

Gyrokinetic turbulence and transport in the Mega Ampere Spherical Tokamak



G. J. Colyer

Exeter College

University of Oxford

Thesis submitted for the Degree of

Doctor of Philosophy

Trinity Term 2016

Abstract

In this thesis we consider energy transport due to turbulence in the Mega Ampere Spherical Tokamak (MAST), using a multi-scale framework based on local δf gyrokinetics. Transport is modelled globally by means of local flux-tube simulations at each magnetic flux surface. In view of the evidence indicating that turbulent ion transport can be substantially suppressed by rotation shear, our flux tubes span only electron-gyroradius scales. This both simplifies our physics model and reduces the computational resource requirements to a manageable level. Our electrostatic simulations of a MAST H-mode discharge exhibit turbulent electron heat transport comparable with experiment in the outer core region, and are broadly consistent with the paradigm of an electron temperature profile governed by pedestal height and critical electron temperature gradient. Neoclassical transport dominates the ions, and is also comparable with experiment. The two species are decoupled in the sense that the collisional equilibration between ions and electrons is small compared to the sources across most of the plasma.

Focusing on a single flux surface in this region, still using both kinetic electrons and kinetic ions, we find that at early simulation times the heat flux “quasi-saturates” with the turbulence dominated by streamer-like radially elongated structures. However, the zonal fluctuation component continues to grow slowly until much later times, eventually leading to a new saturated state dominated by the zonal component.

Simplifying further to an adiabatic ion model (which shows the same slow evolution behaviour), we find that in the final saturated state (which determines the macroscopic energy transport) the electron heat flux is approximately proportional to the collision rate. We outline an explanation of this effect based on zonal-nonzonal interactions and a scaling of the zonal damping rate with electron-ion collisionality. Improved energy confinement with decreasing collisionality has previously been observed experimentally in STs, and is favourable towards the performance of future devices, which are expected to be hotter and thus less collisional.

In memory of my father, who went ahead of me

Peter John Colyer MA MSt DPhil

1943–2014

Contents

List of Figures	8
1 Overview	13
2 Introduction	19
2.1 Introduction to gyrokinetics	19
2.2 Self-consistent transport modelling	23
2.3 Rotation and flow shear	24
2.4 Electromagnetic modes	26
2.5 The GS2 code	28
3 Simulating global energy transport in MAST using local electron-scale flux tubes	31
3.1 Introduction	31
3.2 Transport solution method	33
3.3 Experimental data	35
3.4 Simulation setup	36
3.5 Trinity simulation results	39
3.6 Turbulence properties	43
3.7 Conclusions and future work	44
4 Long-time and ion modelling effects	45
5 Collisionality scaling of the electron heat flux in ETG turbulence	47
5.1 Introduction	47
5.2 Governing equations and numerical set-up	51
5.2.1 Gyrokinetic equation	51
5.2.2 Adiabatic ions	53
5.2.3 Collisions	54
5.3 Results	56
5.3.1 Time evolution: from quasi-saturation to long-time steady state	56
5.3.2 Collisionality scaling: numerical results and theory	58
5.3.3 Damping of zonal modes	62
5.3.4 Numerical tests of the role of collisions and zonal modes	65

5.3.5	Simplified simulations for extended collisionality range . . .	67
5.3.6	Spatial structure of the saturated state	69
5.4	Summary and discussion	71
5.4.1	Previous work on ETG and the collisionality scaling	74
5.4.2	ETG vs. ITG turbulence near and far from threshold	75
6	Conclusions and further work	79
A	Correspondence between energy confinement time and heat flux scalings	83
B	Plasma parameters	85
C	Numerical details	87
C.1	GS2 spatial coordinates	87
C.2	Numerical grids, resolution and boundary conditions	88
D	Linear growth rates	91
E	Issues of numerical convergence in nonlinear simulations	95
E.1	Convergence of the heat flux	95
E.2	Background flow shear	100
F	Long-time linear damping of the electron zonal modes	103
F.1	Zonal evolution equation	103
F.2	Long-wavelength ordering and time scales	104
F.3	Zeroth order	105
F.4	First order	106
F.5	Second order	107
F.6	Damping rate	108
F.7	Zonal damping by same-species collisions	109
F.8	Case of ion zonal modes	110
	Acknowledgements	113
	Bibliography	117

List of Figures

1.1	MAST plasma (courtesy B. Lloyd and MAST Team, CCFE)	15
2.1	Gyroaveraging: the gyrokinetic distribution function describes rings that experience the ring-averaged fields. The black arrows indicate the position vectors of the particle, \mathbf{r} , and of the gyrocentre, \mathbf{R} , relative to the same (arbitrary) origin.	21
2.2	Linear growth rates calculated with various codes (named in legend) using numerical geometry (first three curves) and Miller geometry (last two curves) for the collisionless, electrostatic benchmark case (based on MAST shot 8500 at $\rho_n = 0.6$) used in Roach et al. [13] (data courtesy W. Guttenfelder and Y. Camenen). (Only the numerical geometry curves were published in [13].)	29
3.1	Radial profiles of (a) n_e and (b) q at $t = 0.289$ s for MAST shot 8500	35
3.2	Linear growth rate spectra without flow shear at the four flux surfaces specified in Table 3.1. Circles indicate stability.	37
3.3	Schematic showing four of the nested flux surfaces used for the transport calculation. Each square represents (not to scale – formally smaller in the multiscale expansion) the cross-section at the outboard midplane through the corresponding flux tube.	38
3.4	(a) T_e after 0.01 s showing rapid core temperature increase; (b) T_e at end of run with modified χ_e where dotted; (c) T_i compared with neoclassical-only run	40
3.5	(a) T_e gradient profile; (b) integrated sources and exchange heating	41
3.6	Time evolution of T_e (solid and dotted, to match Figure 3.4(b)) and T_i (dashed). Each symbol shows a transport time step taken by Trinity.	42
3.7	Electron heat flux at $r/a = 0.65$: (a) spectrum averaged over final transport time step; (b) total (circle), and total from flux tube with perturbed electron temperature gradient (triangle): the dashed line between the two symbols gives an indication of stiffness (its slope) and critical gradient (its intercept with $Q_e = 0$).	43

4.1	Time traces of (a) spectral intensity of the electrostatic potential, summed over all k_x , $\sum_{k_x} \phi_{\mathbf{k}_\perp} ^2$, for the zonal $k_y = 0$ component (black) and for multiple individual non-zonal $k_y \neq 0$ components (all yellow); (b) electron heat flux Q_e	46
5.1	Non-dimensionalised electrostatic potential $e\phi/T\rho_*$ (where $\rho_* = \rho_e/a$) at the outboard midplane, for $\nu = 0.2\nu_{\text{nom}}$ (here ν_{nom} is the “nominal”, i.e., experimental value of collisionality), $a/L_T = 3.3$: (a) quasi-saturated state at $t = 1200.3 a/v_{te}$, (b) saturated state at $t = 7835.8 a/v_{te}$, for large-box simulations. See Appendix B and Appendix C for the meaning of symbols.	49
5.2	Evolution in time of (a) the turbulent electron heat flux in electron gyroBohm units, and (b) the square of the zonal velocity, $(k_Z\phi_Z)^2$, for two simulations with electron collisionality $\nu = 0.2\nu_{\text{nom}}$, electron temperature gradient $a/L_T = 3.3$ (green: small-box simulation; yellow: large-box simulation; see Appendix C.2 for details).	57
5.3	Variation of (a) the time-averaged normalised electron heat flux Q/Q_{GB} , and (b) the rms zonal velocity $k_Z\phi_Z$, defined by (5.17), versus normalised electron collisionality $\nu a/v_{te}$, at the nominal (experimental) value of the temperature gradient $a/L_T = 3.42$, and at $a/L_T = 3.3$. Symbol shapes indicate the simulation series (“small box” or “large box”), as explained in Appendix E. The dot-dashed line shows the experimental power-law scaling, and the dotted line shows the theoretical linear scaling, $Q/Q_{\text{GB}} \propto \nu_*$, equation (5.31).	58
5.4	Zonal damping rate normalised to collisionality, γ_Z/ν , versus $k_x\rho_e$, spanning the range of collisionalities shown in Figure 5.3 (solid colors). The final states of various saturated nonlinear simulations were used as initial conditions, with the nonlinearity switched off. Also shown (black crosses) are the corresponding damping rates for a simulation at $\nu = \nu_{\text{nom}}$ in which electron-ion collisions were turned off (formally by setting $Z_{\text{eff}} = 0$; see section 5.2.3); and a simulation at $\nu = \nu_{\text{nom}}$ in which electron-ion collisions were retained but magnetic drifts were turned off (blue open circles). The solid line with slope k_x^2 corresponds to the scaling (5.30); the dashed line with slope k_x^4 corresponds to the scaling expected when $\nu_{ei} = 0$ (see Appendix F.7).	64
5.5	The large-box simulation shown in Figure 5.2 (yellow) was restarted at $t = 8397.7 a/v_{te}$ without zonal interactions in the nonzonal evolution equation (cyan). The heat flux in this modified simulation returns to a level that is close to the high early “quasi-saturated” level. For direct comparison, the same modified simulation was also rerun from initial noise, giving the same heat flux level (purple).	65

5.6	The two simulations corresponding to the red squares in Figure 5.3 are shown: the large-box simulations with $a/L_T = 3.42$ and collisionalities $\nu = \nu_{\text{nom}}/2$ (cyan) and $\nu = \nu_{\text{nom}}/3$ (red). The higher-collisionality case is restarted at $t = 10003.6 a/v_{te}$ with certain terms in the collision operator reduced to the lower collisionality value, whilst others are retained at the higher value: (a) only electron-electron collisions reduced (blue); (b) only electron-ion collisions reduced (purple); (c) only nonzonal collisions reduced (green); (d) only zonal collisions reduced (yellow).	66
5.7	Variation of (a) the time-averaged electron heat flux Q/Q_{GB} , (b) the rms zonal velocity $k_Z \phi_Z$, adding to the points from Figure 5.3 (now circles) further points (red crosses) obtained by varying only the electron-ion collisionality ν_{ei} , but keeping the electron-electron collisionality at the nominal value, $\nu_{ee} = \nu_{\text{nom}}$, all at the nominal temperature gradient $a/L_T = 3.42$. The dot-dashed line shows the experimental scaling (5.1) and the dotted line shows the theoretical linear scaling, $Q/Q_{\text{GB}} \propto \nu_*$, equation (5.31).	68
5.8	Zonal electrostatic potential ϕ_Z in a flux-tube cross-section in the outboard midplane as a function of the radial spatial coordinate x and time, for the same “large-box” simulation as shown in Figure 5.2 ($\nu = 0.2 \nu_{\text{nom}}$, $a/L_T = 3.3$). The black line shows the start of the time-averaging window used to characterise the saturated state.	69
5.9	For the same case as Figure 5.2: (a) the spectrum of the zonal potential; (b) the heat-flux spectrum; both are averaged over the time window of the saturated state, which is shorter for the large-box simulation (yellow) than for the small-box simulation (green).	70
5.10	2D spectra of the nonzonal electrostatic potential for the same large-box simulation as in Figure 5.2: (a) the spectrum averaged over the saturated time-window $t \geq 7502.3 a/v_{te}$; (b) instantaneous spectrum at the same time as Figure 5.1(b); (c) & (d) instantaneous spectra at later times during the saturated state.	71
D.1	Linear growth rates versus k_y , at $k_x = 0$, with zero flow shear $\gamma_E = 0$, for various collisionalities as indicated. All other parameters are the nominal ones given in Table B.1 (including $a/L_T = 3.42$).	91
D.2	Linear growth rates versus a/L_T , with zero flow shear $\gamma_E = 0$, for various collisionalities as indicated, showing that the linear critical gradient (where the growth rate crosses zero) is insensitive to collisionality.	92

E.1	Variation of (a) the time-averaged normalised electron heat flux Q/Q_{GB} and (b) the rms zonal velocity $k_Z\phi_Z$, defined by (5.17), versus normalised electron collisionality $\nu a/v_{te}$ for our entire simulation series using both small and large boxes, and two temperature gradients, including unresolved runs (hollow symbols; infilled symbols are the resolved points shown in Figure 5.3).	96
E.2	Evolution in time of (a) the normalised turbulent electron heat flux Q/Q_{GB} , (b) the square of the zonal electrostatic potential, $\phi_Z^2 = \sum_{k_x} \phi_{k_x,0} ^2$, (c) the square of the zonal velocity, $(k_Z\phi_Z)^2 = \sum_{k_x} k_x^2 \phi_{k_x,0} ^2$, (d) the square of the zonal shear, $\sum_{k_x} k_x^4 \phi_{k_x,0} ^2$, and (e) the rms zonal wavenumber k_Z , for large-box simulations at the nominal $a/L_T = 3.42$ and at various collisionalities (colours: same as Figure D.1, plus yellow for $\nu = 2\nu_{\text{nom}}$).	97
E.3	Spectrum of the zonal perpendicular temperature perturbation, for $\nu = 0.1\nu_{\text{nom}}$, $a/L_T = 3.3$, for a small-box simulation (green) and a large-box simulation (yellow), at the final simulation time in each case. . . .	98
E.4	Evolution in time of (a) the square of the zonal electrostatic potential, $\phi_Z^2 = \sum_{k_x} \phi_{k_x,0} ^2$, (b) the square of the zonal shear, $\sum_{k_x} k_x^4 \phi_{k_x,0} ^2$, and (c) the rms zonal wavenumber k_Z , for $\nu = 0.2\nu_{\text{nom}}$, $a/L_T = 3.3$ (green: small-box simulations; yellow: large-box simulations). This complements Figure 5.2.	99
E.5	Electrostatic potential ϕ at the outboard midplane, for $\nu = 0.2\nu_{\text{nom}}$, $a/L_T = 3.3$: (a) quasi-saturated state at $t = 1204.2 a/v_{te}$, (b) long-time saturated state at $t = 7841.5 a/v_{te}$, for small-box simulations showing 2×2 copies of each periodic domain. The large-box version of this simulation is shown in Figure 5.1.	101

CHAPTER 1

Overview

In his book *Sustainable Energy – without the hot air* [1], David MacKay makes a compelling case that the UK can only reach a level of energy production and consumption that is sustainable in the long term (whether this means an environmentally acceptable level of carbon emissions, or simply matching available energy sources) either by drastic reductions in demand, or by a step change upward in new supply technology: to advanced forms of nuclear fission, country-sized installations for capturing renewable energy, or nuclear fusion. Any one of these options is an enormous challenge to achieve. This thesis is concerned with the last one listed, nuclear fusion, and more specifically with the role played by plasma turbulence in tokamak device performance.

Understanding turbulence is a scientific challenge in its own right, and we will not be concerned in this work with the social or technological issues related to fusion, important though they are to energy policy and for scientists to consider in their role as citizens. We will start by touching briefly on the various elements making up the title of this thesis.

Plasma turbulence and device performance are connected through the linked concepts of transport and confinement. Transport refers to the transfer across a medium of a physical property. For example, heat conduction and convection are two forms of transport involving the transfer of thermal energy. Confinement is the

opposite: it refers to the degree to which transport is lacking; so higher transport means less confinement, and lower transport means more confinement. Confinement is typically quantified in terms of a confinement time, which is the characteristic time over which the transport would remove the property concerned from the system. For a system in steady state it is equal to the quantity held by the system divided by its replenishment (or equivalently loss) rate. For example, the energy confinement time for a system in steady state, with heating power (and loss power) P and internal energy W , is $\tau_E = W/P$. The well-known Lawson criterion for ignition in fusion reactors [2] is that the product $nT\tau_E$ must exceed a certain threshold value, where n is density and T is temperature. The pressure $p = nT$ is bounded by engineering constraints, so this imposes a lower limit on τ_E . Plasma turbulence can be thought of as small-scale convection that adds to the collisional conduction to produce a much higher overall conductivity at large scales, reducing the confinement time. With a better understanding of this so-called anomalous transport, it should be possible to improve confinement in devices.

Plasmas consist of electrically charged particles, as a result of which electromagnetic fields play an important role in their dynamics, and the evolution of the plasma and the electromagnetic fields are interdependent. (Maxwell's equations are coupled to the plasma equations.) In a sufficiently strong magnetic field, the Lorentz force dominates charged particle motion perpendicular to the field, so that this motion is circular to lowest order, the so-called Larmor or cyclotron motion. Cross-field motion on scales larger than these Larmor orbits is therefore inhibited, leading to the concept of a 'magnetic bottle' for holding the hot plasma away from the solid (and therefore cold) vacuum vessel walls. Motion parallel to a straight magnetic field is unaffected by the Lorentz force and therefore not inhibited in this way, but the tokamak (Russian for 'toroidal magnetic chamber') attempts to circumvent this problem by means of a toroidal topology in which one 'end' of the 'bottle' joins back onto the other. The magnetic field is axisymmetric, and may be written in

terms of the poloidal flux function ψ , labelling nested flux surfaces [3]:

$$\mathbf{B} = I(\psi)\nabla\zeta + \nabla\psi \times \nabla\zeta \quad (1.1)$$

where ζ is the toroidal angle. The first term is the toroidal component and the second term the poloidal component of the magnetic field. $I(\psi) = B_T R$ is constant on each flux surface, where B_T is the toroidal magnetic field and R the major radius coordinate; $R\nabla\zeta$ is the unit vector in the toroidal direction.

The Mega Ampere Spherical Tokamak (MAST) is a plasma experiment at the Culham Centre for Fusion Energy (CCFE) near Oxford. Spherical tokamaks, also called spherical tori, and hereafter STs, differ from conventional ones in that the toroidal topology is achieved with a reduced aspect-ratio geometry: more like a cored apple than an onion-ring (Figure 1.1).




Figure 1.1: MAST plasma (courtesy B. Lloyd and MAST Team, CCFE)

Very high temperatures are required for fusion, similar to those at the centre of the Sun. (Deuterium-tritium fusion peaks above 10 keV [4]. Although MAST is not equipped to handle tritium, which is radioactive, nor to cope with the heating that would result from significant fusion, it is designed to operate with deuterium plasmas under fusion-relevant conditions.) The aforementioned constraints on pressure in terrestrial devices then imply a very low plasma density, one that in other contexts

would be regarded as a good vacuum. Both the low density and the high temperature reduce the role of the Coulomb collisions between the charged particles, and this means that fusion plasmas are far away from the collisional continuum (fluid) limit. Kinetic theory is still a continuum theory, but expressed in terms of a probability distribution function over phase space, including velocity as well as real-space coordinates. For magnetized plasmas, gyrokinetics is a simplified kinetic theory valid in certain limits, that averages over the Larmor orbits and thereby reduces the phase space from six dimensions to five. This is one of the ways in which it makes simulation more tractable.

The following chapter introduces gyrokinetics in a little more detail, and includes some other relevant background physics. But first, we conclude this initial overview by summarising the overall structure of the work contained in this thesis.

In chapter 3, we model a MAST plasma using a multi-scale framework in which the plasma turbulence occurs on space and time scales much smaller than those of the large-scale transport. This scale-separation permits a large fraction of the whole device to be simulated, yet still in sufficient detail, a feat that would be computationally intractable without such scale-separation. Features of the plasma allow us to consider electrons and ions essentially independently, and to concentrate on electron turbulence: we assume that ion turbulence is suppressed by background flow shear, so we model only electron gyroradius scales; and we exclude magnetic perturbations. To our knowledge this is the first such multi-scale simulation of an ST or at electron scales. We find that the experimental level of transport is consistent with our simplified modelling.

We also find that long transients can occur in the turbulence simulations, due to slowly increasing zonal modes; that is, modes with only radial small-scale variation. We introduce this phenomenon in a short chapter 4.

In chapter 5, we simplify our model of the electron physics further, including only the adiabatic response of the ions, and investigate in more detail the role of

zonal modes. We consider a single flux surface in this chapter and therefore perform only the turbulence modelling (small-scale part of the multi-scale framework): this tells us the heat flux at that flux surface; the temperature and temperature gradient are given. We find that the zonal modes are weakly damped by electron-ion collisions, and themselves regulate the level of turbulence and the heat flux, which is due to nonzonal modes. As a result, the heat flux increases roughly linearly with collisionality, which is in agreement with previously observed experimental results. We sketch a theoretical explanation of this scaling.

Each of chapters 3 and 5 is written in the style of a paper, with its own self-contained introductory material applicable to that chapter. Chapter 3 is joint work with M. A. Barnes, C. M. Roach and A. A. Schekochihin. Chapter 5 and the appendices at the end of the thesis are joint work with A. A. Schekochihin, C. M. Roach, M. A. Barnes, Y.-c. Ghim, W. Dorland and F. I. Parra.

Finally in chapter 6, overall conclusions are drawn and some pointers towards further work are considered.

CHAPTER 2

Introduction

2.1 Introduction to gyrokinetics

We start from the general kinetic equation

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \frac{\partial f_s}{\partial \mathbf{x}} + \mathbf{a}_s \cdot \frac{\partial f_s}{\partial \mathbf{v}} = C_s(f), \quad (2.1)$$

which describes the evolution of the probability distribution function $f_s(\mathbf{x}, \mathbf{v}, t)$, where \mathbf{x} is position, \mathbf{v} is velocity, and t is time, subject to ‘external’ accelerations \mathbf{a}_s . Each species has its own distribution function indicated by the subscript s . The ‘collision’ operator on the right-hand side represents ‘internal’ interactions both within and between species, so in general C_s depends on $f_{s'}$ for all s' ; hence the omission of a subscript on f in its argument.

In a plasma, the species consist of: electrons, one or more ion species, and zero or more neutral species. We will consider here only simple, two-species plasmas (electrons and ions). The acceleration will be given by the Lorentz force law

$$\mathbf{F}_s = m_s \mathbf{a}_s = q_s \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right), \quad (2.2)$$

where \mathbf{F}_s is the force, m_s is the mass, q_s is the electric charge, c is the speed of light, and \mathbf{E} and \mathbf{B} are the ensemble-averaged electric and magnetic fields respectively. The collision operator will then model the ensemble-averaged effect of the residual Coulomb interactions between individual particles. In this way, the kinetic equation

describes the evolution of an ensemble of macroscopically identical plasmas, and f is an ensemble-averaged distribution function.

The kinetic equations are closed by coupling them to Maxwell's equations for the ensemble-averaged fields.¹ These equations include integrals of f as source terms, namely the total charge and current densities:

$$\rho = \sum_s q_s \int d^3\mathbf{v} f_s, \quad \mathbf{j} = \sum_s q_s \int d^3\mathbf{v} f_s \mathbf{v}. \quad (2.3)$$

Kinetic theory is a simplified description, but f still covers a six-dimensional phase space and evolves in time, and solving a full kinetic system is typically still too resource-intensive for today's computers. One approach to further simplification is suggested by the fact that Maxwell's equations involve only the above first and second velocity moments of f . The idea is to take velocity moments of the kinetic equation and rewrite the theory in terms of these, which are functions over real space only. This leads to the so-called hierarchy problem: in general, the resulting set of equations is infinite because the equation for the n th moment involves other moments including the $(n+1)$ th. Fluid theories deal with this problem by assuming certain orderings that allow the hierarchy to be truncated in some way; the typical one is the collisional limit (used for example to obtain the Navier-Stokes equation from the kinetic equations for neutral fluids), in which the collisional mean free path is smaller than the other length scales of interest. However, in plasmas, including fusion plasmas, the collisional mean free path can be longer than macroscopic lengths such as the device size: the plasma is approaching the collisionless rather than the collisional limit. In these circumstances, processes such as phase mixing

¹Using the exact physical (rather than ensemble-averaged) electromagnetic fields in (2.2), and still requiring (2.3) to relate them self-consistently to f , would do away with the collision term in (2.1); but solving this system would then be equivalent to solving the coupled equations of motion of all the individual particles, and this is intractable. Equivalently, one can derive (2.1) (without the collision term) from the individual particle equations of motion (2.2) (using the exact physical electromagnetic fields), by taking f to be a sum of δ -functions, one for each particle [5]. If, however, f is a smoothed one-particle distribution function, as is the case here, then the collision term incorporates the effect of two-particle correlations, which are not captured by such f .

can give rise to small-scale structure in the velocity-dependence of f , which is not captured by a small number of moments.

Gyrokinetics is another approach to simplifying kinetic theory. Abel et al. [3] provide a systematic exposition of gyrokinetics in tokamaks. More detail will be found in the present thesis in section 5.2. In gyrokinetics, the distribution function is still a function of both real and velocity space, but the dimensionality is reduced from six to five. The gyrokinetic formalism is based on an expansion in powers of $\rho_* = \rho/L$, the ratio of the gyroradius ρ to the background gradient scale length L , and this parameter is assumed to be small. It is also assumed that the perturbed distribution function δf is small, that the fluctuation frequencies ω associated with the perturbations are small compared to the gyrofrequency Ω , and the parallel gradients are small compared to the perpendicular gradients (or wavenumbers, k_{\parallel} and k_{\perp} respectively):

$$\frac{\omega}{\Omega} \sim \frac{k_{\parallel}}{k_{\perp}} \sim \frac{\delta f}{f} \sim \frac{\rho}{L} = \rho_*. \quad (2.4)$$

With these orderings, to lowest order the Larmor orbits are circular and the fields




Figure 2.1: Gyroaveraging: the gyrokinetic distribution function describes rings that experience the ring-averaged fields. The black arrows indicate the position vectors of the particle, \mathbf{r} , and of the gyrocentre, \mathbf{R} , relative to the same (arbitrary) origin.

do not change during the period of a single orbit. It is then possible to integrate out the dependence of the perturbed distribution function δf upon the gyrophase angle φ around the circle, leaving only two remaining dimensions of velocity space. Figure 2.1 illustrates this gyroaveraging process. $\delta f_s(\mathbf{r}, v_{\parallel}, v_{\perp}, \varphi)$ is replaced by $g_s = h_s(\mathbf{R}, v_{\parallel}, v_{\perp}) - q_s \langle \phi \rangle F_{Ms}/T_s$, where $\mathbf{r} = \mathbf{R} + \boldsymbol{\rho}_s$, $\boldsymbol{\rho}_s = \mathbf{b} \times \mathbf{v}/\Omega_s$ is the vector gyroradius for species s , and \mathbf{b} is the unit vector in the direction of \mathbf{B} . h_s is referred to as the ring distribution function.² The second term in g_s is the Boltzmann response to the electrostatic potential $\phi(\mathbf{r})$, where $\langle \cdot \rangle$ denotes the gyroaverage at constant \mathbf{R} , F_{Ms} is the lowest-order Maxwellian background distribution function ($f_s = F_{Ms} + \delta f_s$), and T_s is the species temperature. Furthermore, the orderings imply that time resolution is not required down to the gyrofrequency scale, spatial resolution perpendicular to the magnetic field is not required up to background scales, nor is spatial resolution required far below the gyroradius scale. Together, these improvements make solution just tractable on today's (highest performance) computers.

