Spatial distribution of zonal flow drive: from tokamaks via novel hybrids to stellarators

Tobias Mohandas Schuett

Doctor of Philosophy

University of York

School of Physics, Engineering and Technology

November 2025

Abstract

Zonal flows are of central importance for the confinement and therefore performance of future fusion reactors that are based on the idea of toroidal magnetic confinement. By taking energy out of turbulent fluctuations and by shearing turbulent eddies apart, these flows reduce the radial turbulent transport which is responsible for the majority of the heat and particle losses in tokamaks and optimised stellarators.

Turbulent transport properties are known to be strongly affected by plasma shaping. Yet not all origins are identified, and one key missing gap in our understanding is the spatial distribution of the nonlinear drive of zonal flows due to Reynolds stress. In this work we show that this spatial distribution is generally non-trivial and generally not correlated with the distribution of turbulent fluctuations, for both axisymmetric and non-axisymmetric shaping of the nested toroidal surfaces. In particular, it is found that the zonal flow drive peaks away from the outboard midplane, where turbulent fluctuations are usually the strongest.

In pursuit of combining the advantages of tokamaks and stellarators and to bridge the gap between them, a new method was developed to optimise three-dimensionally perturbed tokamak equilibria for quasi-axisymmetry, strong enough for key stabilising benefits of stellarators to be expected. We examined the general turbulent properties in this continuous parameter space from tokamaks to stellarators, and showed that the zonal flow drive envelope here also is non-trivial, developing additional local maxima due to the localised three-dimensional shaping. Further, by extending the analysis to more traditional stellarators, the zonal flow drive is also shown to be non-trivial in the most unstable flux tube of the largest stellarator experiment, Wendelstein 7-X.

Finally, we derive and discuss how the Reynolds stress analysis can be generalised to the gyrokinetic nonlinearity, describing the zonal flow drive with kinetic effects.

Contents

List of Figures

List of Tables

Acknowledgments

This Thesis would not have been possible without the many people who helped me along the way.

First of all I would like to thank my advisor, István Cziegler, for introducing me to plasma physics and fusion research during my Erasmus year at York, for taking me on as a graduate student, for always being supportive of me following my curiosity, and for the many detailed and often frame-tale style discussions that we have had.

I would also like to thank David Dickinson and Sophia Henneberg for their mentorship and collaboration. Thanks to David for answering my many questions on gyrokinetics and `GS2`, for all the help with implementation and debugging, and for the perhaps most efficient email exchanges I have ever had. Thanks to Sophia for bringing the potential of stellarators to my attention during her visit to York, for hosting me in Greifswald, and for teaching me about stellarators, optimisation and so much else.

Thanks also to Simon Freethy for mentorship in Culham and for our insightful discussions. Thanks to my progression advisors Howard Wilson and Jon Graves for providing additional guidance and direction. Thanks to Per Helander and Gabe Plunk for helpful discussions and for detailed feedback on manuscripts.

Thanks to everyone from the York Plasma Institute, Culham Centre for Fusion Energy, the wider Fusion Centre for Doctoral Training cohort, and those I met during my visit at the Max-Planck Institute for Plasma Physics in Greifswald.

In particular, thanks to Daniel Greenhouse, Felicity Maiden, Joe Bryant, Nick Osborne, and Nicola Lonigro for the camaraderie in Culham. To Steven Thomas, Yorick Enters and Lena Howlett for our heroic efforts in hunting for zonal flows in MAST-U. To Paul Costello, Yann Narbutt, Alan Goodman, Bob Davies and many others for welcoming me in Greifswald. To Stephen Biggs-Fox for writing the initial code which laid the foundations for some of this work. And to Bailey

Cook for our collaboration on a kinetic version of the zonal flow drive analysis.

This project was funded by the UK Engineering and Physical Sciences Research Council under grant number EP/S022430/1. The gyrokinetic simulations in this work were made possible by the Archer2 cluster (grant numbers EP/R034737/1 and EP/X035336/1) and the Viking2 cluster at York. Stellarator optimisations and calculations were made possible by the Max Planck Computing and Data Facility.

