passes from $\theta = 0$ (bad curvature region) to $\theta = \pi$ (good curvature region) it experiences a varying growth rate, resulting in a so-called Floquet mode [37] (analogous to Bloch states in condensed matter physics). The Floquet oscillation period of the growth rate is $T_F = 2\pi\hat{s}/\gamma_E$.

To highlight the effects on linear physics of the NGP approximation for radial wave numbers, we consider the time evolution of the fluctuating electrostatic potential φ , averaged over θ and summed along one particular twist-and-shift chain. In Figure 2.7, $\sum_{k_x} |\langle \hat{\varphi}_{\mathbf{k}} \rangle_{\theta}|^2$ is plotted for several values of the grid spacing in radial wave number, Δk_x , with $\langle \cdot \rangle_{\theta}$ denoting an average over the poloidal angle. Figure 2.7a shows the

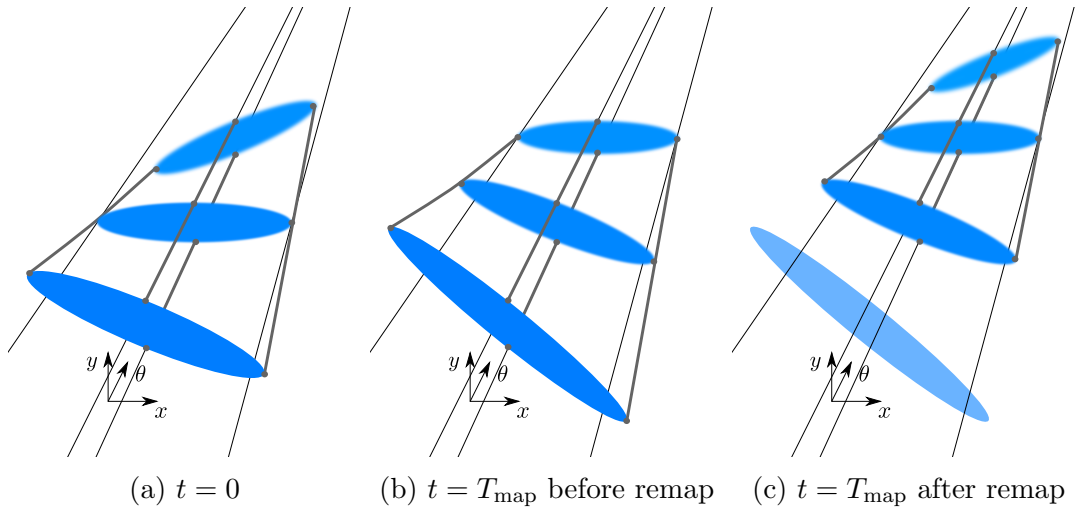
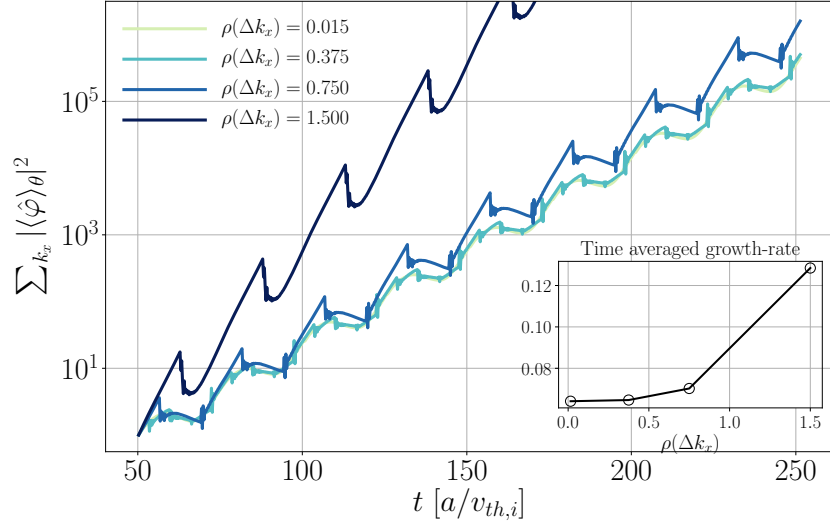


Figure 2.6: Cartoon illustrating how the wave-number re-mapping amounts to moving the computational domain in θ . Initially (2.6a), the structure is twisted by magnetic shear. Over a period of time T_{map} , flow shear contributes to the twisting, enhancing it at one end and decreasing it at the other (2.6b). In order to always include radially extended structures in the simulation, the computational domain is shifted in θ via a wave-number re-map (2.6c).


results obtained with the discrete-in-time approach to flow shear, which uses the NGP approximation. Discrete jumps are generated every time a wave number re-mapping occurs for this twist-and-shift chain. For fine-enough grid spacings, the simulation recovers from the jumps, and the time-averaged growthrate remains approximately correct. For larger Δk_x however, the simulation fails to capture the correct evolution of the Floquet mode. Figure 2.7b shows the results with the new continuous-in-time approach: it does not generate any discrete jumps at re-maps, and reproduces the correct Floquet behaviour over a much larger range of Δk_x than the discrete-in-time algorithm. Note that typical nonlinear gyrokinetic simulations currently use $\rho_i(\Delta k_x) \sim 0.1$. Section 2.4.3 will show that our approach to capture the linear physics correctly by avoiding the NGP approximation can lead to substantial computational savings for nonlinear simulations.

2.4.2 Test cases for Fourier transforms

We now show two simplified cases that test the changes made to the Fourier transforms in the continuous-in-time algorithm.



(a) Discrete-in-time



(b) Continuous-in-time

Figure 2.7: Time evolution of the normalised electrostatic potential for a single twist-and-shift chain, summed over the parallel coordinate. Panel 2.7a shows the results obtained with the widely used discrete-in-time approach for flow shear. Panel 2.7b shows the results obtained with our new approach, which treats flow shear continuously over time. We point out that the inserts have a different scale for their y -axis. Here, $\rho_i k_y = 0.3$, $\gamma_E = 0.2$ and all other physical parameters are set to their Cyclone Base Case values.

The first case highlights the effect on real-space quantities when the NGP approximation is being made in Fourier transforms. Given a field $\Psi(t, x, y)$ with the initial condition $\Psi(0, x, y) = \cos [2(\Delta k)y]$, we solve a simple shearing motion in the lab frame


$$\left. \frac{\partial \Psi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Psi}{\partial y} = 0, \quad (2.56)$$

with the shearing rate $\gamma_E = 0.2$ and $\Delta k_x = \Delta k_y = \Delta k = 0.05$. To show the effect of the NGP approximation on Fourier transforms, this is solved numerically for the Fourier-transformed Ψ in the shearing frame (x, y^*) . In this frame,


$$\left. \frac{\partial \hat{\Psi}_{\mathbf{k}}}{\partial t} \right|_{x,y^*} = 0, \quad (2.57)$$

so the only operation we have to apply at every time step is the wave number re-mapping described in Section 2.3.4, followed by an inverse Fourier transform to visualise Ψ in real space. In the discrete-in-time approach, the transform is given by (2.54) without the first phase factor. In the continuous-in-time approach, the phase factor is included. The time evolution obtained with each approach is shown in Figure 2.8 for $t \in [0, T]$, where T is the time required for $k_x^*(t)$ to change from 0 to $-\Delta k_x$. The correct shearing motion is indicated by the line $y^* = 0$. Until the first re-map, i.e. $t < T/2$, the discrete-in-time approach can be seen to lag behind the correct shearing motion (\bar{k}_x remains 0). After the re-map (\bar{k}_x has now jumped to $-\Delta k_x$), it is ahead of the correct motion, until both coincide again at $t = T$ (i.e. when $\bar{k}_x = k_x^*$). In the continuous-in-time approach, the exact time evolution of k_x^* is included in the Fourier transform, and the numerical result coincides with the correct shearing motion at all times.

Our second simplified case illustrates the effect of the NGP approximation in Fourier transforms, in the presence of a nonlinear term. In this test case, the NGP approximation becomes exact when $t = n/\gamma_E$, with $n = 0, 1, 2, \dots$; i.e. at those times, $\bar{k}_x(T, k_x, k_y) = k_x^*(T, k_x, k_y)$ for every k_x and k_y in the simulation, and Fourier transforming is equivalent in both the discrete- and continuous-in-time approaches. At $t \neq n/\gamma_E$ however, Fourier transforming in the discrete-in-time approach will produce errors in the nonlinear term. We will now show that those errors can accumulate over



(a) Discrete-in-time



(b) Continuous-in-time

Figure 2.8: Time evolution of Ψ in the lab frame, according to equation (2.56). The discrete-in-time approach (2.8a) does not include the phase factor appearing in Fourier transforms, whereas the continuous-in-time approach (2.8b) does. Both methods become equivalent at $t = T$, when the NGP approximation is exact. The dashed black line corresponds to $y^* = 0$ and shows the tilt that the mean flow shear should be producing. The numerical parameters for this first test case are given in Section 2.4.2.

time, and alter the time evolution of the simulation. For our test, we consider two fields, $\Psi(t, x, y)$ and $\Phi(t, x, y)$, with initial conditions

$$\Psi_0(x, y) = \cos [2(\Delta k)y], \quad (2.58)$$

$$\Phi_0(x, y) = \frac{C}{\sigma^2} e^{-(x^2+y^2)/(2\sigma^2)}, \quad (2.59)$$

and with $\Delta k_x = \Delta k_y = \Delta k = 0.025$, $N_x = 467$, $N_y = 234$, $C = -10^5$, and $\sigma = 5\pi$. We suppose that their respective time evolution is given by

$$\left. \frac{\partial \Psi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Psi}{\partial y} + \{\Psi, \Phi\}_{x,y} = 0, \quad (2.60)$$

$$\left. \frac{\partial \Phi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Phi}{\partial y} = 0, \quad (2.61)$$

where $\{\cdot, \cdot\}$ is the Poisson bracket defined by equation (2.26). The analytical solution for Ψ can be computed as being

$$\Psi(t, x, y) = \cos [2(\Delta k)Y(t, x, y)], \quad (2.62)$$

where we defined $Y(t, x, y) = y^* \cos(\Phi_0 t / \sigma^2) - x \sin(\Phi_0 t / \sigma^2)$. Numerically, equations (2.60) and (2.61) are solved in Fourier space, and the Poisson bracket is computed pseudo-spectrally, as explained in Section 2.2.5. Note that the maximum amplitude of the analytical solution is $|\Psi| = 1$.

In Figure 2.9a, we plot the analytical solution (2.62) over a time period $T = 1/\gamma_E$. We indicate with dots the areas of the box where analytical and numerical results should not be compared, because the numerical solution is periodic in x and y , while the analytical one is not. In Figure 2.9b, we show the difference $|\Delta \Psi| = |\Psi_{\text{num}} - \Psi|$ between the analytical solution and the numerical result, obtained with the discrete-in-time algorithm. Although the numerical solution looks more correct at $t = T/2$ (when the NGP approximation is exact for half of the values of k_y) and $t = T$ (when it is exact for all values of k_y), errors in the pseudo-spectral computation of the Poisson bracket at $0 < t < 1/\gamma_E$ have altered the time-evolution of the simulation. In Figure 2.9c, we see that the continuous-in-time approach deviates from the analytical solution


by less than 0.01 at any time, compared to differences of up to 1.4 observed with the discrete-in-time scheme.

Finally, Figure 2.10 illustrates how the error $|\Delta\Psi|$ is affected when Δk is reduced. A smaller Δk (i.e. a larger simulation domain) does not alter the pattern or the maximum amplitude of the error produced by the missing phase factor. It does however reduce the error within the original simulation domain. This simple test case therefore suggests that the error from the missing phase factor in local gyrokinetic codes would not vanish when simply reducing Δk_x , which agrees with the observations made in [51]. This would also suggest that to make this error vanish, only the original portion of the simulation domain should be considered, while simultaneously reducing Δk_x . However, we report in Section 2.4.3 that nonlinear simulations do not seem to agree with this last point.


2.4.3 Nonlinear simulations

Having validated the implementation of the continuous-in-time approach in Sections 2.4.1 and 2.4.2, we now present nonlinear, electrostatic gyrokinetic simulations performed with the local code GS2. They show that our new approach to flow shear can lead to computational savings compared to the discrete-in-time approach, and that differences in fluxes can persist if Fourier transforms are not evaluated correctly when computing the nonlinear term of the gyrokinetic equation. The plasma parameters are identical to those of the JET shot #68448 (presented in Table B.1 of Appendix B), apart from $a/L_{T_i} = 1.756$ and $\gamma_E = -0.063$. The simulations include two kinetic species (deuterium ions and electrons), numerical hyperviscosity [80], and collisions. The numerical resolutions are the same as the ones listed as ‘values used’ in Table D.1 of Appendix D, apart from $K_x = 3.49/\rho_i$ and $K_y = 0.99/\rho_i$.


Figure 2.11 shows how the turbulent ion heat flux Q_i differs in simulations using the new continuous-in-time approach to flow shear, compared to simulations with the discrete-in-time approach. In Figure 2.11a, Δk_x is the only parameter that is being changed (along with the number of k'_x s so as to keep the maximum wave number fixed). It can be seen that the two approaches to flow shear do not converge to the



(a) Analytical time evolution of Ψ .




(b) $|\Delta\Psi| = |\Psi_{\text{num}} - \Psi|$, using the discrete-in-time approach.




(c) $|\Delta\Psi| = |\Psi_{\text{num}} - \Psi|$, using the continuous-in-time approach.


Figure 2.9: Analytical solution to equations (2.60) and (2.61) for Ψ , and comparison to numerical solutions obtained with the discrete- and continuous-in-time approaches. At $t = T$, the NGP approximation is exact, making both approaches equivalent for that time step. The numerical parameters for this second test case are given in Section 2.4.2.




(a) Discrete-in-time, $\Delta k_x = 0.025$.
Identical to $t/T = 0.25$ in Figure 2.9b.



(b) Discrete-in-time, $\Delta k_x = 0.0125$.



(c) Continuous-in-time, $\Delta k_x = 0.025$.
Identical to $t/T = 0.25$ in Figure 2.9c.



(d) Continuous-in-time, $\Delta k_x = 0.0125$.

Figure 2.10: $|\Delta\Psi| = |\Psi_{\text{num}} - \Psi|$ for the nominal Δk_x and for Δk_x halved, with both discrete- and continuous-in-time approaches.

same heat flux. In this particular case, the value from the continuous-in-time method is lower by about 20%. Simulations were performed by only adding the phase factor to the discrete-in-time approach (blue crosses), and other simulations were performed by removing the phase factor from the continuous approach (red dots): both clearly show that it is this phase factor that explains the difference in convergence, and that the changes to the linear parts of the code have little impact at large L_x (as one would expect from the discussion in Section 2.4.1). At lower values of L_x in Figure 2.11a, the main result is that the continuous-in-time approach stays closer to its converged value for smaller box sizes than the discrete approach. This agrees with the results of Figure 2.7, which showed that the continuous algorithm is able to reproduce the linear physics correctly with a larger Δk_x (smaller L_x) than the discrete-in-time approach. At the smallest L_x considered here, simulations where the phase factor was added to the discrete-in-time approach (blue cross at $L_x \sim 20\rho_i$) or where the phase factor removed from the continuous approach (red dot at $L_x \sim 20\rho_i$) indicate that, at small L_x , the changes to the linear parts of the code can play a bigger role than the error from the missing phase factor. Surprisingly, we also note that the phase factor appears to have almost no impact in the simulations with $L_x \sim 20\rho_i$, unlike at $L_x \sim 150\rho_i$. At $L_x \sim 20\rho_i$, we observe that the size of the simulation domain is becoming comparable to the radial extent of the eddies. And from Figure 2.10, we know that the error from the missing phase factor is strongest at the radial edges of the simulation. When the box is small enough, we therefore expect the missing phase factor to be unable to substantially affect a whole eddy: only the parts of the eddy that are close to the radial edges of the box are strongly impacted. In contrast, when the box becomes larger, entire eddies are affected by the phase factor. However, additional work would be needed to determine if this could explain why simulations with smaller boxes are less affected by the phase factor. Figure 2.11b is meant to highlight the cases with smaller L_x . The fluxes are this time plotted against $T_{\text{map}}(k_y)/T_F$, i.e., the number of wave-number re-mappings that occur per Floquet period (meaning that the smallest L_x is now on the right of the Figure). T_{map} is evaluated at the wave number k_y with the largest amplitudes in the simulation. All the points are identical to those on Figure 2.11a, except for those in square boxes for which both Δk_x and Δk_y were changed

(again, keeping the maximum wave numbers fixed). The conclusions are the same as for Figure 2.11a.

It is interesting to note that the exact mechanism through which the added phase factor decreases the converged heat flux remains unclear. The only notable difference when adding the phase is that the amplitude of the box-scale zonal mode ($k_x = \Delta k_x$, $k_y = 0$) decreases, as is illustrated by Figure 2.12. We previously showed in Figures 2.10a and 2.10b that omitting the phase factor produces errors that are larger at the radial edges of the simulation domain. However it is not yet clear how this error relates to the box-scale zonal mode, and how this mode might affect the level of transport observed in simulations. Future work could for example determine whether the larger zonal flows in the absence of phase factor might, over a portion of the simulation domain, oppose the background flow shear enough to enable Floquet modes to grow faster and hence lead to higher turbulent fluxes.

Contrary to what was suggested at the end of Section 2.4.2, our simulations seem to indicate that considering only a smaller subset of a large simulation domain does not reduce the error from the missing phase factor. One could have imagined turbulent fluctuations in the absence of the phase factor setting up a region in the middle of the box where, on average over time, the temperature gradient is substantially enhanced. However we do not observe this. Further work would be needed to confirm whether a different turbulent heat diffusivity $\chi_i = -Q_i/(dT_i/dr_\psi)$ is indeed obtained with and without phase factor. In particular, the diffusivity would have to be averaged over time and space for small subsets of the computational domain.


The implementation of the continuous-in-time approach can have an impact on computational costs in two ways. Firstly, the additional flexibility in the choice of Δk_x with the continuous approach can lead to substantial computational savings. In the nonlinear simulations presented here, the discrete-in-time approach needs $T_{\text{map}}/T_{\text{F}} \sim 0.12$ to remain converged, whereas the continuous approach is still approximately correct at $T_{\text{map}}/T_{\text{F}} \sim 0.18$. This corresponded to a 46% decrease in computing time. Secondly, the continuous approach requires three large matrices to be inverted at initialisation or when changing the time-step. The discrete approach only requires a single matrix inversion. In theory, this could increase the computing time by up to a factor of three

if a simulation were to require its time-step to change very often. In practice however, the time-step typically changes only a few times throughout a simulation. In the many simulations we have carried out with the continuous approach, the need for three matrix inversions has never increased the computing time by more than 25%, and the typical increase is around 15%.

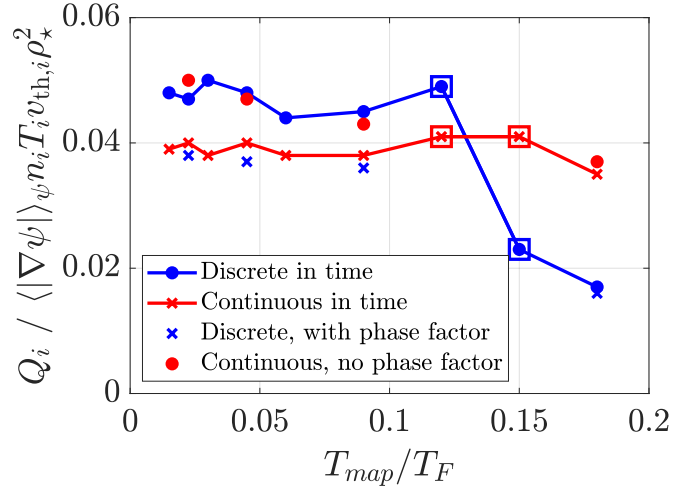
Lastly, while the converged fluxes for this JET discharge differ significantly between the discrete and continuous approaches, it is important to stress that this would by no means be the case for every nonlinear simulation that includes flow shear. In fact, the majority of simulations that we have performed so far yield the same fluxes with both approaches, provided Δk_x is small enough.