The first-order corrections to the circular Larmor orbits are the so-called drift motions which appear in the gyrokinetic equation (see section 5.2.1). These are due to inhomogeneity in the magnetic field (the magnetic drifts) and to perpendicular electric fields (the $\mathbf{E} \times \mathbf{B}$ drift). The magnetic drifts (both curvature and ∇B drifts) are enhanced in STs owing to the small aspect ratio. In turbulence dominated by the $\mathbf{E} \times \mathbf{B}$ drift, plasma circulates perpendicular to the magnetic field along contours of constant electrostatic potential. The region of clockwise or anticlockwise circulation around each extremum of the fluctuating potential is referred to informally as an eddy.

Background equilibrium gradients (such as in temperature) can drive insta-

²Since $v_{\parallel} \sim v_{\perp} \sim v_t = \sqrt{2T/m}$ (the thermal speed), the 'rings' correspond to particle trajectories that are actually helices in true 3d space. However, because $k_{\parallel}/k_{\perp} \sim \rho_*$ is small, the ring-average is equal to the 'helix-average' at leading order. Equivalently, the helices look like rings if one imagines that space is compressed in the parallel direction so as to make the scales of parallel and perpendicular variation look similar.

bilities that cause these fluctuations to grow. However, on length scales much smaller than the gyroradius they are suppressed by the gyroaveraging: gyroaveraged quantities are multiplied by Bessel functions $J_0(k_\perp \rho)$, arising from factors like $\langle \exp(i\mathbf{k} \cdot \mathbf{r}) \rangle_{\mathbf{R}} = \exp(i\mathbf{k} \cdot \mathbf{R}) J_0(k_\perp \rho)$ in the Fourier-space representation of the gyroaverage, so they vanish when $k_\perp \rho \gg 1$. On length scales larger than the gyroradius, however, the fastest-growing modes have the highest k_\perp , since the growth rate $\gamma \sim \omega \sim (\rho_* v_t) k_\perp$ (by ordering). Saturated gyrokinetic turbulence might therefore be expected to be dominated by fluctuations in the region $k_\perp \rho \sim 1$.

Within the gyrokinetic ordering, which has already eliminated the fastest (gyrofrequency) time scales, the drift fluctuations are now the fastest in time and shortest in length. At second order in the gyrokinetic expansion appear the transport equations governing the interaction between these and the slower, longer variations in background fields (see the next section and section 3.2).

Further detail about gyrokinetics and its application to the present work will be found in section 5.2. This introduces the gyrokinetic equation (section 5.2.1), the simplified modelling of ions (section 5.2.2) that is used in chapter 5, and the gyrokinetic collision operator (section 5.2.3).

2.2 Self-consistent transport modelling

The Trinity code [6, 7] solves the 1-d transport equations that are derived at second order in the gyrokinetic ordering, self-consistently with the first-order equations solved by the GS2 code [8, 9]. More detail is given in section 3.2. The background magnetic field topology consists of nested magnetic flux surfaces, and we are interested in cross-field transport in the radial direction orthogonal to these surfaces. In the present work, only energy transport is modelled; density and momentum (rotation) profiles have been fixed at their experimental values. This is a reasonable first step, because we will be looking to verify that the modelled energy transport

agrees with experiment. If it does, then it is also consistent with experiment to use the other experimental profiles and the experimental magnetic equilibrium. If the energy transport does not agree, then the simulation may not be self-consistent, but one can still (therefore) conclude that the energy transport model does not adequately explain the experiment.

2.3 Rotation and flow shear

Plasma rotation has an important effect on gyrokinetic instabilities and hence on confinement, especially in small devices with neutral beam injection (NBI) such as MAST. NBI is intended mainly for heating: high-energy neutral atoms are directed towards the centre of the plasma, where they ionize and become confined. Clearly NBI acts as a particle source and an energy source, but it also acts as a momentum source, injecting toroidal angular momentum into the plasma and causing it to rotate. In fact, equilibrium flows are nearly purely toroidal [3], because friction with trapped particles (which are not free to circulate poloidally) damps poloidal flows in the passing population. The effect is enhanced in STs, which have a low moment of inertia and a high trapped particle fraction.

If the equilibrium magnetic field has non-zero components in both the toroidal and poloidal directions (in other words it is neither purely toroidal nor purely poloidal), then a toroidal equilibrium flow has non-zero components both parallel and perpendicular to this field. The perpendicular component is produced by drift motion, primarily the $\mathbf{E} \times \mathbf{B}$ drift. Considered in the frame co-rotating with a given flux surface, the perpendicular flow shear, denoted γ_E , is proportional to the radial gradient of the angular velocity of the plasma, Ω (not to be confused with the cyclotron frequency), and inversely proportional to the magnetic safety factor q [10, 11]:

$$\gamma_E = \frac{r}{q} \frac{d\Omega}{dr} \quad (2.5)$$

Flow shear can suppress microturbulence via an effect similar to phase mixing (which is the shearing of structure in phase space such that velocity moments in real space are damped as $t \rightarrow \infty$). Consider a solution that is a single spatial Fourier mode (in the perpendicular direction) at any particular time; a general solution is a superposition of such modes. The mode is characterized by plane waves with wavenumber $\mathbf{k}_\perp = (k_x, k_y)$. A constant perpendicular flow shear (which means a shear in the x direction of a flow in the y direction, where x is the radial coordinate and y the other coordinate perpendicular to the magnetic field) will cause the wavefronts to tilt progressively in time, such that the k_x of the solution evolves according to

$$k_x(t) = k_x(0) + \gamma_E k_y t \quad (2.6)$$

Thus, the $k_\perp = \sqrt{k_x^2 + k_y^2}$ of such a solution grows unboundedly as $t \rightarrow \pm\infty$, and its drive is suppressed via the Bessel functions $J_0(k_\perp \rho)$ except during the time window when k_x is small. If $\gamma_E \gg \gamma$ (the growth rate), then during this window the mode will not have time to grow significantly, and if this is true for all modes then turbulence should be quenched. The Waltz rule [12], which is often used in ad-hoc models of the nonlinear saturation level, hypothesizes a stronger condition: that $\gamma_E \gtrsim \gamma_{\max}$ is sufficient for suppression of turbulence.

Roach et al. [13] and Field et al. [14, 15] found that perpendicular flow shear stabilized low- k_y modes, in the MAST plasmas they considered, when γ_E exceeded the maximum linear growth rate, in accordance with the Waltz rule. (Roach et al. [13] also found, for plasma parameters based on the Cyclone reference case [16], that at even higher levels of γ_E the destabilizing parallel flow shear eventually overcame the perpendicular shear stabilization, but this phenomenon has not yet been obviously manifested in MAST simulations.) In the present work we consider a MAST shot (number 8500) for which we show linear stabilization of low- k_y modes across most of the plasma, and it is therefore expected that neoclassical transport

accounts for most of the ion heat flux. We wish to model a shot for which this is the case across a broad radial range, not confined to a narrow ‘internal transport barrier’ (ITB). (Field et al. [14, 15] were investigating ITBs and so their shot selection was made on that basis.) Transport analysis of the experimental data for 8500 [17] found $\chi_i \sim \chi_{nc}$ in $r/a \gtrsim 0.5$, where r is the minor radius, a the minor radius of the last closed flux surface (LCFS), χ_i the ion heat diffusivity, and χ_{nc} the neoclassical ion heat diffusivity. Transport modelling using the TGLF and XPTOR codes [18] found $\chi_i \sim \chi_{nc}$ in $r/a \lesssim 0.7$, and that the steady-state profiles were sensitive to the details of the neoclassical model. Though not in perfect agreement, these results are broadly consistent with low- k_y suppression by flow shear: the anomalous energy transport is predominantly carried by electrons. The present work is the first to carry out a fully first-principles coupled gyrokinetic and transport calculation for a MAST plasma.

2.4 Electromagnetic modes

Gyrokinetic calculations are often done in the ‘electrostatic’ limit, meaning that there are no fluctuations in the magnetic field. This is to ignore the magnetic field arising from fluctuating currents in the plasma: fluctuations in the parallel motion generate $\delta\mathbf{B}_\perp$ (usually represented by δA_\parallel) and fluctuations in the perpendicular motion generate δB_\parallel . Inclusion of these perturbations (adding Ampère’s law to the dynamical system) makes the calculation electromagnetic, and can be important in STs, for which the plasma β (ratio of plasma pressure to magnetic field pressure) can be higher than in conventional tokamaks.

Modes present electrostatically can persist when electromagnetic effects are included [19]. There are also distinctively electromagnetic modes: notably ‘microtearing’ modes, which have odd parity in electrostatic potential and even parity in A_\parallel , along the field line.

Applegate et al. [20] performed electromagnetic simulations of MAST shot 6252, assuming $T_i = T_e$, and found unstable microtearing modes, especially on the inner flux surfaces, that were highly extended in the parallel direction. This makes resolved simulation very challenging. Applegate et al. [21] attribute this prevalence of microtearing in STs to the higher plasma β (which is highest in the core, where the pressure is highest) and higher curvature drift. Roach et al. [22] performed electromagnetic simulations of MAST shot 8500 – which is the shot that we consider throughout the present work – and also found unstable microtearing modes except on the outermost flux surface. These microtearing modes were all at ion scales,³ and had lower growth rates than the experimental level of flow shear (which was not included in the simulations). These results therefore help to motivate our simplification to electrostatic simulation, of this shot at electron scales.

Applegate et al. [20] and Roach et al. [22] compared electromagnetic with electrostatic simulations for inner flux surfaces, where they found that including electromagnetic effects decreased the growth rates of the electrostatic electron-scale modes (but did not stabilize them completely). Applegate et al. [20] also performed such a comparison for an outer flux surface, where they found essentially no difference between electrostatic and electromagnetic modes at electron scales. This is also supporting evidence that, at electron scales, whilst electromagnetic effects can have a quantitative effect on growth rates for inner flux surfaces, they may not be qualitatively significant in the cases that we consider later.

³The term microtearing is typically used to refer to such scales, with $k_y \rho_i \lesssim 1$, where ρ_i is the ion gyroradius. More generally, including electron scales, the term tearing-parity may be used. It was only for a different equilibrium with a much higher β than shot 8500, that Roach et al. [22] found tearing-parity instabilities (and no electrostatic instabilities) at electron scales as well as ion scales, for the mid-radius flux surface they considered.

2.5 The GS2 code

GS2 solves the local gyrokinetic equation as an initial value problem. It supports a variety of geometries for the background magnetic field. In all the work described here, we have used the ‘Miller’ local parameterization of toroidal flux-surface geometry [23]. Each flux surface is assumed to be axisymmetric (independent of toroidal angle), and its poloidal cross-section is described by the following expressions for the cylindrical coordinates R and Z (relative to the toroidal symmetry axis), with the poloidal angle θ as a parameter:

$$\begin{aligned} R &= R_0 + r \cos [\theta + (\sin^{-1} \delta) \sin \theta] \\ Z &= \kappa r \sin \theta \end{aligned}$$

Radial profiles of the best-fit values of the Miller parameters – the major radius R_0 , the minor radius r , the triangularity δ , and the elongation κ – are obtained with the experimental data, having been calculated during the magnetic equilibrium reconstruction. Their radial derivatives are also required. In later chapters the cylindrical coordinates are not used, and R_0 will be denoted simply by R .

GS2 may also be operated using a numerically specified equilibrium flux-surface geometry. As part of the code comparison undertaken for Roach et al. [13] we found (at $\rho_n = 0.6$) that Miller geometry produced linear growth rates deviating from the numerical equilibrium result by no more than a factor of 1.4 across the range $0.1 \leq k_y \rho_i \leq 36$ (Figure 2.2).

For the nonlinear simulations necessary to calculate energy fluxes, GS2 is operated in flux-tube mode. The simulation domain – the flux tube – winds around a particular flux surface following the equilibrium magnetic field. It is long in the parallel direction, covering a change in poloidal angle from $-\pi$ to π . At a poloidal angle of zero, where it crosses the outboard midplane, the flux tube has a rectangular cross-section that is assumed to be small: flux-surface quantities such as density,

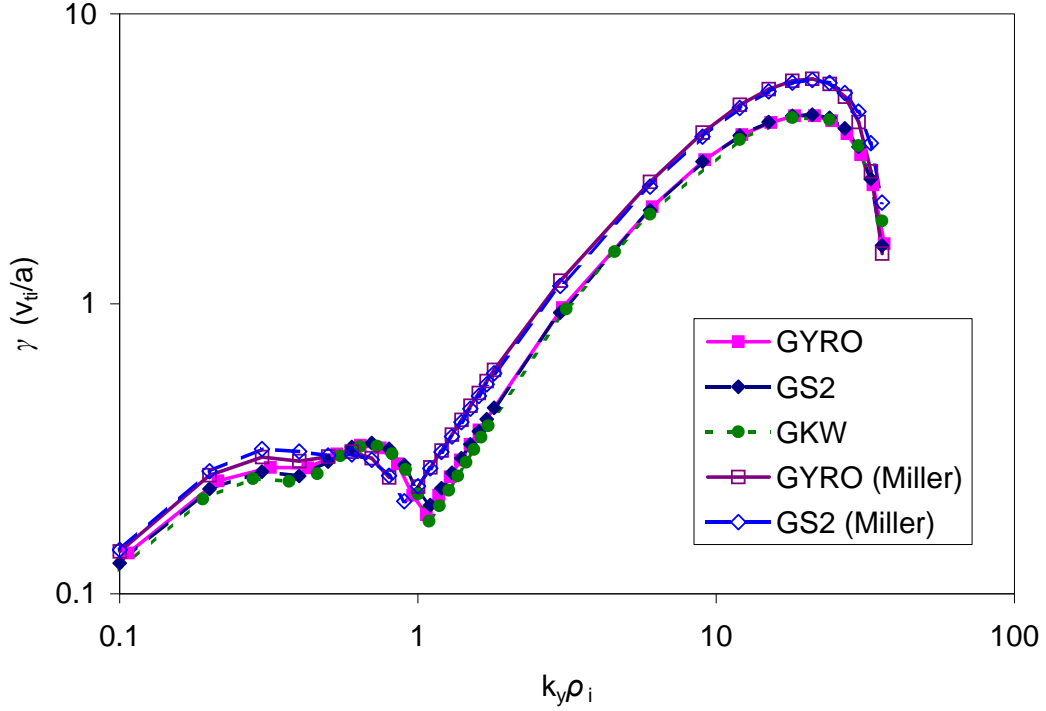


Figure 2.2: Linear growth rates calculated with various codes (named in legend) using numerical geometry (first three curves) and Miller geometry (last two curves) for the collisionless, electrostatic benchmark case (based on MAST shot 8500 at $\rho_n = 0.6$) used in Roach et al. [13] (data courtesy W. Guttenfelder and Y. Camenen). (Only the numerical geometry curves were published in [13].)

temperature and the Miller parameters are represented everywhere in the flux tube by the value and the (radial) gradient at the centre of this rectangle. This is what is meant by ‘local’, and it requires that the perpendicular box size L should be small compared to the gradient scale lengths L_X of these quantities (for all $X = n_s, T_s$, etc.) In order to capture the turbulent structure around the gyroradius scale, the box must also be larger than the species gyroradius ρ_s . Thus we have the double inequality for the validity of the flux-tube approach, $\rho_s \ll L \ll L_X$. In particular, this then implies $\rho_{*s} \equiv \rho_s/L_X \lll 1$, where \lll expresses the requirement that an intermediate scale must fit between the two, separated by \ll on both sides. In MAST, the relatively low magnetic field means that $\rho_{*i} \gtrsim 1/100$ can occur even when L_X is no smaller than the overall size of the plasma (i.e. than the minor radius of the

last closed flux surface). Thus, the flux tube approximation is not robust for ions. However, it is reasonable for electrons, provided the electrons are not significantly coupled to the ion scales, since $\rho_i \approx 60\rho_e$ for deuterium ions.

The flux tube covers all poloidal angles, but only a small fraction, say $1/n_0$, of 2π in toroidal angle. The entire flux surface can be considered to be modelled by n_0 periodic copies of the flux tube, with a toroidal angle of $2\pi/n_0$ between each and the next, so we are restricting ourselves to toroidal mode numbers n that are multiples of n_0 . (The above flux tube approximation implies toroidal mode numbers $n \gg 1$. In practice we are interested in statistical properties that do not require us actually to assemble the whole flux surface.) In accordance with this idea, GS2 uses periodic boundary conditions in y , and also in x . Along the field line, ‘twist-and-shift’ boundary conditions are typically used at the ends of the flux tube [24]. The amplitude of the (j, k) th mode at one end of the flux tube (where j and k are integers labelling the k_x and k_y grid points) is set equal to the amplitude of the $(j + Jk, k)$ th mode at the other end, where J is an integer. J determines the aspect ratio of the flux-tube cross-section at the outboard midplane $\theta = 0$, given the magnetic shear \hat{s} :

$$JL_y/L_x = 2\pi N\hat{s}, \quad (2.7)$$

where L_x and L_y denote the perpendicular box sizes in x and y respectively, and $2\pi N$ is the length of the flux tube in poloidal angle.⁴

⁴We simplified to a single box size L in the description above. We also tacitly assumed a restriction that is present in the GS2 code itself, namely that J/N (which is the input parameter `jt`), and not merely J (which is called `js`), be an integer. In GS2, N is also always an *odd* integer: the ends of the flux tube, where the parallel boundary conditions are applied, are always at the inboard midplane.

CHAPTER 3

Simulating global energy transport in MAST using local electron-scale flux tubes

3.1 Introduction

In this chapter we consider the applicability to the MAST experiment of a transport model in which the ion thermal energy transport is dominated by the neoclassical contribution and the electron thermal energy transport is obtained from local nonlinear gyrokinetic simulations across the plasma. This approach is motivated by both physical and computational considerations. MAST has tangential neutral beam injection (NBI); there is experimental evidence that NBI-heated MAST plasmas have high levels of sheared toroidal flow (with flow velocities of the order of the ion thermal speed), and that anomalous ion transport in MAST can be substantially suppressed [14, 15]. Previous gyrokinetic simulations [13] support the paradigm that flow shear can stabilize otherwise unstable linear modes at ion gyroradius scales. Nonlinear gyrokinetic simulations spanning both ion and electron gyroradius scales are very computationally intensive, even at artificially reduced mass ratio m_i/m_e [25], and it is not yet feasible to run many such simulations in parallel, as would be required for incorporation within a global transport code. Local gyrokinetics at the

ion gyroradius scale may also be problematic for STs because there is only weak separation between this scale and the background gradient length scales. However, if because of flow-shear suppression ion scales need not be modelled gyrokinetically, then there is the possibility to run gyrokinetic simulations at electron scales only, for which the local approximation is much more robust, and it becomes feasible to incorporate such simulations within a transport code.

In this chapter we apply the Trinity multiscale transport framework [6], using GS2 as its embedded flux code, to MAST discharge number 8500, which is held in the International Multi-Tokamak Confinement Profile Database [26].

In previously published work [7] Trinity has been used to study the JET and ASDEX tokamaks; in these cases only ion scales were modelled gyrokinetically, and electron-scale anomalous transport was taken to be negligible. A similar study of the DIII-D tokamak has been undertaken using the TGYRO code coupled to GYRO [27].

Individual flux surfaces of MAST plasmas have been modelled nonlinearly at electron-gyroradius scales using GS2 by Joiner et al. [28] and by Roach et al. [13], who found saturated fluxes consistent with the experimental level of electron heat transport. Using an adiabatic ion model, runaway growth in the heat flux was observed at large times but was found to be healed by including either flow shear or collisions; it was believed that this is due to the stabilization of weak trapped-electron modes at $k_y \rho_i \sim 1$. Roach et al. [13] found that in their simulation kinetic ions without collisions only delayed the runaway rather than eliminating it. In the present chapter, we include both kinetic ions and collisions, and do not find such runaway growth. We usually take the flow shear to be zero, as it is not expected to affect the energy transport significantly at electron scales.

Our gyrokinetic simulations are done in the ‘electrostatic’ limit, meaning that there are no fluctuations in the magnetic field. It is anticipated that in future work the effect of magnetic field fluctuations would need to be included: allowing $\delta\mathbf{B}$

can introduce additional instability, especially in the core [20, 21, 22], and can also affect the growth rate of modes present electrostatically. It will be seen below that indeed our present approach does fail towards the centre of the plasma.

Further details of the transport solution method (3.2) and of the experimental data (3.3) are given in the next two sections. We then present our simulation setup including the results of linear studies (section 3.4), followed by our results (in sections 3.5 and 3.6). At the end of this chapter we draw conclusions and point to some possible further work (section 3.7).

3.2 Transport solution method

The Trinity transport framework [6] is based on a separation of dynamical scales between mean fields and fluctuations. The transport equations, which govern the evolution of macroscopic mean plasma quantities such as density and pressure, operate on device-size length scales and confinement time scales. The nonlinear gyrokinetic equations, which govern the dynamics of fluctuating particle distribution functions and electromagnetic fields, operate on gyroradius length scales and drift-wave time scales. These scales are self-consistently coupled via averaging on an intermediate scale. The Trinity framework is implemented as an open-source Fortran code of the same name, which calls the gyrokinetic code GS2 (as a library) at each point in its coarse space-time grid.¹ Because of the scale separation, GS2's finer space-time (and velocity space) grid is point-like on the transport scales, and time- and space-averaged outputs over the entire GS2 domain provide the coupling terms in the transport equations solved by Trinity. The energy transport equation for each species $s = i, e$ is as follows:

$$\frac{3}{2} \frac{\partial p_s}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial \psi} (V' Q_s) + \mathcal{E}_s = S_s - \frac{3}{2} \left(\frac{\partial p_s}{\partial t} \right)_{\text{exp}}. \quad (3.1)$$

¹Both codes are hosted at the same Subversion repository; see <http://gyrokinetics.sourceforge.net>.

Here all quantities are functions only of the flux-surface label ψ – the poloidal magnetic flux, which is the radial coordinate – and time. $p_s = n_s T_s$ is the scalar pressure, n_s is the density, T_s is the temperature (in energy units), $V(\psi)$ is the volume enclosed by the flux surface, $V' \equiv dV/d\psi$, and Q_s is the energy flux including both turbulent and neoclassical contributions. The collisional equilibration term

$$\mathcal{E}_s = \frac{3}{2} n_s \sum_{s'} \nu_{ss'}^\varepsilon (T_s - T_{s'}) \quad (3.2)$$

exchanges energy between species (its species sum is zero: $\sum_s \mathcal{E}_s = 0$), where

$$\nu_{ss'}^\varepsilon = \frac{8\sqrt{2\pi} (m_s m_{s'})^{\frac{1}{2}} q_s^2 q_{s'}^2 n_{s'} \ln \Lambda_{ss'}}{3 (m_s T_{s'} + m_{s'} T_s)^{\frac{3}{2}}}, \quad (3.3)$$

the $m_{s,s'}$ are species masses, the $q_{s,s'}$ are species charges, and $\ln \Lambda_{ss'}$ is the Coulomb logarithm.

S_s is the prescribed energy source; it is in principle a function of ψ and t , but for the work described here we are looking at a single snapshot in time for which the profiles (including the external sources) are therefore functions of ψ only. However, the experimental $(\partial p_s / \partial t)_{\text{exp}}$ is also available; in the modelling we subtract this from S_s to create a new effective source (the right-hand side of (3.1)), and then aim to evolve the system to steady-state so that the modelled $\partial p_s / \partial t$ (on the left-hand side) vanishes. The modelled steady-state temperature profile will then agree with the experimental profile obtained at the snapshot time if the energy flux divergence + collisional exchange profile agrees; that is, by comparing the temperature profiles we see whether our transport model accounts for the experimental level of transport.

S_s should in principle include a contribution \mathcal{S}_s due to turbulent heating (averaged over the intermediate length and time scales) [6]. Like the collisional equilibration term, the turbulent heating term exchanges energy between species: $\sum_s \mathcal{S}_s = 0$. The turbulent heating consists of the dissipation of turbulence due to collisions (which is a positive-definite heating contribution) minus the turbulent drive due to instabilities (which is generally cooling): the average sink and source respectively

of the turbulent energy cascade; its species sum is zero because the turbulent cascade is local to each flux surface [3]. Although \mathcal{S}_s is not constrained to be zero for each species individually, we expect it to be small when the species interact only weakly, and it must be small when, as we will ultimately find here, the turbulence is dominated by electrons (because the species sum is zero, and in such a case the ion contribution to the sum is small). From a computational perspective, inclusion of the turbulent heating term makes it harder to find a steady state, because its fluctuation level can be large even though its time-average is small. So we defer any further consideration of turbulent heating to further work, setting $\mathcal{S}_s = 0$ here.

As our boundary conditions we fix T_s to their experimental values at the outermost flux surface of the radial transport grid, and we set the fluxes $Q_s = 0$ at the magnetic axis $\psi = 0$.

3.3 Experimental data




Figure 3.1: Radial profiles of (a) n_e and (b) q at $t = 0.289$ s for MAST shot 8500

MAST shot 8500 was previously used for linear gyrokinetic stability studies [22] and as the basis for a comparison of several different gyrokinetic codes [13]. The Profile Database [26] data for this shot are provided at a single time-slice, $t = 0.289$ s (measured from the nominal start of the shot). At this time the plasma is in H-mode, with an edge pedestal temperature of about 0.1 keV, central ion and elec-

tron temperatures of about 1.2 keV, and a central electron density of $4 \times 10^{19} \text{ m}^{-3}$. Full radial profiles of the electron density n_e and of the safety factor q are shown in Figure 3.1.

The sources in the database are taken from a transport analysis of the experimental data [17]. This analysis found that $\chi_i \sim \chi_{\text{nc}}$ in $r/a \gtrsim 0.5$, where r is the minor radius and a the minor radius of the last closed flux surface (LCFS). A transport simulation for this shot using the TGLF model [18] found $\chi_i \sim \chi_{\text{nc}}$ in $r/a \lesssim 0.7$, and that the steady-state profiles were sensitive to the details of the neo-classical model. Both of these previous results are broadly consistent with low- k_y suppression by flow shear and with anomalous transport predominantly carried by electrons, which makes the shot a good experimental candidate for our approach.

3.4 Simulation setup

Initially we carried out linear simulations without flow shear, using the Miller [23] equilibrium parameters indicated in Table 3.1. We looked at the two outer flux surfaces considered in [22], and two new surfaces closer to the magnetic axis. Figure 3.2 shows the linear growth rates at these four flux surfaces, and indicates that the ion-scale modes have positive growth rates only on the outer two surfaces of the four.

ρ_n	r/a	R/a	δ	κ	\hat{s}
0.2	0.248	1.666	0.144	1.48	-0.58
0.4	0.481	1.626	0.186	1.52	0.76
0.6	0.693	1.577	0.223	1.59	2.05
0.8	0.867	1.521	0.268	1.70	3.67

Table 3.1: Some flux surface parameters for MAST shot 8500. ρ_n is the square-root of the toroidal flux enclosed by the flux surface, normalized to $\rho_n = 1$ at the last closed flux surface (LCFS). $a = 0.550 \text{ m}$ is the minor radius (half the midplane diameter) of the LCFS.