My greatest thanks go to my mum and dad, for being great parents and for always supporting me in my decisions. Finally, thanks to Emer, for constant love and laughs that have always kept the spirits high, especially during this final writing-up period.

Declaration

I declare that this work has not previously been presented for a degree or other qualification at this University or elsewhere. This Thesis is a presentation of original work and I am the sole author, except where appropriately stated. All sources are acknowledged as references.

Chapter 2 is based on the following publication and code release

> T M Schuett, I Cziegler, and D Dickinson (Nov. 2025a). The poloidal distribution of electrostatic zonal flow drive in strongly shaped tokamaks. In: *Plasma Physics and Controlled Fusion* 67.11, p. 115022
>
> Michael Barnes, David Dickinson, William Dorland, Peter Alec Hill, Joseph Thomas Parker, Colin Malcolm Roach, Maurizio Giacomin, Noah Mandell, Ryusuke Numata, Tobias Schuett, Stephen Biggs-Fox, Nicolas Christen, Jason Parisi, George Wilkie, Lucian Anton, Justin Ball, Jessica Baumgaertel, Greg Colyer, Michael Hardman, Joachim Hein, Edmund Highcock, Gregory Howes, Adrian Jackson, Michael T. Kotschenreuther, Jungpyo Lee, Huw Leggate, Adwiteey Mauriya, Bhavin Patel, Tomo Tatsuno, and Ferdinand Van Wyk (July 2024). GS2 v8.2.0. Version 8.2.0

The initial algorithm that was implemented was supplied by Stephen Biggs-Fox.

Chapter 3 is based on the following two publications

> T. M. Schuett and S. A. Henneberg (Dec. 2024b). Exploring novel compact quasi-axisymmetric stellarators. In: *Phys. Rev. Res.* 6, p. L042052
>
> Tobias M Schuett and Sophia A Henneberg (June 2025). Optimization of compact quasi-axisymmetric stellarators. In: *Plasma Physics and Controlled Fusion* 67.6, p. 065024

The coil solutions shown in section 3.12.1 were designed by Sophia Henneberg.

The work presented in this Thesis was presented at various conferences, including:

> T. M. Schuett, I. Cziegler, and D. Dickinson (2025b). The poloidal distribution of electrostatic zonal flow drive in strongly shaped tokamaks [Oral Contributed]. *29th EU-US Transport Task Force Workshop, Budapest, Hungary*
>
> T. M. Schuett, S. A. Henneberg, D. Dickinson, and I. Cziegler (2025). Ion temperature gradient turbulence in compact quasi-axisymmetric stellarators [Poster]. *Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting, New York, USA*
>
> T. M. Schuett and S. A. Henneberg (2024c). Optimization of compact quasi-axisymetric stellarators [Oral Invited]. *International Stellarator and Heliotron Workshop, Hiroshima, Japan*
>
> T. M. Schuett and S. A. Henneberg (2024a). Exploring novel compact qausi-axisymmetric stellarators [Poster]. *Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting, New York, USA*

1

Introduction

1.1 Toroidally confined plasmas for fusion energy

The economic flourishing of humanity is tightly coupled to the availability of abundant energy. Controlled thermonuclear fusion, when achieved, can provide base load electricity that is carbon free, produces negligible radioactive waste and uses abundant fuel sources. When nuclei of light elements are combined or *fused* to form heavier elements, an excess of energy is released which can be captured[1]. The fusion reaction with the highest likelihood is the one between deuterium and tritium[2], two heavier isotopes of hydrogen. When both fuse, a helium nucleus (alpha particle) and neutron are created as well as an excess of energy,

$$\mathrm{H}_1^2 + \mathrm{H}_1^3 \rightarrow \mathrm{He}_2^4 + \mathrm{n} + 17.6\,\mathrm{MeV}. \tag{1.1.1}$$