2.5 Conclusions

In this Chapter, we presented the system of equations solved by local gyrokinetic codes such as **GS2** in the electrostatic approximation. We defined shearing-frame coordinates that allow the equations to be Fourier transformed across the magnetic field, and we gave a brief overview of the key boundary conditions that are employed. We then presented the linearly implicit algorithm used in **GS2**. We formulated a new approach to flow shear for linearly implicit, local δf gyrokinetic codes. This method treats flow shear continuously over time, unlike the widely-used, discrete-in-time approach developed by Hammett *et al* [49]. The main complication resides in the response-matrix approach adopted by codes with an implicit-in-time linear algorithm: when $\gamma_E \neq 0$, this matrix becomes time dependent and it is prohibitively expensive to re-compute it at every time step. In the discrete-in-time approach to flow shear, this time dependence was approximated by combining a nearest-grid-point approximation (NGP) for radial wave numbers, together with a wave-number re-mapping procedure. In the new approach to flow shear, a decomposition of terms allows the time dependences in the linear gyrokinetic equation to be treated explicitly in time, without further restricting the time step size. Additionally, time dependences arising in the quasineutrality equation are approximated by linear interpolation, and Hammett's re-mapping in wave number is applied without NGP approximation. Following these steps leads




(a)



(b)

Figure 2.11: Turbulent heat flux, plotted for different Δk_x in 2.11a, and for different $T_{map}(k_y)/T_F$ in 2.11b (here we choose $\rho_i k_y = 0.5$ which is the wave number with the largest amplitudes). All points in 2.11a and 2.11b are identical, except for the ones highlighted by square boxes, for which both Δk_x and Δk_y were modified. Here, $a/L_{T_i} = 1.756$, $\gamma_E = -0.063$, and all other parameters are set to their values presented in Appendix B.



(a) *Discrete approach, no phase factor.* (b) *Continuous approach with phase factor.*

Figure 2.12: Amplitude of the electrostatic potential for the box-scale zonal mode in solid black and for the dominant non-zonal mode in dashed red. Here, $a/L_{Ti} = 1.756$, $\gamma_E = -0.063$, and all other parameters are set to their values presented in Appendix B.

to an implicit-in-time linear algorithm, where flow shear is included continuously in time, and for which response matrices only need to be computed during initialisation. Finally, the pseudo-spectral evaluation of the nonlinear term also requires modifications to take into account the time dependence of the radial wave number, as has been argued in [51].

We implemented this new approach in the GS2 code, and showed numerical simulations that validate the changes made to both linear and nonlinear parts of the code. Linearly, the discrete-in-time approach produces discontinuous jumps at every re-map in wave number, and it fails to reproduce Floquet oscillations when the grid-spacing in radial wave number becomes too large. The continuous-in-time approach produces no such jumps, and is able to capture the correct Floquet behaviour for a wider range of resolutions in wave number. For the nonlinear term, a first test case illustrates the errors in real space quantities when an NGP approximation is made in Fourier transforms, and how the continuous-in-time approach gets rid of those errors by including a phase factor. A second test case shows the impact of these errors when the equations for time evolution include a pseudo-spectrally-computed Poisson bracket: the discrete-in-time approach accumulates errors from the NGP approximation, while the continuous-in-time approach correctly reproduces the expected analytical result. Finally, nonlinear gyrokinetic simulations of a JET discharge demonstrate that the

continuous and discrete approaches can, in some cases, lead to a different converged value for the turbulent heat flux (20% difference in this example). However for the majority of nonlinear simulations that we have performed with flow shear, the two methods yield essentially identical fluxes. It was also shown that the continuous approach can lead to substantial computational savings by allowing for radially narrower boxes (46% decrease in computing time for this example). When using the same box size, the continuous implementation is somewhat slower than the discrete algorithm (about 15% slower for cases in the example presented here, and at most 25% in all simulations performed so far).

Chapter 3

Bistable turbulence in the presence of flow shear

3.1 Introduction

According to the most common paradigm for magnetic-confinement-fusion plasmas, a unique stationary turbulent state can be identified given a certain stirring mechanism and a set of equilibrium plasma parameters. Multistable solutions – where identical parameters admit distinct turbulent states – are known to occur in neutral fluids [52–54] where they are associated with bifurcations and hysteretic behaviour [55, 56]. Multistability has also been reported in weakly magnetised systems of charged fluids [57, 58]. In this Chapter, we report the discovery of bistable turbulence in a strongly magnetised plasma, using direct numerical simulations. We find that bistability arises in such a plasma through the interplay of two crucial mechanisms: an externally imposed mean flow shear and self-generated “zonal” flows. Our observations are made in a toroidal geometry typically encountered in magnetic-confinement-fusion experiments, though the results could likely be generalised to other systems.

As was mentioned in Section 1.4, previous studies have already established that sheared flows play an important role in regulating turbulence in fusion plasmas. When no mean flow shear is imposed via external sources, it was found that the plasma spontaneously generates sheared “zonal” flows, which contribute to the saturation of turbulence [30–36]. In consequence, it has since then been assumed that the effect of zonal flows is always to reduce turbulent transport – i.e. that zonal flows are beneficial for the plasma confinement in tokamaks. When a mean flow shear is externally imposed

in the plasma, it provides an additional mechanism for saturation. It was found that the shear in the mean flow perpendicular to the magnetic field suppresses turbulent fluctuations [37–48]. It was also shown that the effect of the mean flow shear on the saturation of turbulence weakens away from marginal stability [35, 81].

Here, we find that the weakening suppressive effect of mean shear on turbulence is in fact characterised by an abrupt jump from turbulent states that are strongly affected by mean shear to states that are weakly affected by mean shear. Most importantly, we show that a region of parameter space near this transition allows for the existence of two distinct turbulent states, with identical equilibrium parameters but significantly different levels of transport. Surprisingly, we find that the presence of zonal flows leads to high-transport states, and that zonal flows can thus be detrimental to the confinement in fusion devices – the opposite effect to the one that is usually assumed.

This Chapter is organised in the following way. In Section 3.2, we first explain the context that led us to discover two distinct states in a fusion plasma. We then characterise each of the two states in Section 3.3. Finally, we discuss some of the consequences of our work on turbulent transport in Section 3.4.

3.2 Running nonlinear simulations with a mean flow shear

We choose to study turbulent transport in the presence of an externally imposed flow shear by focusing on the JET discharge #68448 [79], which we described in Appendix B. The turbulence is simulated using the local, δf gyrokinetic code GS2 and the continuous-in-time algorithm for flow shear presented in Chapter 2. We now present two procedures that can be followed to run nonlinear gyrokinetic simulations with a mean flow shear.

The standard procedure to run a simulation with flow shear is the following. First, a simulation is run starting from small fluctuation amplitudes, with the mean flow-shear rate set to zero; once the system has reached a stationary turbulent state, flow shear is introduced in the simulation; this simulation is then run until the system settles in a new stationary state. The final state is considered to be the turbulent state

produced by the plasma in the presence of flow shear. This practice is commonly used for nonlinear gyrokinetic simulations in the presence of flow shear (see for example [82]). It is especially useful when exploring so-called subcritical turbulence [43, 44, 83], where the mean flow shear suppresses linear instabilities on average (i.e. $\langle \gamma \rangle_t \leq 0$) but where the transient growth of Floquet modes is sufficient to sustain turbulent transport. For cases of subcritical turbulence, a stationary turbulent state cannot be reached in simulations including flow shear from the start with arbitrarily small initial fluctuation amplitudes. Hence the usefulness of initialising the system from a zero-shear stationary state.

Surprisingly, when following this standard procedure, we find that our simulations across the core region of the discharge #68448 produce fluxes that are almost an order of magnitude larger than those found experimentally. From here on, we choose to characterise this discrepancy by focusing on a single radius at $r_\psi/a = 0.51$, located well inside the core of the plasma. The choice of this radial location is made to limit the computational cost of simulations. Indeed, the magnetic shear \hat{s} is weaker in the central region of the plasma, leading to less stringent requirements for the numerical resolution in the radial coordinate of flux-tube simulations.

Linear simulations of the JET discharge #68448 show that $\langle \gamma \rangle_t \leq 0$ across most of the core. However, $\langle \gamma \rangle_t > 0$ in the inner region of the core where flow shear is weaker. This allows nonlinear simulations with flow shear to be carried out at these radii following a different approach than the standard procedure that we described above. The alternative approach is to start from small fluctuation amplitudes and directly include flow shear from the beginning of the simulation. The algorithm that we developed and presented in Chapter 2 is particularly well suited for this second approach, as it is efficient at correctly reproducing linear instabilities in the presence of a mean flow shear (see for example the results shown in Section 2.4.1). The radius $r_\psi/a = 0.51$ at which we wish to characterise the discrepancy between simulations and experiment has $\langle \gamma \rangle_t > 0$ so that both the standard and the alternative procedures to include flow shear can be followed at this location.

We hence consider nonlinear simulations with small initial amplitudes and flow shear included from the start at the inner radius $r_\psi/a = 0.51$. Remarkably, we discover

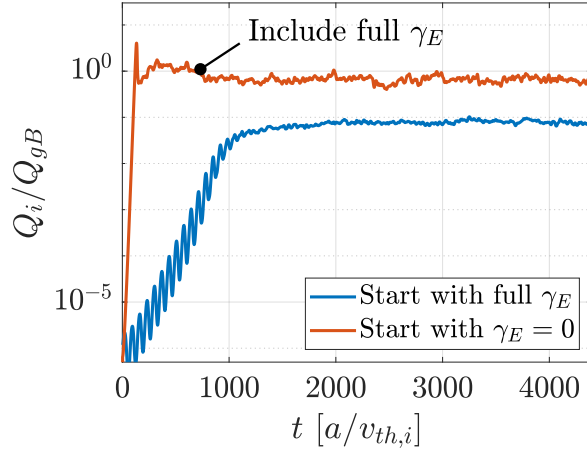


Figure 3.1: Two distinct saturated turbulent states are obtained depending on whether amplitudes are small (low-transport state in blue) or large (high-transport state in red) when flow shear is included in the simulation. Here, $a/L_{T_i} = 1.76$ and all other parameters are set to their values presented in Appendix B.

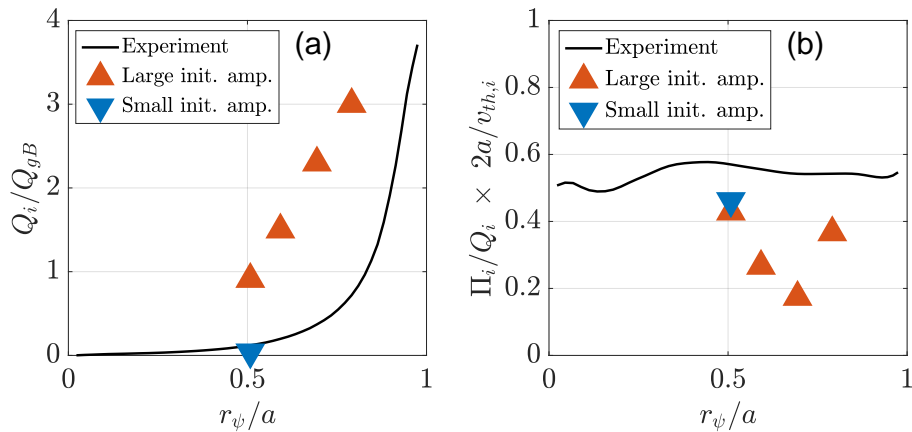


Figure 3.2: Radial profiles of the experimentally observed fluxes, compared to the turbulent fluxes obtained from simulations performed by either including flow shear from the start (i.e., small initial amplitudes), or by restarting from a zero-shear saturated state (i.e., large initial amplitudes). Panel (a) shows the heat flux, and panel (b) the momentum-to-heat-flux ratio. Here, apart from $a/L_{T_i} = 1.76$ and the varying value for r_ψ/a , all parameters are set to their values presented in Appendix B.

that this produces a distinct stationary turbulent state compared to simulations that are initialised from a zero-shear turbulent state. The turbulent fluxes associated with this distinct state appear to lie much closer to the experimental observations. The time traces of the heat flux for both types of simulations are shown in Figure 3.1. The radial profiles for the turbulent fluxes of heat and toroidal-angular momentum obtained from simulations are compared to the experimental observations in Figure 3.2. To our knowledge, this is the first instance of two distinct stationary turbulent states being obtained with identical equilibrium plasma parameters in the context of magnetic-confinement fusion.

The original objective of our work was to use local gyrokinetic simulations to reproduce and predict the profiles of toroidal rotation driven by neutral-beam injection in fusion plasmas. Our approach to that problem is presented in Appendix E. For this approach to successfully reconstruct rotation profiles, turbulent fluxes obtained from gyrokinetic simulations in the presence of flow shear need to match reasonably well the levels of transport reported in the experiment. The existence of two distinct turbulent states makes this more challenging to achieve. Instead of pursuing our original objective, we choose to focus on characterising these two states. In the next Sections, we describe the distinguishing features of the two states, as well as the mechanisms that regulate the saturation of turbulence in each of them.

3.3 Existence of two distinct turbulent states

In the following Sections, we characterise the two distinct, stationary turbulent states that we obtain with identical equilibrium parameters in local gyrokinetic simulations. In Section 3.3.1, we first show the main distinguishing features of the two states. In Section 3.3.2 we then underline the important role of zonal modes. In Section 3.3.3, we highlight the differences between the nonlinear decorrelation times in each of the states, and we relate this to linear instabilities. Finally, in Section 3.3.4, we present two arguments that can be used to describe the region of parameter space in which two distinct states are observed.

3.3.1 Identifying two distinct turbulent states

The first, and perhaps most striking difference between the two states is the level of turbulent transport associated with each of them: the fluxes obtained by initialising the system with a zero-shear state are considerably larger than the fluxes obtained from small initial amplitudes. Hence, we refer to the former as “high-transport” states, and to the latter as “low-transport” states.

By varying the equilibrium temperature-gradient length scale a/L_{T_i} , we observe that the two states can be obtained across a range of temperature gradients close to marginal stability. This is shown in Figure 3.3 where saturated values of Q_i/Q_{gB} are plotted against a/L_{T_i} for a particular value of γ_E . The branch of low-transport states is indicated by downward-facing blue triangles, and the branch of high-transport states is indicated by upward-facing red triangles. The dashed vertical line marks the temperature gradient below which the system is linearly stable on average with the mean flow shear rate $\gamma_E = -0.079$. Figure 3.4 confirms our observation for a range of γ_E . We note that both low-transport and high-transport states can exist above and below the threshold for average linear instability. Previous work has already established that turbulence can either be sustained by a linear instability (supercritical turbulence) or by transient linear growth (known as subcritical turbulence [43, 44, 83]).

Within the region where the two distinct states can occur, simulations with large initial amplitudes saturate in the high-transport state, while simulations with small initial amplitudes settle in the low-transport state. The two states are stationary, and no back-and-forth transitions in time between the two states are observed when simulations are adequately resolved (we discuss an important effect of numerical resolution in Section 3.3.5). Hence, we follow the nomenclature employed in previous work [55–57] and refer to the existence of two distinct, stationary turbulent states for identical equilibrium parameters as “bistability”. The parameters corresponding to $r_\psi/a = 0.51$ in the JET discharge considered here are indicated by a black star in Figure 3.4. As we have mentioned in Section 3.2, we find that the fluxes computed in the low-transport states match the levels of transport observed in the experiment, while the fluxes computed in the high-transport states differ from it by an order of magnitude. From Figure




Figure 3.3: Dependence of the turbulent ion heat flux on the inverse ion-temperature-gradient scale length. In the simulations labelled by green ‘+’ signs, the externally imposed mean flow shear was set to zero. For all other simulations, $\gamma_E = -0.079$ was used. Zonal modes are artificially zeroed out in simulations labelled by black crosses. The black circle denotes a simulation where amplitudes decay in time and no saturated turbulent state is observed. The dashed line marks the temperature gradient below which there is no effective linear instability with $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

3.4, we also notice that the region of bistability broadens with stronger shear. The fact that no low transport states appear below a certain value of γ_E will be discussed in Section 3.3.4. Interestingly, we observe that the experiment appears to lie at the narrow tip of the bistable region at low γ_E . Based on the local simulations that we present here, it may appear counter-intuitive for the experiment to lie so close to the threshold between low transport and high transport, given that a slight perturbation might then trigger a transition to the high-transport branch. However, we are not yet in a position to draw conclusions from this observation: we have so far focused on a single comparison between simulations and experiment, and error bars in experimental measurements imply that the experiment might actually be lying at a slightly different point in the $(\gamma_E, a/L_{T_i})$ plane.

We have established that the two states differ dramatically in their associated turbulent fluxes. We now present results showing that they are also distinguished by striking differences in the real-space structure of the turbulence. In Figure 3.5, we show typical snapshots of turbulence in a low-transport state. The contours of the fluctuating elec-

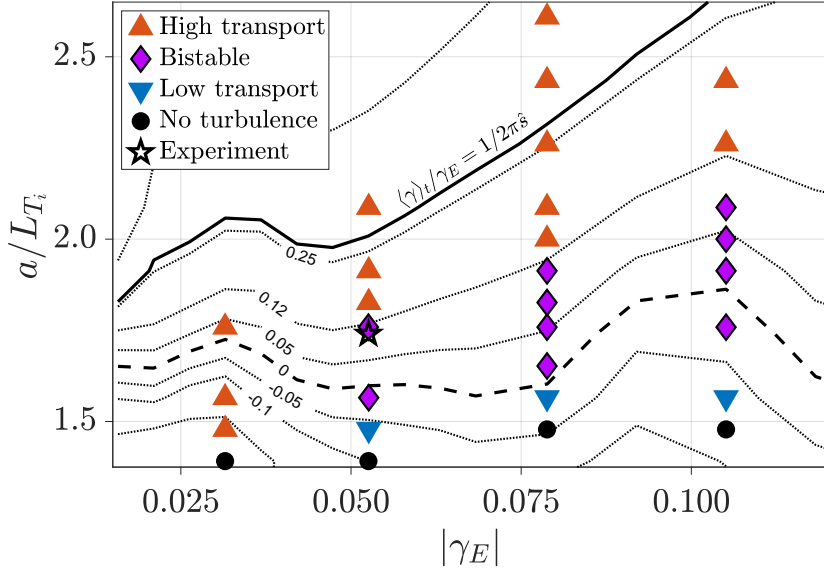


Figure 3.4: Occurrences of the two distinct turbulent states. Purple diamonds indicate parameters at which both a high- and a low-transport state can be obtained. Black circles denote simulations where amplitudes are decaying in time and no saturation of the turbulence is observed. Lines indicate contours of constant $\langle \gamma \rangle_t / \gamma_E$, where $\langle \gamma \rangle_t$ is the effective linear growth rate. The dashed line marks the temperature gradient below which $\langle \gamma \rangle_t < 0$. The solid line indicates $T_F \langle \gamma \rangle_t = 1$ (recalling that $T_F = 2\pi \hat{s} / \gamma_E$). Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. $\langle \gamma \rangle_t$ is obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$.

trostatic potential are plotted in the plane perpendicular to the magnetic field at the outboard of the torus. We recall from Section 2.2.3 that the x coordinate measures the distance along the normal to a magnetic-flux surface and the y coordinate labels the magnetic-field lines within a magnetic-flux surface. We also recall that the simulation is done in the frame moving with the mean flow at $x = 0$. In the present case, the y -component of the mean flow has the opposite sign to x . The turbulent eddies feature a clear tilt as they are being sheared by the mean flow, similarly to previous observations in [35, 46, 81, 84]. In Figure 3.6, we show consecutive snapshots of turbulence with the same equilibrium parameters as in Figure 3.5, but in the high-transport state: bands of high-amplitude eddies propagate radially across the simulation domain, and eddies do not feature any clear tilt across the box. This intermittent high-transport state is reminiscent of the advecting structures reported in [72, 85].