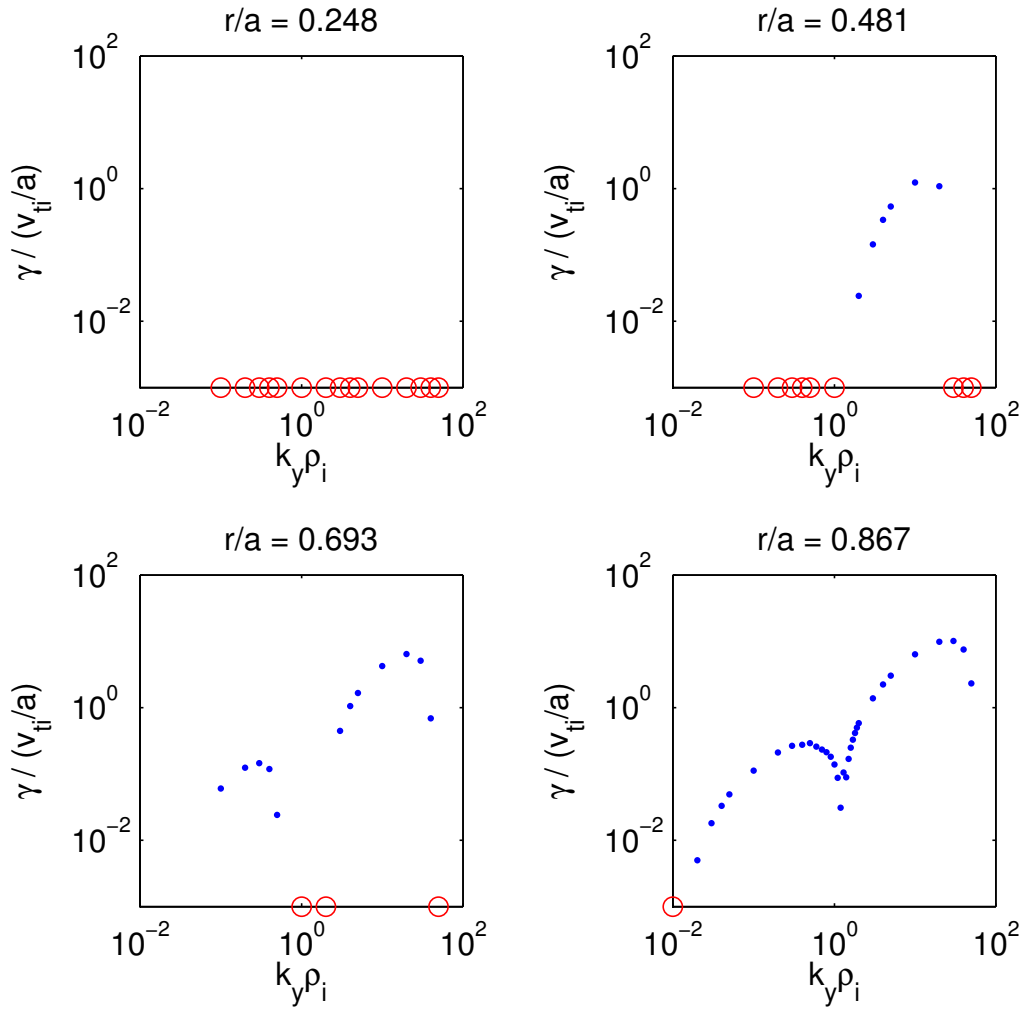


Figure 3.2: Linear growth rate spectra without flow shear at the four flux surfaces specified in Table 3.1. Circles indicate stability.

We also carried out linear simulations with flow shear, which verified – as has been confirmed for other MAST shots [13, 14, 15] – that the positive growth rates for ion-scale modes found without flow shear are then eliminated. This does not preclude transient initial growth of finite duration, which we found may be important nonlinearly in the outermost region (at the $r/a = 0.867$ flux surface).

These results were used to select the code parameters for the full simulation. In particular, $k_y \rho_i = 1$ was chosen as the lowest nonzero k_y mode (corresponding to the box size in y) at the bottom end of the upper (electron scale) branch, and $k_y \rho_i = 17$ as the highest mode at approximately the peak growth rate of that branch.

This range excludes the ion-scale modes, and zero flow shear was used for the full simulation. Note that even at the low- k_y cutoff here, $k_y \rho_i = 1$, we are still at a much smaller scale (higher wavenumber) than those at the low- k_y end of the ion-scale spectrum, which must be captured in a full nonlinear simulation of ion-scale modes, and which are correspondingly closer to the background gradient length scales.




Figure 3.3: Schematic showing four of the nested flux surfaces used for the transport calculation. Each square represents (not to scale – formally smaller in the multiscale expansion) the cross-section at the outboard midplane through the corresponding flux tube.

For the full simulation 8 flux surfaces were used, with a Miller equilibrium representation and $r/a = 0.05, 0.15, \dots, 0.75$. This range excludes the outermost of the above four linearly simulated flux surfaces. The number of flux surfaces is constrained to be small by computational resource limits, and sets a lower limit on the macroscopic gradient length scale (which however is always formally large compared to the flux tube cross-section). How to choose the number of flux surfaces optimally is an interesting topic for future investigation.


There were two kinetic species: electrons and deuterium ions. Within each flux tube there were 40 grid points in the x (radial) direction and 54 in the y direction; after dealiasing, the minimum (strictly positive) and maximum values of k_y were

$1/\rho_i$ and $17/\rho_i \approx 0.3/\rho_e$ respectively. The parallelization across species energy and y coordinates used 144 processors per flux tube, with three flux tubes per flux surface. The setup is shown schematically in Figure 3.3. At each transport iteration, each flux tube simulation proceeded for at least 20,000 time steps (the time step being determined adaptively within GS2, as described immediately below). Evidence that this was sufficiently long can be seen in later Figures 3.6 and 4.1, but the determination of convergence in time is still an area for further investigation, and we return to this point in the Conclusions chapter (chapter 6).


GS2's adaptive time-stepping, which applies only to nonlinear simulations, is designed to ensure that the Courant-Friedrichs-Lewy (CFL) criterion for numerical stability is always satisfied, but only marginally: that the time step is small enough, but not too much smaller. It does this by estimating how well the condition is satisfied, using the current maximum value of the (nonlinearly advecting) drift velocity, and adjusting the time step up or down accordingly. Further details of the same scheme are documented for GS2's sister code AstroGK [29], which was derived from GS2.

3.5 Trinity simulation results


In the inner part of the plasma, consistent with the lack of electron-scale linear instability and with the weak exchange heating between ions and electrons, we found no significant turbulent electron heat flux, and the core electron temperature showed no sign of ceasing to rise (Figure 3.4(a)). This is problematic for steady-state prediction even in the outer part of the plasma, because it does not allow the source heating power in the inner part to reach the outer part. This issue was addressed by setting the electron thermal diffusivity artificially within the region $r < r_c$ to be constant in space and equal to the computed value at the boundary: $\chi_e(r, t) = \chi_e(r_c, t)$. The chosen $r_c = 0.45a$ lies within the unstable outer part,



(a)




(b)




(c)

Figure 3.4: (a) T_e after 0.01 s showing rapid core temperature increase; (b) T_e at end of run with modified χ_e where dotted; (c) T_i compared with neoclassical-only run



(a)



(b)

Figure 3.5: (a) T_e gradient profile; (b) integrated sources and exchange heating

and so $\chi_e(r_c, t) > 0$. This then allows the source heating power in the inner part to flow outward, providing the correct flux boundary condition for the outer part and allowing steady state to be approached after longer simulation times. We are then able to compare the final profiles in the region $r > r_c$ (where the diffusivity is still determined from first principles) with experiment. The agreement with the experimental temperature profile is shown in Figure 3.4(b); Figure 3.5(a) shows the corresponding gradient.

In addition to the turbulent contribution to the heat flux obtained by time- and space-averaging over each flux tube domain, the Chang-Hinton approximation to

the neoclassical energy transport [30] was added to the ions and completely dominated the ion energy transport, as shown in Figure 3.4(c), where the T_i profile at the end of the same run is compared against a run with neoclassical T_i evolution only (no turbulence simulation, and T_e fixed to the experimental profile). Note that for both species the exchange heating is small compared to the source, as shown in Figure 3.5(b), and so the transport problem largely decouples into independent ion and electron channels. We therefore need not be very concerned, in the context of the electron transport, by the questionable applicability of the Chang-Hinton approximation to MAST ions [31]. However, the critical electron temperature gradient for linear (ETG) instability is expected to depend on T_e/T_i [32], so there is an indirect link between the two channels, and it would be possible in future to improve the ion model (especially for STs) by implementing Roach’s re-evaluation of the Chang-Hinton approximation for tight-aspect-ratio, noncircular plasmas [31].

Figure 3.6 shows the temperature evolution against time at all plasma radii for both species, in the run discussed above. The adaptation of the transport time step is clearly seen. Note that this run corresponds to several sequential simulations, each starting from the final result of the previous one, totalling about 180,000 core hours of CPU time.




Figure 3.6: Time evolution of T_e (solid and dotted, to match Figure 3.4(b)) and T_i (dashed). Each symbol shows a transport time step taken by Trinity.

3.6 Turbulence properties

Figure 3.7(a) shows the electron heat flux spectrum at $r/a = 0.65$ averaged over the final transport time step. The nonlinear peak in k_y lies below the peak of the linear spectrum (cf. Figure 3.2), and the nonlinear range of both k_x and k_y lies well within the spectral boundaries of the nonlinear simulation domain.

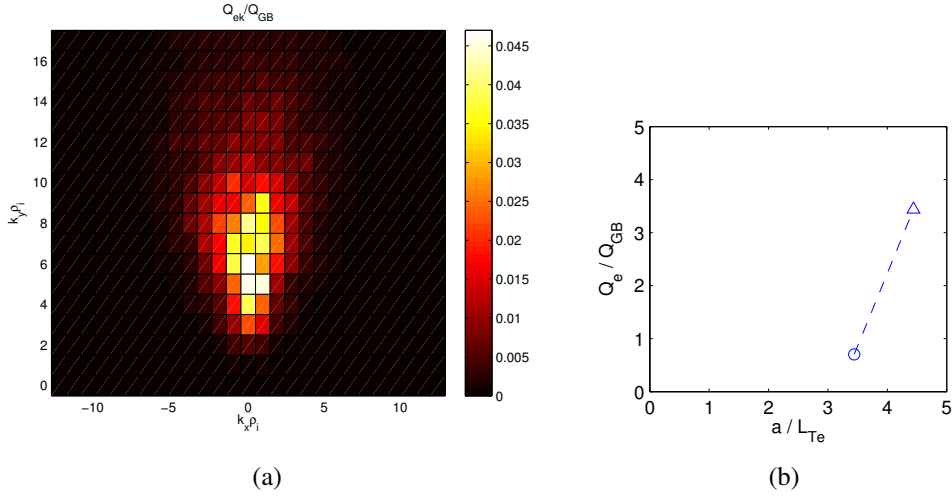


Figure 3.7: Electron heat flux at $r/a = 0.65$: (a) spectrum averaged over final transport time step; (b) total (circle), and total from flux tube with perturbed electron temperature gradient (triangle): the dashed line between the two symbols gives an indication of stiffness (its slope) and critical gradient (its intercept with $Q_e = 0$).

As part of its Newton algorithm for solving the transport equation, at each iteration Trinity runs GS2 flux tubes both with nominal parameters and with perturbed drive gradients. This automatically provides an indication of critical gradient and stiffness, as shown in Figure 3.7(b).

Jenko et al. [32] derive the following formula (cf. their equation (7)) for the ETG linear critical gradient:

$$\begin{aligned} (R/L_{Te})_{\text{crit}} = \max\{ & (1 + Z_{\text{eff}} T_e/T_i)(1.33 + 1.91\hat{s}/q) \\ & \times (1 - 1.5\epsilon)[1 + 0.3\epsilon(d\kappa/d\epsilon)], 0.8R/L_n\}, \end{aligned} \quad (3.4)$$

where $\epsilon = r/R$ and Z_{eff} is the effective ion charge. This formula gives a critical gradient of $(a/L_{Te})_{\text{crit}} = 1.7$ for our parameters (Table B.1). The coefficients in

(3.4) are based on a best fit, but a qualitative theoretical argument is also provided by Jenko et al. [32] that leads (at the very beginning of their page 4102) to a factor $(1 - \epsilon)$ where the factor $(1 - 1.5\epsilon)$ appears instead above. If this replacement is made, our critical gradient becomes $(a/L_{Te})_{\text{crit}} = 2.9$. This is consistent with Figure 3.7(b); the nonlinear critical gradient (the projection of the slope to intercept $Q_e = 0$) may be upshifted relative to the linear critical gradient [16].

3.7 Conclusions and future work

Our results are broadly compatible with an electron temperature profile governed by edge pedestal height (which we have taken as a given; it simply sets the outer boundary condition here) and ETG critical gradient, with the ions decoupled. The electrostatic turbulent electron heat transport is comparable to experiment in the region $r/a > 0.45$ where the magnetic shear is positive (that is, q is increasing with r ; see Figure 3.1(b)) and the neoclassical ion transport is also comparable to experiment.

Additional instabilities are required to match experiment in the negative magnetic shear region. $\delta\mathbf{B}$ and flow shear terms are already implemented in the code and could be investigated. Additional transport physics is also already in the code, for modelling density and momentum transport and for the inclusion of turbulent heating. It would also be possible to improve the ion model by implementing Roach's extension [31] of the Chang-Hinton approximation [30], as discussed in section 3.5.

CHAPTER 4

Long-time and ion modelling effects

In this short bridging chapter we will highlight one surprising feature found in the simulations described in chapter 3, replicating it in a single flux tube. This will motivate our subsequent study of the long-term behaviour of ETG turbulence in a single flux tube, outside of the much more computationally challenging context of multiscale transport simulations. In chapter 5 we will also simplify the model further to an adiabatic ion response (section 5.2.2), no longer modelling the ions kinetically. The present chapter indicates that similar effects, though not investigated as fully, are apparent with kinetic ions, as were used in the Trinity simulations.

During a Trinity run, at each transport time step the flux-tube simulations are continued from the previous state, with updated parameters to reflect the updated radial profiles. Where these parameters do not much change, as in the outer part of the plasma in this case, Trinity effectively conducts a very long flux tube simulation. It is apparent in our results that the zonal component can grow slowly throughout the run. Figure 4.1(a) shows a time trace of the spectral intensity of the electrostatic potential, $\sum_{k_x} |\phi_{\mathbf{k}_\perp}|^2$, for a corresponding standalone GS2 run with fixed parameters; this isolates the behaviour from any global features of the full transport simulations in chapter 3. The black curve is the zonal component. Figure 4.1(b) shows that simultaneously the turbulent electron heat flux gradually falls. (After $v_{it}/a = 220$, the limit shown here, there is no appreciable further change, at least as far as the end




Figure 4.1: Time traces of (a) spectral intensity of the electrostatic potential, summed over all k_x , $\sum_{k_x} |\phi_{\mathbf{k}\perp}|^2$, for the zonal $k_y = 0$ component (black) and for multiple individual non-zonal $k_y \neq 0$ components (all yellow); (b) electron heat flux Q_e .

of the simulation at $v_{ti}t/a = 500$.) Roach et al. [13] carried out similar flux-tube simulations that were shorter in duration (up to $v_{ti}t/a \approx 50$). At these earlier times, the simulation in Figure 4.1 indeed appears to have reached steady state, because the later slow trend is not yet established. We will refer to this apparent earlier steady state as “quasi-saturated”.

Slowly growing zonal and falling non-zonal turbulence levels have been reported by Guttenfelder and Candy [33] in ETG simulations with adiabatic ions (AI), and low AI transport is also mentioned by Peterson et al. [34]. Mantica et al. [35] exhibit slow zonal and turbulence changes in ion temperature gradient (ITG) simulations using drift-kinetic electrons [36]. To our knowledge our own results are the first to show such behaviour in ETG simulations with fully kinetic ions (KI).

In the next chapter, we investigate this phenomenon in more detail for the simplified AI model. It will turn out that in the final saturated state at the end of the slow growth phase, the zonal modes can strongly regulate the turbulence, and that the weak damping of the zonal modes by electron-ion collisions results in a simple linear scaling of the heat flux with collisionality.

CHAPTER 5

Collisionality scaling of the electron heat flux in ETG turbulence

5.1 Introduction

Experiments on MAST in which heat transport is dominated by electrons find that the thermal energy confinement time τ_E varies with the normalized electron collisionality ν_* according to the scaling [37]

$$B\tau_E \propto \nu_*^{-0.82 \pm 0.1}, \quad (5.1)$$

where B is the magnetic field. This scaling is favourable towards improved confinement in the hotter, lower collisionality plasmas anticipated in future devices.¹

In this chapter, we investigate how the electron heat flux Q varies with electron collisionality in simulations of electron-scale plasma turbulence in MAST, using the local gyrokinetic flux-tube code GS2 [8, 9]. At constant geometry and $\rho_* = \rho_e/a$ (ρ_e is the electron gyroradius, a is the equilibrium length scale), it can be shown (see Appendix A) that

$$B\tau_E \propto \left(\frac{Q}{Q_{GB}} \right)^{-1}, \quad (5.2)$$


¹Although B and ν_* are not necessarily independently controllable for a given device design, it is anticipated that this scaling with ν_* is strong enough to influence τ_E in the same direction as $B\tau_E$. For example, in the collisionality scan undertaken by Valovič et al. [37], $\nu_* \propto 1/B^4$, and this strong B^4 dependence will not be overturned by the single factor of B in (5.1). Further details concerning this scaling, including the other control parameters held constant here, will be found in Appendix A.

where $Q_{\text{GB}} = nT v_{te} \rho_*^2$ is the electron gyroBohm heat flux. It therefore ought to be possible to recover the scaling (5.1) from a local calculation of the electron heat flux.


We wish to discover whether this experimental MAST scaling may be understood in terms of electron temperature gradient (ETG) driven turbulent transport. With this goal in mind, we carry out local gyrokinetic simulations restricted to electrostatic perturbations at electron-gyroradius scales, ignoring both electromagnetic (e.g., “microtearing” [38, 39]) modes and ion-gyroradius-scale effects. The restriction to electrostatic perturbations is a matter of considering as simple a model as possible for a plasma where β is low and magnetic perturbations are relatively small. The neglect of the ion-gyroradius-scale turbulence is justifiable on the grounds that, in typical MAST plasmas, ion turbulence is considerably suppressed by radial shear in the background flow, and ion transport is close to the neoclassical level [17, 22, 14, 15]. This is fortunate, as spanning both ion- and electron-gyroradius scales requires prohibitively large computational resources. We therefore limit ourselves here to electron-gyroradius scales. We further limit ourselves here to a Boltzmann-ion model, treating only the electrons kinetically. Again, we wish to adopt as simple a model as possible.² Whilst greatly simplified, the equations we solve (which are described in more detail in the next section) do reproduce the experimental scaling of electron heat flux with collisionality.

The reference simulation parameters are based on experiment, which is close to the threshold for the onset of turbulence. In this region of parameter space, we reproduce the experimental scaling of electron heat flux with collisionality. This scaling is only revealed, however, if the simulation times are sufficiently long to reach a true steady state, which requires them to be much longer than the electron-

²Abel and Cowley [40] remark that modelling electrons kinetically at only ion scales is formally invalid – deliberately unresolved in some sense. Modelling ions kinetically at only electron scales could also be questioned in a similar fashion. We have circumvented this issue by modelling only electrons kinetically at only electron scales.



(a)



(b)

Figure 5.1: Non-dimensionalised electrostatic potential $e\phi/T\rho_*$ (where $\rho_* = \rho_e/a$) at the outboard midplane, for $\nu = 0.2\nu_{\text{nom}}$ (here ν_{nom} is the “nominal”, i.e., experimental value of collisionality), $a/L_T = 3.3$: (a) quasi-saturated state at $t = 1200.3 a/v_{te}$, (b) saturated state at $t = 7835.8 a/v_{te}$, for large-box simulations. See Appendix B and Appendix C for the meaning of symbols.

collision time scale. At earlier times, there is a transient “quasi-saturated” state with higher heat flux, in which the zonal modes (which do not themselves contribute to the heat flux) are small but slowly growing. When they have grown to a sufficient level, the nonzonal modes and the heat flux are significantly suppressed. Figure 5.1 illustrates these two regimes by showing the electrostatic potential (which is proportional to the density perturbation) in an outboard-midplane cross-section of a flux tube in MAST both in the earlier quasi-saturated state and the later long-time saturated state, based on one of the simulations reported below. In the quasi-saturated state, the zonal modes do not appear to play a special role, and radially extended “streamers” can be seen, as is usually expected for ETG turbulence [41, 42]. In contrast, in the long-time saturated state, a strong zonal component comes to dominate, structuring the turbulence into “vortex streets” and dramatically weakening radial transport. In this final saturated state, the nonlinear drive of the zonal modes is balanced by their weak collisional damping, dominated by electron-ion collisions. Scans in collisionality reveal that the saturated heat flux increases with increasing collisionality, in rough proportionality: $Q/Q_{GB} \propto \nu_*$.

A brief outline of the rest of the chapter is as follows. In section 5.2, we describe the equations that are solved and the simulation set-up. In section 5.3, we present our main results, including the long-time evolution of the turbulence, the dependence of the saturated heat flux on collisionality, and the structure of the saturated turbulent state. We also sketch a simple theoretical argument that explains the collisionality scaling of the heat flux (section 5.3.2). In section 5.4, a summary of our findings is given, our results are put in the context of some earlier work, and a discussion is given of the apparent differences and possible similarities between the ETG and ITG turbulent states in light of the conclusions of the present study.

5.2 Governing equations and numerical set-up

Our study is based on the MAST shot 8500, for which data are available from the International Tokamak Profile Database [26]. This shot was analysed by Field et al. [17], and a linear gyrokinetic study was performed by Roach et al. [22]. In the present work, we consider a single flux surface, for which the detailed plasma parameters are given in Appendix B (they are referred to as “nominal” parameters); these were kept fixed throughout our study, except for varying collisionality and electron temperature gradient where indicated.

5.2.1 Gyrokinetic equation

We use the GS2 continuum gyrokinetic code [9] to obtain the perturbed distribution function and electrostatic field in local flux-tube geometry. The electron distribution function is written

$$f = F + \delta f, \quad (5.3)$$

where

$$F = n \left(\frac{m}{2\pi T} \right)^{3/2} \exp \left(-\frac{mv^2}{2T} \right) \quad (5.4)$$

is the equilibrium Maxwellian background distribution. We will omit species subscripts when this will cause no confusion. n , m , T and \mathbf{v} are the electron density, mass, temperature and velocity, respectively, and the perturbed distribution function is split into a Boltzmann response associated with the perturbed electrostatic potential ϕ and a gyrotropic non-Boltzmann, generally non-Maxwellian part:

$$\delta f = \frac{e\phi(\mathbf{r})}{T} F + h(\mathbf{R}, v_{\perp}, v_{\parallel}), \quad (5.5)$$

where e is the absolute value of the electron charge. Note that ϕ is a function of the particle position \mathbf{r} , whereas h is a function of the gyrocentre position $\mathbf{R} = \mathbf{r} - \mathbf{b} \times \mathbf{v}/\Omega_e$, where \mathbf{b} is the unit vector along the magnetic field and $\Omega_e = -eB/mc$ is the

electron cyclotron frequency (its sign is the negative sign of the electron charge). The gyrocentre distribution h is otherwise independent of the gyroangle, being a function of the parallel $v_{\parallel} = \mathbf{v} \cdot \mathbf{b}$ and perpendicular $v_{\perp} = (v^2 - v_{\parallel}^2)^{1/2}$ velocities. Equivalently, its velocity-space variables can be (and are) chosen to be the energy $\mathcal{E} = v^2/2$ and magnetic moment $\mu = v_{\perp}^2/2B$, so $v_{\parallel} = \pm(2\mathcal{E} - 2\mu B)^{1/2}$. In this representation, the gyrocentre distribution h satisfies the electrostatic gyrokinetic equation (GKE) [43] (see review [3])

$$\frac{\partial}{\partial t} \left(h + \frac{e\langle\phi\rangle}{T} F \right) + (v_{\parallel} \mathbf{b} + \mathbf{v}_B) \cdot \nabla h + \langle \mathbf{v}_E \rangle \cdot \nabla h + \langle \mathbf{v}_E \rangle \cdot \nabla F = \langle C[h] \rangle, \quad (5.6)$$

where $\langle \dots \rangle$ denotes gyroaveraging at constant gyrocentre position \mathbf{R} ,

$$\mathbf{v}_B = \frac{\mathbf{b}}{\Omega_e} \times \left(v_{\parallel}^2 \mathbf{b} \cdot \nabla \mathbf{b} + \frac{v_{\perp}^2}{2} \frac{\nabla B}{B} \right) \quad (5.7)$$

is the magnetic drift velocity, and

$$\mathbf{v}_E = \frac{c}{B} \mathbf{b} \times \nabla \phi \quad (5.8)$$

is the $\mathbf{E} \times \mathbf{B}$ drift velocity. The energy is injected into the system via the last term on the left-hand side of (5.6), which contains the radial gradients of the equilibrium distribution:

$$\langle \mathbf{v}_E \rangle \cdot \nabla F = \langle v_{Ex} \rangle \frac{\partial F}{\partial x} = \langle v_{Ex} \rangle \left[\frac{1}{L_n} + \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right) \frac{1}{L_T} \right] F, \quad (5.9)$$

where $v_{te} = (2T/m)^{1/2}$ is the electron thermal speed. We have defined $L_n = -d \ln n / dx$ and $L_T = -d \ln T / dx$, the gradient scale lengths of the equilibrium electron density and temperature profiles, respectively. The natural normalisation of these lengths is the tokamak minor radius a . In the above definitions, x is the spatial coordinate transverse to the flux surface and y is a coordinate within the flux surface transverse to the magnetic field (in the outboard midplane, it is approximately the arc length perpendicular to the field) — these curvilinear, non-orthogonal coordinates, as used by GS2, are explained in Appendix C.1. Note that, in these coordinates,

$$v_{Ex} = \mathbf{v}_E \cdot \nabla x = \frac{cq}{B_0^2 r} \frac{d\psi}{dr} \frac{\partial \phi}{\partial y} \quad (5.10)$$

(see Appendix C.1 for the definition of all symbols).