The excess energy of $17.6\,\mathrm{MeV}$ per reaction is released in the form of kinetic energy, with $14.1\,\mathrm{MeV}$ going to the neutron and $3.5\,\mathrm{MeV}$ to the alpha particle. Deuterium is naturally occurring in water, and tritium can be produced (or bred) within the fusion reactor by capturing the neutron in a so-called breeder blanket that surrounds the volume in which the reaction takes place. There the neutron first reacts with lithium to produce tritium before heating up water which generates the electricity through a conventional steam cycle. The fusion reaction 1.1.1 requires high temperatures since nuclei generally repel each other via the Coulomb force, unless they come very close to each other and attract each other via the strong nuclear force. Such incidental proximity can be achieved through high temperatures and thanks to the tunnelling effect. Thus, for the likelihood of the deuterium-tritium fusion reaction to be maximised both species must be in the temperature range around $10\,\mathrm{keV}$ which corresponds to approximately 100 million degrees Celsius. At these temperatures, electrons are no longer bound to their nuclei and a fluid of electrons and ions is formed. This

---

[1] Heavy elements on the other hand must be split for energy to be released, corresponding to the more well-known method of obtaining energy from the reaction of nuclei: nuclear fission.

[2] Strictly speaking this reaction is only the most likely up to temperatures of $\sim 500\,\mathrm{keV}$. This, however, is such a large value that most designs aim to operate well below it.

loosely[3] defines a *plasma*, a ‘gas of charged particles’, often referred to as the ‘fourth state of matter’. Further, these extreme temperatures require a method to confine this plasma which is not based on material boundaries, motivating the use of magnetic fields to push back against the plasma’s urge to expand, and to create so-called *magnetic confinement*.[4]

Toroidal magnetic confinement is based on the idea that a charged particle in a magnetic field mostly follows the magnetic field lines, resulting in a spatially confined motion if the field lines were to close on themselves. A magnetic field which exhibits many such circular field lines is topologically toroidal in nature, with each circular field line being centred on the same location but their radii and vertical location, Z, varying, see the left illustration in figure 1.1.1. The fact that the field lines bend around in a circle, however, does impact the particle motion to a degree that matters. So-called magnetic drifts arise which depend on the charge of the particle and thus separate ions and electrons from each other along the vertical Z-direction (see figure 1.1.1 for the definition of this axis with respect to the torus traced out by the magnetic field). The resulting vertical electric field results in rapid radial motion due to the excellent electrical conductivity of the plasma along the magnetic field, and ends confinement too quickly for fusion conditions to be set up in time. This charge separation can be easily avoided by adding *twist* to the magnetic field lines, such that they move around the central field line in a helical motion, consisting of a toroidal and a *poloidal* component, as illustrated in figures 1.1.1. Poloidal refers to the short way around the torus. The twist prevents the charge separation for the same reason that a spinning honey spoon prevents the honey from running off the spoon: even though the drift is constant in its (vertical) direction, to the degree that it takes particles away from their field line while on the top half of their helical trajectory, it brings them back on the bottom half.

The helical field lines now trace out toroidal surfaces which are nested within each other and particles mostly stay on these nested surfaces. A consequence of this is that both pressure and the *magnetic flux* are constant on these surfaces. Any pressure differences that might arise are shorted out quickly since, in general, any point on the surface is connected to any other point on the surface via a magnetic field line. The magnetic flux is constant since no field line inside the nested surface leaves that volume such that the number of field lines that pass through a poloidal cross-section of the volume is constant across all toroidal locations.

---

[3] For a stricter definition see (Freidberg 2008, Chapter 7).

[4] Another approach is to aim for enough reactions to happen fast enough such that confinement is secondary, motivating the idea of inertial fusion energy.