Figure 3.5: Consecutive snapshots of the turbulence in real space for the low-transport state where $a/L_{T_i} = 1.76$ and $\gamma_E = -0.079$. In the top panels, the fluctuating electrostatic potential is plotted at three successive times at the outboard of the torus, in the plane perpendicular to \mathbf{B} . In the bottom panels, the zonal flow is plotted at the same times. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

The differences in real space between the two states are also reflected in quantities that are Fourier analysed across the magnetic field (using the spectral representation that we presented in Section 2.2.4). Here, we choose to focus on the energy spectrum, averaged in time and θ , $E_{\mathbf{k}} = (|k_x| + |k_y|) \langle |\hat{\varphi}_{\mathbf{k}}|^2 \rangle_{t,\theta}$. Recalling that the electrostatic potential is a real quantity ($\varphi \in \mathbb{R}$), we follow the convention of the GS2 code and only consider modes with $k_y \geq 0$. For this discussion, we also exclude modes with $k_y = 0$ for reasons that will be addressed in Section 3.3.5.

In Figure 3.7, we plot the energy spectra for typical high-transport and low-transport states obtained with identical equilibrium parameters. In the absence of a mean flow shear, the radial-wave-number spectrum is symmetric. When $\gamma_E \neq 0$, however, this symmetry in the radial direction is broken, and asymmetric k_x -spectra can be obtained. Comparing panels (a) and (b) of Figure 3.7, we observe that the asymmetry in k_x is noticeably stronger in the low-transport state than it is in the high-transport state. In anticipation of the discussion that we propose in the following Sections, this observation suggests that the mean flow shear plays a more important role in the saturation of

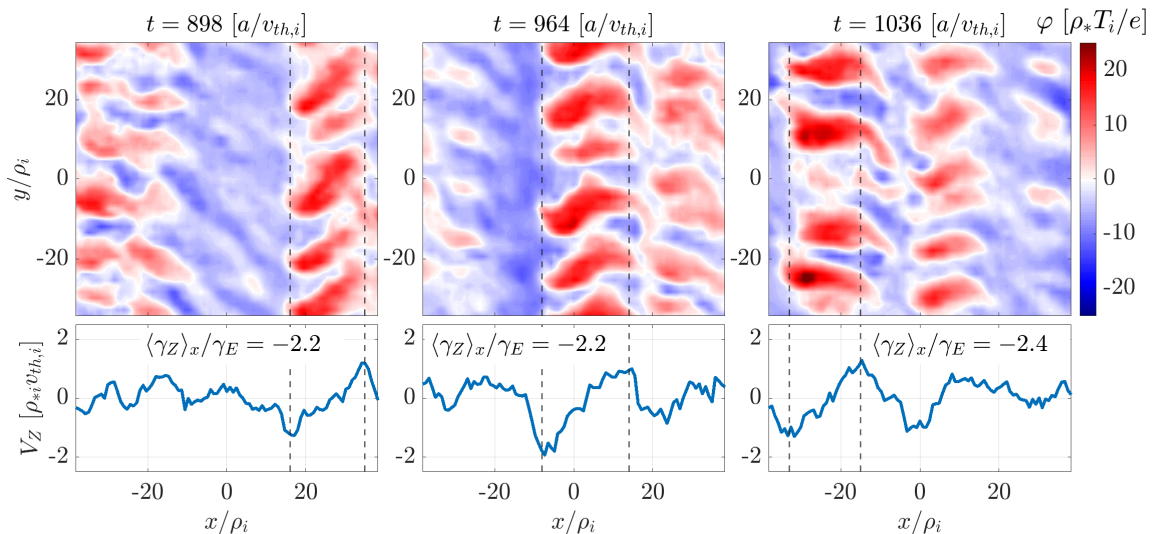


Figure 3.6: Consecutive snapshots of the turbulence in real space for the high-transport state where $a/L_{T_i} = 1.76$ and $\gamma_E = -0.079$. In the top panels, the fluctuating electrostatic potential is plotted at three successive times at the outboard of the torus, in the plane perpendicular to \mathbf{B} . In the bottom panels, the zonal flow is plotted at the same times. The zonal shear averaged between the two vertical dashed lines is compared to the externally imposed mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

turbulence for low-transport states than for states on the high-transport branch. We point out that the same observation applies to the k_x spectra of other quantities, such as the fluctuating particle density or the fluctuating temperature.

Finally, panels (c-f) of Figure 3.7 compare the energy spectra in k_x and k_y to the scaling law obtained via a critical-balance argument proposed in the absence of a mean flow shear [65]. In that argument, the time scales associated with the parallel streaming and the perpendicular decorrelation of turbulent eddies are comparable at all spatial scales, leading to the scaling for the energy spectrum $E_{\mathbf{k}} \sim |\mathbf{k}|^{-7/3}$. We observe good agreement with this scaling for the k_x -spectrum in both states. In panel (f) of Figure 3.7, we notice that the k_y -spectrum in the low-transport state appears to be steeper than what would be expected from the critical balance argument. However the spectrum seems to become less steep at large k_y , which might indicate that a slightly higher maximum k_y would be needed for the simulations to fully agree with critical balance. Alternatively, this steeper k_y spectrum could be the result of flow shear

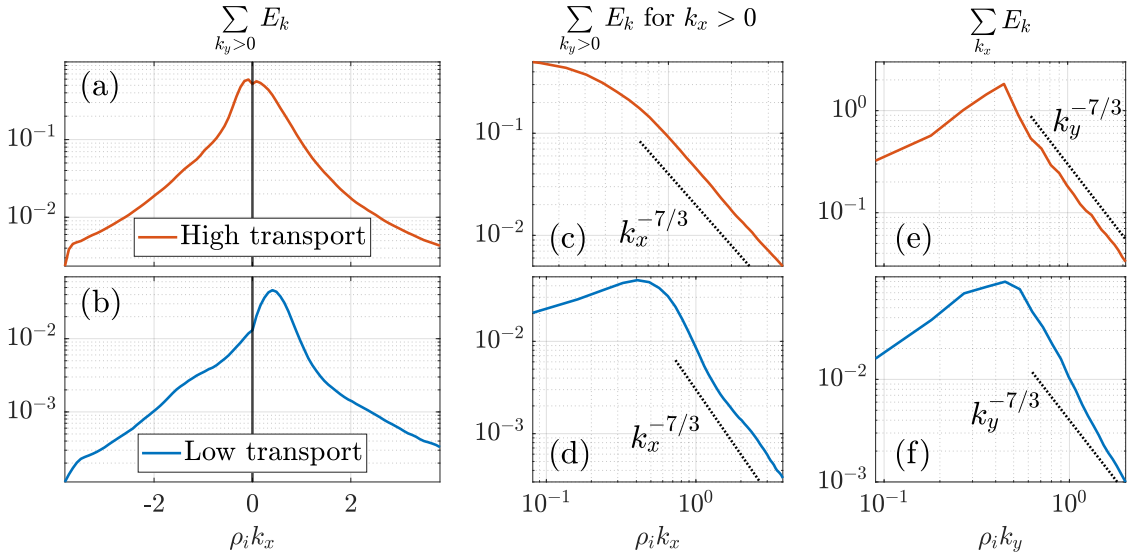


Figure 3.7: Energy spectra for a high-transport (a,c,e) and a low-transport (b,d,f) state obtained with identical equilibrium parameters. Panels (c) to (f) compare the spectra to predictions of the critical-balance argument, which does not take into account the effect from mean flow shear. Here, $a/L_{T_i} = 2.0$, $\gamma_E = -0.105$, and all other parameters are set to their values presented in Appendix B.

advecting structures along the magnetic field, where the dissipation by finite-Larmor-radius effects is stronger. This would suggest that the critical-balance argument may need to be extended to consider turbulent states in which the mean flow shear plays an important role.

3.3.2 Importance of zonal modes

We find that the presence of large “zonal” modes is a crucial distinguishing feature between the two states. Modes are called zonal when they have no spatial variation other than in the radial (x) direction, corresponding to modes with $k_y = 0$ in the spectral representation that we defined in Section 2.2.4. They are linearly stable and cannot be sheared by a toroidal mean flow, but they can exchange energy with non-zonal modes via nonlinear interactions. Zonal modes include zonal flows characterised by the velocity

$$V_Z = -\frac{c}{B} |\hat{\mathbf{b}} \times \nabla r_\psi| \frac{\partial \varphi_Z}{\partial r_\psi}, \quad (3.1)$$

where φ_Z is the zonal part of the electrostatic potential. The zonal flows are known to develop through a secondary instability of the modes driven unstable by the temperature gradient [32, 33]. The associated zonal shearing rate $\gamma_Z = \partial V_Z / \partial r_\psi$ can affect the rest of the turbulence in a manner analogous to the mean flow shear γ_E . When the amplitudes of zonal modes become large enough, we find that the zonal shear can compete with the mean flow shear. In the lower panels of Figure 3.6, we plot the zonal flow V_Z : the zonal shear γ_Z opposes the background shear γ_E within bands propagating radially, and both have comparable magnitudes. In such regions where the zonal and mean shears oppose each other, unsheared non-zonal fluctuations are allowed to grow faster, feeding the zonal modes nonlinearly until the system settles in the high-transport state. We also show in Figure 3.3 that the transport obtained in the complete absence of mean flow shear (green ‘+’ symbols) is much closer to the high-transport states than to the low-transport states.

Meanwhile, zonal modes do not seem to play a crucial role for the turbulent dynamics of low-transport states: unlike in the high-transport states, no long-lived structures with γ_Z cancelling γ_E are observed, as we show in the lower panels of Figure 3.5. Further evidence of this weaker impact of zonal modes in low-transport states can be seen in simulations where we artificially set zonal modes to zero at every time step, indicated by black crosses in Figure 3.3. Despite this unphysical truncation introduced in the system, and independently of the initial condition, a saturated state is obtained that is – apart from a slight change in the flux – indistinguishable from the low-transport state. As a result of the interplay between the zonal modes and the mean flow, high-transport states are only obtained when the initial fluctuation amplitudes are sufficiently large, or when the amplitudes in the low-transport state become large enough for the zonal shear to compete with γ_E ; e.g., when a/L_{T_i} is increased or γ_E is decreased past a certain threshold.

Effectively this shows that, in the presence of a mean flow shear, zonal modes can help sustain high levels of transport. This is the opposite effect to what is commonly expected from zonal modes. We point out that this transport-enhancing role of zonal modes is only observed close to marginal stability. In Figure 3.8, we present simulations where zonal modes are artificially zeroed out beyond the region of existence of low




Figure 3.8: Dependence of the turbulent ion heat flux on the inverse ion-temperature-gradient scale length. Here, $\gamma_E = -0.079$. Zonal modes are artificially zeroed out in simulations labelled by black crosses. The dashed line marks the temperature gradient below which there is no linear instability in the presence of mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

transport states. As can be seen from the Figure, the more widely expected effect of zonal modes is recovered at higher a/L_{T_i} , where their presence is shown to suppress turbulent transport.

3.3.3 Nonlinear decorrelation times

As we have discussed in Sections 3.3.1 and 3.3.2, the turbulent amplitudes in the two distinct states are significantly different, with the mean flow shear regulating the saturation of low-transport states, and zonal shear partially cancelling the effect of mean flow shear in high-transport states. In this Section, we consider the typical time scale associated with the nonlinear decorrelation of turbulent eddies, which we denote by τ_{NL} .

We define this nonlinear decorrelation time scale to be

$$\tau_{\text{NL}}^{-1} = \frac{c}{B} \frac{[\varphi_{\text{NZ}}]_{\text{rms}}}{\ell_x \ell_y}, \quad (3.2)$$

where we denote the eddy correlation length in the x and y direction by ℓ_x and ℓ_y respectively, and where $[\varphi_{\text{NZ}}]_{\text{rms}} = \sqrt{\langle \varphi_{\text{NZ}}^2 \rangle_{t,x,y}}$ is the root mean square of the nonzonal part of the electrostatic potential φ_{NZ} , averaged over time, x and y . The definition

(3.2) is motivated by the fact that the time derivative and the nonlinear term in the gyrokinetic equation (2.4) must have comparable sizes in a saturated turbulent state.




Figure 3.9: Two-point spatial correlation function for a low-transport state (a), a high-transport state (b) and a state with no mean flow shear (c). The correlation lengths ℓ_x and ℓ_y correspond to the e -folding of $\text{Cor}[\varphi_{\text{NZ}}]$ along the solid and dashed lines, respectively. In the three states shown here, $a/L_{T_i} = 1.76$. In (a) and (b), $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

Furthermore, we define the two-point spatial correlation function as

$$\text{Cor}[\Phi](\delta_x, \delta_y) = \frac{\langle \Phi(t, x, y) \Phi(t, x + \delta_x, y + \delta_y) \rangle_{t,x,y}}{\langle \Phi^2(t, x, y) \rangle_{t,x,y}^{1/2} \langle \Phi^2(t, x + \delta_x, y + \delta_y) \rangle_{t,x,y}^{1/2}}, \quad (3.3)$$

for any fluctuating scalar quantity Φ . The correlation lengths ℓ_x and ℓ_y are chosen to correspond to the e -folding of $\text{Cor}[\varphi_{\text{NZ}}]$ along the δ_x direction (adjusted to match the tilt induced by the flow shear) and the δ_y direction. Typical examples of $\text{Cor}[\varphi_{\text{NZ}}]$ are shown in Figure 3.9. Note that the exact definition of the correlation lengths is somewhat arbitrary. Another, choice that is commonly found in the literature is to define ℓ_x and ℓ_y as the integral of $\text{Cor}[\varphi_{\text{NZ}}]$ along the δ_x and δ_y axes – which, in our case, yields similar results to the e -folding lengths.

In Figure 3.10, we compare the nonlinear decorrelation times obtained in low-transport states, high-transport states, and states with no mean flow shear. We observe that, while high-transport states yield similar results to cases with $\gamma_E = 0$, the decorrelation times of low-transport states differ significantly, both in absolute value and in their dependence on the temperature-gradient length scale a/L_{T_i} .

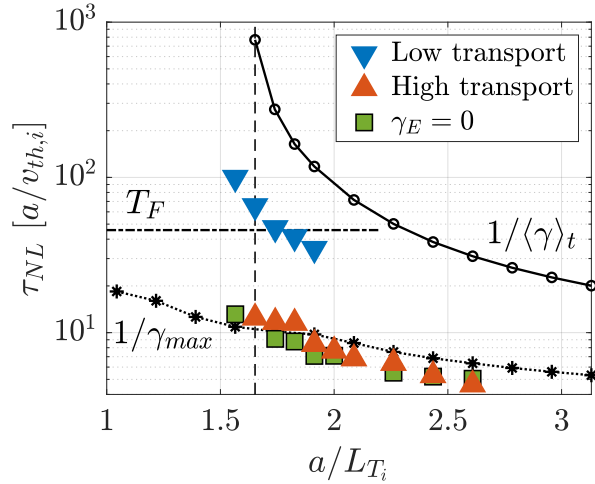


Figure 3.10: Comparison between linear time scales and nonlinear decorrelation time scales. Here $\gamma_E = -0.079$, except for the cases indicated by green squares which have no mean flow shear. Simulations to the left of the dashed vertical line are linearly stable on average with $\gamma_E = -0.079$. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. Linear growth rates are obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest instantaneous and average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$.

We recall from Section 2.2.6 that, linearly, mean flow shear leads to the existence of Floquet modes, which are described by an effective, time-averaged growth rate $\langle \gamma \rangle_t$, and an instantaneous growth rate whose maximum we denote by γ_{\max} . In Figure 3.10, we show how the time scales associated with these two linear growth rates depend on the ion temperature gradient. Comparing the linear time scales with the nonlinear decorrelation time scales, we notice that $\tau_{\text{NL}} \sim 1/\gamma_{\max}$ in the high-transport and no-shear states. In low-transport states, however, τ_{NL} lies between $1/\gamma_{\max}$ and $1/\langle \gamma \rangle_t$, and its dependence on the temperature gradient is best described by $\tau_{\text{NL}} \sim 1/\langle \gamma \rangle_t$. This is a clear indication that there are two significantly different linear time scales that can play an important role in turbulence sheared by a mean flow: the time scale associated with the average growth of Floquet modes, and the time scale associated with the instantaneous growth of the modes.

3.3.4 Identifying the region of bistability

In this Section, we present two arguments that identify a region of the parameter space $(\gamma_E, a/L_{T_i})$ in which bistability is allowed to occur, and we characterise the threshold in temperature gradient above which low-transport states cease to exist.

For the time-averaged growth rate $\langle\gamma\rangle_t$ to correspond to the nonlinear decorrelation time scale in turbulent low-transport states, we argue that one Floquet period T_F must be allowed to pass before the turbulence decorrelates. Indeed, the average linear growth rate $\langle\gamma\rangle_t$ can only be properly defined for time scales that are longer than T_F . Therefore, for $\langle\gamma\rangle_t$ to be able to affect a turbulent plasma, the nonlinear decorrelation time τ_{NL} must be comparable to or longer than the Floquet period. Following this argument, a necessary condition for the existence of low-transport states is that $\tau_{\text{NL}} \gtrsim T_F$. We evaluate τ_{NL} through equation (3.2) and find that this condition is approximately satisfied for low-transport states, as shown on Figure 3.10. By estimating $\tau_{\text{NL}} \sim 1/\langle\gamma\rangle_t$ for low-transport states, we can express the condition necessary for the existence of low-transport states as:

$$T_F \langle\gamma\rangle_t \lesssim 1. \quad (3.4)$$

This condition can be checked by running computationally inexpensive linear simulations. In Figure 3.11, we show the dependence of $T_F \langle\gamma\rangle_t$ on the ion temperature gradient and the mean flow-shear rate. The solid black line indicates the temperature gradient below which the condition (3.4) is met, allowing low-transport states to exist. As can be seen from Figure 3.11, all instances of the turbulent low-transport states are situated below the black line, and thus comply with the condition (3.4). This is also shown in Figure 3.4, where linear simulations are interpolated to obtain the curve along which $T_F \langle\gamma\rangle_t = 1$, indicated by a solid black line. In Section 3.3.1, we observed that no low transport states exist below a certain threshold value in γ_E . Following the argument we just presented, it would appear that this is a result of the $T_F \langle\gamma\rangle_t = 1$ curve getting closer to the $\langle\gamma\rangle_t = 0$ curve. We also note at this point that the condition (3.4) can only provide an upper bound for the threshold in temperature gradient above which low-transport states cease to exist: we showed in Figure 3.10 that τ_{NL}

does not exactly follow $1/\langle\gamma\rangle_t$, but that it rather lies between $1/\gamma_{\max}$ and $1/\langle\gamma\rangle_t$. In consequence, the actual threshold in a/L_{T_i} is lower than predicted by (3.4).




Figure 3.11: Dependence of the product of the Floquet period and the time-averaged linear growth rate on the mean flow shear rate and the ion-temperature-gradient-length scale. We indicate the region where the low-transport state is observed by the dotted line with downward-pointing triangles (nonlinear simulations were only performed for $|\gamma_E| \leq 0.105$). The necessary condition (3.4) for the existence of low-transport states is satisfied below the solid black line. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B. Linear growth rates $\langle\gamma\rangle_t$ is obtained with $\rho_i k_y \in \{0.25, 0.5, 0.75\}$ and the plot shows the fastest average growth rates within this set – typically this corresponds to $\rho_i k_y = 0.5$.