5.2.2 Adiabatic ions

In a simple (two-species, hydrogenic, i.e., with ion charge = $+e$) plasma, the quasineutrality condition for the perturbations is

$$\delta n_e = \delta n_i, \quad (5.11)$$

where $\delta n = \int d^3\mathbf{v} \delta f$ is the density perturbation for each species. In terms of h , this density perturbation is

$$\frac{\delta n}{n} = -\frac{Ze\phi}{T} + \frac{1}{n} \int d^3\mathbf{v} \langle h \rangle_{\mathbf{r}}, \quad (5.12)$$

where $\langle \cdot \rangle_{\mathbf{r}}$ denotes the gyroaveraging operator at constant particle position \mathbf{r} (needed here because the velocity integral must be performed at fixed particle position \mathbf{r} , while h is a function of \mathbf{R}), and $Z = 1$ for the ions and -1 for the electrons. We have assumed a decomposition of the perturbed ion distribution function analogous to (5.5). If we consider only perturbations at scales perpendicular to the magnetic field that are much smaller than the ion gyroradius, $k_{\perp} \rho_i \gg 1$, the non-Boltzmann part of the ion density response in (5.12) can be neglected because $\langle h_i \rangle_{\mathbf{r}}$ is suppressed by the averaging over large ion Larmor orbits. In formal terms, this approximation is the lowest-order expansion in the electron-ion mass ratio. Therefore, in the present study, we will assume that the ion distribution is entirely described by the Boltzmann (sometimes called adiabatic) response:

$$\frac{\delta n_i}{n_i} = -\frac{e\phi}{T_i}. \quad (5.13)$$

Combining this with the quasineutrality equation (5.11) and the full gyrokinetic electron density response given by (5.12) with $Z = -1$, we get the following equation for ϕ :

$$\frac{e\phi}{T_e} \left(1 + \frac{1}{\tau} \right) = -\frac{1}{n} \int d^3\mathbf{v} \langle h \rangle_{\mathbf{r}}, \quad (5.14)$$

where h satisfies (5.6) and $\tau = T_i/T_e$.

Equations (5.6) and (5.14) constitute a closed system, which is solved numerically by the version of the GS2 code that we use here.

5.2.3 Collisions

The electron collision operator used in GS2, appearing on the right-hand side of (5.6), consists of the electron-ion pitch-angle scattering operator (to lowest order in the mass-ratio expansion) and a simplified model of the electron-electron collision operator [22, 44, 45]: in \mathbf{k}_\perp space, it is

$$\langle C[h] \rangle_{\mathbf{k}_\perp} = \langle C_{ee}[h] \rangle_{\mathbf{k}_\perp} + \nu_{ei} \frac{v_{te}^3}{v^3} \left[\frac{\partial}{\partial \xi} \frac{(1 - \xi^2)}{2} \frac{\partial h_{\mathbf{k}_\perp}}{\partial \xi} - \frac{(1 + \xi^2)}{4} \frac{v^2}{v_{te}^2} k_\perp^2 \rho_e^2 h_{\mathbf{k}_\perp} \right], \quad (5.15)$$

where $\xi = v_\parallel/v$. The electron-electron operator is defined fully in [44]; it has the form

$$\begin{aligned} \langle C_{ee}[h] \rangle_{\mathbf{k}_\perp} = & \nu_D \frac{\partial}{\partial \xi} \frac{(1 - \xi^2)}{2} \frac{\partial h_{\mathbf{k}_\perp}}{\partial \xi} + \frac{1}{v^2} \frac{\partial}{\partial v} \left(\frac{1}{2} v^4 \nu_\parallel F \frac{\partial}{\partial v} \frac{h_{\mathbf{k}_\perp}}{F} \right) \\ & - \frac{1}{4} [\nu_D(1 + \xi^2) + \nu_\parallel(1 - \xi^2)] \frac{v^2}{v_{te}^2} k_\perp^2 \rho_e^2 h_{\mathbf{k}_\perp} + \dots, \end{aligned}$$

where the additional correction terms indicated by ‘ \dots ’ are constructed so as to ensure the conservation of particle number, momentum and energy:

$$\int d^3\mathbf{v} \langle C_{ee}[h_{\mathbf{k}_\perp}] \rangle = 0, \quad \int d^3\mathbf{v} v_\parallel \langle C_{ee}[h_{\mathbf{k}_\perp}] \rangle = 0, \quad \int d^3\mathbf{v} v^2 \langle C_{ee}[h_{\mathbf{k}_\perp}] \rangle = 0.$$

ν_D and ν_\parallel are also defined in [44]. The numerical implementation of the whole collision operator is detailed in [45].³

³As originally written, GS2 was not configured to simulate kinetic electrons with adiabatic (Boltzmann) ions and also to include electron-ion collisions. This is because GS2 always formally requires at least one kinetic ion species, and so single-species KE+AI simulations were actually performed as single-species kinetic-ions simulations with the sign of the charge in the Boltzmann response appropriately reversed. The equations solved are then correct for simulations with kinetic electrons and adiabatic ions without collisions or with same-species collisions only. For the present work, we added the capability to include the electron-ion collision term in such simulations; previously it was included only for electrons in multi-species simulations with both kinetic ions and kinetic electrons.

The electron-ion collision rate is $\nu_{ei} = Z_{\text{eff}}\nu_{ee}$, where $\nu_{ee} \equiv \nu = \sqrt{2}\pi n e^4 T^{-3/2} m^{-1/2} \ln \Lambda$ is the electron-electron collision rate and $\ln \Lambda$ is the Coulomb logarithm. In a plasma with multiple ion species, $Z_{\text{eff}} = \sum_i n_i Z_i^2 / \sum_i n_i Z_i$ arises from summing the individual electron-ion collision operators over all ion species ($Z_i e$ is that species' charge); in a hydrogenic plasma, $Z_{\text{eff}} = 1$. Note, however, that GS2 treats Z_{eff} as an independent parameter in equation (5.15), formally allowing one to vary ν_{ei} relative to ν_{ee} without affecting the quasineutrality equation (5.11). This is normally done to include the effect of electron collisions with impurity ion species that are not being modelled kinetically, but are present in real experiments. In this work, it will also be useful to us in section 5.3.5 as a method of varying this ratio artificially.

Electron-ion collisions relax the parallel electron flow towards the stationary ion background, resistively dissipating the current associated with an electron flow below the ion gyroscale.⁴ As a result, the electron-ion collision operator, and hence the electron collision operator overall, does not conserve (electron) momentum. The last term in the bracket in (5.15) is a spatial diffusion — it is a finite-gyroradius (finite-Larmor-radius: FLR) effect arising from the gyroaveraging of the collision operator [44].

In Appendix C, we provide the details about our numerical simulation parameters, grids and resolution.

⁴The electron flow is also a current because, in the limit of $k_{\perp}\rho_i \gg 1$, ion gyromotion averages out over the electron scales leaving, to lowest order, no ion-flow response in either parallel or perpendicular direction; cf. the adiabatic-ion approximation (5.13), which arises from the averaging out of the gyrokinetic ion-density response.

5.3 Results

5.3.1 Time evolution: from quasi-saturation to long-time steady state

Figure 5.2(a) shows the turbulent electron heat flux as a function of time for two simulations that have identical plasma parameters and differ only in numerical grid parameters (their coincidence is some evidence of numerical convergence; for more details of the numerics and of the various issues of convergence, see Appendix E, which also contains plots for other values of ν and a/L_T and of other zonal quantities). The heat flux Q is calculated from the solution of equations (5.6) and (5.14) according to

$$Q = \overline{\int d^3\mathbf{v} \frac{mv^2}{2} \delta f v_{Ex}}, \quad (5.16)$$

where the overbar indicates a flux-surface average. We will normalise Q to the gyroBohm value $Q_{\text{GB}} = nT v_{te} \rho_*^2$, where $\rho_* = \rho_e/a$. Besides its physical meaning as the heat flux, Q is also a good proxy for the turbulent fluctuation level of the non-zonal modes ($k_y \neq 0$; the $k_y = 0$ components of ϕ or h do not contribute to Q).

After a short exponential transient during the linear growth phase, the system reaches a “quasi-saturated” turbulent state, which, however, is not the final steady state. The final saturated state is reached much later, after a slow decline in heat flux accompanied by slow growth of the zonal ($k_y = 0$) component of the turbulence: Figure 5.2(b) shows the evolution of the zonal velocity squared,

$$(k_Z \phi_Z)^2 = \sum_{k_x} k_x^2 |\phi_{k_x,0}|^2 \quad (5.17)$$

(this equation also defines the characteristic zonal scale k_Z ; see equation (5.18) for ϕ_Z separately). Snapshots of ϕ , shown in Figure 5.1 and corresponding to early and late times in one of these simulations, are a vivid illustration of the different structure of the “quasi-saturated” and the final steady state: the former resembles the




Figure 5.2: Evolution in time of (a) the turbulent electron heat flux in electron gyroBohm units, and (b) the square of the zonal velocity, $(k_Z \phi_Z)^2$, for two simulations with electron collisionality $\nu = 0.2 \nu_{\text{nom}}$, electron temperature gradient $a/L_T = 3.3$ (green: small-box simulation; yellow: large-box simulation; see Appendix C.2 for details).

streamer-dominated state usually associated with ETG turbulence [41, 42], whereas the latter is a zonal-mode-dominated state, which we will now proceed to analyse. Note that it is this long-time saturated state that matters for determining the level of transport because, even though the time for it to emerge is long by the standards of electron-gyroscale dynamics, it is still much shorter than the transport time scale in a tokamak.⁵ This is because the self-consistent combination of δf local (gradient-driven) gyrokinetics at the gyroradius scale, and radial transport evolution at the system scale, implies a scale separation in both space and time [7, 3].

⁵Using MAST data from [37], the energy confinement time τ_E is of order 10 ms, whereas a/v_{te} is of order 10 ns, which is a factor of 10^6 smaller. The longest runs in this chapter evolve for times of order $10^4 a/v_{te}$, so there is still good scale separation.

5.3.2 Collisionality scaling: numerical results and theory

The two clearest results from the final saturated state of these simulations are summarised in Figure 5.3, which is based on a parameter scan in electron collisionality for two values of the electron temperature gradient, at two different numerical resolutions. The saturated heat flux increases roughly proportionally to the collisionality, whereas the zonal velocity is essentially independent of it. The heat-flux scaling is in remarkable agreement with the experimental scaling (5.1).

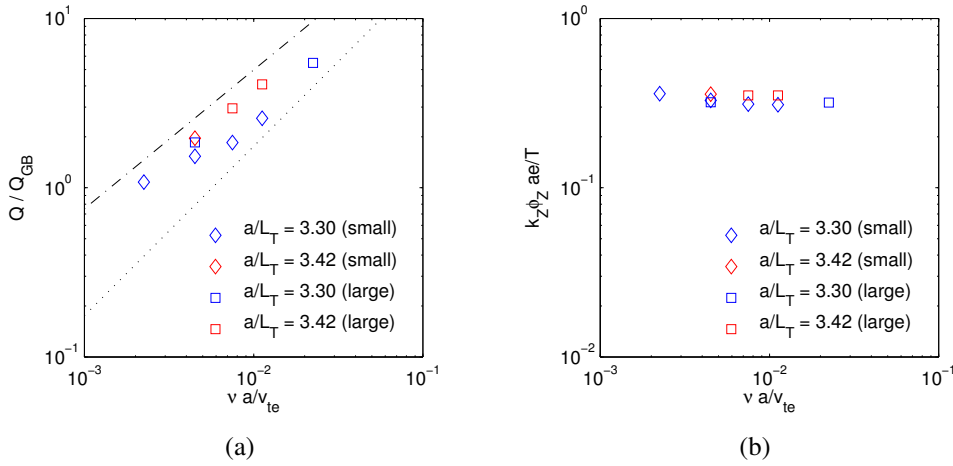


Figure 5.3: Variation of (a) the time-averaged normalised electron heat flux Q/Q_{GB} , and (b) the rms zonal velocity $k_Z \phi_Z$, defined by (5.17), versus normalised electron collisionality $\nu a/v_{te}$, at the nominal (experimental) value of the temperature gradient $a/L_T = 3.42$, and at $a/L_T = 3.3$. Symbol shapes indicate the simulation series (“small box” or “large box”), as explained in Appendix E. The dot-dashed line shows the experimental power-law scaling, and the dotted line shows the theoretical linear scaling, $Q/Q_{GB} \propto \nu_*$, equation (5.31).

Let us outline a simple explanation of these results, which we will then follow up with a series of numerical experiments designed to test its plausibility.

Let us split the gyrokinetic equation (5.6) explicitly into equations governing the evolution of the nonzonal and zonal components of the distribution function and the associated electrostatic potential (cf. [46]),

$$h = h_{NZ} + h_Z, \quad \phi = \phi_{NZ} + \phi_Z, \quad (5.18)$$

where the subscripts NZ, Z denote nonzonal ($k_y \neq 0$) and zonal ($k_y = 0$) modes, respectively:

$$\begin{aligned}
& \frac{\partial}{\partial t} \left(h_{\text{NZ}} + \frac{e\langle\phi_{\text{NZ}}\rangle}{T} F \right) + (v_{\parallel} \mathbf{b} + \mathbf{v}_B) \cdot \nabla h_{\text{NZ}} - \langle C[h_{\text{NZ}}] \rangle \\
& \quad + \underbrace{\langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{Z}} + \langle \mathbf{v}_E \rangle_{\text{Z}} \cdot \nabla h_{\text{NZ}}}_{\text{Z-NZ interaction (I)}} \\
& \quad + \overline{\langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{NZ}} - \langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{NZ}}} \\
& = \underbrace{-\langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla F}_{\text{energy injection (II)}}, \tag{5.19}
\end{aligned}$$

$$\begin{aligned}
& \frac{\partial}{\partial t} \left(h_{\text{Z}} + \frac{e\langle\phi_{\text{Z}}\rangle}{T} F \right) + (v_{\parallel} \mathbf{b} + \mathbf{v}_B) \cdot \nabla h_{\text{Z}} - \underbrace{\langle C[h_{\text{Z}}] \rangle}_{\text{damping (III)}} \\
& = \underbrace{-\overline{\langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{NZ}}}}_{\text{energy injection (IV)}}. \tag{5.20}
\end{aligned}$$

The overline denotes spatial averaging over y , i.e., the $k_y = 0$ component. Equation (5.20) is the y average of the gyrokinetic equation (5.6); then (5.19) is the result of subtracting (5.20) from (5.6).

We conjecture that the dominant balance governing the saturated state of the nonzonal modes is between the zonal-nonzonal interaction terms (I) and the linear drive (energy-injection) term (II) in (5.19).⁶ We estimate these terms as

$$\text{(I)} \sim \langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{Z}} \sim \frac{c}{B} k_y \phi_{\text{NZ}} k_{\text{Z}} h_{\text{Z}}, \tag{5.21}$$

and

$$\text{(II)} \sim \langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla F \sim \frac{c}{B} k_y \phi_{\text{NZ}} \frac{F}{L_T} \tag{5.22}$$

where k_{Z} is the typical zonal wavenumber and k_y the typical nonzonal wavenumber (we have used $v_{E_x} \sim (c/B)k_y\phi$; see (5.10)). The second zonal-nonzonal interaction

⁶We are thus treating the collision term in the nonzonal equation as subdominant, or at least as unimportant to this aspect of the dynamics. Numerically we find that it cannot be neglected as it regularises the fine velocity-space structure arising due to the phase-mixing of h_{NZ} [47].

The parallel streaming and magnetic-drift terms likely play a part in determining the spatial structure of the turbulence [48, 49, 47], but we shall see that we do not need to determine k_y , k_{Z} or k_{\parallel} . In [46], a split between zonal and nonzonal components is performed for integrated entropy balance equations in which these terms do not appear. One could base a similar balance argument on these equations. Whilst our argument presented here, based on the underlying dynamics, may be more transparent, the fact that the key terms are also the ones appearing in the integrated equations is a hint that perhaps the argument has a deeper thermodynamic significance.

term, $\langle \mathbf{v}_E \rangle_Z \cdot \nabla h_{\text{NZ}}$, is of the same order as (5.21) if we assume that

$$\frac{h_{\text{NZ}}}{F} \sim \frac{e\phi_{\text{NZ}}}{T}, \quad \frac{h_Z}{F} \sim \frac{e\phi_Z}{T} \quad (5.23)$$

for both nonzonal and zonal modes. Balancing (5.21) and (5.22), we find, after cancellation of $k_y\phi_{\text{NZ}}$, that

$$k_Z h_Z \sim \frac{F}{L_T}. \quad (5.24)$$

This is a form of mixing-length hypothesis [50], suggesting that the perturbed zonal gradients $\sim k_Z h_Z$ compensate the background equilibrium gradients associated with F .⁷ It follows from (5.23) and (5.24) that

$$k_Z \frac{e\phi_Z}{T} \sim \frac{1}{L_T}. \quad (5.25)$$

Thus, the gradients of the zonal modes (e.g., the zonal velocity or the zonal temperature gradient) are independent of collisionality. This independence is indeed seen in Figure 5.3(b).

The only nonlinearity present in the zonal equation (5.20) is the nonzonal-nonzonal interaction term (IV); zonal modes are not directly driven by background gradients because ∇F is in the x direction. One can think of (5.20) as a kind of Langevin equation for zonal modes, which are excited by coupling between nonzonal modes and damped by collisions (at long times, the only damping mechanism is collisional; see Appendix F).⁸ Therefore, the dominant balance in (5.20) is be-

⁷They need not necessarily flatten the background gradient completely or everywhere. From Figure 5.2(b), the normalized zonal velocity $k_Z \rho_e e\phi_Z / T \rho_* \approx 0.3$ (we also find $k_Z \rho_e \delta T_Z / T \rho_* \approx 0.3$), which should be compared with the background gradient $a/L_T = 3.3$. For the (random-noise) initial conditions used in our simulations, the minimum a/L_T required to sustain turbulence is between 3.0 and 3.3, so the perturbation levels correspond roughly to the distance away from this nonlinear critical gradient (of which we have not made a precise determination because the system appears to be strictly subcritical in the presence of even small flow shear; the linear critical gradient *without* flow shear is about 2.4, as shown in Appendix D). Furthermore, the perturbed gradient is not constant (since $k_Z \neq 0$); the actual magnitude of the gradient at any particular x can exceed 0.3.

⁸How precisely the zonal flows are generated is a matter for future research. We do find that disconnecting the zonal feedback on the nonzonal modes but keeping the nonzonal-nonzonal interactions in the evolution equation for the zonal modes in the “quasi-saturated” state (as is done in the numerical experiments of section 5.3.4) still leads to growth of zonal modes, but they grow at a different (larger) rate than the ones in simulations where their feedback is preserved. This means

tween the nonlinear energy injection (IV) and collisional damping (III):

$$\frac{c}{B} k_Z k_y \phi_{\text{NZ}} h_{\text{NZ}} \sim \gamma_Z h_Z, \quad (5.26)$$

where γ_Z is the collisional damping rate of the zonal modes. Note that in conjecturing such a balance, we are assuming the zonal modes are not subject to a strong nonlinear instability that would break them down back into nonzonal perturbations, thus resulting in a purely collisionless saturated state. This possibility (which, for example, appears to be realised for zonal flows in certain regimes of ITG turbulence, where it is known as the tertiary instability [52]) would amount to a dominant balance in (5.20) between the energy-injection and energy-removal effects within the nonlinear term (IV). We are expressly assuming this does not happen in our near-marginal ETG turbulence.

Combining equations (5.25), (5.26), and (5.23) again, we get

$$\frac{h_{\text{NZ}}^2}{h_Z^2} \sim \frac{\phi_{\text{NZ}}^2}{\phi_Z^2} \sim \frac{\gamma_Z e B L_T}{c k_y T} \sim \frac{1}{k_y \rho_e} \frac{\gamma_Z}{v_{te}/L_T}. \quad (5.27)$$

Therefore, estimating the heat flux (5.16), we find

$$\frac{Q}{Q_{\text{GB}}} \sim \frac{n \delta T_{\text{NZ}} v_{Ex}}{n T v_{te} \rho_*^2} \sim \frac{\delta T_{\text{NZ}}}{T} \frac{c k_y \phi_{\text{NZ}}}{B v_{te} \rho_*^2} \sim k_y \rho_e \left(\frac{e \phi_{\text{NZ}}}{T \rho_*} \right)^2, \quad (5.28)$$

where we have used (5.23) to estimate $\delta T/T \sim e\phi/T$. Using (5.27) to relate ϕ_{NZ}^2 to ϕ_Z^2 and (5.25) to estimate ϕ_Z^2 , we get

$$\frac{Q}{Q_{\text{GB}}} \sim \frac{\gamma_Z}{v_{te}/L_T} \left(\frac{e \phi_Z}{T \rho_*} \right)^2 \sim \frac{\gamma_Z}{k_Z^2 \rho_e^2 (v_{te}/a)} \frac{a}{L_T}. \quad (5.29)$$

Thus, we are able to estimate the electron heat conductivity entirely in terms of the linear damping rate and characteristic scale of the zonal flows. It is possible to show analytically (Appendix F) and confirm numerically (section 5.3.3) that in the long term, zonal modes are damped by collisional (Ohmic) resistivity:

$$\gamma_Z \sim \nu_{ei} k_Z^2 \rho_{pe}^2, \quad (5.30)$$

that, even when the zonal modes are small, their excitation by the nonzonal modes still depends on the small modifications (subdominant as far as the heat flux is concerned) that they produce in the latter — and thus the mechanism of this excitation must be some form of “negative viscosity” or a similar effect [51], rather than just stochastic coupling of nonzonal modes into $k_y = 0$.

where ν_{ei} is the electron-ion collision rate and $\rho_{pe} = \rho_e B/B_p$ is the “poloidal gyro-radius” of the electrons (B_p is the poloidal magnetic field; see Appendix F.4 for a more precise definition of ρ_{pe}). Using (5.30) in (5.29), we finally obtain

$$\frac{Q}{Q_{\text{GB}}} \sim \frac{\nu_{ei}}{v_{te}/a} \left(\frac{B}{B_p} \right)^2 \frac{a}{L_T} \propto \nu_{ei}. \quad (5.31)$$

Thus, a relatively simple theoretical argument has produced a linear scaling of the heat flux with collisionality. Note that all dependence on k_Z or any other wave numbers has cancelled in the final expression (5.31), and so, in order to obtain the heat flux, we need not know the spatial scales of either zonal or nonzonal modes. Considering the simplicity of the argument, and the level of agreement between it, our numerical results (Figure 5.3), and the experimental MAST scaling (5.1), we find it quite compelling.

5.3.3 Damping of zonal modes

In the theoretical argument of section 5.3.2, a crucial step was to use the expression (5.30) for the collisional damping of the zonal modes, which allowed us to estimate the heat flux according to (5.29) and avoid having to theorise about the characteristic scale of the zonal modes (a nontrivial question, possibly with a nonuniversal answer). In Appendix F, the damping rate (5.30) is derived analytically. Physically, the situation can be summarised as follows.

Consider a zonal perturbation with some perpendicular wave number k_x satisfying (5.20) with zero right-hand side — a linear equation. In the absence of collisions, this perturbation will decay quickly (on the time scale $\sim a/v_{te}$), but not to zero, leaving a finite residual zonal field [53]. With collisions present, after a period of a few collision times, which is still much shorter than the damping time, $\nu^{-1} \ll \gamma_Z^{-1} \sim (\nu k_x^2 \rho_e^2)^{-1}$ in the long-wavelength limit $k_x \rho_e \ll 1$, it is possible to show (see Appendix F.3) that, to lowest order in $k_x \rho_e$, the remaining perturbation is a perturbed Maxwellian with a density (or, equivalently, ϕ) and a temperature

perturbation, both constant on each flux surface. These perturbations then decay diffusively due to perpendicular particle diffusion (equivalently, resistivity) arising from the collision operator. We already saw in section 5.2.3 that the gyrokinetic collision operator (5.15) contains FLR terms that have the form of a spatial diffusion. These terms correspond to the displacement of gyrocentres by distances $\sim \rho_e$ due to collisions during Larmor rotation. Solving the “zonal transport” problem more carefully, one can show (see Appendix F.4) that collisions also displace the centres of banana (and corresponding passing) orbits by distances of order the poloidal gyroradius $\rho_{pe} = (B/B_p)\rho_e$, which is larger. This leads to the damping rate (5.30). It is proportional to ν_{ei} (rather than ν_{ee} , on which it depends weakly) because it is essentially the Ohmic resistive damping of electron currents (both parallel and perpendicular; see Appendix F.4), and it is due to electron-ion friction (cf. [54]).

The calculation of Appendix F, where this is demonstrated more carefully, can be checked in our numerical simulations, to ascertain that it is indeed this effect that is responsible for the zonal damping. Figure 5.4 shows the zonal damping rate normalised by the collision frequency, γ_Z/ν , for a number of simulations that were restarted with the nonlinearity (the right-hand side) in equation (5.20) turned off. We see that, for a range of collisionalities ν and in a broad range of wave numbers $k_x\rho_e$, the scaling (5.30), $\gamma_Z/\nu \sim k_x^2\rho_{pe}^2$, is followed quite well. As a further successful test, we find that if we turn off magnetic drifts ($\mathbf{v}_B \cdot \nabla h$ in equation (5.20)), thus removing the banana orbits, the zonal damping rates drop by close to an order of magnitude (blue open circles in Figure 5.4). This is roughly consistent with a reduction of γ_Z by a factor of $(B/B_p)^2 \approx (qR/r)^2 \approx 18$ in our geometry, to $\gamma_Z \sim \nu k_x^2\rho_e^2$, with the dominant diffusion due in this case to finite Larmor orbits, as explained above.

Finally, we test the theoretical expectation that the dominant contribution to the damping of the electron zonal flows, which are also currents, is proportional specifically to the electron-ion collision frequency ν_{ei} . Figure 5.4 (black crosses) shows that when the electron-ion collisions are turned off, $\nu_{ei} = 0$, leaving only the




Figure 5.4: Zonal damping rate normalised to collisionality, γ_Z/ν , versus $k_x \rho_e$, spanning the range of collisionalities shown in Figure 5.3 (solid colors). The final states of various saturated nonlinear simulations were used as initial conditions, with the nonlinearity switched off. Also shown (black crosses) are the corresponding damping rates for a simulation at $\nu = \nu_{\text{nom}}$ in which electron-ion collisions were turned off (formally by setting $Z_{\text{eff}} = 0$; see section 5.2.3); and a simulation at $\nu = \nu_{\text{nom}}$ in which electron-ion collisions were retained but magnetic drifts were turned off (blue open circles). The solid line with slope k_x^2 corresponds to the scaling (5.30); the dashed line with slope k_x^4 corresponds to the scaling expected when $\nu_{ei} = 0$ (see Appendix F.7).

(momentum conserving) electron-electron collisions ν_{ee} , the damping rates drop dramatically and scale as $\gamma_Z \sim \nu_{ee} k_x^4 \rho_{pe}^4$, as indeed expected theoretically (see Appendix F.7).

Thus, the scaling (5.30) and the theory that leads to it (Appendix F) appear to be sound and successfully reproduced in our simulations.⁹

⁹It is perhaps worth pointing out that such an agreement is only possible in simulations that use a sufficiently realistic electron gyrokinetic collision operator (see section 5.2.3), an indispensable property being momentum conservation by the electron-electron collisions and a correct capturing of Ohmic resistivity by the electron-ion ones. See, however, section 5.3.5 for certain simplifications that are allowed.