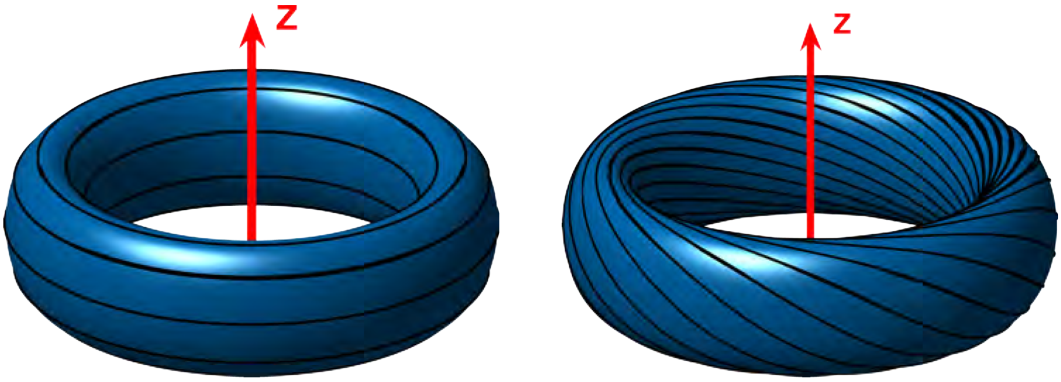

Figure 1.1.1. – Purely toroidal field versus helical field. **Left:** a purely toroidal field resulting in net particle drifts that end confinement. **Right:** a helical field, i.e. one with a toroidal and a *poloidal* component, where particle drifts cancel out on average, guaranteeing confinement (in the single-particle picture). Image taken from (Imbert-Gérard, Paul & Wright 2024).

Note that the neutrons produced in the fusion reaction 1.1.1 leave the plasma due to their neutral charge, and are used within the reactor walls to breed tritium and heat the water to generate the electricity. In contrast, the alpha particles with their non-zero charge now also stay confined to the nested surface that they were born on and heat up the plasma[5]. The goal of magnetic confinement fusion is to achieve sufficiently high temperatures and densities in the core of the plasma efficiently enough such that the plasma *ignites*, i.e. such that the fusion reactions heat the plasma enough to sustain themselves and no external heating is required. More specifically, it is the triple product of temperature T, density n, and energy confinement time τ_E which must exceed some critical value for ignition. T must be at an optimal point to maximise the fusion reactivity, thereby leaving $n\tau_E$, the Lawson parameter, as the quantity to be optimised for (Freidberg 2008).

The challenge thus becomes to find a toroidal confinement equilibrium where the central pressure is maximised in an economical way, while guaranteeing that the plasma is stable to perturbations. Here economically refers to the requirements on recirculating power: if too much of the generated power is needed to sustain the confinement equilibrium, too little is fed into the grid.

1.1.1 Tokamaks and stellarators

Arguably the broadest distinction one can make between toroidal confinement equilibria is whether or not they feature continuous rotational symmetry around

---

[5] This is the case in an ideal design. In reality, collisional transport, turbulent transport and imperfections in the magnetic field would make alphas leave the plasma at some point.

their central Z-axis, which is marked in figure 1.1.1. This toroidal symmetry is referred to as *axisymmetry*. The two most developed magnetic confinement fusion devices are the *tokamak* and the *stellarator*. Tokamaks are axisymmetric, whereas stellarators break axisymmetry and are inherently three-dimensional. Examples of both are shown in figure 1.1.2.