To better characterise this threshold above which no lower states exist (i.e., the transition from purple diamonds to red triangles in Figure 3.4), we expect that the role of zonal modes must be taken into account. Indeed, we presented evidence in Section 3.3.2 that the zonal modes are crucial in sustaining high-transport states. Here, we propose a simple scaling argument that focuses on the partial cancellation of the effect of mean flow shear by the zonal shear. As will become clear, the argument presented here involves order-of-magnitude estimates, which prevent us from making a quantitative prediction for the threshold. Nevertheless, our result characterises the shape of the curve defining the transition threshold in the parameter space $(\gamma_E, a/L_{T_i})$.

We estimate the zonal shearing rate to scale as $\gamma_Z \propto \partial^2 \varphi_Z / \partial x^2 \sim \varphi_Z / \ell_x^2$, where ℓ_x is the radial correlation length of the turbulent eddies. In doing so, we implicitly assume

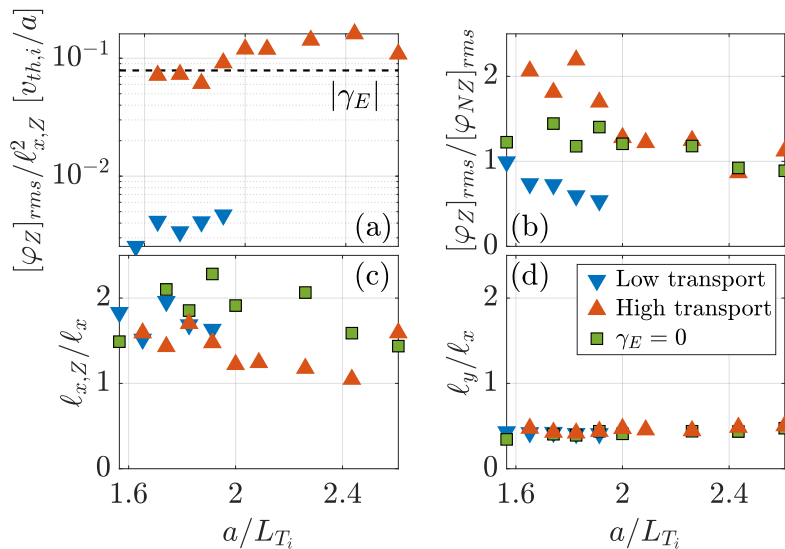


Figure 3.12: Estimates used for the scaling argument at the end of Section 3.3.4. Panel (a) shows the estimated zonal flow shear rate. Panel (b) shows the relative magnitudes of zonal and non-zonal modes. Panel (c) compares the non-zonal radial correlation length ℓ_x to the zonal radial correlation length $\ell_{x,Z}$. Panel (c) compares the non-zonal correlation lengths in x and y . Here, $\gamma_E = -0.079$, except for simulations indicated by green squares which are obtained without mean flow shear. Apart from a/L_{T_i} and γ_E , all parameters are set to their values presented in Appendix B.

that $\ell_x \sim \ell_{x,Z}$, where $\ell_{x,Z}$ is the zonal correlation length computed from $\text{Cor}[\varphi_Z]$. Our simulations indicate that, in the low-transport state, zonal modes have amplitudes comparable to the rest of the turbulence, so we assume $\varphi_Z \sim \varphi$. As we argued in the previous Section, we estimate that $\tau_{\text{NL}} \sim 1/\langle\gamma\rangle_t$ in supercritical low-transport states. Note that a different estimate of τ_{NL} would be needed for subcritical states where $\langle\gamma\rangle_t < 0$. Finally, we find the correlation lengths to be comparable in x and y , i.e., $\ell_x \sim \ell_y$.

Combining these estimates, we obtain $\gamma_Z \sim \langle\gamma\rangle_t$ for low-transport states. We argue that low-transport states transition to the high-transport branch when γ_Z becomes of the order of γ_E , partially cancelling the effect of the mean shear. The transition line in the $(\gamma_E, a/L_{T_i})$ plane should thus follow a curve of constant $\langle\gamma\rangle_t/\gamma_E$. We find reasonable agreement with this statement in Figure 3.4, where contours of constant $\langle\gamma\rangle_t/\gamma_E$ are shown as dotted lines.

In general, we find that the estimates used for this scaling argument are reasonably

well met, as shown in Figure 3.12. One discrepancy appears to be the magnitude of the zonal flow shear that is required to prevent the saturation of low-transport states. In panel (a) of Figure 3.12, the low-transport states denoted by downward-pointing blue triangles show that the zonal shear (estimated by $[\varphi_Z]_{\text{rms}}/\ell_x^2$) increases with a/L_{T_i} but that it remains smaller than the mean flow shear rate γ_E . This indicates that even a moderate magnitude of the zonal shear can disrupt the low-transport state. A potential explanation could be the existence of a positive feedback loop between the zonal shear and the average linear growth rate $\langle\gamma\rangle_t$: even a moderate zonal shear can reduce the stabilising effect of the mean shear, which increases $\langle\gamma\rangle_t$, leading to larger fluctuation amplitudes and a stronger zonal shear, until the system saturates in the high-transport state. Another explanation could lie in the role played by mode-rational flux surfaces, which we discuss in the next Section. In any case, this moderate zonal shear explains why low-transport states in Figure 3.4 cease to exist around $\langle\gamma\rangle_t/\gamma_E \simeq 0.12$ rather than around $\langle\gamma\rangle_t/\gamma_E$ of order unity. We also point out that panel (a) of Figure 3.12 shows that the zonal flow shear is comparable to γ_E in the high-transport state, so that the two shearing rates are able to partially cancel each other out. This confirms the results that we presented in Section 3.3.2. The lower boundary of the bistability region appears to be due to the amplitudes in the high-transport states becoming too low to sustain a zonal shearing rate that is able to compete with γ_E .

3.3.5 On the role of mode-rational surfaces

In Section 2.2.7, we explained how the “twist-and-shift” boundary condition used in codes such as GS2 leads to the existence of mode-rational surfaces. On a mode-rational surface of order m , magnetic field lines in a simulation are forced to artificially close onto themselves after m poloidal turns. The layers of constant zonal flow shear that are commonly associated with mode-rational surfaces are of particular concern in the context of turbulence with a mean flow shear: including more mode-rational surfaces in the computational domain might fictitiously interfere with the externally imposed (and physically correct) mean flow shear. Since mode-rational surfaces are a numerical artefact, it must be ensured that their presence does not significantly affect the bistability reported in this work.

To test the effect of mode-rational surfaces in our simulations, we consider two ways to change the number of mode-rational surfaces that are included in a simulation. From the definition (2.23) of j_{twist} and equation (2.33), the radial distance between two mode-rational surfaces of order m can be expressed as $\Delta_{x,m} = 1/(m\hat{s}\Delta k_y)$. Thus, the largest m that is included in a simulation is determined by the maximum radial wave number K_x :

$$\frac{2\pi}{\Delta_{x,m}} \leq K_x \Rightarrow m \leq \frac{K_x}{2\pi\hat{s}\Delta k_y}. \quad (3.5)$$

By increasing K_x , one can therefore include higher-order mode-rational surfaces in a simulation. By doing so, one has to ensure that the maximum wave number in the y direction K_y is as large as $K_x/(2\pi\hat{s})$. On the other hand, the radial density of mode-rational surfaces for a given order m is $1/\Delta_{x,m} = m\hat{s}\Delta k_y$, so that the density of low-order mode-rational surfaces in a simulation can be reduced by decreasing Δk_y , as has been argued in [71].

Near the threshold temperature gradient above which no low-transport states exist, we find that the existence of the lower state can be compromised by the presence of mode-rational surfaces in our simulations. In Figure 3.13a, we show that the low-transport states near the threshold in a/L_{T_i} survive for an ever shorter time as K_x is increased (i.e., as higher-order mode-rational surfaces are included in the simulation). In Figure 3.13b, we show that these low-transport states become long-lived again when Δk_y is reduced (i.e., when the radial density of low-order mode-rational surfaces is reduced).

In order to better understand the influence of mode-rational surfaces on the saturation of sheared turbulence, we focus on the dynamics of passing electrons. Here, we refer to particles as “passing” when their parallel velocity is large enough for them to travel beyond a single $-\pi < \theta < \pi$ portion of the torus. Since $v_{\text{th},e}/v_{\text{th},i} \sim \sqrt{m_i/m_e} \gg 1$, electrons are typically much faster than ions, and they can therefore stream further along magnetic field lines before the turbulence decorrelates. Thus, passing electrons contribute significantly to the effect of mode-rational surfaces in simulations, as was already been identified in [68]. We now compare simulations obtained with fully kinetic electrons, to results obtained by imposing the passing non-zonal electron response to




Figure 3.13: Time traces of the turbulent ion heat flux, showing the fragility of the low-transport state close to the threshold above which low-transport states cease to exist. Panel (a) shows a scan in the maximum radial wave number K_x , and panel (b) shows how these results change when the spacing Δk_y for wave numbers in the y direction is halved. Here, $\rho_i K_x = 1.7$, $\rho_i \Delta k_y = 0.09$, $\rho_i K_y = 0.9$, $a/L_{T_i} = 2.087$, $\gamma_E = -0.105$ and all other parameters are set to their values presented in Appendix B.

be Boltzmann (which corresponds to considering the limit $m_e/m_i \rightarrow 0$, as was shown in [86–88]). The results are presented in Figure 3.14.

As can be seen in panel (a) of Figure 3.14, the low-transport turbulent states are found to be robust with the Boltzmann response for the passing electrons. The presence of mode-rational surfaces is clearly visible in the radial-wave-number spectrum of the zonal electrostatic potential shown in panel (b): every $2\pi/\Delta_{x,m}$ (highlighted by dots) corresponds to a sharp, single-wave-number spike in the spectrum. We note that the spikes are significantly less prominent with the modified electron response, in agreement with our expectation that passing electrons are major contributors to the adverse effect of mode-rational surfaces. Lastly, the panel (c) in Figure 3.14 shows that mode-rational surfaces can lead to substantial levels of zonal flow shear that remains constant over time. One should note that the fast oscillations of $\langle \gamma_Z \rangle_t$ for the case with kinetic electrons and the doubled K_x may be due to the fact that, in this case, K_y is slightly too small (see remark below equation 3.5). We therefore attribute the fragility of the low-transport states near the temperature-gradient threshold to the disruption of mean flow shear by the large, fictitious zonal shear associated with mode-rational surfaces.

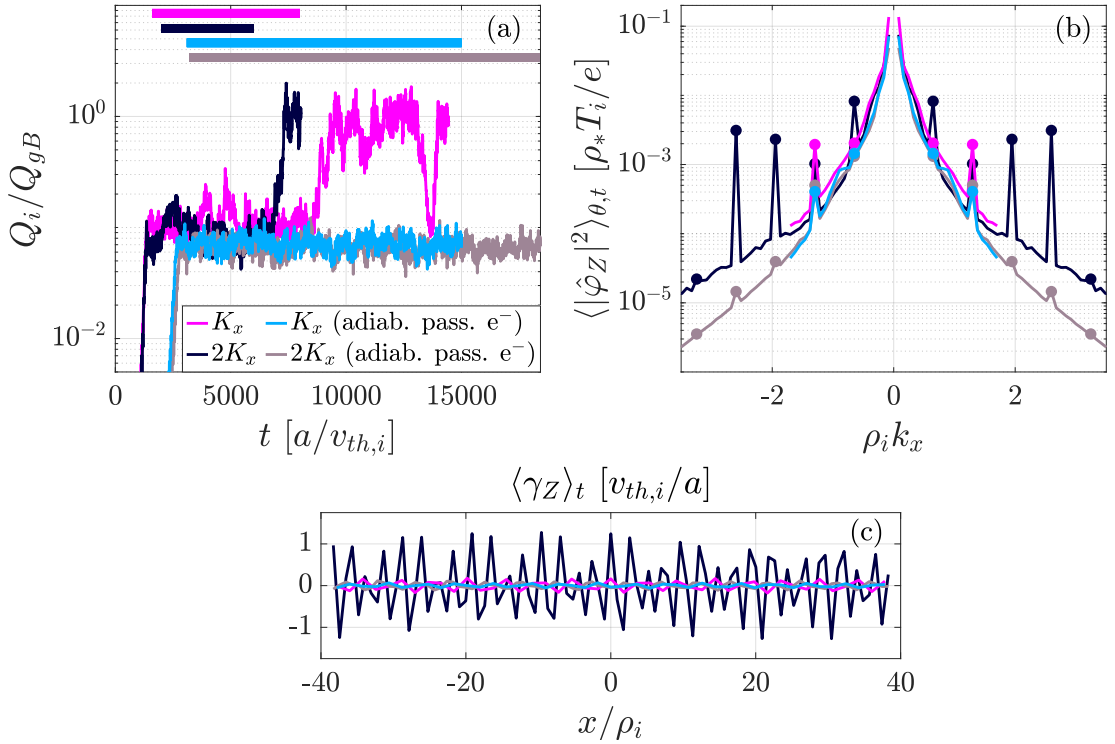


Figure 3.14: Effect of mode-rational surfaces in simulations with fully kinetic electrons, and in simulations with an imposed Boltzmann response for passing non-zonal electrons (labelled as “adiabatic passing electrons”). Panel (a) shows the time traces of the turbulent ion heat flux, where the horizontal bars indicate the windows used for time averages. Panel (b) shows the time-averaged radial-wave-number spectrum of the zonal electrostatic potential. The dots indicate $2\pi/\Delta_{x,m}$, where $\Delta_{x,m}$ is the distance between two mode-rational surfaces of order m . Panel (c) shows the corresponding time-averaged radial profiles of the zonal flow shear. Here, $\rho_i K_x = 1.7$, $\rho_i \Delta k_y = 0.09$, $\rho_i K_y = 0.9$, $a/L_{T_i} = 2.087$, $\gamma_E = -0.105$ and all other parameters are set to their values presented in Appendix B.

We point out that this sensitivity to the presence of mode-rational surfaces was only observed in the low-transport state, close to the threshold in temperature gradient above which we reported that the low-transport state ceases to exist (i.e., the transition from purple diamonds to red triangles in Figure 3.4). While this shows that mode-rational surfaces do not affect the main conclusions of our work, it suggests that the region of bistability may become broader if the effect of mode-rational surfaces could be further reduced. This is illustrated in Figure 3.15, where simulations with the full effect of mode-rational surfaces (fully kinetic electrons) produce a lower threshold in a/L_{T_i} than simulations with a reduced effect of the mode-rational surfaces (obtained

with the modified passing electron response).

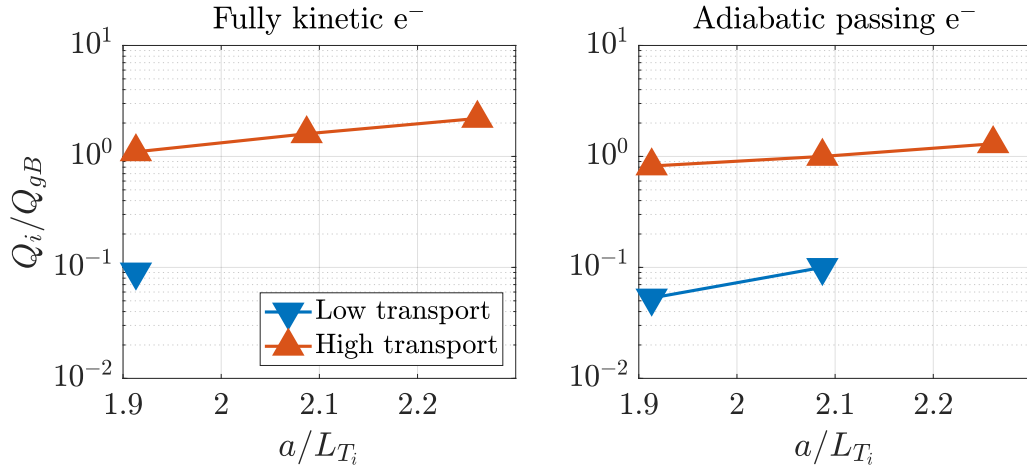


Figure 3.15: Threshold in temperature gradient above which low-transport states cease to exist, for simulations with fully kinetic electrons, and simulations with an imposed Boltzmann response for passing non-zonal electrons (labelled as “adiabatic passing electrons”). Here, apart from $\gamma_E = -0.079$ and the varying a/L_{T_i} , all other parameters are set to their values presented in Appendix B.

3.4 Consequences of bistability

In this Section, we present some remarkable consequences that follow from the existence of bistable turbulence. We then discuss the challenge that our results might pose to the numerical modelling of turbulent transport in fusion devices.

3.4.1 Bifurcations in turbulent transport

A consequence of the bistability reported in this work is the possible existence of bifurcations in turbulent transport. To illustrate this, we consider the simplified situation described in Figure 3.16.

As we have argued in previous Sections, the low-transport states cease to exist when the fluctuation amplitudes increase past a threshold value. If we now consider a plasma in which, instead of being fixed, the local temperature gradient is slowly increasing in time (for example because the externally injected power exceeds the power leaving the plasma), a discontinuous jump will be triggered from a low-transport state to the high-transport branch. This corresponds to moving from point A to point B in Figure 3.16.

The same jump can be triggered by decreasing the mean flow shear at fixed a/L_{T_i} in a low-transport state.

As shown in Figure 3.4, we also observe that the subcritical low-transport states exist closer to marginal stability in a/L_{T_i} than the subcritical high-transport states. We attribute this to the intermittent nature of the high-transport state associated with the radially propagating bands shown in Figure 3.6. Previous work with neutral fluid flows [89] and fusion plasmas [35, 44, 90] has indeed shown that the survival in time of subcritical turbulence is compromised by rare, large fluctuations. Similarly to the transition from low to high transport, we argue that subcritical high-transport states can drop to the low-transport branch if the temperature gradient decreases. This corresponds to moving from point C to D in Figure 3.16. From the bottom left of Figure 3.4, we expect that the same transition could be achieved by increasing the mean flow shear in a subcritical high-transport state, while keeping a/L_{T_i} fixed.




Figure 3.16: Illustration of bifurcations and hysteresis in the turbulent heat flux as a/L_{T_i} is modified. The different stages, labelled as A-B-C-D on the sketch in panel (a), each correspond to a stage in the sequence of simulations shown on the right. Panel (b) indicates the sequence of temperature gradients that are given as inputs to the code. Each change in the gradient requires a new simulation, where the final state of the previous simulation is used as an initial condition. Panel (c) shows the corresponding turbulent heat flux that is obtained. Note that step D has been run for a longer time (not shown in this Figure) to ensure that it does indeed correspond to a saturated low-transport state. Here, apart from $\gamma_E = -0.079$ and the varying a/L_{T_i} , all other parameters are set to their values presented in Appendix B.

3.4.2 Mean-gradient relaxation cycles

The existence of bifurcations in turbulent transport opens up the possibility for relaxation cycles of the mean gradients to develop. This hinges on two findings that we show in Figure 3.17. First, we observe a significant gap between the highest heat flux obtained in low-transport states and the lowest heat flux obtained in high-transport states. Second, we observe that the ratio of the turbulent momentum flux to heat flux, Π_i/Q_i , is almost identical in the low-transport and high-transport states. This indicates that the Prandtl number is approximately the same in the two states.