5.3.4 Numerical tests of the role of collisions and zonal modes

With the theoretical argument presented in section 5.3.2 in mind, let us now build up the evidence that the long-time steady state of the ETG turbulence in our simulations is controlled by zonal-nonzonal interactions and by the electron-ion collisional damping of the zonal modes. First, we remark that whilst the parameter values are such that the linear instability growth rates are numerically similar to the nominal (experimental) collisionality, these growth rates vary only weakly with reducing collisionality (see Figure D.1). It does not seem plausible that such insensitive linear physics – the linear instability is not primarily collisional – can explain the strong collisionality dependence of the nonlinearly saturated state. Clues to the actual (nonlinear) origin of this dependence can be obtained via nonlinear simulations with modified dynamics, as described below.




Figure 5.5: The large-box simulation shown in Figure 5.2 (yellow) was restarted at $t = 8397.7 a/v_{te}$ without zonal interactions in the nonzonal evolution equation (cyan). The heat flux in this modified simulation returns to a level that is close to the high early “quasi-saturated” level. For direct comparison, the same modified simulation was also rerun from initial noise, giving the same heat flux level (purple).

First let us show that the zonal component regulates the amplitude of the rest of the turbulence, which determines the heat flux. The cyan and purple curves in Figure 5.5 show the time evolution¹⁰ of the heat flux in simulations with identical parameters to one of the simulations in Figure 5.2, but with the nonlinear term

¹⁰In each of Figures 5.5 and 5.6, the restart time indicated is simply the time at the end of the previous run, and its precise value has no physical significance. The number of time steps in a run is typically a round number, but the normalized time step and hence the final time is not.

artificially adjusted in such a way that the zonal modes no longer affect the evolution of the nonzonal modes: namely, the zonal components have been zeroed out in the calculation of the nonlinear term, so that in equation (5.6), $\langle \mathbf{v}_E \rangle \cdot \nabla h$ is replaced by $\langle \mathbf{v}_E \rangle_{\text{NZ}} \cdot \nabla h_{\text{NZ}}$. The zonal modes are still allowed to be nonlinearly driven by the nonzonal modes, but the zonal modes do not then feed back on the nonzonal modes; the nonzonal evolution is entirely independent of the zonal evolution. This eliminates the nonlinear terms used to obtain the dominant balance (5.24). The heat-flux collapse occurring in the full simulation is prevented by this change, confirming that the collapse is indeed mediated by the effect of the zonal modes (which was turned off) on the nonzonal modes (which carry the heat flux). The “quasi-saturated” streamer-dominated state is thus just the saturated state that would have emerged had the zonal flows been prohibited or suppressed.




Figure 5.6: The two simulations corresponding to the red squares in Figure 5.3 are shown: the large-box simulations with $a/L_T = 3.42$ and collisionalities $\nu = \nu_{\text{nom}}/2$ (cyan) and $\nu = \nu_{\text{nom}}/3$ (red). The higher-collisionality case is restarted at $t = 10003.6 a/v_{te}$ with certain terms in the collision operator reduced to the lower collisionality value, whilst others are retained at the higher value: (a) only electron-electron collisions reduced (blue); (b) only electron-ion collisions reduced (purple); (c) only nonzonal collisions reduced (green); (d) only zonal collisions reduced (yellow).

Collisions damp the zonal modes, and in this context are important precisely *because* they are small, as this means that the finite zonal fields that emerge from any fast collisionless evolution [53] are damped very weakly and so can grow to

dynamically significant amplitudes and regulate the turbulence. Let us show that it is the electron-ion collisions that affect the zonal modes in the crucial way; these are momentum non-conserving (for electrons), as they act to relax the electron flow and thereby dissipate the associated current — in other words, they give rise to Ohmic resistivity (see discussion in section 5.2.3). In Figure 5.6, we show time evolution of the heat flux corresponding to two different collisionalities in the same series of simulations (the two shown as red squares in Figure 5.3 — the large-box simulations with $a/L_T = 3.42$). The heat flux in the higher-collisionality simulation is shown by the cyan curve, the heat flux in the lower-collisionality one by the red curve. If we rerun the higher-collisionality simulation with ν_{ei} unchanged but ν_{ee} reduced to the lower value (blue curve), or with the zonal collisionality unchanged but the nonzonal collisionality reduced to the lower value (green curve), there is no significant change in the saturated heat flux. By contrast, if we reduce only ν_{ei} , leaving ν_{ee} unchanged (purple curve), or if we reduce only the zonal collisionality, leaving the nonzonal collisionality unchanged (yellow curve), the heat flux drops to a value consistent with the lower-collisionality case (red). Thus it is the electron-ion collisions on the zonal component that primarily determine the heat-flux collisionality dependence.

5.3.5 Simplified simulations for extended collisionality range

We have argued that it is the effect of electron-ion collisions on the zonal modes that sets the collisionality dependence of the saturated heat flux. Collisions in the nonzonal gyrokinetic equation (5.19) are regularising, in that collisions are needed to dissipate the fine structure that the distribution function develops in velocity space. This means that, while collisions cannot be dropped for the nonzonal modes, the particular value of ν has a weak effect — for example on linear growth rates (see Appendix D). This is similar to the well-known situation in fluid turbulence,

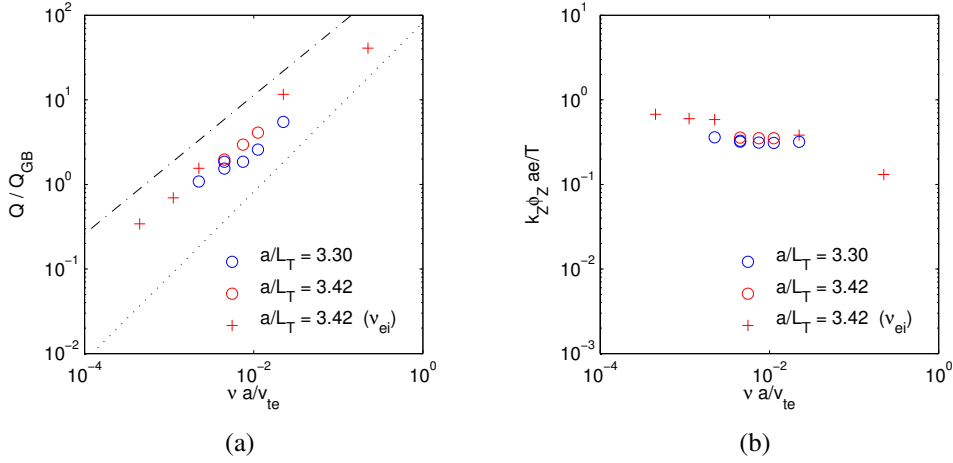


Figure 5.7: Variation of (a) the time-averaged electron heat flux Q/Q_{GB} , (b) the rms zonal velocity $k_Z \phi_Z$, adding to the points from Figure 5.3 (now circles) further points (red crosses) obtained by varying only the electron-ion collisionality ν_{ei} , but keeping the electron-electron collisionality at the nominal value, $\nu_{ee} = \nu_{nom}$, all at the nominal temperature gradient $a/L_T = 3.42$. The dot-dashed line shows the experimental scaling (5.1) and the dotted line shows the theoretical linear scaling, $Q/Q_{GB} \propto \nu_*$, equation (5.31).

where a small viscosity is needed to provide dissipation but the saturated state is independent of the exact value of this viscosity. This understanding of the underlying physics opens up an opportunity to probe the collisionality dependence of electron transport without paying the high price of increased velocity-space resolution that reducing collisionality would exact. The strategy is to vary only the collisionality affecting the zonal modes or only ν_{ei} (although in the latter case, the subdominant part of the zonal damping rate $\sim \nu_{ee} k_x^4 \rho_e^4$ — see section 5.3.3 and Appendix F.7 — will eventually take over).¹¹

Figure 5.7 shows the collisionality dependence of the electron heat flux over a wider range of values of ν_{ei} than in Figure 5.3, accessed using the latter strategy: varying ν_{ei} only while keeping ν_{ee} at its nominal value. We see that the result of

¹¹We further stress that one should be careful in the interpretation of such simplified simulations outside the range of collisionalities used for the full simulations. The simplified simulations could become unrepresentative of the original system if the fixed electron-electron collisions suppress new modes that would otherwise have emerged and dominated the dynamics at either low or high collisionality. We have not investigated whether this is the case for the simulations reported here.

this extended collisionality scan is to confirm the general plausibility of our picture of the ETG transport: Q/Q_{GB} stays approximately proportional to ν_{ei} across a wider range of its values, and the independence of the zonal flow velocity of ν_{ei} also approximately persists over this wider range. Note that these simulations were performed with reduced perpendicular spatial resolution and without flow shear; see Appendix C.2 and Appendix E for details.

5.3.6 Spatial structure of the saturated state

Finally, let us provide some details about the structure of the zonal and nonzonal components of the turbulence in its final saturated state.




Figure 5.8: Zonal electrostatic potential ϕ_Z in a flux-tube cross-section in the out-board midplane as a function of the radial spatial coordinate x and time, for the same “large-box” simulation as shown in Figure 5.2 ($\nu = 0.2\nu_{nom}$, $a/L_T = 3.3$). The black line shows the start of the time-averaging window used to characterise the saturated state.

The zonal component of the saturated turbulent field is long-lived, with a spatial structure that can persist, once established, for as long as the total simulation time. Figure 5.8 is a Hovmöller (space-time) diagram of the zonal potential for the same

case as shown in Figure 5.2. The zonal pattern barely changes during the time window used for time averaging the heat flux and other quantities in the saturated state of this simulation.

The spatial spectrum of the zonal potential is almost monochromatic in this particular case, as shown in Figure 5.9(a). In general, we find that the zonal spectra typically have only a small number of sharp peaks. A single extra harmonic is just visible in Figure 5.9(a); another example is given in Figure E.3.




Figure 5.9: For the same case as Figure 5.2: (a) the spectrum of the zonal potential; (b) the heat-flux spectrum; both are averaged over the time window of the saturated state, which is shorter for the large-box simulation (yellow) than for the small-box simulation (green).

As an aside, one need not be too concerned here that $k_x \rho_e$ is sufficiently small as to be approaching ion scales. First, ion scales do not actually appear in our simplified model: the ion mass, and hence gyroradius scale, are effectively infinite. Of course that does not preclude their relevance in reality. Second, however, one expects ion physics to occur similarly at small values of $k_x \rho_i$, and so the separation between the two species scales is not necessarily reduced just because these values are small.

By contrast with the monochromatic zonal spectrum, the nonzonal spectrum is broadband when averaged over time. Figure 5.9(b) shows the heat-flux spectrum against k_y , and Figure 5.10(a) shows the nonzonal ϕ spectrum for the large-box simulation of Figure 5.2 against both k_x and k_y . Again by contrast with the static




Figure 5.10: 2D spectra of the nonzonal electrostatic potential for the same large-box simulation as in Figure 5.2: (a) the spectrum averaged over the saturated time-window $t \geq 7502.3 a/v_{te}$; (b) instantaneous spectrum at the same time as Figure 5.1(b); (c) & (d) instantaneous spectra at later times during the saturated state.

zonal component, the nonzonal component is rapidly fluctuating. Figures 5.10(b)-(d) show the same 2D spectrum of ϕ at particular instants of time. Thus, the nonzonal spectrum has a “flickering” appearance: at any given time, a small number of modes are much more intense than the others, but the dominant modes change over time, giving rise to a smooth time-averaged spectrum. This high “k-space intermittency” is perhaps natural in the near-marginal saturated state at experimentally relevant parameters, with only a small number of modes excited at any given time.

5.4 Summary and discussion

The prevailing view of the structure of ETG turbulence in tokamaks and the associated levels of transport has its origin in the first gyrokinetic simulations, which did not include collisions [41, 42] (and did not correspond to near-threshold con-

ditions in a spherical tokamak). Whereas the ITG turbulent state has long been believed to be zonal-flow dominated [51], the ETG fluctuations were characterized by long radial eddies (“streamers”) that enhanced the transport to a level comparable with ITG turbulence, overcoming the reduction by a factor of $(m_e/m_i)^{1/2}$ expected from the relationship between the electron and ion gyroscopes (the scales at which the two types of fluctuations were driven).

The present study differs from the more traditional approach to modelling ETG turbulence in three respects: collisions are included; simulations are run for a much longer time; and the equilibrium parameters are relevant to the experimental situation in a real device, namely MAST, and therefore place the system close to a marginal state with respect to the ETG drive (note also that MAST is a spherical tokamak, so has a somewhat different magnetic geometry compared to the more prevalent large fusion devices such as TFTR, JET or ITER).

As a result, we find that, in application to the physical regime we have considered, the standard picture of ETG transport is in need of substantial revision. The high-transport, streamer-dominated nonlinear state does indeed emerge, but persists only transiently, over relatively short simulation times (short compared to the energy confinement time but still long compared to a typical eddy turnover time, and long compared to typical simulation times used for ETG in the past). It turns out that this state is not entirely steady — while the heat flux might appear to be statistically stationary, there is a slow growth of the zonal component of the fluctuations, which eventually (after $tv_{te}/a \sim$ a few thousand) reaches dynamical strength compared to the transport-setting nonzonal modes and proceeds to change the character of the turbulence. A new, long-time, zonal-dominated saturated state emerges, whose structure is more reminiscent of what is traditionally expected of ITG, rather than ETG, turbulence (see Figure 5.1 and further discussion in section 5.4.2). We emphasize that it is the *final* saturation level of the heat flux in gyrokinetic flux-tube simulations, averaged over the turbulent fluctuation scales in length and time, that

is physically relevant for transport calculations. We have found that the turbulent heat flux supported by the new long-time nonlinear state can be much lower than in the “quasi-saturated” streamer-dominated state if the collisionality of the plasma is low (see Figures 5.2(a), 5.3(a) and 5.7(a)). The (roughly linear) collisionality dependence of the heat flux found in our simulations turns out to be in remarkably good agreement with the experimental scaling (5.1) [37].¹²

We have proposed a phenomenological argument (section 5.3.2) whereby this collisionality scaling can be understood if one assumes that the saturation of the nonzonal modes is governed by the zonal gradients coming into approximate balance with the equilibrium gradients (to be more precise, the zonal-nonzonal interactions balancing the linear drive), while the saturation of the zonal modes is set by a balance between their nonlinear excitation by the nonzonal interactions and their damping by Ohmic resistivity. The latter is operative because electron-scale flows are also currents — and so it is the electron-ion collisions that play the defining role in setting the zonal damping rate. We have supported our view by a series of numerical experiments that confirmed the crucial role of the zonal modes in enabling the emergence of the new saturated state (section 5.3.4) and the crucial role of the electron-ion collisions on the zonal modes in setting its collisionality dependence (sections 5.3.4 and 5.3.5). The Ohmic damping of the zonal flows is an analytical result (Appendix F), but we have also systematically confirmed that it is captured correctly in our simulations (section 5.3.3). We have also documented some key qualitative features of the long-time saturated state: the long-time coherence and approximate monochromaticity of the zonal modes (cf. [55, 56]) and, in contrast, the dominant individual nonzonal modes “flickering” with time in and out of existence to give rise to a broad-band time-averaged spectrum (section 5.3.6).

¹²Earlier simulations of ETG in MAST [28, 13] did report heat fluxes roughly consistent with the estimated experimental level of electron heat transport in the cases that were simulated. Whilst we obtain a difference in heat flux between the “quasi-saturated” state and the final state across the entire range of collisionality shown in Figure 5.3, including at the nominal collisionality, the most dramatic differences occur when the collisionality is well below its nominal value (see Figure E.2(a)).

5.4.1 Previous work on ETG and the collisionality scaling

Long-time changes in the saturated state of gyrokinetic simulations have previously been reported in other numerical studies. Mantica et al. [35] reported a change in the zonal-nonzonal balance at long times in an ITG simulation. Guttenfelder and Candy [33], in their collisionless ETG simulations with adiabatic ions using the GYRO code, found a long-time reduction in transport associated with an increased level of zonal perturbations, which occurred at a low but not at a higher level of flow shear. We have not investigated here the dependence on flow shear, except via the simplified simulations with zero flow shear reported in section 5.3.5. It is not impossible that the transport scaling described in this chapter exists within a window in which the flow shear is not too large (although it should perhaps still be large enough to suppress ion-scale transport in order for a heat-flux calculation restricted to electron scales to make sense).

Perhaps most relevantly for comparisons with our work, Nakata et al. [57], using a slab ETG model, adiabatic ions, a Lorentz collision operator (i.e., the pitch-angle scattering operator in (5.15)), and comparing two different sets of parameters, reported that transport at long times in their linearly more unstable case could be suppressed below the level of their linearly more stable case, owing to the formation of zonal flows that collimated the turbulence into “vortex streets” and acted as a barrier to radial transport (see also earlier studies of ETG zonal flows [58, 59]). This suppression appears to be consistent with our findings.

An alternative explanation for the experimental collisionality scaling of heat transport in spherical tokamaks that has been previously suggested relies on the transport associated with microtearing turbulence [37, 60, 61]. Since, in the present work, we have limited ourselves to electrostatic perturbations, the microtearing instability is excluded and so it is clear that the collisionality scaling we have found does not require it. Generally speaking, the arguments that we apply to ETG

are fairly generic and may well apply in some form to other instabilities. We have, of course, also not excluded the possibility that other modes may produce a similar scaling for different reasons. The microtearing contribution in particular to the overall electron transport in real fusion devices remains a live and pressing research subject (made challenging, however, by the particular difficulty of obtaining well-resolved simulations of electromagnetic turbulence in tokamaks).

It is an interesting question how general our picture of ETG turbulent transport might prove to be. In this connection it should be noted that conventional tokamaks, as opposed to spherical ones, have not been found to exhibit the same strong scaling (5.1) of confinement time with collisionality: the scaling exponent between $B\tau_E$ and ν_* is essentially zero [62, 63]. Clearly, identifying the reason or reasons for this difference is an important issue for further study.

5.4.2 ETG vs. ITG turbulence near and far from threshold

The idea that the ETG turbulent state is dominated by zonal modes leads one naturally to the question of whether the saturated states of ETG and ITG turbulence are essentially similar, at least qualitatively.

Our ETG turbulence model, equations (5.6) and (5.14), describes kinetic electrons with Boltzmann (adiabatic) ions. The simplest possible model for ITG turbulence would involve kinetic ions with Boltzmann electrons — the latter physically justified by fast streaming of electrons along field lines. The main mathematical difference between these two models is that the Boltzmann electron response must be restricted to perturbations that have variation along magnetic-field lines. Namely, the density perturbation is, in contrast to the ETG case (5.13),

$$\frac{\delta n}{n} = \frac{e(\phi - \bar{\phi})}{T_e}, \quad (5.32)$$

where $\bar{\phi}$ is the flux-surface average [64, 65, 40]. It is this latter feature that is normally believed to be responsible for the difference between the zonal-flow-

dominated ITG state and the streamer-dominated ETG state: radial variations in the zonal ϕ do not affect δn as electrons cannot respond to radially varying zonal modes at ion scales; these can, therefore, grow large enough to break up the primary ITG-driven streamers via a secondary instability [66, 52] and isotropise the turbulence, whereas for ETG turbulence, the latter effect was believed to be too weak to destroy the streamers [41, 42] (although the physics that determined the radial scale of the streamers perhaps remained unclear [67] — possibly again a (weaker) secondary instability [66, 68]).

What we have found in our simulations is that, whereas there is indeed no trace of a fast onset of a secondary instability similar to one that breaks up the ITG streamers (and is indeed seen in most ITG simulations), the ETG streamers' lease on life granted by a stronger density response to zonal perturbations is nevertheless only temporary: the zonal component does find a way to grow slowly, until it is large enough to take control over the nonzonal fluctuations.

In addition to the ITG instability giving rise to streamer-like modes and the secondary instability breaking them up into zonal flows, the third pillar of the standard picture of ITG saturation is a “tertiary instability” whereby the zonal modes that have grown to a certain critical amplitude break up, returning energy to nonzonal perturbations [52]. This mechanism of regulating the zonal component of the turbulence does not rely on collisional damping and thus should give rise to heat fluxes that are mostly independent of collisionality — as indeed it appears to do, at least when the ITG turbulence is simulated in regimes that are far from the instability threshold [48]. However, it has long been known that at temperature gradients that are close to the linear threshold (and thus, arguably, most relevant experimentally), the tertiary instability is ineffective, and a state of strongly suppressed transport ensues, called the Dimits upshift of the critical temperature gradient [16, 52]. Clearly, in the case of ETG turbulence, should a tertiary instability of the zonal modes appear, their collisional damping would cease to matter and consequently the colli-

sionality dependence of the heat flux should flatten off (cf. [69]). It is then tempting to suppose that our ETG state, in which the zonal modes are stable, collisionally damped and the heat flux scales with collisionality, is simply the ETG version of the Dimits-shift regime. Ascertaining whether this is indeed the case requires a larger parameter scan (in particular, in the temperature gradients) than has been undertaken here and has to be left for future work. By analogy with the ETG case, another tempting conjecture would be that the Dimits regime for the ITG turbulence should also have an (ion) collisionality scaling, by way of an argument similar to the one presented in section 5.3.2.¹³ This in fact is not a new idea [70, 71, 51, 69], although the situation remains murky as including more realism (in particular, kinetic electron response) in gyrokinetic simulations of ITG turbulent transport can nullify or even reverse its collisionality dependence [72]. Since the collisionality scaling of the ETG transport that we report at least has some experimental support [37], we feel more secure about the relevance of our conclusions — although “stress-testing” them by relaxing various assumptions (e.g., perhaps most urgently, allowing kinetic ions [67]) and by extending the equilibrium parameter ranges is an essential direction for future research.

Thus, the relationship between ETG and ITG is complicated and indeed within each of these two regimes, there are likely to be several sub-regimes depending on the part of the parameter space that one is interested in — in particular, how far from the threshold the turbulence is. This said, our results suggest that there are perhaps more similarities between ETG and ITG physics than previously

¹³In porting this argument to the ion scales, one must not forget that an important difference between the ETG and ITG cases is the specially pivotal role played by electron-ion collisions for the ETG zonal modes. For ions, ion-electron collisions are always subdominant in the mass-ratio expansion because the ions carry nearly all of the momentum, and collisions with the lighter electrons have a negligible impact on the ion flow. By contrast, the relative flow of electrons (which constitutes a current) is strongly affected by collisions with the heavier, sluggish ions. Therefore, for electrons, collisions are not momentum-conserving, whereas for ions, to lowest approximation they are. How this difference plays out in the calculation of the rate of damping of ion zonal flows is outlined in Appendix F.8. In the end, it turns out, however, that the damping rate scales in the same way as for the electron case, $\gamma_{Zi} \sim \nu_{ii} k_x^2 \rho_{pi}^2$, although for a different reason.

thought. Seeking a unified and universal description of various types of drift-wave turbulence may therefore be a worthwhile and reasonable aspiration.

In conclusion, let us summarise again what has been achieved here. We have found that in the saturated state of ETG turbulence at driving gradients close to experimental levels in MAST, the electron heat transport decreases with collisionality, in agreement with experimental evidence [37]. This behaviour points toward improved confinement in future devices. We have explained it based on a simple theoretical picture (backed up by numerical tests) of an ETG turbulent state dominated by the interaction between nonzonal and zonal modes, with the collisionality dependence of the heat flux originating from the collisionality dependence of the resistive damping of the zonal modes. On a practical note, ETG simulations with generic initial conditions must be run to very long times to capture the (crucial) effect of zonal modes on the saturated state.

CHAPTER 6

Conclusions and further work

In this thesis we have seen that, for a typical MAST plasma, the turbulent electron transport level in the outer core – as obtained via a multi-scale approach, using local gyrokinetic flux-tubes to calculate fluxes within a large-scale transport model, over a macroscopic range of the plasma minor radius – is broadly consistent with that arising due to electrostatic ETG turbulence. We have also seen that a crude model for the neoclassical level of ion transport may be adequate for the electron modelling, because interspecies equilibration is low compared to sources. During these simulations we found that the zonal fluctuation components of the electron turbulence could slowly grow and cause the saturated heat flux to reduce. We then investigated this phenomenon in more detail for a single flux surface, simplifying our physics model further to a pure Boltzmann ion response. We found that, for the parameters at this flux surface, the final saturated state is determined by a balance between nonlinear drive and weak zonal damping rate $\sim \nu_{ei} k_x^2 \rho_{pe}^2$, where $k_x^2 \rho_{pe}^2 \ll 1$, mediated by the effect of the zonal component on the nonzonal turbulence. This leads to an electron heat flux $Q_e \propto \nu_{ei}$, which agrees with the experimentally obtained scaling of confinement time with collisionality.

In concluding this thesis, we will now consider some further questions arising from these results, and pointers towards potential further work.

The first question arising is to what extent the local turbulence behaviour investi-

gated in chapter 5 is relevant to global transport simulations of the type in chapter 3? This is a subtle issue because the transport solver should automatically adjust the parameters to obtain and then maintain the correct level of heat flux for power balance. So whilst within the turbulence simulation at a single transport iteration we might see a fall in the heat flux as the zonal component grows, over multiple transport time steps we should see a change in the temperature gradient that ultimately keeps the turbulence far enough above threshold. Moreover, this change could be quite small, owing to the stiffness exhibited in Figure 3.7(b). The evolution of the zonal component would therefore be different to that in the fixed-gradient simulations of chapter 5, and it may be necessary to develop new diagnostics for convergence. To assist with convergence to the long-time steady state, it should be beneficial always to restart flux tubes from their current state, rather than reinitialising.

We have not investigated fully the range of plasma parameters over which our local collisionality scaling regime holds. This will affect to what extent it can determine the global confinement scaling, an issue that is certainly relevant to future devices, including MAST Upgrade, which may be able to access lower collisionalities. It would be interesting to run model simulations for different plasma conditions. An initial attempt could be made to model hotter, less collisional plasmas by increasing the outer boundary temperature in the Trinity simulation.

Even a single Trinity simulation requires significant high-performance computing resources, and this is another issue for future multi-scale transport simulations. In chapter 5 we presented evidence that more spatial resolution and more spectral resolution are required at the two ends of the collisionality range, but also that it may be possible to avoid this requirement by varying only electron-ion collisionality. The shortcut is worth investigating further, as otherwise even crude global simulations are too expensive for routine use given current computing capabilities.

For practical and especially predictive modelling purposes, another important avenue for further work is to investigate any effect on the collisionality scaling of

a more complete physics model, in particular including kinetic ions and/or electromagnetic perturbations. To capture the ion physics fully, however, by spanning both ion and electron gyroradius scales, would require much greater computational resources, even for a single flux-tube simulation, than is routinely available. Electromagnetic perturbations may be important in real devices with high enough plasma β , especially for electron modes – as we have noted, electron flows are also currents – but modelling these is again very computationally challenging.