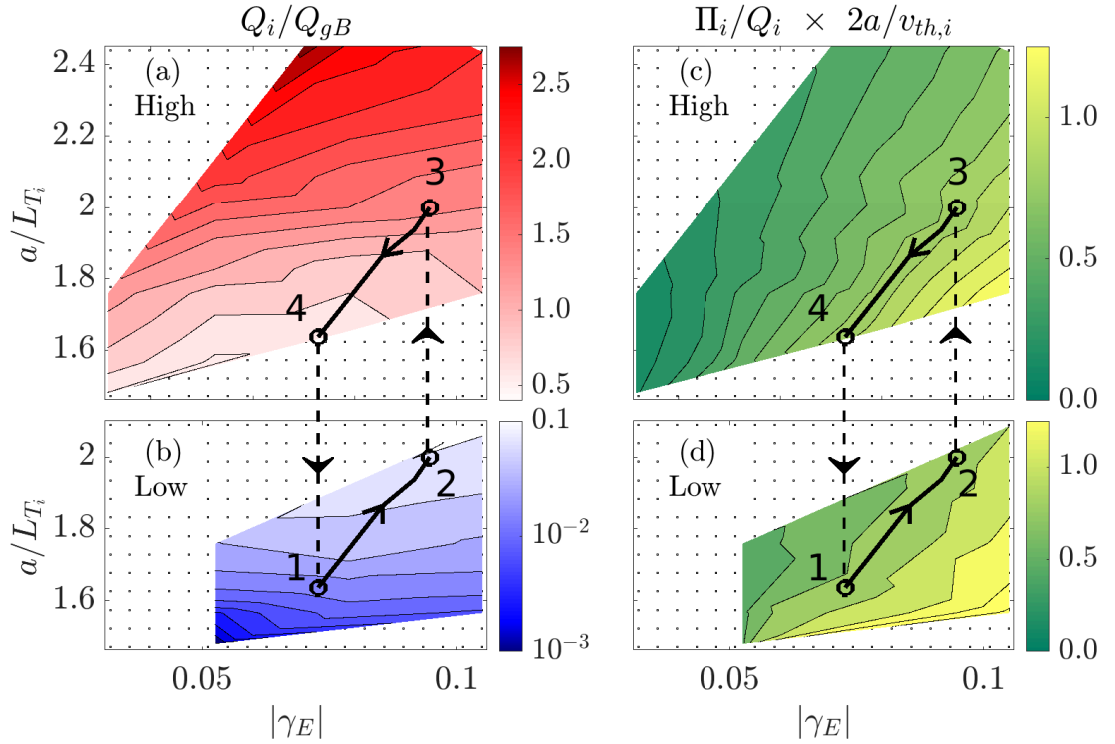


Figure 3.17: Dependence of the ion heat flux (a,b) and the momentum-to-heat-flux ratio (c,d) on the imposed flow shear and the inverse ion-temperature-gradient scale length. The top panels show results for the high-transport states, and the bottom panels show the low-transport states. We point out that there is a gap between the values of the heat flux obtained in (a) and those obtained in (b). The constant Π_i/Q_i path defined by points 1, 2, 3 and 4 gives an example of the successive stages of a gradient-relaxation cycle, when the heat injected into the plasma corresponds to a steady-state flux within the aforementioned gap. Here, apart from γ_E and a/L_{T_i} , all other parameters are set to their values presented in Appendix B.

We now consider a thought experiment with an external power P injected into the

volume bounded by a given magnetic-flux surface of area S , via a beam of neutral atoms with energy E . As is argued in [91], the turbulent heat flux exiting the magnetic-flux surface is $Q_i \sim P/S$, and $\Pi_i/Q_i \sim E^{-1/2}v_{\text{th},i}/a$. Thus, a given beam configuration corresponds to a unique $(Q_i, \Pi_i/Q_i)$. We consider an initial situation where the input power is such that (P, E) corresponds to the levels of turbulent fluxes of a low-transport state (point 1 in Figure 3.17). From this initial stationary state, we increase P by small successive increments, keeping E fixed. In response, the plasma equilibrium will evolve through a succession of low-transport stationary states with an ever larger Q_i , but with Π_i/Q_i staying constant (along the solid arrow between points 1 and 2 in Figure 3.17). Above a certain threshold, we find that there is a range in input power with no corresponding stationary solutions for the turbulent fluxes: in Figure 3.17, these are the powers too high to match the low-transport state at point 2, but too low to match the high-transport state at point 3. It is interesting to ask what will happen to a plasma where the input power falls into this gap. One possibility would be for the temperature gradient to continue increasing until the plasma transitions to a high-transport state (jumping from 2 to 3 in Figure 3.17). In this state, the outgoing heat flux is larger than what can be sustained by the external power input, and the temperature gradient starts to flatten. As a/L_{T_i} decreases (from 3 to 4), the turbulent fluctuations remain too large to allow a transition back to a low-transport state. Eventually, a/L_{T_i} becomes too small for the high-transport state to survive, and the system transitions back to the lower state (from 4 to 1). The flux is now too low compared to the power input, so the temperature gradient builds up again, and the cycle repeats itself. In this scenario, no proper steady state is reached when the input power falls within a “forbidden” gap, and the temperature gradient and mean flow shear would experience periodic relaxation cycles.

3.4.3 Implications for the numerical modelling of turbulence

Our work has important implications for research in nuclear fusion. Indeed, the design of future fusion devices relies on quantitative modelling of turbulent transport. The turbulent fluxes exiting the core of the plasma determine the equilibrium profile

gradients that can be sustained, and thus limit the fusion energy yield of a device. Furthermore, computing the outgoing fluxes is crucial to ensure that the power deposited onto the wall of fusion devices remains within limits that plasma-facing materials can tolerate.

In experiments, turbulent fluxes are set by the external injection of particles, heat and momentum into the plasma, with profile gradients evolving until a stationary state is reached. However, because of computational cost, direct numerical simulations – such as the ones presented in this work – follow the local approximation (see Section 2.2.3) and only consider a given radial location in the device, at which they solve the inverse problem: for given local equilibrium quantities (e.g., profile gradients), the simulations determine the associated turbulent fluxes. The flux-to-gradient problem and its inverse can be considered equivalent if a one-to-one correspondence exists between the turbulent transport and the equilibrium parameters.

The existence of bistable turbulence in the presence of a mean flow shear implies that this correspondence is not always one-to-one. This poses a new challenge for modelling transport: as our own experience discussed in Section 3.2 has shown, it can become very difficult to reach a reasonable agreement between the levels of transport observed in the experiment and the fluxes obtained from simulations. Furthermore, when predicting transport levels for new experiments, it is not clear how one can be sure that the stationary state obtained in simulations is the ‘true’ state that would be observed experimentally.

3.5 Conclusions

In this Chapter, we reported the existence of bistable turbulence in fusion plasmas with an externally imposed mean flow shear. Two distinct turbulent states are observed in local gyrokinetic simulations with identical equilibrium parameters. We found that the level of transport observed in the experiment was matched by the lower-transport state, in which the mean flow shear plays a crucial role in regulating the turbulence. In the high-transport state, the zonal modes were found to partially cancel the mean flow shear, allowing turbulence to saturate at higher amplitudes. This is the opposite

of the usual turbulence-suppressing effect attributed to zonal modes. We presented the main distinguishing features of the two observed states, and characterised the region of parameter space in which bistability is observed. Finally, we discussed some of the consequences following from the existence of bistability, and the challenge that this poses for transport modelling in the future.

Our findings shed new light on near-marginal turbulence in the presence of a mean flow shear, and on the role of zonal modes in regulating this turbulence. The existence of bistability, with two distinct states obtained from identical equilibrium parameters, suggests a new way to approach long-standing questions around bifurcations and gradient relaxation cycles observed in fusion devices [92–95]. This work also presents a new challenge for a research area where the prevailing assumption has been a one-to-one correspondence between plasma parameters and turbulent transport. Further work could determine how the extent of the bistable region might be modified, for example by exploring the effect of collisions on the saturation of zonal modes [34, 96]. Experiments could test the existence of bistability in fusion devices, following a scenario similar to the one described in Figure 3.17. Lastly, we note that the details of the plasma analysed here, such as toroidicity or the exact nature of the drive for turbulence, do not appear to be crucial to our understanding of bistable states: the only requirements we have identified so far are an applied flow shear and the ability of the plasma to generate zonal flows. Therefore, we expect that similar bistability could be observed in a variety of systems.

Chapter 4

Summary and conclusions

Following the δf gyrokinetic approach, a novel algorithm was developed to simulate mean rotational flow shear in the core of magnetic-confinement fusion plasmas. Using this algorithm, it was determined that two distinct, stationary turbulent states can be obtained with identical mean plasma parameters in the presence of a mean flow shear – a phenomenon known as bistability.

In Chapter 2, a novel algorithm was developed to include a mean rotational flow shear in linearly implicit gyrokinetic codes. The algorithm was implemented and validated in the **GS2** code. This new, continuous-in-time algorithm for flow shear was shown to have several advantages compared to the discrete-in-time approach [49] that has been widely used until now. First, the new algorithm allows fast, linearly implicit codes to reproduce linear instabilities in the presence of flow shear without any discrete jumps in the simulations. An important consequence is that the instabilities are reproduced correctly with much sparser radial-wave-number resolution than with the discrete-in-time approach. The same observation was made for nonlinear simulations: using the new algorithm, the turbulent fluxes remain close to their converged value with sparser resolution. The second advantage of the new algorithm is that it corrects the implementation of nonlinear interactions in presence of flow shear, similarly to the work in [51]. In the present thesis, it was shown that this correction can lead to a different value of the converged fluxes compared to the discrete-in-time approach.

As we discussed in Section 2.4.3, the exact reason for this difference in the turbulent fluxes remains an open question. One possibility for future work would be to determine if the larger zonal modes produced with the discrete-in-time algorithm may

be responsible for this discrepancy. The associated sheared zonal flows might oppose the mean flow shear and produce slightly higher fluxes, similarly to the mechanism sustaining high-transport states that we presented in Chapter 3. A first step could be to compare the continuous and discrete-in-time algorithms for simulations where the zonal modes are artificially zeroed out.

So far, the new algorithm was not applied to collisions in GS2. Instead, collisions are still treated using the discrete-in-time approach for flow shear, which may become a limitation in cases where flow shear and collisional effects are both important. However, it should be possible to extend the algorithm to also include collisions. Another limitation of the algorithm is that its numerical scheme is only stable for sufficiently small grid spacing in the radial wave number. However, it was found that the grid-spacing required to make the scheme unstable is much larger than any reasonable choice for simulations of turbulent transport. Nevertheless, it might be useful, albeit difficult, to develop a mechanism ensuring that the scheme is always stable.

Using the new algorithm, it was found in Chapter 3 that bistable turbulence can be obtained in the presence of a mean rotational flow shear. Two distinct stationary states feature significantly different levels of turbulent transport, with the low-transport state matching the levels of fluxes observed experimentally at the JET tokamak. The occurrences of the two states appear to be regulated by the competition between the mean, externally imposed flow shear, and spontaneously generated sheared zonal flows. This work found that the zonal flows can oppose the mean flow shear and help sustain the high-transport turbulent state. This result stands in contrast to the usual expectation that zonal flows always act to suppress turbulent transport. The existence of bistability poses a new challenge for modelling turbulent transport in tokamaks. It shows that the prevailing paradigm of a one-to-one correspondence between the mean plasma parameters and the level of turbulent transport is not always correct. In consequence, it can be challenging to match the level of transport in simulations with the ones observed in experiments. When designing future experiments, it also becomes unclear how to ensure that the transport predicted by simulations will correspond to the transport produced in the actual experiment.

It is essential to point out that the bistability discussed in this work is obtained from local simulations, in which the equilibrium parameters are assumed to be identical across the radial extent of the simulation domain. In particular, the radial dependence of the profile gradients driving the turbulence is not taken into account. In this local approximation, the existence of bistability implies that, for a given equilibrium profile, two radial locations that lie arbitrarily close to each other (i.e., with slightly different values for their respective local profile gradients) can produce two distinct states with vastly different turbulent fluxes (e.g., the two neighbouring radial locations could correspond to the points 2 and 3 in Figure 3.17 that we discussed above).

Presumably, this would not occur in an actual experiment: the turbulent eddies at those two radial locations would partly overlap and mix with each other. Modelling this would require simulations that include so-called “global” effects, such as the radial dependence of profile gradients. Several gyrokinetic codes are capable of including global effects, and could hence be used to address these questions. Similar considerations have previously been discussed for simulations near marginal stability in the absence of flow shear [97]. In global simulations with a mean flow shear, radii located between two neighbouring lower-transport and higher-transport states would most likely feature an intermediate level of turbulent transport. Future work will be needed to clarify how such a smooth transition might modify the discontinuous, hysteretic relationship between the turbulent fluxes and the temperature gradient that we observe in local simulations (see for example Figure 3.3, where Q_i/Q_{gB} is plotted against a/L_{T_i}).

It would be particularly interesting to determine how the profiles of temperature and mean toroidal flow might look in plasmas where certain radial locations are in the high-transport state while other radii are in the low-transport state. In order to simulate this, one could consider following a similar approach to the one developed in the multiscale gyrokinetic code TRINITY [98]: the turbulent fluxes are computed given initial equilibrium profiles; the fluxes are then used to solve the transport equations of particle density, momentum and energy on the slower equilibrium time scale; the updated equilibrium profiles are used to compute the new turbulent fluxes; and the procedure is repeated until converged profiles are obtained. In order to avoid the

radial discontinuity in fluxes that we discussed above, one might need to include global effects when computing the turbulent fluxes at radii located between neighbouring low-transport and high-transport states. This last point may not be needed if the length scale associated with the mixing of turbulence is much shorter than the radial spacing used to solve the transport equations.

Another possibility for future work could be to further test the influence of mode-rational surfaces on bistability. This could be based on simulations that do not suffer from the numerical limitations of the twist-and-shift parallel boundary condition, e.g. “flux-tube train” simulations where several flux tubes are connected along the parallel direction [99].

In this work, we limited our analysis to an experimentally relevant tokamak plasma in which turbulent transport is driven by the ion temperature gradient instability (ITG). An important point for future research would be to simplify as much as possible the system that is required to observe bistable turbulence, so as to better isolate the key physical parameters. A related question would be to explore whether the existence of bistability prevails when other instabilities such as electron temperature gradient (ETG) or trapped electron modes (TEM) dominate. Given the importance of sheared zonal flows for bistable turbulence, one would for instance expect that bistability might not be observed for certain TEM regimes where the saturation of turbulence does not appear to be regulated by zonal modes [100, 101].

Lastly, it might be of interest to develop a heuristic model for the suppression of turbulent transport by flow shear that accounts for the existence of bistability. Such a model would for instance include an estimate of the nonlinear correlation time scale based on linear results, similarly to the results presented in Section 3.3.3.

We now consider the broader goal of fusion to maximise the energy yield of tokamaks and to mitigate the heat load on the plasma-facing components. Following the results presented in this thesis, there are several avenues that could be explored from a theoretical and computational perspective. For example, it would seem advantageous to aim for tokamaks to reach high core temperatures in the low-transport turbulent state, without ever transitioning to the high-transport state.

A path towards such a scenario might be to maximise the width of the bistable region in parameter space, allowing the low-transport state to survive at higher temperature gradients. A necessary condition for the existence of low-transport states was found in Section 3.3.4 to be $T_F \langle \gamma \rangle_t < 1$. Given that $T_F = 2\pi\hat{s}/\gamma_E$, it would for instance be interesting to determine what happens to bistability when the magnetic shear \hat{s} is varied. For example, a question could be whether a lower magnetic shear (as proposed for the so-called “hybrid” scenario for ITER [102]) can lead to a wider region of bistability.

Another possibility would be to increase the gap in the turbulent transport level between low-transport states and high-transport states. One way to achieve this might be to vary the average linear growth rate independently of the maximum instantaneous linear growth rate, for example by modifying the magnetic geometry and making field lines trace a longer distance on the inboard side of the torus, where the ITG instability is suppressed.

Experimentally, the existence of bistability could be tested in fusion devices by following a procedure similar to the thought experiment presented in Figure 3.17. Since this work found that zonal flows can sometimes be detrimental to the plasma confinement, another possibility for experimental work might be to explore whether the confinement in tokamaks may be enhanced by somehow disrupting the zonal modes.

Finally we note that this thesis has focused on one particular experiment. A further opportunity for future work would be to determine if bistability occurs for various experimental conditions. This last point may also suggest a new way to explore long-standing questions around relaxation cycles and bifurcations, see e.g. [92–95], which are commonly observed in tokamaks but are not well understood yet.

Appendix A

Using JETPEAK equilibria with the GS2 code

Codes and experimental datasets can often vary in their choices of coordinates and normalisations. In this Appendix, we summarise some of the aspects to consider when making combined use of the local, δf gyrokinetic code GS2 and the experimental database JETPEAK [79].

A.1 Sign conventions

The JETPEAK database provides access to a substantial set of well-diagnosed discharges conducted at the JET tokamak. It includes output from EFIT [103] and TRANSP [104] for the magnetic geometry, and from ASCOT [105] for particle, momentum and heat deposition profiles. From here on, we use the subscript ‘J’ to refer to quantities defined in JETPEAK.

The toroidal magnetic geometry in tokamaks can be described by the right-handed cylindrical coordinate system (R, ϕ, Z) , with R the major radius, ϕ the toroidal angle and Z the vertical direction. Another set of coordinates is useful when describing nested toroidal magnetic flux surfaces, namely (ψ, θ, ϕ) with ψ the poloidal magnetic flux, θ the poloidal angle and ϕ the toroidal angle. Using these coordinates, the magnetic field can be expressed as

$$\mathbf{B} = I(\psi)\nabla\phi + \nabla\phi \times \nabla\psi. \quad (\text{A.1})$$

In tokamaks, the poloidal component of the magnetic field is generated by driving an electric current \mathbf{I}_p in the plasma along the toroidal direction. In GS2, the convention

is that ϕ grows in the direction of the plasma current, ie:

$$\mathbf{I}_p \cdot \nabla\phi > 0, \quad (\text{A.2})$$

and that the coordinate system (ψ, θ, ϕ) is right-handed. From this we get that

$$\mathbf{B} \cdot \nabla\theta > 0. \quad (\text{A.3})$$

A length denoted by R_{geo} in GS2 and the normalising magnetic field B_r are defined by the relation

$$I(\psi) = B_r R_{\text{geo}}. \quad (\text{A.4})$$

Typically, the user fixes R_{geo} arbitrarily and this then fixes B_r . Here, we additionally choose the convention

$$B_r > 0, \quad (\text{A.5})$$

which then implies that

$$\text{sgn}(R_{\text{geo}}) = \text{sgn}(I) = \text{sgn}(\mathbf{B} \cdot \nabla\phi). \quad (\text{A.6})$$

In JETPEAK, the toroidal direction $\nabla\phi_J$ is defined to be clockwise when looking down from above the tokamak. The sign of R_{geo} can then be computed in the following way:

$$\begin{aligned} \text{sgn}(R_{\text{geo}}) &= \text{sgn}(\mathbf{B} \cdot \nabla\phi) \\ &= \text{sgn}(\mathbf{B} \cdot \mathbf{I}_p) \\ &= \text{sgn}(\mathbf{B} \cdot \nabla\phi_J) \text{sgn}(\mathbf{I}_p \cdot \nabla\phi_J) \\ &= \text{sgn}(\text{BASIC.BT}) \text{sgn}(\text{BASIC.IP}), \end{aligned} \quad (\text{A.7})$$

where BASIC.BT and BASIC.IP are the JETPEAK variables corresponding to $\mathbf{B} \cdot \nabla\phi_J$ and $\mathbf{I}_p \cdot \nabla\phi_J$, respectively.