The dominant role of zonal-nonzonal interactions in the nonzonal electron GKE in this work suggests that it may be possible to model gyrokinetic turbulence without the nonzonal-nonzonal term in the nonzonal GKE (i.e. without NZ-NZ-NZ triad interactions). This has been done in the geophysical fluid context [73] and for fluid-like models in the plasma context [55, 56], and it would be interesting to investigate it in gyrokinetics. The approach may not be any more efficient than the full physics if implemented straightforwardly in GS2, but may eventually lead to resource savings (if indeed successful as a physics model), by enabling either a drastic reduction in number of k_y 's, or a fully spectral code without the overhead cost of performing Fourier Transforms.

Finally, we draw attention to the slow radial propagation of the zonal waveform that is evident in Figure 5.8. Preliminary investigations indicate that this is due to the small background flow shear $\gamma_E \neq 0$. An analogous atmospheric phenomenon, the poleward migration of zonal jets, has been seen in the geophysical context [74], and it would be interesting to pursue this connection.

APPENDIX A

Correspondence between energy confinement time and heat flux scalings

In the experimental scaling (5.1), τ_E is the energy confinement time, i.e., the time it takes the power loss $P \sim QA$ to carry away the thermal energy $W \sim nTV$, where A and V scale with flux surface area and volume respectively, Q is the energy flux, n is density, and T is temperature. Therefore,

$$\tau_E \sim \frac{W}{P} \sim \frac{nTa}{Q}, \quad (\text{A.1})$$

where $a \sim V/A$ is the macroscopic length scale.

The appropriate normalising quantity for Q is the ‘‘gyro-Bohm’’ heat flux $Q_{\text{GB}} = nTv_{te}\rho_*^2$, where $v_{te} = (2T/m)^{1/2}$ is the thermal speed, $\rho_* = \rho_e/a$, $\rho_e = v_{te}/\Omega_e$ is the gyroradius and Ω_e is the cyclotron frequency (since we are interested in electron transport, we are using electron quantities, but the argument for ions is exactly the same). In the gyrokinetic ordering, the heat flux is $Q = \mathcal{O}(\rho_*^2)$ automatically (see, e.g., [3]), so the normalised heat flux Q/Q_{GB} must be order unity and independent of ρ_* . This normalised heat flux is

$$\frac{Q}{Q_{\text{GB}}} = \frac{Q}{nTv_{te}\rho_*^2} = \frac{Q}{nTa\Omega_e\rho_*^3} \sim \frac{1}{\Omega_e\tau_E\rho_*^3}, \quad (\text{A.2})$$

where we have used (A.1) to obtain the last expression. At constant ρ_* , (A.2) implies

$$\frac{Q}{Q_{\text{GB}}} \propto \frac{1}{B\tau_E}, \quad (\text{A.3})$$

so we are probing the same dependence on collisionality by measuring Q/Q_{GB} in gyrokinetic simulations as the experiments do by measuring $B\tau_E$ with ρ_* fixed.

Let us now explain what is meant when scaling (5.1) is claimed to hold experimentally [37]. The starting point is to assume that the following parametrisation of the energy confinement time holds:

$$\Omega_e \tau_E = K \nu_*^{x_\nu} \rho_*^{x_\rho} q^{x_q} \beta^{x_\beta} \kappa^{x_\kappa}, \quad (\text{A.4})$$

where K is a constant, κ is elongation [23], β is plasma beta, q is the safety factor and $\nu_* = \nu / (\varepsilon^{3/2} v_{te} / qR)$ is the effective electron collision rate ν / ε divided by the bounce frequency $\varepsilon^{1/2} v_{te} / qR$, where $\varepsilon = a/R$ is the inverse aspect ratio, and R the major radius. This is the dimensionless quantity characterising collisional detrapping. The experiments looking for the ‘‘collisionality scaling’’ strive to hold all these parameters constant apart from ν_* and measure the scaling exponent x_ν of $B\tau_E$ with respect to ν_* . Since in our numerical study, we do not vary the equilibrium geometry, the same exponent x_ν will apply to the normalised collisionality in our gyrokinetic simulations, which is simply $\nu a / v_{te}$.

Note finally that, in view of (A.2), if Q/Q_{GB} is independent of ρ_* (as it is hard-wired to be within the gyrokinetic ordering), it must be the case in (A.4) that $x_\rho = -3$, which is indeed consistent with the experimental findings [37].

APPENDIX B

Plasma parameters

The reference, or “nominal”, plasma parameters for our simulations represent MAST shot 8500 at $t = 0.289$ s. They are specified in Table B.1. We use a Miller [23] representation of the flux-surface geometry, which is assumed to be axisymmetric.

minor radius	r/a	0.65
major radius	R/a	1.46
safety factor	q	1.9
magnetic shear	\hat{s}	1.8
Shafranov shift	R'/a	-0.25
pressure gradient	$p'\beta/p$	-0.12
elongation	κ	1.57
elongation gradient	κ'	0.40
triangularity	δ	0.22
triangularity gradient	δ'	0.16
density gradient	a/L_n	-1.2
electron temperature gradient	a/L_T	3.42
effective ion charge	Z_{eff}	1.15
species temperature ratio	$T_e/T_i = 1/\tau$	1.06
electron collisionality	$\nu_{\text{nom}}a/v_{te}$	0.02
flow shear	$\gamma_E a/v_{te}$	-0.003

Table B.1: Nominal dimensionless local equilibrium parameters based on MAST shot 8500 at $t = 0.289$ s. Here $a = 0.55$ m is the half diameter of the last closed flux surface at the elevation of the magnetic axis and r is the same for the particular flux surface we have chosen. The prime symbol denotes the derivative with respect to normalized minor radius: $' \equiv a(d/dr)$. Scale lengths are defined by $a/L_n \equiv -n'/n$ (density) and $a/L_T \equiv -T'/T$ (temperature).

In our simulations, we varied the collisionality ν relative to its nominal value ν_{nom} given in Table B.1. We carried out such a collisionality scan for two values of a/L_T , the nominal one given in Table B.1 and a slightly smaller (slightly more marginal) $a/L_T = 3.3$. All other parameters were held fixed.

A small amount of flow shear γ_E , given in Table B.1, was included in the non-linear simulations. As in [13], this was found to assist convergence and is of the same order as the experimental level in MAST (see Appendix E.2). The sign of γ_E is unimportant because the equilibrium is up-down symmetric [75].

APPENDIX C

Numerical details

C.1 GS2 spatial coordinates

GS2 solves the gyrokinetic equation (5.6) in a flux tube that winds its way toroidally and poloidally around the toroidal flux surface designated by a label r (see Table B.1), which is defined as its half diameter at the elevation of the magnetic axis of the tokamak. The position along the flux tube is determined by the poloidal angle θ , which can be thought of as the coordinate parallel to the magnetic field. The flux-tube domain that we use in GS2 extends over $-\pi \leq \theta \leq \pi$, where $\theta = 0$ is the outboard midplane and $\theta = \pm\pi$ the inboard midplane.

If we demand that

$$\mathbf{B} = \nabla\alpha \times \nabla\psi, \quad (\text{C.1})$$

where ψ is the poloidal magnetic flux and α is chosen to satisfy (C.1), each field line is labelled by the pair of Clebsch coordinates (ψ, α) . Then (ψ, α, θ) can be used as the three independent (but not orthogonal!) curvilinear coordinates. Designating the central field line of a flux tube by (ψ_0, α_0) , the GS2 coordinates (x, y) transverse to the magnetic-field direction are related to (ψ, α) in the following manner [76, 77]:

$$x = \frac{q(\psi_0)}{B_0 r(\psi_0)} (\psi - \psi_0), \quad (\text{C.2})$$

$$y = \frac{1}{B_0} \left. \frac{d\psi}{dr} \right|_{r=r(\psi_0)} (\alpha - \alpha_0), \quad (\text{C.3})$$

where the safety factor $q(\psi)$ and the radial flux-surface label $r(\psi)$ are both flux functions and B_0 is a reference (constant) magnetic field defined to be the toroidal magnetic field at the mid-radius of the flux surface containing our flux tube. The x coordinate is “radial” in the sense that it is transverse to the flux surface and increases away from the magnetic axis. The second coordinate y is often called “binormal” because it effectively measures distances perpendicular to the field line but within the flux surface (at constant ψ). In the outboard midplane, (x, y) form an approximately orthogonal grid in the cross-section of the flux tube perpendicular to the magnetic field, with x and y approximately measuring true distances from the centre of the flux tube. Following the field line around the flux surface, the flux tube twists due to magnetic shear, with ∇x , ∇y and $\nabla\theta$ generally not orthogonal to each other.

GS2 is pseudospectral with respect to (x, y) but not θ ; that is, except for the evaluation of the nonlinear term, its data grid has the coordinates $(k_x, k_y) = \mathbf{k}_\perp$ and θ . For the evaluation of the nonlinear term, the Fast Fourier Transform is used to go from (k_x, k_y) to (x, y) (and then back again after the evaluation). In the context of the GS2 grids, the terms “parallel” and “perpendicular” refer to θ and to the other two coordinates (either (x, y) or (k_x, k_y) as appropriate), respectively.

C.2 Numerical grids, resolution and boundary conditions

In our largest nonlinear simulations, the spatial grid has 48 cells in θ and 108×144 points in $x \times y$, the range of positive perpendicular wavenumbers being $(k_x \rho_e, k_y \rho_e) \in [0.012, 0.40] \times [0.012, 0.57]$. We refer to these as “large-box simulations” in the text. The simulations referred to as “small-box simulations” are the same except for the perpendicular grid parameters: the box size in real

space is halved, but extended to somewhat higher perpendicular wavenumbers (i.e., higher spatial resolution): 80×108 in $x \times y$, spanning the spectral range $(k_x \rho_e, k_y \rho_e) \in [0.023, 0.60] \times [0.024, 0.84]$. The parameter regimes in which these two classes of simulations are useful are discussed in Appendix E. The wavenumber ranges are intended to capture electron physics; the Boltzmann ion response ensures that we need not worry about where they lie in relation to ion scales.

Dealiasing is used for the nonlinear term: zero-padding with higher-wavenumber modes is introduced before the Fourier step into real space; any nonlinear interactions that wrap around unphysically in reciprocal space will be nullified when these modes are thrown away after the inverse Fourier step. The number of $x \times y$ points indicated here (applicable only to nonlinear runs) is the same as the total number of modes including the dealiasing modes, but the ranges of wavenumbers given above do not include the dealiasing modes. All data shown here, including plots in real space, derive solely from the evolved, non-dealiasing modes: the additional dealiasing modes are used only transiently during the nonlinear part of each time step, and are not included in the code’s output.

Twist-and-shift parallel boundary conditions [24, 76] link together some of the k_x modes at low k_y , whereas the $k_y = 0$ modes are periodic in θ . At the remaining (open) domain ends, the parallel boundary condition is $h = 0$ for incoming particles. The boundary conditions are automatically periodic in the perpendicular directions (x, y) by virtue of the spectral representation.

The velocity-space grid has 18 energies $\mathcal{E} = v^2/2$, and, for passing particles, 16 grid-points in $\lambda = \mu/\mathcal{E}$ for each sign of $v_{\parallel} = \pm(2\mathcal{E} - 2\mu B)^{1/2}$, where $\mu = v_{\perp}^2/2B$ is the magnetic moment. Trapped particles are accommodated by additional λ grid points, one for particles bouncing at each θ .

Linear simulations (Appendix D) use the same setup as nonlinear ones, but evolve each perpendicular wavenumber independently.

For the “simplified” simulations in section 5.3.5, in which only the electron-ion

collisionality was varied, the grid sizes were: 24 cells in θ ; 40×54 points in $x \times y$, spanning the spectral range $(k_x \rho_e, k_y \rho_e) \in [0.023, 0.30] \times [0.024, 0.41]$; 12 energies \mathcal{E} ; and 8 λ 's for passing particles in each direction. This version of the code applied the zero incoming boundary condition to $g = \langle \delta f \rangle = h + e \langle \phi \rangle F/T$, rather than to h .¹

¹ The option to use h was added to the code during the course of the present work. The difference between the two options vanishes as $\phi \rightarrow 0$ at $\theta \rightarrow \pm\infty$, but it has been found in linear simulations [78] that the use of h leads to faster convergence as the actual, finite θ domain size is increased.

APPENDIX D

Linear growth rates




Figure D.1: Linear growth rates versus k_y , at $k_x = 0$, with zero flow shear $\gamma_E = 0$, for various collisionalities as indicated. All other parameters are the nominal ones given in Table B.1 (including $a/L_T = 3.42$).

Figure D.1 shows the growth rate of the most unstable linear mode at $k_x = 0$ for each k_y in the numerical domain (which was described in Appendix C.2). These linear simulations have zero flow shear, $\gamma_E = 0$.

At the given a/L_T , the effect of decreasing collisionality is to decrease slightly the maximum growth rate, and to extend to somewhat lower k_y the unstable wavenumber interval. At $k_y \rho_{te} \approx 0.2$ where the saturated heat flux is maximal (see Figure 5.9(b)), the dependence of the linear spectrum on collisionality is weak.

Note that this scan is consistent with the previous observation [13] that collisions lead to a stability gap between ETG- and ITG-driven modes (the range of k_y 's unstable to ITG being lower than the wave numbers shown here).




Figure D.2: Linear growth rates versus a/L_T , with zero flow shear $\gamma_E = 0$, for various collisionalities as indicated, showing that the linear critical gradient (where the growth rate crosses zero) is insensitive to collisionality.

Figure D.2 shows the dependence of the maximum growth rate on the temperature gradient (still with $\gamma_E = 0$). We see that the linear critical gradient varies only between 2.3 and 2.5 throughout this whole range spanning a factor of 50 in collisionality. This is in agreement with the earlier findings [32] that collisions have very little effect on the linear critical gradient.

Note that while these linear simulations were done at zero flow shear, our non-linear runs had $\gamma_E = -0.003 v_{te}/a$, a value of the same order as the measured rotational shear in the MAST shot on which the rest of our equilibrium parameters were based. The ETG instability in the presence of flow shear can become subcritical in the sense that growth is transient and all modes eventually decay (cf. [79, 80, 81, 82, 83, 84]). Given a finite initial perturbation, such transient growth still leads to a saturated nonlinear state (because long enough transient growth can be as good as a formally unstable mode as far as maintaining turbulence is concerned

[80, 81, 82, 84]), although we have found the minimum value of a/L_T required for this to be higher than the critical gradient implied by Figure D.2: it is typically between 3 and 3.3. The mapping out of a “zero-turbulence manifold” (cf. [85, 86]) for the ETG problem has been left outside the scope of this study — and would be a formidable computational challenge, requiring very many very-long-time simulations. It appears plausible that the key parameters in such a study would be the electron temperature gradient a/L_T and collisionality $\nu_{ei}a/v_{te}$, rather than the ion temperature gradient and the perpendicular flow shear γ_E , which has a profound effect on the turbulence threshold in ITG [85, 86, 84]. The magnetic geometry factor $q/\varepsilon \approx B/B_p$, which appears in the scaling (5.31), may play a similar role — improving confinement as this quantity is reduced — in both cases.

APPENDIX E

Issues of numerical convergence in nonlinear simulations

We discuss here the selection of runs used to construct Figure 5.3, and the associated numerical convergence issues. We also include additional evolution plots, complementing the two that were shown in Figure 5.2. This appendix concludes with a short discussion on the role of flow shear.

E.1 Convergence of the heat flux

Figure 5.3 was based on a set of simulations picked from a larger parameter scan that we have carried out. The choice of simulations was based on whether we deemed them to be numerically converged. Here we explain how this was decided.

Figure E.1 shows the results of both small- and large-box simulations (see Appendix C.2 for the explanation of what this means), over a range of collisionalities and for two values of the temperature gradient. Runs excluded from Figure 5.3 are shown as hollow points. They have been excluded because one or more important physical quantity measured in these runs was found to be dominated by either the box scale (the lowest k) or by grid scales (the highest k 's), and, therefore, the convergence of the simulation with respect to box size and/or spatial resolution could

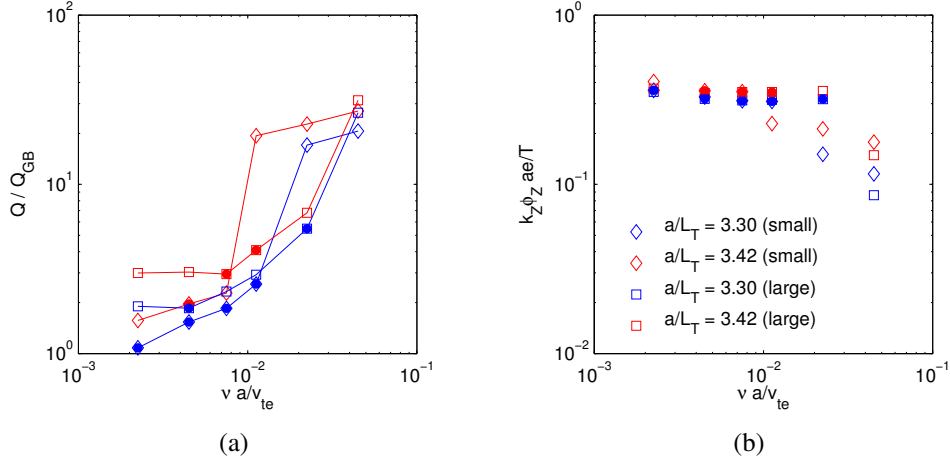


Figure E.1: Variation of (a) the time-averaged normalised electron heat flux Q/Q_{GB} and (b) the rms zonal velocity $k_Z \phi_Z$, defined by (5.17), versus normalised electron collisionality $\nu a/v_{te}$ for our entire simulation series using both small and large boxes, and two temperature gradients, including unresolved runs (hollow symbols; infilled symbols are the resolved points shown in Figure 5.3).

not be relied upon.

Two particular pathologies have been identified, and are also visible in Figure E.2, which shows the time evolution of the full set of large-box simulations at the nominal temperature gradient ($a/L_T = 3.42$).

(i) As the collisionality is reduced, it can be seen from the growth of the zonal velocity (Figure E.2(c)) that the increasing zonal ϕ^2 (Figure E.2(b)) more than compensates initially for the gradually falling zonal wavenumber (Figure E.2(e)), until the final level of zonal velocity is reached. However, the zonal shear (Figure E.2(d)) remains relatively constant in time. For the highest collisionality (yellow), the heat flux never collapses, and we believe that in this case it is because the required zonal wavenumber is too close to the simulation box scale; see Figure E.2(e). Such uncollapsed runs are visible as the high points in Figure E.1(a).

(ii) At the low end of the collisionality scale in Figure E.2 (blue, red and green), it can be seen that there is no significant difference in the heat fluxes as the collisionality is varied, leading to the plateau of hollow red squares in Figure E.1(a) (and a

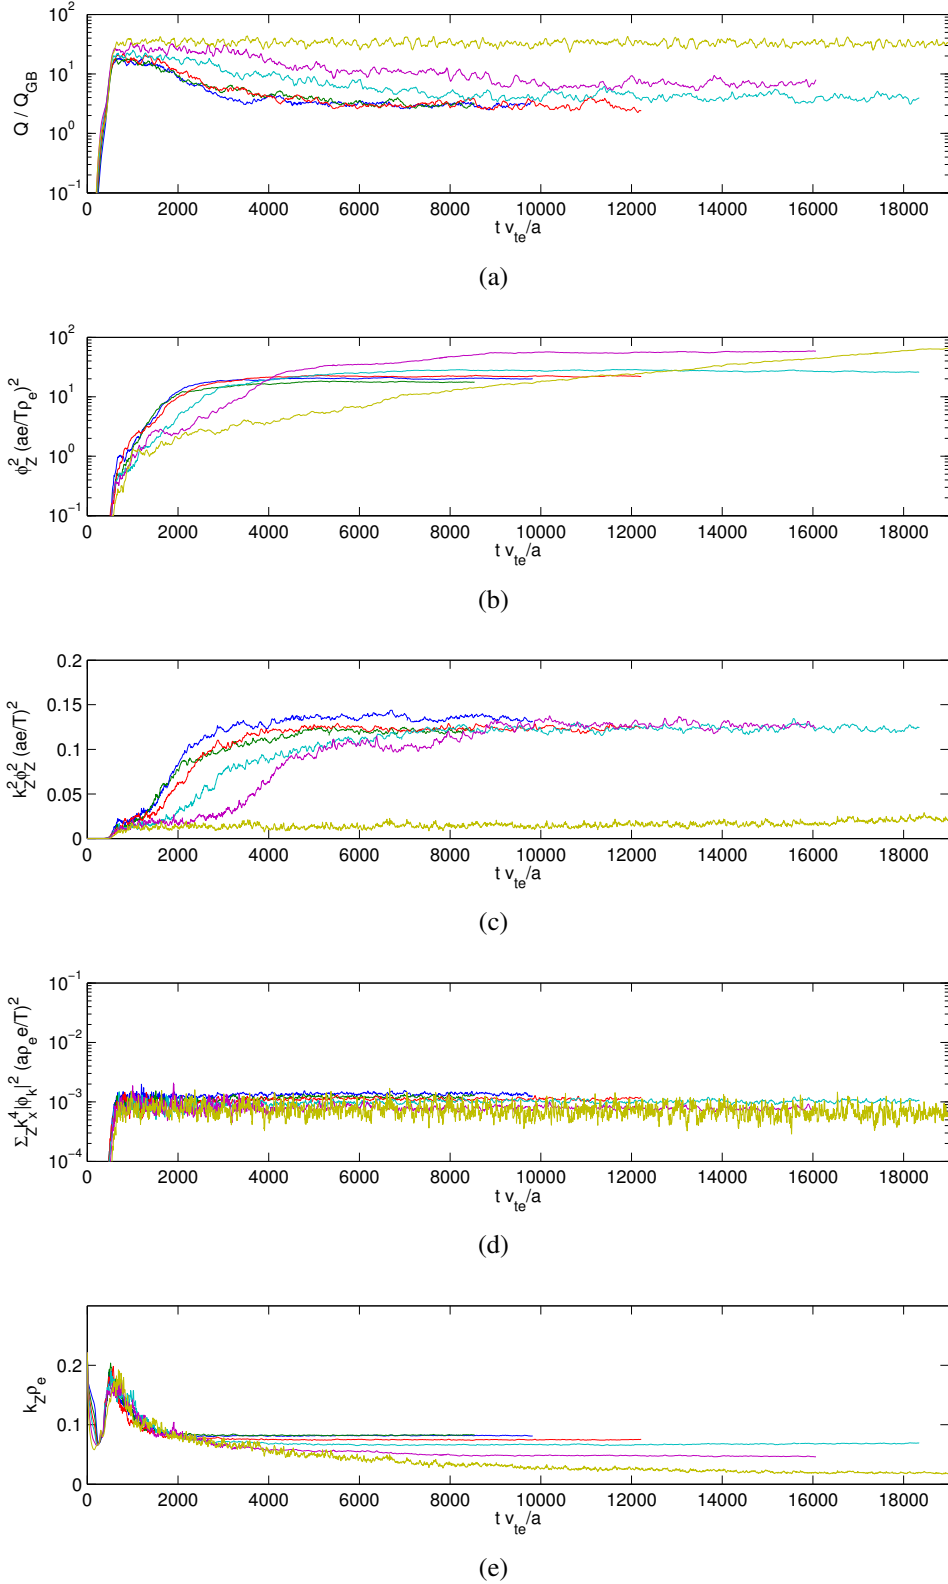


Figure E.2: Evolution in time of (a) the normalised turbulent electron heat flux Q/Q_{GB} , (b) the square of the zonal electrostatic potential, $\phi_Z^2 = \sum_{k_x} |\phi_{k_x,0}|^2$, (c) the square of the zonal velocity, $(k_Z \phi_Z)^2 = \sum_{k_x} k_x^2 |\phi_{k_x,0}|^2$, (d) the square of the zonal shear, $\sum_{k_x} k_x^4 |\phi_{k_x,0}|^2$, and (e) the rms zonal wavenumber k_Z , for large-box simulations at the nominal $a/L_T = 3.42$ and at various collisionalities (colours: same as Figure D.1, plus yellow for $\nu = 2 \nu_{nom}$).

similar plateau for the other temperature gradient, blue squares). These runs are also excluded because of box-scale effects. In these cases, somewhat counterintuitively, the small-box runs, which, however, extend to higher k 's, appear to be resolved and do not exhibit the plateau. Figure E.3 shows an example (the lowest-temperature-gradient and lowest-collisionality case), in which the perpendicular temperature is dominated by the box scale for the large-box run, which is excluded.




Figure E.3: Spectrum of the zonal perpendicular temperature perturbation, for $\nu = 0.1 \nu_{\text{nom}}$, $a/L_T = 3.3$, for a small-box simulation (green) and a large-box simulation (yellow), at the final simulation time in each case.

Eliminating both pathologies gives a fairly clean collisionality scaling when the remaining resolved runs are combined. Given that the zonal damping, according to the scaling demonstrated in section 5.3.3, is lowest for the box-scale mode, it is not surprising that there are convergence problems at this end of the spectrum. Nonetheless, all the simulations that we consider resolved have saturated with the nonlinear state dominated by zonal modes with wavenumbers above the box-scale mode.

Thus, to summarise, at higher collisionalities, we need larger boxes but can sometimes get away with less spatial resolution, whereas at low collisionalities, we need higher spatial resolution, but can sometimes get away with smaller boxes.