Similarly, we have to adjust the signs carried by quantities related to the safety factor or to the toroidal angular frequency. The flux-surface averaged safety factor q is defined as

$$q(\psi) = \frac{1}{2\pi} \int_0^{2\pi} d\theta \left| \frac{\mathbf{B} \cdot \nabla\phi}{\mathbf{B} \cdot \nabla\theta} \right|_{\psi}, \quad (\text{A.8})$$

and therefore:

$$\text{sgn}(q) = \text{sgn}(R_{\text{geo}}). \quad (\text{A.9})$$

The toroidal angular frequency associated with the background plasma flow is defined as

$$\Omega_\phi = \mathbf{u} \cdot \nabla\phi, \quad (\text{A.10})$$

to lowest order in ρ_* . Following our conventions, its sign is hence given by:

$$\begin{aligned} \text{sgn}(\Omega_\phi) &= \text{sgn}(\mathbf{u} \cdot \mathbf{I}_p) \\ &= \text{sgn}(\mathbf{u} \cdot \nabla\phi_J) \text{sgn}(\mathbf{I}_p \cdot \nabla\phi_J) \\ &= \text{sgn}(\text{ION.ANGF}) \text{sgn}(\text{BASIC.IP}), \end{aligned} \quad (\text{A.11})$$

where `ION.ANGF` is the `JETPEAK` variable corresponding to $\mathbf{u} \cdot \nabla\phi_J$.

A.2 Normalisation for the magnetic field

The reference magnetic field strength B_r used in `GS2` normalisations is not completely trivial to compute from the `JETPEAK` dataset. As was detailed in Section [A.1](#), the reference magnetic field is defined by:

$$B_r = \frac{I(\psi)}{R_{\text{geo}}}, \quad (\text{A.12})$$

where the user is free to choose $|R_{\text{geo}}|$ and where the sign of R_{geo} is set by [\(A.7\)](#). In this work, we set $|R_{\text{geo}}|$ to be the major radius of the flux surface's middle point at the height of the magnetic axis (i.e. equivalent to the `GS2` variable `Rmaj` before normalisation). The function $I(\psi)$ is computed in terms of the safety factor as

$$I(\psi) = 2\pi q \left/ \int_0^{2\pi} d\theta \right|_{\psi} \frac{J_r}{R^2 \frac{d\psi}{dr} \Big|_{\theta}}, \quad (\text{A.13})$$

where r is the minor radius and

$$\begin{aligned} J_r &= [\nabla\theta \cdot (\nabla\phi \times \nabla r)]^{-1} \\ &= rR \end{aligned} \quad (\text{A.14})$$

is the Jacobian of the transformation to the (r, θ, ϕ) coordinate system. In JETPEAK, the magnetic geometry is provided by the output of the TRANSP code. The magnetic flux ψ is specified as a function of space in two different discretised grids: one follows the shape of flux surfaces but is coarse, while the other is a finer, evenly-spaced grid in (R, Z) . The integral in (A.13) is computed on the flux-surface following grid, while the radial derivative of ψ is evaluated at the nearest grid point from the finer (R, Z) grid via:

$$\left. \frac{d\psi}{dr} \right|_{\theta} = \left. \frac{\partial\psi}{\partial R} \right|_Z \left. \frac{\partial R}{\partial r} \right|_{\theta} + \left. \frac{\partial\psi}{\partial Z} \right|_R \left. \frac{\partial Z}{\partial r} \right|_{\theta}. \quad (\text{A.15})$$

Appendix B

JET discharge

Many thanks to Henri Weisen and Paula Sirén for providing the experimental data discussed in this Appendix. In this work, we focus on the plasma discharge #68448 carried out at the JET tokamak. This discharge is of interest because part of its heating stems from injecting neutral beams into the plasma, which drives a sheared mean toroidal flow. The discharge is documented in the JETPEAK database [79]. Figure B.1 shows the radial profiles of the temperature, the density, and the angular frequency of the mean plasma flow. The associated deposition profiles of torque and energy calculated with the ASCOT code [105] are shown in Figure B.2. The main plasma parameters at $r_\psi/a = 0.51$ are presented in Table B.1. The main ion species is deuterium. We note that non-hydrogenic impurities are neglected throughout this work, so that the main ion density is effectively set to be equal to the electron density. Electromagnetic effects are neglected so that the thermal-to-magnetic-pressure ratio β is approximated to zero.




Figure B.1: Radial profiles of the JET discharge #68448. Panel (a) shows the temperature, panel (b) the particle density and panel (c) the mean rotation frequency.




Figure B.2: Radial profiles of the energy deposition (a) and of the torque deposition (b) in the discharge #68448. $S_{\epsilon,i}$ is the external source of energy, and $S_{\Delta,i}$ is the energy transfer with electrons (see equation (E.8)). The deposition is computed based on the mean plasma parameter profiles, using the ASCOT code. Section E.2 explains how the turbulent fluxes in the experiment can be computed from the profiles (a) and (b).

I_p	2.6MA	Plasma current
B_T	2.9T	Vacuum toroidal field on axis
P_{NBI}	17MW	Neutral beam heating power
R_ψ	$3.06a$	$[\max(R) + \min(R)]/2$ for this flux surface
r_ψ	$0.508a$	$[\max(R) - \min(R)]/2$ for this flux surface
$ q_0 $	1.43	flux-surface averaged safety factor
\hat{s}	0.574	flux-surface averaged magnetic shear
κ	1.36	Miller elongation [106]
$d\kappa/dr_\psi$	$0.146/a$	elongation gradient
δ	0.0571	arcsin of Miller triangularity [106]
$d\delta/dr_\psi$	$0.129/a$	gradient of GS2 triangularity
γ_E	$-0.0553v_{\text{th},i}/a$	background flow shear rate
Ω_ϕ	$-0.08v_{\text{th},i}/a$	background flow angular frequency
n_i/n_e	1.0	ion to electron density ratio
$1/L_{n_i}$	$0.602/a$	inverse ion density gradient length
$1/L_{n_e}$	$0.602/a$	inverse electron density gradient length
T_e/T_i	0.855	electron to ion temperature ratio
$1/L_{T_i}$	$1.7392/a$	inverse ion temperature gradient length
$1/L_{T_e}$	$1.551/a$	inverse electron temperature gradient length
ν_{ii}	$2.6 \times 10^{-4}v_{\text{th},i}/a$	ion collisionality
ν_{ee}	$0.02v_{\text{th},i}/a$	electron collisionality
β	0.0125	$2e\mu_0n_iT_i/B_r^2$

Table B.1: Parameters for the JET discharge #68448 at $r_\psi/a = 0.51$. The gradient length of any quantity ξ is given by $L_\xi = 1/[d \log(\xi)/dr_\psi]$. μ_0 denotes the vacuum permeability.

Appendix C

Discretisation in the GS2 code

In this Section, we present how the system of gyrokinetic and quasineutrality equations is discretised in **GS2**. The discretisation in time and poloidal angle θ can be modified via the parameters r_t and r_θ respectively. The user can set r_t between zero (fully-implicit linear scheme) and one (fully-explicit scheme). Similarly r_θ can be set between zero (centred in θ) and one (upwinding in θ). The discrete indices used in this Appendix are defined in Table C.1. With these definitions, the time derivative of any fluctuating quantity Φ is approximated by

$$\frac{\partial\Phi}{\partial t} \approx \frac{1 + \sigma r_\theta}{2} \frac{\Phi_{n+1, j_\theta+1} - \Phi_{n, j_\theta+1}}{\Delta t} + \frac{1 - \sigma r_\theta}{2} \frac{\Phi_{n+1, j_\theta} - \Phi_{n, j_\theta}}{\Delta t}, \quad (\text{C.1})$$

where σ is the sign of $\mathbf{v} \cdot \hat{\mathbf{b}}$. Similarly, derivatives in θ are approximated by

$$\frac{\partial\Phi}{\partial\theta} \approx (1 - r_t) \frac{\Phi_{n+1, j_\theta+1} - \Phi_{n+1, j_\theta}}{\Delta\theta} + r_t \frac{\Phi_{n, j_\theta+1} - \Phi_{n, j_\theta}}{\Delta\theta} \quad (\text{C.2})$$

and all other terms are approximated by

$$\begin{aligned} \Phi \approx \frac{1 - \sigma r_\theta}{2} \left(r_t \Phi_{n, j_\theta} + (1 - r_t) \Phi_{n+1, j_\theta} \right) \\ + \frac{1 + \sigma r_\theta}{2} \left(r_t \Phi_{n, j_\theta+1} + (1 - r_t) \Phi_{n+1, j_\theta+1} \right). \end{aligned} \quad (\text{C.3})$$

We now use bold-font vector notation for the potential and distribution function, to denote discrete arrays whose indices span every (θ, k_x) of a single twist-and-shift chain. Matrices with two such indices are denoted using bold upper case letters. With this notation, discretising the gyrokinetic equation (2.24) according to (C.1), (C.2) and (C.3) yields:

$$\mathbf{A}_{n+1}^* \hat{\mathbf{g}}_{n+1} + \mathbf{B}_n^* \hat{\mathbf{g}}_n = \mathbf{C}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} + \mathbf{D}_n^* \hat{\boldsymbol{\varphi}}_n, \quad (\text{C.4})$$

Time	$0 \leq n \leq N_t$	$t_n = n\Delta t$
Poloidal angle	$-(N_\theta - 1)/2 \leq j_\theta \leq (N_\theta - 1)/2$	$\theta_{j_\theta} = j_\theta \Delta\theta$
Radial wave number	$-(N_x - 1)/2 \leq j_x \leq (N_x - 1)/2$	$\bar{k}_{x,0} = j_x \Delta k_x$
Wave number in y^*	$0 \leq j_y < N_y$	$k_{y,j_y} = j_y \Delta k_y$
Twist-and-shift chain	$0 \leq j_c < 2\pi \hat{s} k_y / (\Delta k_x)$	chain $\leftrightarrow (j_y, j_c)$
Along a chain	$j = j_\theta + N_\theta(j_x - j_c)/(2\pi \hat{s} k_y)$	all (θ, k_x) in a chain

Table C.1: Indices used in discrete notations.

where we have defined the matrices

$$\mathbf{A}_{n+1,j,j'}^* = \frac{1}{\Delta t} \left(\frac{1 + \sigma r_\theta}{2} \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \delta_{j,j'} \right) \quad (\text{C.5})$$

$$+ (1 - r_t) w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (\delta_{j+1,j'} - \delta_{j,j'}) \quad (\text{C.6})$$

$$+ i(1 - r_t) \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta+1} \cdot \mathbf{k}_{n+1,j+1}^* \delta_{j+1,j'} \right) \quad (\text{C.7})$$

$$+ \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n+1,j}^* \delta_{j,j'} \Big), \quad (\text{C.8})$$

$$\mathbf{B}_{n,j,j'}^* = - \frac{1}{\Delta t} \left(\frac{1 + \sigma r_\theta}{2} \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \delta_{j,j'} \right) \quad (\text{C.9})$$

$$+ r_t w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (\delta_{j+1,j'} - \delta_{j,j'}) \quad (\text{C.10})$$

$$+ i r_t \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta+1} \cdot \mathbf{k}_{n,j+1}^* \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n,j}^* \delta_{j,j'} \right), \quad (\text{C.11})$$

$$\mathbf{C}_{n+1,j,j'}^* = - \frac{e Z_s F_{0,s}}{T_s} w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (1 - r_t) (J_{0s,n+1,j+1}^* \delta_{j+1,j'} - J_{0s,n+1,j}^* \delta_{j,j'}) \quad (\text{C.12})$$

$$- i \frac{e Z_s F_{0,s}}{T_s} (1 - r_t) \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta+1} \cdot \mathbf{k}_{n+1,j+1}^* J_{0s}^* n + 1, j + 1 \delta_{j+1,j'} \right) \quad (\text{C.13})$$

$$+ \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n+1,j}^* J_{0s}^* n + 1, j \delta_{j,j'} \Big) \quad (\text{C.14})$$

$$- i k_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_\psi} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) (1 - r_t) \quad (\text{C.15})$$

$$\cdot \left(\frac{1 + \sigma r_\theta}{2} J_{0s,n+1,j+1}^* \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} J_{0s,n+1,j}^* \delta_{j,j'} \right), \quad (\text{C.16})$$

$$\mathbf{D}_{n,j,j'}^* = -\frac{eZ_s F_{0,s}}{T_s} w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} r_t \left(J_{0s,n,j+1}^* \delta_{j+1,j'} - J_{0s,n,j}^* \delta_{j,j'} \right) \quad (\text{C.17})$$

$$-i \frac{eZ_s F_{0,s}}{T_s} r_t \left(\frac{1 + \sigma r_{\theta}}{2} \mathbf{V}_{d,s,j_{\theta}+1} \cdot \mathbf{k}_{n,j+1}^* J_{0s,n,j+1}^* \delta_{j+1,j'} \right. \quad (\text{C.18})$$

$$\left. + \frac{1 - \sigma r_{\theta}}{2} \mathbf{V}_{d,s,j_{\theta}} \cdot \mathbf{k}_{n,j}^* J_{0s,n,j}^* \delta_{j,j'} \right) \quad (\text{C.19})$$

$$-ik_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_{\psi}} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) r_t \left(\frac{1 + \sigma r_{\theta}}{2} J_{0s,n,j+1}^* \delta_{j+1,j'} \right. \quad (\text{C.20})$$

$$\left. + \frac{1 - \sigma r_{\theta}}{2} J_{0s,n,j}^* \delta_{j,j'} \right). \quad (\text{C.21})$$

Finally, the quasineutrality equation (2.31) is evaluated at time step $n+1$ and written in the discrete form:

$$\mathbf{E}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} = \mathbf{W}_{n+1}^* \hat{\boldsymbol{g}}_{n+1}, \quad (\text{C.22})$$

where we have defined the matrix

$$\mathbf{E}_{n+1,j,j'}^* = \sum_s \frac{e^2 Z_s n_s}{T_s} (1 - \Gamma_{s,n+1,j}^*) \delta_{j,j'}. \quad (\text{C.23})$$

The velocity-space integral operator \mathbf{W} is defined such that, for any function Φ ,

$$\mathbf{W}_{n+1,j,j'}^* \Phi = \sum_s Z_s \int d^3 w J_{0s,n+1,j}^* \delta_{j,j'} \Phi. \quad (\text{C.24})$$

Appendix D

Source code and input parameters

In this Appendix, we indicate how to access the various source codes used for this thesis. We then list the parameter scans that were carried out to ensure that the conclusions of this work are not affected by numerical resolutions. Finally, we provide a typical example of a GS2 input file that was used for this thesis.

D.1 Source code

The version of the GS2 code used for this work is available at https://bitbucket.org/gyrokinetics/gs2/branch/ndc_branch, with the newest commit at the time of writing being 0abdcda. The associated version of “Makefiles” necessary for compilation is available at https://bitbucket.org/gyrokinetics/makefiles/branch/ndc_branch under the commit ba24979, and the additional “utils” files required to run the code are available at https://bitbucket.org/gyrokinetics/utils/branch/ndc_branch under the commit 8e41f9a.

The scripts used to generate GS2 input files based on experimental data stored in the JETPEAK database [79] are available at https://github.com/nicolaschristen/jetpeak_v_gs2, with the newest commit at the time of writing being 16c0198. The scripts also allow comparisons between fluxes obtained from GS2 simulations, and fluxes observed in the experiment.

D.2 GS2 input parameters

In order to ensure that the results presented in this thesis are not affected by insufficient numerical resolution, scans were carried out for various numerical parameters (see Table D.1). In addition, it was checked that the two distinct saturated states are also obtained with the widely used, discrete-in-time algorithm for flow shear that we describe in Section 2.3.4.

Parameter	Values tested	Type of scan	Value used	Units
Δk_x	0.04 – 0.08	nonlinear	0.08	$1/\rho_i$
K_x	1.7 – 15.2	nonlinear	3.8	$1/\rho_i$
Δk_y	0.045 – 0.18	nonlinear	0.09	$1/\rho_i$
K_y	0.99 – 1.98	nonlinear	1.98	$1/\rho_i$
ntheta	16 – 128	linear	32	-
negrid	6 – 48	linear	16	-
ngauss	3 – 20	linear	5	-
vcut	2.5 – 4.5	linear	2.5	-
Δt	0.025 – 0.1	linear	linearly: 0.1	$a/v_{th,i}$

Table D.1: Range of numerical parameters that were tested with GS2. Here, **ntheta** is roughly the number of grid points in θ , **negrid** is the number of energy grid points, $4 \times$ **ngauss** is the number of untrapped pitch angles, and **vcut** is the number of standard deviations from the Maxwellian distribution of velocities above which the fluctuating distribution function is set to zero

We now provide a typical GS2 input file used for this thesis, based on the experimental data of the JET discharge #68448. This particular input file corresponds to steps A and C in the panel (c) of Figure 3.16. To obtain the low-transport state (step A in Figure 3.16), the simulation should be started with the mean flow shear turned on, i.e., with **g_exb** set to the value given here. To obtain the high-transport state (step C in Figure 3.16), the simulation can be started with no flow shear until a saturated state is reached, i.e., with **g_exb** set to zero. This simulation should then be restarted (instructions are given in the comments below) with the flow shear turned on.

```
!! ----- !!  
!! GS2 input file based on the JET discharge #68448 !!  
!! ----- !!
```

```
&species_knobs  
  nspec = 2  
/  
  
&species_parameters_1  
  z = 1.0  
  mass = 1.0  
  dens = 1  
  temp = 1  
  tprim = 1.91312 ! corresponds to a/LTi  
  fprim = 0.60228  
  uprim = 0.0  
  vnewk = 0.00026042  
  type = 'ion'  
/  
  
&species_parameters_2  
  z = -1.0  
  mass = 2.7e-4  
  dens = 1  
  temp = 0.85478  
  tprim = 1.5509  
  fprim = 0.60228  
  uprim = 0.0  
  vnewk = 0.019972  
  type = 'electron'  
/  
  
&dist_fn_species_knobs_1  
  fexpr = 0.45  
  bakdif = 0.05  
/  
  
&dist_fn_species_knobs_2  
  fexpr = 0.45  
  bakdif = 0.05  
/  
  
&collisions_knobs  
  collision_model='default'  
/
```

```

&parameters
  beta = 0.012575
  zeff = 1
/

&theta_grid_parameters
  ntheta = 32
  nperiod = 1
  rhoc = 0.50825
  shat = 0.57383
  qinp = -1.4253
  Rmaj = 3.0642
  R_geo = -3.0642
  shift = -0.10502
  akappa = 1.3594
  akappri = 0.1458
  tri = 0.057107
  tripri = 0.12938
/

&dist_fn_knobs
  adiabatic_option = "iphi00=2"
  gridfac = 1.0
  boundary_option = "linked"
  mach = -0.079881
  g_exb = -0.0788355 ! corresponds to gamma_E
/

&theta_grid_knobs
  equilibrium_option = 'eik'
/

&theta_grid_eik_knobs
  itor = 1
  iflux = 0
  irho = 2
  ppl_eq = F
  gen_eq = F
  efit_eq = F
  local_eq = T
  eqfile = 'dskeq.cdf'
  equal_arc = T
  bishop = 4
  s_hat_input = 0.57383
  beta_prime_input = -0.052589
  delrho = 1.e-3
  isym = 0

```

```

writelots = F
/

&kt_grids_knobs
  grid_option = 'box'
/

&kt_grids_box_parameters
  ! total number of ky's: naky = (ny-1)//3 + 1
  ! total number of kx's: nakx = 2*(nx-1)//3 + 1
  ny = 72                ! i.e. kymax = 2.0
  nx = 144               ! i.e. >= 1 twist-and-shift links for kymax
  y0 = 11.111            ! i.e. dky = 1/y0 = 0.09
  jtwist = 4             ! i.e. dkx = 2*pi*shat*dky/jtwist = 0.0811
  mixed_flowshear = .true. ! turns on continuous-in-time algo for flow shear
/

&fields_knobs
  field_option = 'implicit'
  force_maxwell_reinit = .false.
/

&le_grids_knobs
  ngauss = 5
  negrid = 16
  vcut = 2.5
/

&init_g_knobs
  chop_side = F
  phiinit = 1.e-3
  ! location to save/read restart file (overwritten when restarting)
  restart_file = "nc/run.nc"
  ginit_option = "noise"    ! FOR RESTARTS : set to "many"
  clean_init = .true.
  read_many = .true.
/

&knobs
  fphi = 1.0
  fapar = 0.0
  faperp = 0.0
  delt = 0.025
  nstep = 200000
  avail_cpu_time = 86400    ! 24hrs, adapt to available resources
  delt_option = "default"  ! FOR RESTARTS : set to "check_restart"
/

```

```

&nonlinear_terms_knobs
  nonlinear_mode = 'on'
  cfl = 0.25
/

&reinit_knobs
  delt_adj = 2.0
  delt_minimum = 1.e-4
  delt_cushion = 10000
/

&layouts_knobs
  ! consider layout = 'lxyes' for better performance
  layout = 'xyles'
  local_field_solve = F
/

&hyper_knobs
  hyper_option = 'visc_only'
  const_amp = .false.
  isotropic_shear = .false.
  D_hypervisc = 0.05
/

&gs2_diagnostics_knobs
  write_fluxes = .true.
  print_flux_line = T
  write_nl_flux = T
  print_line = F
  write_line = F
  write_omega = F
  write_final_fields = T
  write_g = F
  write_verr = T
  nwrite = 50
  navg = 50
  nsave = 3000
  omegainst = 500.0
  save_for_restart = .true.
  omegatol = -1.0e-3
  save_many = .true.
/

!! ----- !!
!                               !
!               End of input file               !
!! ----- !!