For completeness, we include as Figure E.4 the time traces of various zonal quantities for the simulations used in Figure 5.2. When these two figures are combined, the complete set analogous to Figure E.2 is obtained. Note that these demon-

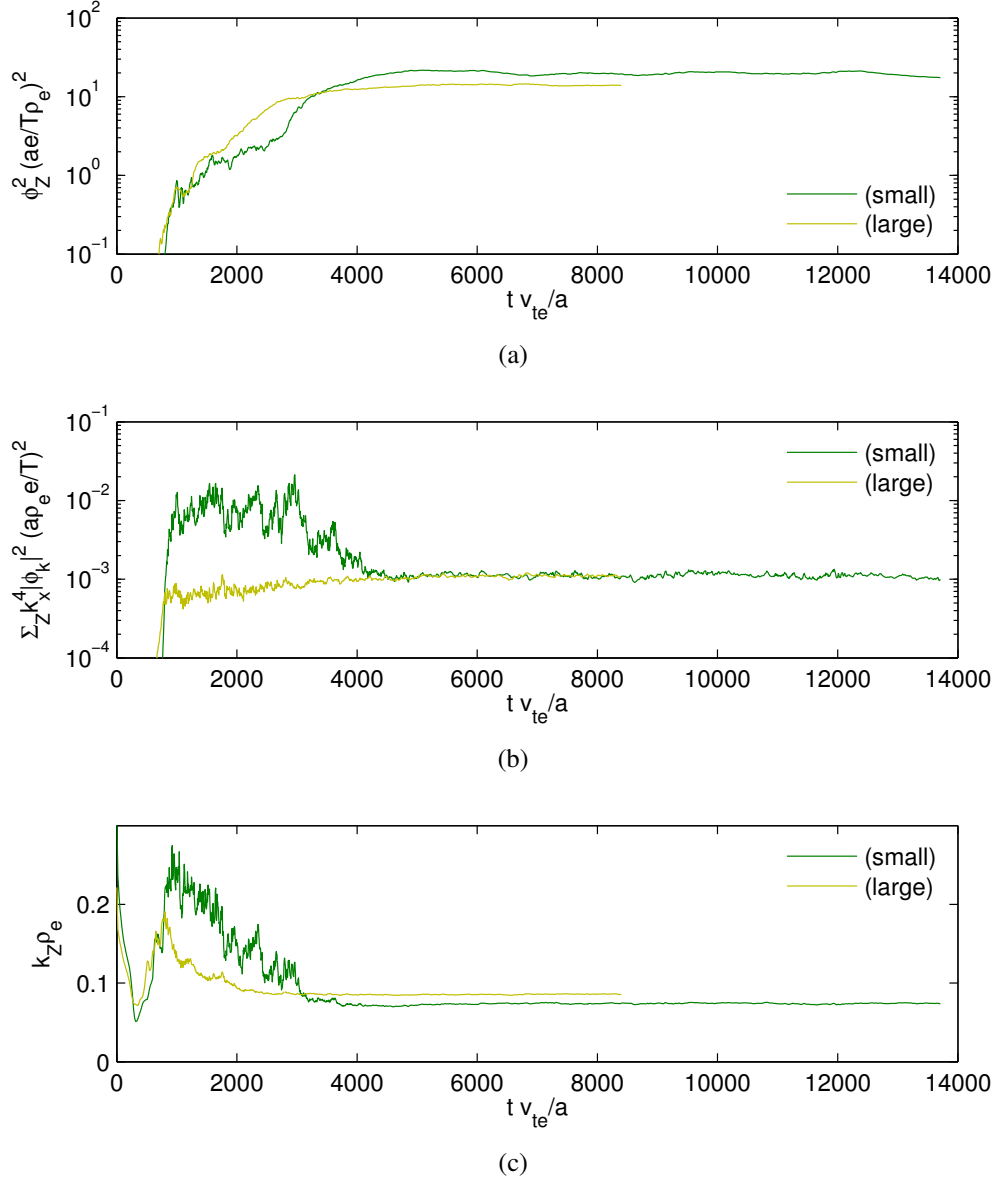



Figure E.4: Evolution in time of (a) the square of the zonal electrostatic potential, $\phi_Z^2 = \sum_{k_x} |\phi_{k_x,0}|^2$, (b) the square of the zonal shear, $\sum_{k_x} k_x^4 |\phi_{k_x,0}|^2$, and (c) the rms zonal wavenumber k_Z , for $\nu = 0.2 \nu_{nom}$, $a/L_T = 3.3$ (green: small-box simulations; yellow: large-box simulations). This complements Figure 5.2.

strate some lack of convergence in the “quasi-saturated” state, but the long-time state is resolved. Spatial cross-sections of the electrostatic potential for these two simulations are shown in Figure 5.1 (large-box) and Figure E.5 (small-box). The morphology of the quasi-saturated state is qualitatively distinguishable between the two simulations. This difference is perhaps associated with the clear difference in the zonal shear shown in in Figure E.4(b). The small-box simulation has higher spatial resolution, i.e., a higher maximum k , and the higher power of k_x in the zonal shear (compared to the zonal velocity) emphasises the high end of the spectrum. The long-time saturated states, on the other hand, appear qualitatively the same. Note that in the large-box simulation in Figure 5.1(b), there are seven periods of the dominant zonal wavelength across the box, whereas in the doubled small box in Figure E.5(b) there can only be an even number, which is six (two times the three across each smaller periodic box). This level of discrepancy between the zonal wavenumbers must be expected given the finite domain size.


E.2 Background flow shear

Various authors have reported runaway growth in ETG turbulence simulations. To aid saturation, Guttenfelder and Candy [33], following earlier work [87], added background flow shear to their simulations. Roach et al. [13], extending the simulations of Joiner et al. [28], observed runaway growth that was healed by adding either flow shear or sufficient collisions (attributed by them to the stabilisation of weak trapped-electron modes close to the box scale at $k_y \rho_e \approx 0.02$).

In our simulations, we have included an experimentally realistic level of flow shear, $\gamma_E = -0.003 v_{te}/a$, and we do not find runaway behaviour. Note that in all of the cases reported above, the normalised perturbed zonal shear (Figure E.2(d) and Figure E.4(b)) is an order of magnitude larger than the background flow shear. In the “simplified” simulations of section 5.3.5, zero flow shear was used, and it is



(a)



(b)

Figure E.5: Electrostatic potential ϕ at the outboard midplane, for $\nu = 0.2\nu_{\text{nom}}$, $a/L_T = 3.3$: (a) quasi-saturated state at $t = 1204.2 a/v_{te}$, (b) long-time saturated state at $t = 7841.5 a/v_{te}$, for small-box simulations showing 2×2 copies of each periodic domain. The large-box version of this simulation is shown in Figure 5.1.

reasonable to suppose that the fixed electron-electron collisions helped regularise these simulations over a wider range of collisionalities than was accessible when varying electron-ion and electron-electron collisions together. The simplified simulations support our view that the scaling of the heat flux with the collision rate that we have found does not depend directly on flow shear. Of course, physically, in a real tokamak, flow shear is still needed to suppress the ion-scale turbulence, which we did not model.

APPENDIX F

Long-time linear damping of the electron zonal modes

In this Appendix, we derive the scaling (5.30) of the zonal damping rate.

F.1 Zonal evolution equation

Consider the linearised form of the zonal gyrokinetic equation (5.20) (i.e., ignore its right-hand side). It is convenient to Fourier transform this equation in the field-perpendicular coordinate x (there is no dependence on y by definition of the zonal modes), while leaving the field-parallel derivatives in position space (the zonal modes here are just ones with $k_y = 0$, but are allowed to have parallel variation). This gives us

$$\frac{\partial}{\partial t} \left[h + J_0(k_\perp \rho) \frac{e\phi}{T} F \right] + v_\parallel \mathbf{b} \cdot \nabla h + ik_x v_{Bx} h = \langle C[h e^{-ik_x \rho x}] e^{ik_x \rho x} \rangle, \quad (\text{F.1})$$

where $J_0(k_\perp \rho)$ with $\rho = v_\perp / \Omega_e$ is the Fourier-space form of the gyroaveraging operator, $k_\perp = k_x |\nabla x|$ (here the curvilinear coordinate x is defined by (C.2) and so $\nabla x = (q/B_0 r) \nabla \psi$), $\rho_x = (\nabla x) \cdot (\mathbf{b} \times \mathbf{v}_\perp) / \Omega_e$, the angle brackets $\langle \dots \rangle$ denote the averaging of the linearised collision operator over the gyroangle [88, 44], and

$$v_{Bx} = \mathbf{v}_B \cdot \nabla x = \frac{q}{B_0 r} \mathbf{v}_B \cdot \nabla \psi, \quad (\text{F.2})$$

where \mathbf{v}_B is given by (5.7). The quasineutrality equation (5.14) becomes

$$\frac{e\phi}{T} \left(1 + \frac{1}{\tau}\right) = -\frac{1}{n} \int d^3\mathbf{v} J_0(k_\perp \rho) h. \quad (\text{F.3})$$

F.2 Long-wavelength ordering and time scales

Assuming that the zonal modes have long wavelengths, $k_x \rho_e \ll 1$, we seek the solution to equations (F.1) and (F.3) in the form of an asymptotic expansion

$$h = h^{(0)} + h^{(1)} + h^{(2)} + \dots, \quad (\text{F.4})$$

where $h^{(n)} \sim (k_x \rho_e)^n h^{(0)}$. In this ordering, $J_0(k_\perp \rho) \approx 1$ to lowest order and clearly, from (F.3), to lowest order, $e\phi/T \sim h^{(0)}/F$. We introduce a formal ordering of time scales in (F.1):

$$v_\parallel \mathbf{b} \cdot \nabla \sim \frac{v_{te}}{a} \sim \nu, \quad k_x v_{Bx} \sim \frac{k_x v_{te}^2}{\Omega_e a} = k_x \rho_e \frac{v_{te}}{a} \sim (k_x \rho_e) \nu, \quad \frac{\partial}{\partial t} \sim \gamma_Z \sim (k_x \rho_e)^2 \nu. \quad (\text{F.5})$$

This ordering anticipates the kind of solution we are expecting. Generally speaking, any initial zonal perturbation will evolve on three successive time scales: first, it will be damped collisionlessly by the streaming and magnetic-drift terms on the time scale $\sim a/v_{te}$, leaving a finite residual perturbation [53]; then this residual, which is non-Maxwellian, will be damped by collisions on the time scale $\sim 1/\nu$ [89, 90, 91, 92], leaving a residual Maxwellian (this perturbation will be our $h^{(0)}$); and finally, this perturbation will be damped diffusively at the rate γ_Z , as we are about to show.

Thus, the zonal perturbation whose evolution we are going to calculate will be close to a steady state, with its evolution essentially governed by the (neoclassical) collisional transport theory.

F.3 Zeroth order

With the ordering (F.5), at zeroth order, (F.1) becomes

$$v_{\parallel} \mathbf{b} \cdot \nabla h^{(0)} = C[h^{(0)}]. \quad (\text{F.6})$$

If we multiply this equation by $h^{(0)}/F$ and flux-surface average, the left-hand side vanishes (see, e.g., [3], §6.1) and we get

$$-\overline{\int d^3\mathbf{v} \frac{h^{(0)} C_{ee}[h^{(0)}]}{F}} + \overline{\int d^3\mathbf{v} \frac{\nu_{ei} v_{te}^3}{v^3 F} \frac{1 - \xi^2}{2} \left(\frac{\partial h^{(0)}}{\partial \xi} \right)^2} = 0, \quad (\text{F.7})$$

where we have denoted the flux-surface average by an overbar, used the collision operator (5.15) taken to lowest order in $k_{\perp} \rho_e$ and integrated by parts in the term involving electron-ion collisions (cf. [88], §4.2). Both terms in (F.7) are non-negative definite and so must vanish individually. The vanishing of the first of these implies that $h^{(0)}$ must be a perturbed Maxwellian, the vanishing of the second that it has no mean parallel flow. Therefore,

$$h^{(0)} = \left[-\frac{e\phi}{T} \left(1 + \frac{1}{\tau} \right) + \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right) \frac{\delta T}{T} \right] F, \quad (\text{F.8})$$

where the density perturbation associated with $h^{(0)}$ has been fixed by the requirement that, to lowest order, $h^{(0)}$ and ϕ must satisfy (F.3). Note that we omit the superscript (0) on ϕ and δT because we will not need to calculate these fields explicitly to any higher order.

Finally, substituting (F.8) back into (F.6) and using the fact that, $h^{(0)}$ being a flow-less perturbed Maxwellian, $C[h^{(0)}] = 0$, we find that both ϕ and δT must be flux functions to zeroth order:

$$\mathbf{b} \cdot \nabla \frac{e\phi}{T} = \mathbf{b} \cdot \nabla \frac{\delta T}{T} = 0. \quad (\text{F.9})$$

The zeroth-order zonal solution (F.8) is precisely the promised quasi-steady state to which any initial zonal perturbations will relax after initial collisionless and collisional transients.

F.4 First order

At first order in $k_x \rho_e$, equation (F.1) is

$$v_{\parallel} \mathbf{b} \cdot \nabla h^{(1)} + ik_x v_{Bx} h^{(0)} = C[h^{(1)}]. \quad (\text{F.10})$$

Note that any gyroaveraging or FLR effects in the collision operator will only appear in the second order.

In dealing with (F.10), it is convenient to take advantage of the fact that in tokamaks [93],

$$\mathbf{v}_B \cdot \nabla \psi = -v_{\parallel} \mathbf{b} \cdot \nabla \left(\frac{v_{\parallel} I}{\Omega_e} \right), \quad (\text{F.11})$$

where $I(\psi) = B_T R$ is a flux function, B_T is the toroidal magnetic field and R the major radius coordinate. On the right-hand side note that ∇ is taken at constant \mathcal{E} and μ , which owing to the variation of B is not at constant v_{\parallel} . Using (F.2) and recalling that the safety factor $q(\psi)$ and the radial flux-surface label $r(\psi)$ are both flux functions and the reference magnetic field B_0 is a constant, we have

$$v_{Bx} = -v_{\parallel} \mathbf{b} \cdot \nabla \left(\frac{v_{\parallel}}{\Omega_e} \frac{qI}{B_0 r} \right) = -v_{\parallel} \mathbf{b} \cdot \nabla \left(\rho_{pe} \frac{v_{\parallel}}{v_{te}} \right), \quad (\text{F.12})$$

where $\rho_{pe} = (qI/B_0 r) \rho_e$ is, by definition, the ‘‘poloidal gyroradius’’, so called because $qI/B_0 r \sim B/B_p$, where B_p is the poloidal magnetic field. Note that Ω_e and, therefore, $\rho_e = v_{te}/\Omega_e$, is calculated using the total, space-dependent field B , and so ρ_{pe} is *not* a flux function.

Since $h^{(0)}$ is a flux function ($\mathbf{b} \cdot \nabla h^{(0)} = 0$), we may then rewrite (F.10) as follows

$$v_{\parallel} \mathbf{b} \cdot \nabla \left(h^{(1)} - ik_x \rho_{pe} \frac{v_{\parallel}}{v_{te}} h^{(0)} \right) = C[h^{(1)}]. \quad (\text{F.13})$$

This equation expresses a problem in neoclassical theory [93, 54], but we shall not require its explicit solution in order to work out the scalings we need. Note from (F.13) that the first-order zonal perturbation has a parallel electron flow (and, therefore, current) of order

$$\frac{u_{\parallel}}{v_{te}} \sim k_x \rho_{pe} \frac{\delta p}{p}, \quad (\text{F.14})$$

where $\delta p/p$ is the zonal pressure perturbation associated with $h^{(0)}$. The perpendicular $\mathbf{E} \times \mathbf{B}$ zonal flow is, of course, just

$$\frac{v_{Ey}}{v_{te}} \sim k_x \rho_e \frac{e\phi}{T}. \quad (\text{F.15})$$

It is the resistive damping of zonal flows (currents) that will lead to the damping of the zonal modes.

F.5 Second order

In order to calculate this damping, we need evolution equations for the zeroth-order fields ϕ and δT . The time derivatives enter at the second order in (F.1). Using (F.8) and (F.12), we have at this order:

$$\begin{aligned} & \frac{\partial}{\partial t} \left[-\frac{1}{\tau} \frac{e\phi}{T} + \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right) \frac{\delta T}{T} \right] F + v_{\parallel} \mathbf{b} \cdot \nabla h^{(2)} - h^{(1)} v_{\parallel} \mathbf{b} \cdot \nabla \left(ik_x \rho_{pe} \frac{v_{\parallel}}{v_{te}} \right) \\ & = k_x^2 \langle C[h^{(0)} \rho_x] \rho_x \rangle + \left\langle C \left[h^{(2)} - \frac{1}{2} k_x^2 \rho_x^2 h^{(0)} \right] \right\rangle, \end{aligned} \quad (\text{F.16})$$

where the leading-order FLR parts of the collision operator acting on $h^{(0)}$ have appeared. We now transit-average this equation, integrate the magnetic-drift term by parts and use (F.13) to express $v_{\parallel} \mathbf{b} \cdot \nabla h^{(1)}$. The result is

$$\begin{aligned} & \frac{\partial}{\partial t} \left[-\frac{1}{\tau} \frac{e\phi}{T} + \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right) \frac{\delta T}{T} \right] F \\ & = \overline{-ik_x \rho_{pe} \frac{v_{\parallel}}{v_{te}} C[h^{(1)}] + k_x^2 \langle C[h^{(0)} \rho_x] \rho_x \rangle + \langle C[\dots] \rangle}. \end{aligned} \quad (\text{F.17})$$

To separate the evolution of ϕ and δT , we take the density and energy moments of (F.17). Integrating it over velocities (we need not evaluate the argument denoted ‘...’ since by conservation of particles $\int d^3\mathbf{v} \langle C[\dots] \rangle = 0$), we get¹

$$-\frac{n}{\tau} \frac{\partial e\phi}{\partial t T} = \overline{-ik_x \rho_{pe} \int d^3\mathbf{v} \frac{v_{\parallel}}{v_{te}} C_{ei}[h^{(1)}] + k_x^2 \rho_e^2 \int d^3\mathbf{v} \frac{C_{ei}[h^{(0)} \rho_x] \rho_x}{\rho_e^2}}. \quad (\text{F.18})$$

¹The last term in (F.18) is related to the spatial FLR diffusion term that we wrote out explicitly in the gyroaveraged collision operator (5.15).

Only the electron-ion collision terms have survived because, to lowest order, conservation of momentum by the electron-electron collisions implies

$$\int d^3\mathbf{v} v_{\parallel} C_{ee}[h^{(1)}] = 0, \quad \int d^3\mathbf{v} C_{ee}[h^{(0)}\rho_x]\rho_x = 0. \quad (\text{F.19})$$

In the same vein, multiplying (F.17) by v^2/v_{te}^2 and integrating over velocities, we find, using energy conservation by collisions, $\int d^3\mathbf{v} v^2 \langle C[\dots] \rangle = 0$,

$$\frac{3n}{2} \frac{\partial}{\partial t} \frac{\delta T}{T} = \frac{-ik_x \rho_{pe} \int d^3\mathbf{v} \frac{v_{\parallel}}{v_{te}} \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right) C[h^{(1)}] + k_x^2 \rho_e^2 \int d^3\mathbf{v} \frac{C[h^{(0)}\rho_x]\rho_x}{\rho_e^2} \left(\frac{v^2}{v_{te}^2} - \frac{3}{2} \right)}{(\text{F.20})}$$

The difference between this equation and (F.18) is that the contributions from electron-electron collisions do not vanish (because same-species collisions can support non-zero heat fluxes).

F.6 Damping rate

We do not need to solve the neoclassical equation (F.13) for $h^{(1)}$ explicitly to see that, this equation being linear, $h^{(1)}$ will be a linear combination of $e\phi/T$ and $\delta T/T$ with velocity-dependent coefficients, all of which are proportional to $k_x \rho_{pe}$; any homogeneous part of $h^{(1)}$ satisfies (F.6) and so can be absorbed into $h^{(0)}$. Therefore, we may schematically represent (F.18) and (F.20) as the following matrix equation at this order:²

$$\frac{\partial}{\partial t} \frac{e\phi}{T} = \nu_{ei} k_x^2 \rho_{pe}^2 \left(a_{11} \frac{e\phi}{T} + a_{12} \frac{\delta T}{T} \right) \quad (\text{F.21})$$

$$\frac{\partial}{\partial t} \frac{\delta T}{T} = \nu_{ei} k_x^2 \rho_{pe}^2 \left(a_{21} \frac{e\phi}{T} + a_{22} \frac{\delta T}{T} \right) + \nu_{ee} k_x^2 \rho_{pe}^2 \left(b_{21} \frac{e\phi}{T} + b_{22} \frac{\delta T}{T} \right) \quad (\text{F.22})$$

²Note that in the absence of magnetic drifts, in a straight field, $h^{(1)} = 0$ and we must replace ρ_{pe} with ρ_e everywhere, with the collisional evolution of the zonal modes now controlled by the FLR spatial-diffusion terms in the gyrokinetic collision operator.

where a_{ij} and b_{ij} are dimensionless coefficients that depend only on equilibrium parameters. Clearly, if $\nu_{ei} = 0$, the matrix is defective (has one zero row corresponding to the upper equation (F.21)) and the damping rate of both the potential and temperature perturbations of the slowest damped eigenmode vanishes. Therefore,

$$\gamma_Z \propto \nu_{ei} k_x^2 \rho_{pe}^2. \quad (\text{F.23})$$

The order-unity numerical prefactor in the exact expression for γ_Z depends on the ratio ν_{ee}/ν_{ei} , which itself is order unity. Thus, we have shown that the long-time damping rate of the zonal modes always satisfies (5.30), a scaling that is indeed well reproduced in our numerical simulations (Figure 5.4).

F.7 Zonal damping by same-species collisions

The result (5.30) survives even if, artificially, the ratio ν_{ee}/ν_{ei} is made large — as was done in the “simplified” simulations of section 5.3.5. In this case, (F.22) simply implies that

$$\frac{\delta T}{T} \approx -\frac{b_{21}}{b_{22}} \frac{e\phi}{T}, \quad (\text{F.24})$$

and, according to (F.21), the zonal perturbation is still damped at the rate $\gamma_Z \sim \nu_{ei} k_x^2 \rho_{pe}^2$.

If we consider an even more artificial situation in which ν_{ei} is made to vanish completely, we must go to next order when calculating the density moment (F.18) of the zonal kinetic equation (F.1) to allow higher-order finite-drift-orbit-width FLR terms to come in. These terms will be small by an extra factor of $(k_x \rho_{pe})^2$ and so technically we would have to reorder the time derivative two powers of $k_x \rho_e$ higher and then work through two extra orders in our expansion. It is clear though³ that the

³The result was also obtained by Connor [94].

result of this amounts to replacing (F.21) and (F.22) with

$$\frac{\partial}{\partial t} \frac{e\phi}{T} = \nu_{ee} k_x^4 \rho_{pe}^4 \left(b_{11} \frac{e\phi}{T} + b_{12} \frac{\delta T}{T} \right), \quad (\text{F.25})$$

$$\frac{\partial}{\partial t} \frac{\delta T}{T} = \nu_{ee} k_x^2 \rho_{pe}^2 \left(b_{21} \frac{e\phi}{T} + b_{22} \frac{\delta T}{T} \right), \quad (\text{F.26})$$

using (F.26) to infer again the coupling (F.24) between the zonal potential and temperature perturbations, and finally arriving at a zonal damping rate

$$\gamma_Z \sim \nu_{ee} k_x^4 \rho_{pe}^4. \quad (\text{F.27})$$

This is the scaling that is indeed recovered numerically when ν_{ei} is switched off (black crosses in Figure 5.4). Such considerations might appear moot, as $\nu_{ei} = 0$ is unphysical and achievable only in numerical experiments. They do, however, help us gain some insight into the way our theoretical arguments would be modified if we applied them to zonal dynamics in ITG, rather than ETG, turbulence, as suggested in section 5.4.2.

F.8 Case of ion zonal modes

The only differences between the gyrokinetic equation for the ion zonal modes and (F.1) are a minus sign in front of the $e\phi/T$ term, accounting for the ion charge, and the absence of interspecies collisions:

$$\frac{\partial}{\partial t} \left[h_i - J_0(k_\perp \rho) \frac{e\phi}{T_i} F \right] + v_\parallel \mathbf{b} \cdot \nabla h_i + ik_x v_{Bx} h_i = \langle C_{ii} [h_i e^{-ik_x \rho_x}] e^{ik_x \rho_x} \rangle. \quad (\text{F.28})$$

Another very important difference between the ion and electron cases is the Boltzmann-electron closure (5.32), which removes the flux-surface-averaged part of the zonal potential from the electron-density response. Combining (5.32) with the expression (5.12) for the ion density perturbation in terms of h_i , we get the following version of the quasineutrality equation:

$$\frac{e\phi}{T_i} + \frac{e(\phi - \bar{\phi})}{T_e} = \frac{1}{n} \int d^3 \mathbf{v} J_0(k_\perp \rho) h_i. \quad (\text{F.29})$$

That the difference between (F.29) and (F.3) is significant becomes obvious if we take the density moment of (F.28), i.e., integrate it over velocities keeping \mathbf{r} constant (equivalently, multiply by $J_0(k_\perp\rho)$ and integrate over \mathbf{v}), and then flux-surface average. To lowest order in $k_\perp\rho_i$, the term in the square brackets becomes, using (F.29),

$$\frac{1}{n} \int d^3\mathbf{v} J_0(k_\perp\rho) \left[h_i - J_0(k_\perp\rho) \frac{e\phi}{T_i} F \right] \approx \frac{1}{2} k_\perp^2 \rho_i^2 \frac{e\bar{\phi}}{T_i}. \quad (\text{F.30})$$

This has two extra powers of k_x compared to the analogous term for electron zonal flows (see the left-hand side of (F.18)).

By a calculation analogous to Appendix F.3, the zeroth-order solution $h_i^{(0)}$ of (F.28) in the long-time limit must again be a quasi-steady (i.e., slow-evolving), constant on flux surfaces, flow-less perturbed Maxwellian. The rest of the calculation proceeds similarly to the electron case with the exception that only momentum-conserving ion-ion collisions are present and so the velocity integral of the collision terms in the density equation will be smaller by an extra factor of $k_x^2 \rho_{pi}^2$, as explained in Appendix F.7. In view of (F.30), the analog of (F.25) for ion zonal modes will then be

$$\frac{\partial}{\partial t} \frac{1}{2} k_x^2 \rho_{pi}^2 \frac{e\phi}{T_i} = \nu_{ii} k_x^4 \rho_{pi}^4 \left(c_{11} \frac{e\phi}{T_i} + c_{12} \frac{\delta T_i}{T_i} \right), \quad (\text{F.31})$$

where we have absorbed all factors accounting for differences between k_\perp and k_x and between ρ_i and ρ_{pi} into the dimensionless coefficients c_{11} and c_{12} . The extra factors of $k_x^2 \rho_{pi}^2$ on the left-hand side (due to the special role of flux-surface-averaged potential in the Boltzmann-electron closure) and on the right-hand side (due to momentum conservation by the ion-ion collision operator) cancel, leaving $e\phi/T_i$ to evolve on the diffusive time scale $\sim 1/\nu_{ii} k_x^2 \rho_{pi}^2$. The evolution equation for $\delta T_i/T_i$ does not have these extra factors on either side and so features the same time scale. Thus, the ion zonal modes, similarly to the electron ones (although for a different reason), will be damped at a rate

$$\gamma_{Zi} \sim \nu_{ii} k_x^2 \rho_{pi}^2. \quad (\text{F.32})$$

Acknowledgements

This work was supported in part by the RCUK Energy Programme under grant EP/I501045, by the European Communities under the contract of Association between EURATOM and CCFE, by the European Union's Horizon 2020 research and innovation programme under grant agreement number 633053, and by the Foreign and Commonwealth Office.

Computations were performed at the UK's HECToR and ARCHER services under EPSRC grants EP/H002081/1 and EP/L000237/1 respectively, on EFDA's HPC-FF system, and on the Helios system at IFERC-CSC.

Support from CCP-Plasma, grant EP/M022463/1, is also acknowledged, as is the generous hospitality of the Wolfgang Pauli Institute in Vienna and of CIEMAT in Madrid.

We thank Tim Hender, Greg Hammett, Walter Guttenfelder, Ian Abel, Steve Cowley, Edmund Highcock, Ferdinand van Wyk and Jack Connor for enlightening and helpful discussions; and thank Anthony Field, Martin Valovič, Martin O'Brien, Francis Casson, Justin Ball and two anonymous referees also for useful comments on drafts of the papers contained here.

★

I doubt whether any thesis of this type is completed without the existence of a substantial support network behind the student whose name appears on the title page; certainly this one was not. Too many people to name have helped me in too

many ways to list; I thank you all. Even if not mentioned individually – indeed even if forgotten (although for this I can only apologise) – your contribution still counted.

First and foremost I would like to thank my wife Margaret. Without her encouragement, I would not even have started this project, let alone finished it. She has made many sacrifices in its favour, and especially in the last few months has provided and organized much practical support. Our son Evan has also borne the brunt of my increasing absence from home with good grace. I hope that one day he will understand some of what is written here; then he can come to his own conclusion about whether it was worthwhile!

My wider family has also been very supportive, especially my parents over many years, and I have been blessed with good friends now in many places around the world. For their particular inspiration and help I would like to thank Clare Carroll, Richard Clark and Elysium Ltd, Jenny Robinson, David Taubman, my old colleague from the other place James Briginshaw; and in this place all those associated with St Matthew's church, St Ebbe's primary school, and the Oxford Suzuki group.

I have had many good teachers, but would like to single out my first actual physics teacher, known to all as JDW, who set me on the path to this point by noticing my interest in a subject I hadn't yet heard of, then nurturing it in his own unique way. How many schoolchildren can say that they flew in a helicopter, visited a battery factory and learned every Queen song thanks to their physics teacher?

I would like to thank three later mentors: Richard Ansorge, whose light-touch supervision of my undergraduate final-year project gave me my first real taste of research; Chippy Thyagaraja, who was a wise elder; and Bill Dorland. On very little acquaintance Bill demonstrated great insight into my motivation and character, and committed his time and expertise to help me far beyond what my efforts deserved.

Michael Barnes (who is Bill's former student) also invested in me, especially at the beginning of this project, and I am privileged to have worked with him. Felix Parra gave me helpful tuition towards the end, and I am particularly indebted

to him for the original draft of Appendix F. All of the students and postdocs during my time at Oxford have been good friends and interesting colleagues to work alongside. I particularly thank Edmund Highcock for passing on the L^AT_EX template for this thesis, and Ferdinand van Wyk for many things GS2-related. Young-chul Kim contributed insight at a key moment and was an important link between theory and experiment. I would also like to thank my College Adviser at Exeter, Andrew Steane, for his crucial help; the Physics IT team for support with numerous computing issues; James Binney and Ard Louis; Michelle Boshier and others in the Theoretical Physics office; Lucy Tipping and Rebecca Reynolds at the Mathematical, Physical and Life Sciences division office; and all the staff in the Exeter College offices.

I started this project through my employment at the Culham laboratory of the UK Atomic Energy Authority, and I am very grateful for its support, not least financially. I would particularly like to thank Martin Cox, Martin O'Brien and Tim Hender, all of whom I worked for first during my gap year; and then again much later, when Rob Akers was my line manager. I am grateful to Rob for his constant encouragement. Later I was in Colin Roach's group: he has been an insightful collaborator, a supportive line manager during a difficult period for the lab, and is now a friend. I am also especially grateful to Jack Connor for his kindness, and to Anthony Field for his diligence and loyalty. My work would not have run as smoothly without the help of many non-scientists at Culham, and I would particularly like to thank Wendy Coombes, Kamariah Samson, Lynda Lee and Jackie Costello.

Towards the end of this project I have been assisted by many people to find and take the next step. I would particularly like to thank Ted Shepherd, Paul Dellar, Nuno Loureiro, Geoffrey Vallis, Jonathan Winter (Career Innovation), Damian Wilson, Claire and Andrew Down, and Barbara Montanari, for their varied contributions to this endeavour.

I am grateful for the moral support that I have received during this project from

past collaborators not directly involved in the present work, whose faith in me has been a great encouragement. In particular I would like to thank Paola Mantica, and colleagues at General Atomics in San Diego, especially Gary Staebler.

Finally I want to thank my supervisor, Alex Schekochihin. Over the course of this project, Alex has helped me to develop as a research physicist in numerous ways. He has many unique insights into plasma physics, and more generally an unmatched capability to see through extraneous issues to the heart of the scientific question. He has been particularly attentive to me in the final stages, suggesting many improvements to the text and plots of chapter 5 and the appendices, especially Appendix F. I look forward to the possibility of further fruitful collaboration with him in the future.

Bibliography

- [1] D. MacKay. *Sustainable Energy – without the hot air*. UIT, 2008.
- [2] J. Wesson. *Tokamaks*, chapter 1.5, page 10. Oxford University Press, Oxford, 3 edition, 2004.
- [3] I. G. Abel, G. G. Plunk, E. Wang, M. Barnes, S. C. Cowley, W. Dorland, and A. A. Schekochihin. Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows. *Reports on Progress in Physics*, 76:116201, 2013.
- [4] J. Wesson. *Tokamaks*, chapter 1.3, page 5. Oxford University Press, Oxford, 3 edition, 2004.
- [5] T. D. Arber. Plasma kinetic theory. Unpublished lecture notes from the 44th Culham Plasma Physics Summer School, July 2007.
- [6] M. A. Barnes. *Trinity: A Unified Treatment of Turbulence, Transport, and Heating in Magnetized Plasmas*. PhD thesis, University of Maryland, 2008.
- [7] M. Barnes, I. G. Abel, W. Dorland, T. Görler, G. W. Hammett, and F. Jenko. Direct multiscale coupling of a transport code to gyrokinetic turbulence codes. *Physics of Plasmas*, 17:056109, 2010.
- [8] M. Kotschenreuther, G. Rewoldt, and W. M. Tang. Comparison of initial value and eigenvalue codes for kinetic toroidal plasma instabilities. *Computer Physics Communications*, 88(2):128, 1995.
- [9] GS2 is an open-source Fortran code, <http://gyrokinetics.sourceforge.net/>.
- [10] C. Roach. Equilibrium flow shear implementation in GS2. Unpublished notes, 2010.

- [11] J. E. Kinsey, R. E. Waltz, and J. Candy. Nonlinear gyrokinetic turbulence simulations of $E \times B$ shear quenching of transport. *Physics of Plasmas*, 12(6):062302, 2005.
- [12] R. E. Waltz, G. D. Kerbel, and J. Milovich. Toroidal gyro-Landau fluid model turbulence simulations in a nonlinear ballooning mode representation with radial modes. *Physics of Plasmas*, 1(7):2229–2244, 1994.
- [13] C. M. Roach, I. G. Abel, R. J. Akers, W. Arter, M. Barnes, Y. Camenen, F. J. Casson, G. Colyer, J. W. Connor, S. C. Cowley, D. Dickinson, W. Dorland, A. R. Field, W. Guttenfelder, G. W. Hammett, R. J. Hastie, E. Highcock, N. F. Loureiro, A. G. Peeters, M. Reshko, S. Saarelma, A. A. Schekochihin, M. Valovic, and H. R. Wilson. Gyrokinetic simulations of spherical tokamaks. *Plasma Physics and Controlled Fusion*, 51:124020, 2009.
- [14] A. R. Field, C. Michael, R. J. Akers, J. Candy, G. Colyer, W. Guttenfelder, Y.-c. Ghim, C. M. Roach, S. Saarelma, and the MAST Team. Plasma rotation and transport in the MAST spherical tokamak. In *Proceedings of 23rd IAEA Fusion Energy Conference, Daejeon, Republic of Korea*, October 2010. EXC/P8-04.
- [15] A. R. Field, C. Michael, R. J. Akers, J. Candy, G. Colyer, W. Guttenfelder, Y.-c. Ghim, C. M. Roach, S. Saarelma, and the MAST Team. Plasma rotation and transport in MAST spherical tokamak. *Nuclear Fusion*, 51(6):063006, 2011.
- [16] A. M. Dimits, G. Bateman, M. A. Beer, B. I. Cohen, W. Dorland, G. W. Hammett, C. Kim, J. E. Kinsey, M. Kotschenreuther, A. H. Kritz, L. L. Lao, J. Mandrekas, W. M. Nevins, S. E. Parker, A. J. Redd, D. E. Shumaker, R. Sydora, and J. Weiland. Comparisons and physics basis of tokamak transport models and turbulence simulations. *Physics of Plasmas*, 7:969, 2000.
- [17] A. R. Field, R. J. Akers, C. Brickley, P. G. Carolan, C. Challis, N. J. Conway, G. Cunningham, H. Meyer, C. M. Roach, M. J. Walsh, and the MAST Team. Core heat transport in the MAST spherical tokamak. In *Proceedings of 20th IAEA Fusion Energy Conference, Vilamoura, Portugal*, November 2004. EX/P2-11.
- [18] G. M. Staebler, G. Colyer, S. Kaye, J. E. Kinsey, and R. E. Waltz. Testing the trapped gyro-Landau fluid transport model with data from tokamaks and spherical tori. 22nd IAEA Fusion Energy Conference. TH/P8-42.

- [19] M. J. Pueschel, M. Kammerer, and F. Jenko. Gyrokinetic turbulence simulations at high plasma beta. *Physics of Plasmas*, 15(10):102310, 2008.
- [20] D. J. Applegate, C. M. Roach, S. C. Cowley, W. D. Dorland, N. Joiner, R. J. Akers, N. J. Conway, A. R. Field, A. Patel, M. Valovic, and M. J. Walsh. Microstability in a spherical tokamak ‘MAST-like’ high confinement mode equilibrium. *Physics of Plasmas*, 11:5085, 2004.
- [21] D. J. Applegate, C. M. Roach, J. W. Connor, S. C. Cowley, W. Dorland, R. J. Hastie, and N. Joiner. Micro-tearing modes in the Mega Ampère Spherical Tokamak. *Plasma Physics and Controlled Fusion*, 49:1113–1128, 2007.
- [22] C. M. Roach, D. J. Applegate, J. W. Connor, S. C. Cowley, W. D. Dorland, R. J. Hastie, N. Joiner, S. Saarelma, A. A. Schekochihin, R. J. Akers, C. Brickley, A. R. Field, M. Valovic, and the MAST Team. Microstability physics as illuminated in the spherical tokamak. *Plasma Physics and Controlled Fusion*, 47:B323, 2005.
- [23] R. L. Miller, M. S. Chu, J. M. Greene, Y. R. Lin-Liu, and R. E. Waltz. Noncircular, finite aspect ratio, local equilibrium model. *Physics of Plasmas*, 5:973, 1998.
- [24] M. A. Beer. *Gyrofluid Models of Turbulent Transport in Tokamaks*. PhD thesis, Princeton University, 1995.
- [25] T. Görler and F. Jenko. Scale separation between electron and ion thermal transport. *Physical Review Letters*, 100:185002, 2008.
- [26] C. M. Roach, M. Walters, R. V. Budny, F. Imbeaux, T. W. Fredian, M. Greenwald, J. A. Stillerman, D. A. Alexander, J. Carlsson, J. R. Cary, F. Ryter, J. Stober, P. Gohil, C. Greenfield, M. Murakami, G. Bracco, B. Esposito, M. Romanelli, V. Parail, P. Stubberfield, I. Voitsekhovitch, C. Brickley, A. R. Field, Y. Sakamoto, T. Fujita, T. Fukuda, N. Hayashi, G. M. D. Hogeweyj, A. Chudnovskiy, N. A. Kinerva, C. E. Kessel, T. Aniel, G. T. Hoang, J. Ongena, E. J. Doyle, W. A. Houlberg, A. R. Polevoi, ITPA Confinement Database and Modelling Topical Group, and ITPA Transport Physics Topical Group. The 2008 Public Release of the International Multi-tokamak Confinement Profile Database. *Nuclear Fusion*, 48:125001, 2008.

- [27] J. Candy, C. Holland, R. E. Waltz, M. R. Fahey, and E. Belli. Tokamak profile prediction using direct gyrokinetic and neoclassical simulation. *Physics of Plasmas*, 16(6):060704, 2009.
- [28] N. Joiner, D. Applegate, S. C. Cowley, W. Dorland, and C. M. Roach. Electron temperature gradient driven transport in a MAST H-mode plasma. *Plasma Physics and Controlled Fusion*, 48:685, 2006.
- [29] R. Numata, G. G. Howes, T. Tatsuno, M. Barnes, and W. Dorland. AstroGK: Astrophysical gyrokinetics code. *Journal of Computational Physics*, 229(24):9347–9372, 2010.
- [30] C. S. Chang and F. L. Hinton. Effect of finite aspect ratio on the neoclassical ion thermal conductivity in the banana regime. *Physics of Fluids*, 25(9):1493–1494, 1982.
- [31] C. M. Roach. Modelling ohmic confinement experiments on the START tokamak. *Plasma Physics and Controlled Fusion*, 38(12):2187, 1996.
- [32] F. Jenko, W. Dorland, and G. W. Hammett. Critical gradient formula for toroidal electron temperature gradient modes. *Physics of Plasmas*, 8(9):4096, 2001.
- [33] W. Guttenfelder and J. Candy. Resolving electron scale turbulence in spherical tokamaks with flow shear. *Physics of Plasmas*, 18(2):022506, 2011.
- [34] J. L. Peterson, G. W. Hammett, D. R. Mikkelsen, H. Y. Yuh, J. Candy, W. Guttenfelder, S. M. Kaye, and B. LeBlanc. An enhanced nonlinear critical gradient for electron turbulent transport due to reversed magnetic shear. 2011. arXiv:1105.2195v1.
- [35] P. Mantica, C. Angioni, B. Baiocchi, M. Baruzzo, M. N. A. Beurskens, J. P. S. Bizarro, R. V. Budny, P. Buratti, A. Casati, C. Challis, J. Citrin, G. Colyer, F. Crisanti, A. C. A. Figueiredo, L. Frassinetti, C. Giroud, N. Hawkes, J. Hobbirk, E. Joffrin, T. Johnson, E. Lerche, P. Migliano, V. Naulin, A. G. Peeters, G. Rewoldt, F. Ryter, A. Salmi, R. Sartori, C. Sozzi, G. Staebler, D. Strintzi, T. Tala, M. Tsalas, D. V. Eester, T. Versloot, P. C. deVries, J. Weiland, and JET EFDA Contributors. Ion heat transport studies in JET. *Plasma Physics and Controlled Fusion*, 53(12):124033, 2011.
- [36] C. Angioni. Private communication, 2011.

- [37] M. Valovič, R. Akers, M. de Bock, J. McCone, L. Garzotti, C. Michael, G. Naylor, A. Patel, C. M. Roach, R. Scannell, M. Turnyanskiy, M. Wisse, W. Guttenfelder, J. Candy, and the MAST Team. Collisionality and safety factor scalings of H-mode energy transport in the MAST spherical tokamak. *Nuclear Fusion*, 51(7):073045, 2011.
- [38] H. Doerk, F. Jenko, M. J. Pueschel, and D. R. Hatch. Gyrokinetic microtearing turbulence. *Physical Review Letters*, 106:155003, 2011.
- [39] W. Guttenfelder, J. Candy, S. M. Kaye, W. M. Nevins, E. Wang, R. E. Bell, G. W. Hammett, B. P. Leblanc, D. R. Mikkelsen, and H. Yuh. Electromagnetic transport from microtearing mode turbulence. *Physical Review Letters*, 106:155004, 2011.
- [40] I. G. Abel and S. C. Cowley. Multiscale gyrokinetics for rotating tokamak plasmas: II. Reduced models for electron dynamics. *New Journal of Physics*, 15(2):023041, 2013.
- [41] W. Dorland, F. Jenko, M. Kotschenreuther, and B. N. Rogers. Electron temperature gradient turbulence. *Physical Review Letters*, 85:5579, 2000.
- [42] F. Jenko, W. Dorland, M. Kotschenreuther, and B. N. Rogers. Electron temperature gradient driven turbulence. *Physics of Plasmas*, 7:1904, 2000.
- [43] E. A. Frieman and L. Chen. Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *Physics of Fluids*, 25:502, 1982.
- [44] I. G. Abel, M. Barnes, S. C. Cowley, W. Dorland, and A. A. Schekochihin. Linearized model Fokker-Planck collision operators for gyrokinetic simulations. I. Theory. *Physics of Plasmas*, 15:122509, 2008.
- [45] M. Barnes, I. G. Abel, W. Dorland, D. R. Ernst, G. W. Hammett, P. Ricci, B. N. Rogers, A. A. Schekochihin, and T. Tatsuno. Linearized model Fokker-Planck collision operators for gyrokinetic simulations. II. Numerical implementation and tests. *Physics of Plasmas*, 16:072107, 2009.
- [46] H. Sugama, T.-H. Watanabe, and M. Nunami. Linearized model collision operators for multiple ion species plasmas and gyrokinetic entropy balance equations. *Physics of Plasmas*, 16(11):112503, 2009.

- [47] A. A. Schekochihin, J. T. Parker, E. G. Highcock, P. J. Dellar, W. Dorland, and G. W. Hammett. Phase mixing versus nonlinear advection in drift-kinetic plasma turbulence. *Journal of Plasma Physics*, 82:905820212, 2016.
- [48] M. Barnes, F. I. Parra, and A. A. Schekochihin. Critically balanced ion temperature gradient turbulence in fusion plasmas. *Physical Review Letters*, 107:115003, 2011.
- [49] Y.-c. Ghim, A. A. Schekochihin, A. R. Field, I. G. Abel, M. Barnes, G. Colyer, S. C. Cowley, F. I. Parra, D. Dunai, and S. Zoletnik. Experimental signatures of critically balanced turbulence in MAST. *Physical Review Letters*, 110:145002, 2013.
- [50] J. Wesson. *Tokamaks*, chapter 4.18, page 209. Oxford University Press, Oxford, 3 edition, 2004.
- [51] P. H. Diamond, S.-I. Itoh, K. Itoh, and T. S. Hahm. Zonal flows in plasma—a review. *Plasma Physics and Controlled Fusion*, 47:35, 2005.
- [52] B. N. Rogers, W. Dorland, and M. Kotschenreuther. Generation and stability of zonal flows in ion-temperature-gradient mode turbulence. *Physical Review Letters*, 85:5336, 2000.
- [53] M. N. Rosenbluth and F. L. Hinton. Poloidal flow driven by ion-temperature-gradient turbulence in tokamaks. *Physical Review Letters*, 80:724, 1998.
- [54] P. Helander and D. J. Sigmar. *Collisional Transport in Magnetized Plasmas*. Cambridge University Press, Cambridge, 2005.
- [55] J. B. Parker and J. A. Krommes. Zonal flow as pattern formation. *Physics of Plasmas*, 20(10):100703, 2013.
- [56] J. B. Parker and J. A. Krommes. Generation of zonal flows through symmetry breaking of statistical homogeneity. *New Journal of Physics*, 16(3):035006, 2014.
- [57] M. Nakata, T.-H. Watanabe, H. Sugama, and W. Horton. Formation of coherent vortex streets and transport reduction in electron temperature gradient driven turbulence. *Physics of Plasmas*, 17:042306, 2010.
- [58] Y. Idomura, S. Tokuda, and Y. Kishimoto. Global profile effects and structure formations in toroidal electron temperature gradient driven turbulence. *Nuclear Fusion*, 45:1571, 2005.

- [59] Y. Idomura. Self-organization in electron temperature gradient driven turbulence. *Physics of Plasmas*, 13:080701, 2006.
- [60] W. Guttenfelder, J. Candy, S. M. Kaye, W. M. Nevins, R. E. Bell, G. W. Hammett, B. P. LeBlanc, and H. Yuh. Scaling of linear microtearing stability for a high collisionality National Spherical Torus Experiment discharge. *Physics of Plasmas*, 19(2):022506, 2012.
- [61] W. Guttenfelder, J. L. Peterson, J. Candy, S. M. Kaye, Y. Ren, R. E. Bell, G. W. Hammett, B. P. LeBlanc, D. R. Mikkelsen, W. M. Nevins, and H. Yuh. Progress in simulating turbulent electron thermal transport in NSTX. *Nuclear Fusion*, 53(9):093022, 2013.
- [62] ITER Physics Expert Groups on Confinement and Transport and Confinement Modelling and Database, and ITER Physics Basis Editors. Chapter 2 of ITER Physics Basis: Plasma confinement and transport. *Nuclear Fusion*, 39:2175, 1999.
- [63] C. Bourdelle, T. Gerbaud, L. Vermare, A. Casati, T. Aniel, J. F. Artaud, V. Basiuk, J. Bucalossi, F. Clairet, Y. Corre, P. Devynck, G. Falchetto, C. Fenzi, X. Garbet, R. Guirlet, Ö. Gürçan, S. Heuraux, P. Hennequin, G. T. Hoang, F. Imbeaux, L. Manenc, P. Monier-Garbet, P. Moreau, R. Sabot, J.-L. Ségui, A. Sirinelli, D. Villegas, and the Tore Supra Team. Collisionality scaling in Tore Supra: detailed energy confinement analysis, turbulence measurements and gyrokinetic modelling. *Nuclear Fusion*, 51(6):063037, 2011.
- [64] W. Dorland and G. W. Hammett. Gyrofluid turbulence models with kinetic effects. *Physics of Fluids B*, 5:812, 1993.
- [65] G. W. Hammett, M. A. Beer, W. Dorland, S. C. Cowley, and S. A. Smith. Developments in the gyrofluid approach to tokamak turbulence simulations. *Plasma Physics and Controlled Fusion*, 35(8):973, 1993.
- [66] S. C. Cowley, R. M. Kulsrud, and R. Sudan. Considerations of ion-temperature-gradient-driven turbulence. *Physics of Fluids B*, 3:2767, 1991.
- [67] W. M. Nevins, J. Candy, S. Cowley, T. Dannert, A. Dimits, W. Dorland, C. Estrada-Mila, G. W. Hammett, F. Jenko, M. J. Pueschel, and D. E. Shumaker. Characterizing electron temperature gradient turbulence via numerical simulation. *Physics of Plasmas*, 13:122306, 2006.

- [68] F. Jenko and W. Dorland. Prediction of significant tokamak turbulence at electron gyroradius scales. *Physical Review Letters*, 89(22):225001, November 2002.
- [69] P. Ricci, B. N. Rogers, and W. Dorland. Small-scale turbulence in a closed-field-line geometry. *Physical Review Letters*, 97:245001, 2006.
- [70] Z. Lin, T. S. Hahm, W. W. Lee, W. M. Tang, and P. H. Diamond. Effects of collisional zonal flow damping on turbulent transport. *Physical Review Letters*, 83:3645, 1999.
- [71] Z. Lin, T. S. Hahm, W. W. Lee, W. M. Tang, and R. B. White. Gyrokinetic simulations in general geometry and applications to collisional damping of zonal flows. *Physics of Plasmas*, 7:1857, 2000.
- [72] D. R. Mikkelsen and W. Dorland. Dimits shift in realistic gyrokinetic plasma-turbulence simulations. *Physical Review Letters*, 101:135003, 2008.
- [73] K. Srinivasan and W. R. Young. Zonostrophic instability. *Journal of the Atmospheric Sciences*, 69(5):1633–1656, 2012.
- [74] R. Chemke and Y. Kaspi. Poleward migration of eddy-driven jets. *Journal of Advances in Modeling Earth Systems*, 7(3):1457–1471, 2015.
- [75] F. I. Parra, M. Barnes, and A. G. Peeters. Up-down symmetry of the turbulent transport of toroidal angular momentum in tokamaks. *Physics of Plasmas*, 18:062501, 2011.
- [76] M. A. Beer, S. C. Cowley, and G. W. Hammett. Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Physics of Plasmas*, 2:2687, 1995.
- [77] E. Highcock. *The Zero-Turbulence Manifold in Fusion Plasmas*. DPhil thesis, University of Oxford, 2012. arXiv:1207.4419.
- [78] G. Hammett and C. Roach. Private communication, 2012.
- [79] S. L. Newton, S. C. Cowley, and N. F. Loureiro. Understanding the effect of sheared flow on microinstabilities. *Plasma Physics and Controlled Fusion*, 52:125001, 2010.

- [80] E. G. Highcock, M. Barnes, A. A. Schekochihin, F. I. Parra, C. M. Roach, and S. C. Cowley. Transport bifurcation in a rotating tokamak plasma. *Physical Review Letters*, 105:215003, 2010.
- [81] M. Barnes, F. I. Parra, E. G. Highcock, A. A. Schekochihin, S. C. Cowley, and C. M. Roach. Turbulent transport in tokamak plasmas with rotational shear. *Physical Review Letters*, 106:175004, 2011.
- [82] E. G. Highcock, M. Barnes, F. I. Parra, A. A. Schekochihin, C. M. Roach, and S. C. Cowley. Transport bifurcation induced by sheared toroidal flow in tokamak plasmas. *Physics of Plasmas*, 18:102304, 2011.
- [83] A. A. Schekochihin, E. G. Highcock, and S. C. Cowley. Subcritical fluctuations and suppression of turbulence in differentially rotating gyrokinetic plasmas. *Plasma Physics and Controlled Fusion*, 54:055011, 2012.
- [84] F. van Wyk, E. G. Highcock, A. A. Schekochihin, C. M. Roach, A. R. Field, and F. I. Parra. Transitions to subcritical turbulence in a tokamak plasma. 2016. In preparation.
- [85] E. G. Highcock, A. A. Schekochihin, S. C. Cowley, M. Barnes, F. I. Parra, C. M. Roach, and W. Dorland. Zero-turbulence manifold in a toroidal plasma. *Physical Review Letters*, 109:265001, 2012.
- [86] Y.-c. Ghim, A. R. Field, A. A. Schekochihin, E. G. Highcock, C. Michael, and the MAST Team. Local dependence of ion temperature gradient on magnetic configuration, rotational shear and turbulent heat flux in MAST. *Nuclear Fusion*, 54:042003, 2014.
- [87] J. Candy, R. E. Waltz, M. R. Fahey, and C. Holland. The effect of ion-scale dynamics on electron-temperature-gradient turbulence. *Plasma Physics and Controlled Fusion*, 49:1209, 2007.
- [88] A. A. Schekochihin, S. C. Cowley, W. Dorland, G. W. Hammett, G. G. Howes, E. Quataert, and T. Tatsuno. Astrophysical gyrokinetics: Kinetic and fluid turbulent cascades in magnetized weakly collisional plasmas. *Astrophysical Journal Supplement Series*, 182:310, 2009.
- [89] F. L. Hinton and M. N. Rosenbluth. Dynamics of axisymmetric and poloidal flows in tokamaks. *Plasma Physics and Controlled Fusion*, 41:A653, 1999.

- [90] E.-J. Kim, C. Holland, and P. H. Diamond. Collisional damping of ETG-mode-driven zonal flows. *Physical Review Letters*, 91:075003, Aug 2003.
- [91] Y. Xiao, P. J. Catto, and K. Molvig. Collisional damping for ion temperature gradient mode driven zonal flow. *Physics of Plasmas*, 14:032302, 2007.
- [92] Y. Xiao, P. J. Catto, and W. Dorland. Effects of finite poloidal gyroradius, shaping, and collisions on the zonal flow residual. *Physics of Plasmas*, 14:055910, May 2007.
- [93] F. L. Hinton and S. K. Wong. Neoclassical ion transport in rotating axisymmetric plasmas. *Physics of Fluids*, 28(10):3082, 1985.
- [94] J. Connor. Private communication, 2014.