```

Appendix E

Reconstruction method for toroidal rotation profiles

In this Appendix, we briefly present a method that relies on local gyrokinetic simulations to reconstruct and predict mean rotation profiles in experiments where neutral beam injection (NBI) generates a weak mean toroidal flow. Initially, an objective of our work was to develop this method and apply it to various experimental scenarios, in order to probe how the NBI affects the rotation profile. As we explained in Section 3.2, the challenges encountered when comparing simulations with experiments led us to the results presented in Chapter 3, and our focus shifted away from the reconstruction method. In Section E.1, we show the transport equations for density, momentum and energy and define the source terms associated with the injection of particles, angular momentum and heat into the plasma. In Section E.2, we explain how the turbulent fluxes present in an experiment are obtained from the radial profiles of equilibrium quantities. We then present the reconstruction method for rotation profiles, together with an initial result for a JET discharge in E.3. Finally, we discuss possible applications of the method to predict rotation profiles for slowly rotating plasmas in Section E.4.

E.1 Transport equations

Within the asymptotic expansion detailed in Section 2.2.1, we consider the kinetic equation up to $O(\rho_{*s}^3 \Omega_s f_s)$, averaged over the turbulent time scale and over the turbulent length scales across the magnetic field. Following the work in [28, 29, 107], we

then consider moments in velocity space of the equation to describe the time evolution of density, momentum and energy.

E.1.1 Particle transport

The zeroth-order moment describes the time evolution of the background density for species s :

$$\frac{\partial n_s}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial \psi} (V' \Gamma_s) = S_{n,s} \quad (\text{E.1})$$

where n_s is the species density, ψ is the poloidal magnetic flux, $V(\psi)$ is the volume enclosed by a magnetic-flux surface, $V' = dV/d\psi$, and $S_{n,s}$ represents the externally imposed sources of particles. The turbulent flux of particles is given by

$$\gamma_s = \left\langle \int d^3\mathbf{v} (\mathbf{V}_E \cdot \nabla \psi) \delta f_s \right\rangle_t \quad (\text{E.2})$$

and we then define $\Gamma_s = \langle \gamma_s \rangle_\psi$. \mathbf{V}_E denotes the fluctuating $\mathbf{E} \times \mathbf{B}$ velocity and δf_s is the fluctuating part of the distribution function. The brackets $\langle \cdot \rangle_t$ indicate averages over the turbulent time scale, and $\langle \cdot \rangle_\psi$ denotes volume averages over a flux tube.

E.1.2 Momentum transport

The first-order moment is obtained by multiplying the kinetic equation by $m_s R^2 (\mathbf{v} \cdot \nabla \phi)$ before integrating over all velocities, and summing over all species. The time evolution for the angular momentum is then described by:

$$\frac{\partial L}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial \psi} \left[V' \left(\Pi + \sum_s m_s \Omega_\phi \langle R^2 \gamma_s \rangle_\psi \right) \right] = S_L, \quad (\text{E.3})$$

where $L = \sum_s n_s m_s \langle R^2 \rangle_\psi \Omega_\phi$ is the toroidal angular momentum density, Ω_ϕ is the mean toroidal angular frequency, R is the major radius, and S_L represents the total external source of momentum. The time-averaged turbulent flux of toroidal angular momentum for species s is given by:

$$\Pi_s = \left\langle \left\langle \int d^3\mathbf{v} (\mathbf{V}_E \cdot \nabla \psi) (\mathbf{v} \cdot \nabla \phi) m_s R^2 \delta f_s \right\rangle_t \right\rangle_\psi, \quad (\text{E.4})$$

and the total turbulent flux is denoted by $\Pi = \sum_s \Pi_s$. We point out that, with the sign conventions from GS2, a positive momentum flux corresponds to a radially outward flux of momentum in the direction of the plasma current.

E.1.3 Energy transport

The second-order moment is obtained by multiplying the kinetic equation by $m_s v^2/2$ before the velocity-space integration, and yields an equation for the evolution of the plasma pressure:

$$\frac{3}{2} \frac{\partial p_s}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial \psi} (V' Q_s) = S_{\varepsilon,s} + H_s + S_{\Delta,s}, \quad (\text{E.5})$$

where p_s is the plasma pressure and $S_{\varepsilon,s}$ is the external source of energy. The time-averaged turbulent flux of energy is given by

$$Q_s = \left\langle \left\langle \int d^3 \mathbf{v} (\mathbf{V}_E \cdot \nabla \psi) \frac{m_s v^2}{2} \delta f_s \right\rangle_t \right\rangle_\psi \quad (\text{E.6})$$

and the turbulent heating term is

$$H_s = \left\langle \left\langle Z_s e \int d^3 \mathbf{v} h_s \frac{\partial \varphi}{\partial t} \right\rangle_t \right\rangle_\psi - \frac{\partial \Omega_\phi}{\partial \psi} (\Pi_s + m_s \Omega_\phi \langle R^2 \gamma_s \rangle_\psi), \quad (\text{E.7})$$

with Z_s the atomic number, e the elementary charge, h_s the non-Boltzmann part of δf_s and φ the fluctuating electrostatic potential. The collisional energy transfer with species s is given by

$$S_{\Delta,s} = \frac{3}{2} n_s \sum_{s'} \nu_{s,s'}^\varepsilon (T_{s'} - T_s), \quad (\text{E.8})$$

where $\nu_{s,s'}^\varepsilon$ is the collisional frequency for temperature equilibration and T_s is the temperature.

E.2 Computing fluxes from equilibrium profiles

In order to compute the fluxes associated with given equilibrium profiles in the experiment, we assume that the system is at steady state and neglect all time dependencies in the transport equations. The fluxes can then be expressed as functions of the source terms that describe the external injection of particles, momentum and energy into the plasma. The source terms are calculated from the equilibrium profiles using the particle-orbit-following code ASCOT [105]. The angular momentum equation (E.3) can be used to express the momentum flux in terms of the external torque applied to the plasma:

$$\Pi = \frac{1}{V'} \int_0^V dV S_L, \quad (\text{E.9})$$

where we used $V'd\psi = dV$ and where we neglected the contribution from the particle flux (by assuming that the electrons are close to adiabaticity so that the particle fluxes remain small). Similarly, by using equation (E.5), the flux can be expressed in terms of the energy sources:

$$Q_s = \frac{1}{V'} \int_0^V dV (S_{\varepsilon,s} + S_{\Delta,s}). \quad (\text{E.10})$$

We recall from Section 2.2.1 that we assumed a slowly rotating plasma at steady state, and we have thus neglected the contribution from turbulent heating to obtain equation (E.10).

We recall that in GS2, the turbulent fluxes are normalised in the following way:

$$\Pi_N = \frac{1}{\langle |\nabla\psi| \rangle_\psi n_r m_r a v_{\text{th},r}^2 \rho_*^2} \Pi, \quad (\text{E.11})$$

$$Q_{s,N} = \frac{1}{\langle |\nabla\psi| \rangle_\psi n_r T_r v_{\text{th},r} \rho_*^2} Q_s, \quad (\text{E.12})$$

where the subscript N indicates a quantity that has been normalised, and n_r , m_r , T_r , $v_{\text{th},r}$ are the density, mass, temperature and thermal velocity of the reference species, respectively. The presence of $\langle |\nabla\psi| \rangle_\psi$ in the normalisation factors facilitates the comparison of turbulent fluxes from the experiment with those obtained from simulations, as it leads to an expression that is independent of the flux coordinate chosen. For example, applying the normalisation from GS2 yields:

$$\Pi_N = \frac{1}{n_r m_r a v_{\text{th},r}^2 \rho_*^2} \frac{1}{A_\psi} \int_0^V dV S_L, \quad (\text{E.13})$$

where A_ψ is the area of the flux surface, which is readily available in the JETPEAK database of JET discharges. To obtain (E.13), we used the following identity:

$$dV = d\psi \int_0^{2\pi} d\theta \int_0^{2\pi} d\phi \frac{1}{\nabla\theta \cdot (\nabla\phi \times \nabla\psi)} = d\psi \frac{A_\psi}{\langle |\nabla\psi| \rangle_\psi}. \quad (\text{E.14})$$

E.3 Reconstruction method

We choose to describe the momentum flux as the sum of a diffusive term and a pinch term,

$$\Pi = -\chi_L n_i m_i R^2 \frac{d\Omega_\phi}{dr_\psi} - V_L n_i m_i R^2 \Omega_\phi, \quad (\text{E.15})$$

which we identify as the first terms in the Taylor expansion of the momentum flux with respect to the rotation and the rotation shear. Indeed, a symmetry argument presented in [107–109] implies that, for up-down symmetric equilibria such as the ones considered in the present work, the momentum flux vanishes in the combined absence of rotation and rotation shear. In equation (E.15), χ_L is the momentum diffusivity, V_L is the velocity associated with momentum pinch, and $m_e/m_i \ll 1$ allows us to neglect the electron contribution. Following the sign conventions in **GS2**, a positive χ_L corresponds to a radially outward flux of co-current momentum, while a positive V_L corresponds to a radially inward flux of co-current momentum.

From equation (E.15), we express the angular frequency associated with the mean toroidal rotation as

$$\Omega_\phi(r_\psi) = \left[\Omega_\phi(r_{\text{out}}) + \int_{r_\psi}^{r_{\text{out}}} d\tilde{r}_\psi \frac{1}{n_i m_i R^2} \frac{\Pi}{\chi_L} e^{-\Lambda(\tilde{r}_\psi)} \right] e^{\Lambda(r_\psi)}, \quad (\text{E.16})$$

where we use r_{out} to denote a radius located in the outer region of the core, and where we defined the integral

$$\Lambda(r_\psi) = \int_{r_\psi}^{r_{\text{out}}} d\tilde{r}_\psi \frac{V_L}{\chi_L}. \quad (\text{E.17})$$

Considering equilibrium profiles obtained from an experiment with neutral beam injection, the momentum flux Π and the transport coefficients χ_L and V_L can be computed from local gyrokinetic simulations by running the **GS2** code. At any given radial location, two simulations are required to determine χ_L and V_L ; we choose to run one of the simulations with the experimental values $(\Omega_\phi, \Omega'_\phi)$, and the other simulation with $(\Omega_\phi = 0, \Omega'_\phi)$. We follow this procedure for a set of different radial locations, and we then apply linear interpolation to obtain the radial profiles of Π , χ_L and V_L . Finally, the profile $\Omega_\phi(r_\psi)$ is obtained by numerically evaluating the integral in equation (E.16). In Figure E.1, we show that this method can successfully be applied to reconstruct experimental rotation profiles: the reconstruction (indicated by red circles) shows a very good agreement with the experiment (solid black line). The experiment considered here is the JET discharge #68448, of which the details are presented in Appendix B.

In the ideal case when the momentum flux obtained from local gyrokinetic simulations matches the experimental momentum-flux profile, it would be possible to substitute Π in equation (E.16) by its experimental value that we compute via relation (E.9)




Figure E.1: Comparison between the experimental rotation profile (black line), and the rotation profile reconstructed by following the method presented in Section E.3 (red circles). The blue ‘+’ symbols are obtained via the same reconstruction method, but by replacing the momentum flux obtained from local gyrokinetic simulations by the actual momentum-flux profile from the experiment. For cases where the simulations match the transport levels of the experiment, the red circles and the blue ‘+’ symbols would give the same result.

– both approaches would produce the same reconstructed profile. However, for reasons that we address in Chapter 3, a mismatch was found between the simulations and the experiment. As a result, using the experimental Π in the reconstruction does not reproduce the correct rotation profile, as we show with blue ‘+’ symbols in Figure E.1: our simulations in the high-transport state over-estimate the momentum flux (and thus the momentum diffusivity χ_L), which leads to an under-estimate for the ratio Π/χ_L appearing in equation (E.16) when the experimental momentum flux is used.

E.4 Predicting rotation profiles

In theory, the reconstruction method presented in the previous Section could be adapted to obtain predictive estimates of rotation profiles in slowly rotating fusion plasmas. In experiments where the mean rotation is slow enough, the associated flow shear can be assumed to be too weak to substantially affect turbulent eddies. In other words, it can be assumed that, to lowest order in Ω_ϕ , the only dependence on Ω_ϕ and $\partial\Omega_\phi/\partial r_\psi$

appears in expression (E.15) for the turbulent momentum flux Π . As a consequence, if the characteristics of the NBI were to be slightly modified in such an experiment (e.g. by changing the energy, the power, or the injection angle of the beam), the resulting difference in momentum flux exiting the plasma would only alter the rotation profile; other equilibrium parameters such as the temperature or density would remain unaffected. This is, in essence, how the reconstruction method could be made predictive:

- consider an experiment with given equilibrium profiles;
- using local gyrokinetic simulations, reproduce the experimental levels of turbulent transport at different radial locations and obtain radial profiles of the momentum diffusivity χ_L and the momentum pinch V_L ;
- alter the experimental momentum-flux profile to reflect any desired change in the characteristics of the NBI;
- using χ_L and V_L together with the modified experimental momentum-flux profile, evaluate expression (E.16) to obtain $\Omega_\phi(r_\psi)$.

References

- [1] Hoegh-Guldberg, O. *et al.* *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*, chap. 3. Impacts of 1.5°C Global Warming on Natural and Human Systems (World Meteorological Organization, Geneva, Switzerland, 2018). URL <https://www.ipcc.ch/sr15/chapter/chapter-3/>.
- [2] United Nations Framework Convention on Climate Change. *Decision 1/CP.21, Adoption of the Paris Agreement*, chap. Annex, Paris Agreement, Article 2.1.a (2016). URL <https://unfccc.int/resource/docs/2015/cop21/eng/10a01.pdf>.
- [3] IPCC. *Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*, chap. 0. Summary for Policymakers (World Meteorological Organization, Geneva, Switzerland, 2018). URL <https://www.ipcc.ch/sr15/chapter/spm/>.
- [4] Washington, DC: World Resources Institute. Climate Watch Historical GHG Emissions. <https://www.climatewatchdata.org/ghg-emissions> (2021).
- [5] International Renewable Energy Agency, Abu Dhabi. Electricity Storage and Renewables: Costs and Markets to 2030. <https://www.irena.org/publications/>

[2017/Oct/Electricity-storage-and-renewables-costs-and-markets](#)
(2017).

- [6] Wesson, J. *Tokamaks*, chap. 1. Fusion (Oxford University Press, 2011).
- [7] Wesson, J. *Tokamaks*, chap. 4. Confinement (Oxford University Press, 2011).
- [8] Keilhacker, M. *et al.* High fusion performance from deuterium-tritium plasmas in JET. *Nuclear Fusion* **39**, 209–234 (1999). URL <https://doi.org/10.1088/0029-5515/39/2/306>.
- [9] ITER Organization. *ITER Research Plan within the Staged Approach (Level III – Provisional Version)*, chap. Executive Summary (ITER Technical Report, 2018). URL https://www.iter.org/doc/www/content/com/Lists/ITER%20Technical%20Reports/Attachments/9/ITER_Research_Plan_within_the_Staged_Approach_levIII_provversion.pdf.
- [10] Creely, A. J. *et al.* Overview of the SPARC tokamak. *Journal of Plasma Physics* **86**, 865860502 (2020).
- [11] Wesson, J. *Tokamaks*, chap. 5. Heating (Oxford University Press, 2011).
- [12] Hemsworth, R. S. *et al.* Overview of the design of the ITER heating neutral beam injectors. *New Journal of Physics* **19**, 025005 (2017). URL <https://doi.org/10.1088/1367-2630/19/2/025005>.
- [13] Toigo, V. *et al.* Progress in the ITER neutral beam test facility. *Nuclear Fusion* **59**, 086058 (2019). URL <https://doi.org/10.1088/1741-4326/ab2271>.
- [14] G. Serianni on behalf of the NBTF team. Status of the ITER Neutral Beam Test Facility and the first beam operations with the full-size prototype ion source (2021). URL https://conferences.iaea.org/event/214/contributions/17034/attachments/9999/14347/SerianniJoshi_IAEA2021_Oral_20210409_2.pdf.

- [15] Hinton, F. L. & Hazeltine, R. D. Theory of plasma transport in toroidal confinement systems. *Rev. Mod. Phys.* **48**, 239–308 (1976). URL <https://link.aps.org/doi/10.1103/RevModPhys.48.239>.
- [16] Hirshman, S. & Sigmar, D. Neoclassical transport of impurities in tokamak plasmas. *Nuclear Fusion* **21**, 1079–1201 (1981). URL <https://doi.org/10.1088/0029-5515/21/9/003>.
- [17] Horton, W. Drift waves and transport. *Rev. Mod. Phys.* **71**, 735–778 (1999). URL <https://link.aps.org/doi/10.1103/RevModPhys.71.735>.
- [18] Doyle, E. *et al.* Chapter 2: Plasma confinement and transport. *Nuclear Fusion* **47**, S18–S127 (2007). URL <https://doi.org/10.1088/0029-5515/47/6/s02>.
- [19] Mantica, P. *et al.* Progress and challenges in understanding core transport in tokamaks in support to ITER operations. *Plasma Physics and Controlled Fusion* **62**, 014021 (2019). URL <https://doi.org/10.1088/1361-6587/ab5ae1>.
- [20] Shimada, M. *et al.* Chapter 1: Overview and summary. *Nuclear Fusion* **47**, S1–S17 (2007). URL <https://doi.org/10.1088/0029-5515/47/6/s01>.
- [21] Braginskii, S. Transport phenomena in a completely ionized two-temperature plasma. *Sov. Phys. JETP* **6**, 358–369 (1958).
- [22] Mikhailovskii, A. & Tsypin, V. Transport equations and gradient instabilities in a high pressure collisional plasma. *Plasma Physics* **13**, 785 (1971).
- [23] Epperlein, E. & Haines, M. Plasma transport coefficients in a magnetic field by direct numerical solution of the fokker–planck equation. *The Physics of fluids* **29**, 1029–1041 (1986).
- [24] Catto, P. J. & Simakov, A. N. A drift ordered short mean free path description for magnetized plasma allowing strong spatial anisotropy. *Physics of Plasmas* **11**, 90–102 (2004).

- [25] Sadler, J. D., Walsh, C. A. & Li, H. Symmetric set of transport coefficients for collisional magnetized plasma. *Phys. Rev. Lett.* **126**, 075001 (2021). URL <https://link.aps.org/doi/10.1103/PhysRevLett.126.075001>.
- [26] Catto, P. J. Linearized gyro-kinetics. *Plasma Physics* **20**, 719–722 (1978). URL <https://doi.org/10.1088%2F0032-1028%2F20%2F7%2F011>.
- [27] Frieman, E. A. & Chen, L. Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *The Physics of Fluids* **25**, 502–508 (1982). URL <https://aip.scitation.org/doi/abs/10.1063/1.863762>.
- [28] Sugama, H. & Horton, W. Nonlinear electromagnetic gyrokinetic equation for plasmas with large mean flows. *Physics of Plasmas* **5**, 2560–2573 (1998). URL <https://doi.org/10.1063/1.872941>.
- [29] Abel, I. *et al.* Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows. *Reports on Progress in Physics* **76**, 116201 (2013).
- [30] Biglari, H., Diamond, P. H. & Terry, P. W. Influence of sheared poloidal rotation on edge turbulence. *Physics of Fluids B: Plasma Physics* **2**, 1–4 (1990). URL <https://doi.org/10.1063/1.859529>.
- [31] Dimits, A. M. *et al.* Comparisons and physics basis of tokamak transport models and turbulence simulations. *Physics of Plasmas* **7**, 969–983 (2000). URL <https://doi.org/10.1063/1.873896>.
- [32] Rogers, B. N., Dorland, W. & Kotschenreuther, M. Generation and stability of zonal flows in ion-temperature-gradient mode turbulence. *Phys. Rev. Lett.* **85**, 5336–5339 (2000). URL <https://link.aps.org/doi/10.1103/PhysRevLett.85.5336>.
- [33] Diamond, P. H., Itoh, S., Itoh, K. & Hahm, T. Zonal flows in plasma—a review. *Plasma Physics and Controlled Fusion* **47**, R35 (2005).

- [34] Colyer, G. J. *et al.* Collisionality scaling of the electron heat flux in ETG turbulence. *Plasma Physics and Controlled Fusion* **59**, 055002 (2017). URL <https://doi.org/10.1088/1361-6587/aa5f75>.
- [35] van Wyk, F. *et al.* Ion-scale turbulence in MAST: anomalous transport, subcritical transitions, and comparison to BES measurements. *Plasma Physics and Controlled Fusion* **59**, 114003 (2017). URL <https://doi.org/10.1088/1361-6587/aa8484>.
- [36] Ivanov, P. G., Schekochihin, A. A., Dorland, W., Field, A. R. & Parra, F. I. Zonally dominated dynamics and dimits threshold in curvature-driven itg turbulence. *Journal of Plasma Physics* **86**, 855860502 (2020).
- [37] Waelbroeck, F. L. & Chen, L. Ballooning instabilities in tokamaks with sheared toroidal flows. *Physics of Fluids B: Plasma Physics* **3**, 601–610 (1991). URL <https://doi.org/10.1063/1.859858>.
- [38] Artun, M. & Tang, W. M. Gyrokinetic analysis of ion temperature gradient modes in the presence of sheared flows. *Physics of Fluids B: Plasma Physics* **4**, 1102–1114 (1992). URL <https://doi.org/10.1063/1.860118>.
- [39] Newton, S. L., Cowley, S. C. & Loureiro, N. F. Understanding the effect of sheared flow on microinstabilities. *Plasma Physics and Controlled Fusion* **52**, 125001 (2010). URL <https://doi.org/10.1088/0741-3335/52/12/2F125001>.
- [40] Dimits, A. M., Williams, T. J., Byers, J. A. & Cohen, B. I. Scalings of ion-temperature-gradient-driven anomalous transport in tokamaks. *Phys. Rev. Lett.* **77**, 71–74 (1996). URL <https://link.aps.org/doi/10.1103/PhysRevLett.77.71>.
- [41] Casson, F. J. *et al.* Anomalous parallel momentum transport due to ExB flow shear in a tokamak plasma. *Physics of Plasmas* **16**, 092303 (2009). URL <https://doi.org/10.1063/1.3227650>.

- [42] Roach, C. M. *et al.* Gyrokinetic simulations of spherical tokamaks. *Plasma Physics and Controlled Fusion* **51**, 124020 (2009). URL <https://doi.org/10.1088/0741-3335/51/12/124020>.
- [43] Highcock, E. G. *et al.* Transport bifurcation in a rotating tokamak plasma. *Phys. Rev. Lett.* **105**, 215003 (2010). URL <https://link.aps.org/doi/10.1103/PhysRevLett.105.215003>.
- [44] Barnes, M. *et al.* Turbulent transport in tokamak plasmas with rotational shear. *Phys. Rev. Lett.* **106**, 175004 (2011). URL <https://link.aps.org/doi/10.1103/PhysRevLett.106.175004>.
- [45] Citrin, J. *et al.* Ion temperature profile stiffness: non-linear gyrokinetic simulations and comparison with experiment. *Nuclear Fusion* **54**, 023008 (2014). URL <https://doi.org/10.1088/0029-5515/54/2/023008>.
- [46] van Wyk, F. *et al.* Transition to subcritical turbulence in a tokamak plasma. *Journal of Plasma Physics* **82**, 905820609 (2016).
- [47] Synakowski, E. J. *et al.* Roles of electric field shear and shafranov shift in sustaining high confinement in enhanced reversed shear plasmas on the TFTR tokamak. *Phys. Rev. Lett.* **78**, 2972–2975 (1997). URL <https://link.aps.org/doi/10.1103/PhysRevLett.78.2972>.
- [48] Mantica, P. *et al.* Experimental study of the ion critical-gradient length and stiffness level and the impact of rotation in the JET tokamak. *Phys. Rev. Lett.* **102**, 175002 (2009). URL <https://link.aps.org/doi/10.1103/PhysRevLett.102.175002>.
- [49] Hammett, G. W., Dorland, W., Loureiro, N. F. & Tatsuno, T. Implementation of large scale $\mathbf{E} \times \mathbf{B}$ shear flow in the gs2 gyrokinetic turbulence code (2006). Poster presented at the DPP meeting of the American Physical Society.

- [50] Candy, J. & Belli, E. A. Spectral treatment of gyrokinetic shear flow. *Journal of Computational Physics* **356**, 448 – 457 (2018). URL <http://www.sciencedirect.com/science/article/pii/S0021999117309063>.
- [51] McMillan, B. F., Ball, J. & Brunner, S. Simulating background shear flow in local gyrokinetic simulations. *Plasma Physics and Controlled Fusion* **61**, 055006 (2019). URL <https://doi.org/10.1088%2F1361-6587%2F61%2F055006>.
- [52] Snedeker, R. S. & Donaldson, C. d. Observation of a bistable flow in a hemispherical cavity. *AIAA Journal* **4**, 735–736 (1966).
- [53] Burggraf, O. R. & Foster, M. R. Continuation or breakdown in tornado-like vortices. *Journal of Fluid Mechanics* **80**, 685–703 (1977).
- [54] Schmucker, A. & Gersten, K. Vortex breakdown and its control on delta wings. *Fluid Dynamics Research* **3**, 268–272 (1988). URL [https://doi.org/10.1016/0169-5983\(88\)90077-9](https://doi.org/10.1016/0169-5983(88)90077-9).
- [55] Shtern, V. & Hussain, F. Hysteresis in a swirling jet as a model tornado. *Physics of Fluids A: Fluid Dynamics* **5**, 2183–2195 (1993).
- [56] Ravelet, F., Marié, L., Chiffaudel, A. & Daviaud, F. Multistability and memory effect in a highly turbulent flow: Experimental evidence for a global bifurcation. *Physical review letters* **93**, 164501 (2004).
- [57] Simitev, R. D. & Busse, F. H. Bistability and hysteresis of dipolar dynamos generated by turbulent convection in rotating spherical shells. *EPL (Europhysics Letters)* **85**, 19001 (2009). URL <https://doi.org/10.1209/0295-5075/85/19001>.
- [58] Latter, H. N. & Papaloizou, J. C. B. Hysteresis and thermal limit cycles in MRI simulations of accretion discs. *Monthly Notices of the Royal Astronomical Society* **426**, 1107–1120 (2012). URL <https://doi.org/10.1111/j.1365-2966.2012.21748.x>.

- [59] Christen, N., Barnes, M. & Parra, F. I. Continuous-in-time approach to flow shear in a linearly implicit local δf gyrokinetic code. *Journal of Plasma Physics* **87**, 905870230 (2021).
- [60] Kotschenreuther, M., Rewoldt, G. & Tang, W. M. Comparison of initial value and eigenvalue codes for kinetic toroidal plasma instabilities. *Computer Physics Communications* **88**, 128 – 140 (1995). URL <http://www.sciencedirect.com/science/article/pii/001046559500035E>.
- [61] Jenko, F., Dorland, W., Kotschenreuther, M. & Rogers, B. N. Electron temperature gradient driven turbulence. *Physics of Plasmas* **7**, 1904–1910 (2000). URL <https://doi.org/10.1063/1.874014>.
- [62] Peeters, A. G. *et al.* The nonlinear gyro-kinetic flux tube code GKW. *Computer Physics Communications* **180**, 2650–2672 (2009). URL <http://www.sciencedirect.com/science/article/pii/S0010465509002112>. 40 years of CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures.
- [63] Catto, P. J., Bernstein, I. B. & Tessarotto, M. Ion transport in toroidally rotating tokamak plasmas. *The Physics of Fluids* **30**, 2784–2795 (1987). URL <https://aip.scitation.org/doi/abs/10.1063/1.866045>.
- [64] Beer, M. A., Cowley, S. & Hammett, G. Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Physics of Plasmas* **2**, 2687–2700 (1995). URL <https://doi.org/10.1063/1.871232>.
- [65] Barnes, M., Parra, F. I. & Schekochihin, A. A. Critically balanced ion temperature gradient turbulence in fusion plasmas. *Phys. Rev. Lett.* **107**, 115003 (2011). URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.115003>.
- [66] Candy, J. Beta scaling of transport in microturbulence simulations. *Physics of Plasmas* **12**, 072307 (2005). URL <https://doi.org/10.1063/1.1954123>.

- [67] Waltz, R. E., Austin, M. E., Burrell, K. H. & Candy, J. Gyrokinetic simulations of off-axis minimum-q profile corrugations. *Physics of Plasmas* **13**, 052301 (2006). URL <https://doi.org/10.1063/1.2195418>.
- [68] Dominski, J. *et al.* How non-adiabatic passing electron layers of linear microinstabilities affect turbulent transport. *Physics of Plasmas* **22**, 062303 (2015). URL <https://doi.org/10.1063/1.4922659>.
- [69] Weikl, A. *et al.* The occurrence of staircases in ITG turbulence with kinetic electrons and the zonal flow drive through self-interaction. *Physics of Plasmas* **25**, 072305 (2018). URL <https://doi.org/10.1063/1.5035184>.
- [70] C. J., A. *et al.* How eigenmode self-interaction affects zonal flows and convergence of tokamak core turbulence with toroidal system size. *Journal of Plasma Physics* **86**, 905860504 (2020).
- [71] Ball, J., Brunner, S. & C.J., A. Eliminating turbulent self-interaction through the parallel boundary condition in local gyrokinetic simulations. *Journal of Plasma Physics* **86**, 905860207 (2020).
- [72] Chandrarajan Jayalekshmi, A. *Studying the effect of non-adiabatic passing electron dynamics on microturbulence self-interaction in fusion plasmas using gyrokinetic simulations*. Ph.D. thesis, École Polytechnique Fédérale de Lausanne (2020).
- [73] Romanelli, F. Ion temperature-gradient-driven modes and anomalous ion transport in tokamaks. *Physics of Fluids B: Plasma Physics* **1**, 1018–1025 (1989). URL <https://doi.org/10.1063/1.859023>.
- [74] Cowley, S. C., Kulsrud, R. M. & Sudan, R. Considerations of ion-temperature-gradient-driven turbulence. *Physics of Fluids B: Plasma Physics* **3**, 2767–2782 (1991). URL <https://doi.org/10.1063/1.859913>.
- [75] Kadomtsev, B. B. & Pogutse, O. P. Reviews of plasma physics **5**, 249 (1970).

- [76] Liu, C. S. Instabilities in a magnetoplasma with skin current. *Phys. Rev. Lett.* **27**, 1637–1640 (1971). URL <https://link.aps.org/doi/10.1103/PhysRevLett.27.1637>.
- [77] Lee, Y., Dong, J., Guzdar, P. & Liu, C. Collisionless electron temperature gradient instability. *The Physics of fluids* **30**, 1331–1339 (1987).
- [78] Horton, W., Hong, B. & Tang, W. Toroidal electron temperature gradient driven drift modes. *The Physics of fluids* **31**, 2971–2983 (1988).
- [79] Siren, P., Varje, J., Weisen, H. & Giacomelli, L. Role of JETPEAK database in validation of synthetic neutron camera diagnostics and ASCOT- AFSI fast particle and fusion product calculation chain in JET. *Journal of Instrumentation* **14**, C11013–C11013 (2019). URL <https://doi.org/10.1088/1748-0221/14/11/c11013>.
- [80] Belli, E. A. *Studies of numerical algorithms for gyrokinetics and the effects of shaping on plasma turbulence*. Ph.D. thesis, Princeton University (2006).
- [81] Fox, M. F. J. *et al.* Symmetry breaking in MAST plasma turbulence due to toroidal flow shear. *Plasma Physics and Controlled Fusion* **59**, 034002 (2017). URL <https://doi.org/10.1088/1361-6587/aa544b>.
- [82] Scott, B. Gyrokinetic study of the edge shear layer. *Plasma Physics and Controlled Fusion* **48**, A387–A392 (2006). URL <https://doi.org/10.1088/0741-3335/48/5a/s39>.
- [83] Schekochihin, A. A., Highcock, E. G. & Cowley, S. C. Subcritical fluctuations and suppression of turbulence in differentially rotating gyrokinetic plasmas. *Plasma Physics and Controlled Fusion* **54**, 055011 (2012). URL <https://doi.org/10.1088/0741-3335/54/5/055011>.
- [84] Shafer, M. W. *et al.* 2D properties of core turbulence on DIII-D and comparison to gyrokinetic simulations. *Physics of Plasmas* **19**, 032504 (2012). URL <https://doi.org/10.1063/1.3691965>.

- [85] McMillan, B. F., Pringle, C. C. T. & Teaca, B. Simple advecting structures and the edge of chaos in subcritical tokamak plasmas. *Journal of Plasma Physics* **84**, 905840611 (2018).
- [86] Hardman, M. R., Barnes, M., Roach, C. M. & Parra, F. I. A scale-separated approach for studying coupled ion and electron scale turbulence. *Plasma Physics and Controlled Fusion* **61**, 065025 (2019). URL <https://doi.org/10.1088/1361-6587/ab1323>.
- [87] Hardman, M. *Multiscale turbulence in magnetic confinement fusion devices*. Ph.D. thesis, University of Oxford (2019).
- [88] Hardman, M. R., Barnes, M. & Roach, C. M. Stabilisation of short-wavelength instabilities by parallel-to-the-field shear in long-wavelength $E \times B$ flows. *Journal of Plasma Physics* **86**, 905860601 (2020).
- [89] Faisst, H. & Eckhardt, B. Sensitive dependence on initial conditions in transition to turbulence in pipe flow. *Journal of Fluid Mechanics* **504**, 343–352 (2004).
- [90] Highcock, E. G. *et al.* Transport bifurcation induced by sheared toroidal flow in tokamak plasmas. *Physics of Plasmas* **18**, 102304 (2011). URL <https://doi.org/10.1063/1.3642611>.
- [91] Parra, F. I., Barnes, M., Highcock, E. G., Schekochihin, A. A. & Cowley, S. C. Momentum injection in tokamak plasmas and transitions to reduced transport. *Phys. Rev. Lett.* **106**, 115004 (2011). URL <https://link.aps.org/doi/10.1103/PhysRevLett.106.115004>.
- [92] von Goeler, S., Stodiek, W. & Sauthoff, N. Studies of internal disruptions and $m = 1$ oscillations in tokamak discharges with soft-X-ray techniques. *Phys. Rev. Lett.* **33**, 1201–1203 (1974). URL <https://link.aps.org/doi/10.1103/PhysRevLett.33.1201>.

- [93] Wagner, F. *et al.* Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak. *Phys. Rev. Lett.* **49**, 1408–1412 (1982). URL <https://link.aps.org/doi/10.1103/PhysRevLett.49.1408>.
- [94] Hastie, R. Sawtooth instability in tokamak plasmas. *Astrophysics and space science* **256**, 177–204 (1997).
- [95] Connor, J. W. A review of models for ELMs. *Plasma Physics and Controlled Fusion* **40**, 191–213 (1998). URL <https://doi.org/10.1088/0741-3335/40/2/003>.
- [96] Weikl, A. *et al.* Ion temperature gradient turbulence close to the finite heat flux threshold. *Physics of Plasmas* **24**, 102317 (2017). URL <https://doi.org/10.1063/1.4986035>.
- [97] Peeters, A. G. *et al.* Gradient-driven flux-tube simulations of ion temperature gradient turbulence close to the non-linear threshold. *Physics of Plasmas* **23**, 082517 (2016). URL <https://doi.org/10.1063/1.4961231>.
- [98] Barnes, M. *et al.* Direct multiscale coupling of a transport code to gyrokinetic turbulence codes. *Physics of Plasmas* **17**, 056109 (2010). URL <https://doi.org/10.1063/1.3323082>.
- [99] Watanabe, T.-H., Sugama, H., Ishizawa, A. & Nunami, M. Flux tube train model for local turbulence simulation of toroidal plasmas. *Physics of Plasmas* **22**, 022507 (2015). URL <https://doi.org/10.1063/1.4907793>.
- [100] Lang, J., Parker, S. E. & Chen, Y. Nonlinear saturation of collisionless trapped electron mode turbulence: Zonal flows and zonal density. *Physics of Plasmas* **15**, 055907 (2008).
- [101] Merz, F. & Jenko, F. Nonlinear saturation of trapped electron modes via perpendicular particle diffusion. *Physical review letters* **100**, 035005 (2008).

- [102] Gormezano, C. *et al.* Hybrid advanced scenarios: perspectives for ITER and new experiments with dominant RF heating. *Plasma Physics and Controlled Fusion* **46**, B435–B447 (2004). URL <https://doi.org/10.1088/0741-3335/46/12b/037>.
- [103] Lao, L., John, H. S., Stambaugh, R., Kellman, A. & Pfeiffer, W. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion* **25**, 1611–1622 (1985). URL <https://doi.org/10.1088/0029-5515/25/11/007>.
- [104] Transp webpage. <http://w3.pppl.gov/~pshare/help/transp.htm> (2021).
- [105] Hirvijoki, E. *et al.* Ascot: Solving the kinetic equation of minority particle species in tokamak plasmas. *Computer Physics Communications* **185**, 1310–1321 (2014).
- [106] Miller, R. L., Chu, M. S., Greene, J. M., Lin-Liu, Y. R. & E, R. E. W. R. Noncircular, finite aspect ratio, local equilibrium model. *Physics of Plasmas* **5**, 973–978 (1998). URL <https://doi.org/10.1063/1.872666>.
- [107] Parra, F. I., Barnes, M. & Peeters, A. G. Up-down symmetry of the turbulent transport of toroidal angular momentum in tokamaks. *Physics of Plasmas* **18**, 062501 (2011). URL <https://doi.org/10.1063/1.3586332>.
- [108] Peeters, A. G. & Angioni, C. Linear gyrokinetic calculations of toroidal momentum transport in a tokamak due to the ion temperature gradient mode. *Physics of Plasmas* **12**, 072515 (2005). URL <https://doi.org/10.1063/1.1949608>.
- [109] Peeters, A. *et al.* Overview of toroidal momentum transport. *Nuclear Fusion* **51**, 094027 (2011).