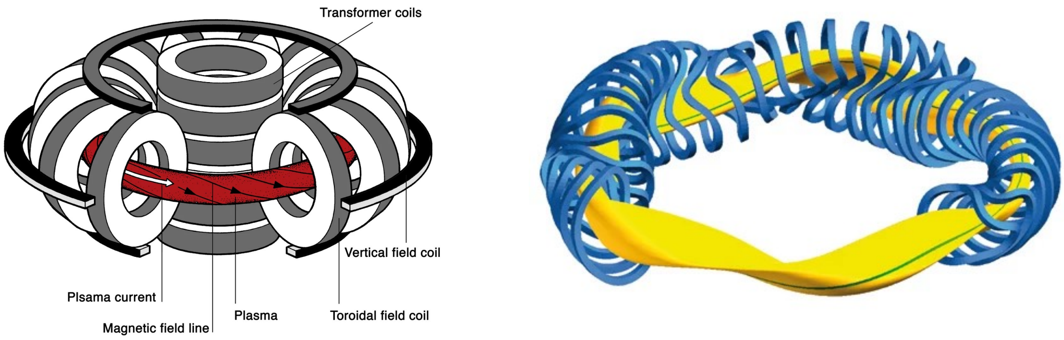

Figure 1.1.2. – Tokamak (left) and stellarator (right), the two concepts for toroidal magnetic confinement fusion, and the two most advanced concepts within magnetic confinement fusion in general. Tokamak figure from https://www.ipp.mpg.de/14869/tokamak and stellarator figure from https://www.ipp.mpg.de/2815232/konzeptentwicklung, both accessed on August 12, 2025.

Both devices create their toroidal magnetic field with external electromagnets, i.e. with currents *outside* the plasma. However, the source of the poloidal field, which makes the field twist, differs. In tokamaks this twist in the magnetic field is generated purely by a toroidal current *inside* the plasma, whereas in stellarators the twist is generated by a mix of currents both *inside* and *outside* the plasma. This ability of external currents to generate twist in the field is only possible due to the broken axisymmetry (Helander 2014). The ratio between the twist from currents inside the plasma versus outside the plasma generally depends on the type of stellarator.

While tokamaks automatically confine particles due to their axisymmetry, and can be built with simple coils, their pure reliance on the plasma current introduces the possibility for instabilities which end confinement quickly (so-called disruptions). Further, this requires tokamaks to have a larger recirculating power than stellarators. Stellarators do not rely on the plasma current for confinement, making them safe against disruptions. However, due to their non-axisymmetric three-dimensional (3D) character the plasma must be carefully optimised to provide confinement to all particles and electromagnets must be optimised to be able to create the plasma. This need for 3D plasmas to generally require optimisation offers another way to broadly distinguish toroidal confinement devices – we will expand on this point later in section 3 when the discussion turns to non-axisymmetric fields.

1.1.2 Nested flux surfaces, magnetic coordinates, and general coordinate definitions

Naturally, toroidal geometry will play an important role in this work. Either toroidal coordinates (r, θ, ϕ) or cylindrical coordinates (R, Z, ϕ) will be used, depending on which are more convenient in a particular setting. Here r refers to the minor radius, θ the (geometrical) poloidal angle, ϕ the (geometrical) toroidal angle, R to the major radius, and Z is the aforementioned vertical coordinate. Their geometrical meaning as well as other important quantities that we will use are depicted in figure 1.1.3.

Many of the basic definitions are motivated by constraints on the confined plasma that follow from magnetohydrodynamic (MHD) analysis. The MHD framework treats the plasma as a single fluid and describes the most macroscopic dynamics of the plasma. For the plasma to be in steady state it must be in (ideal) MHD equilibrium, which is equivalent to the requirement that

$$\mathbf{J} \times \mathbf{B} = \nabla p \tag{1.1.2}$$

is satisfied throughout the plasma volume. Here $\mathbf{B}$ is the magnetic field, $\mathbf{J}$ the current density, and p the pressure. Since $\mathbf{B} \cdot \nabla p = 0$ and $\mathbf{J} \cdot \nabla p = 0$ follows from 1.1.2, the requirement on $\mathbf{B}$ and $\mathbf{J}$ is thus that they lie on surfaces of constant pressure. Thus, as long as the pressure profile is monotonic, the position within the plasma volume can be specified by p and both angles, i.e. $\mathbf{B}(p, \theta, \phi)$ and $\mathbf{J}(p, \theta, \phi)$[6].

From the fact that $\mathbf{B}$ lies on surfaces of constant pressure it follows – after some calculation which we omit here for brevity – that it can be written as (Helander 2014, section 2.2)

$$\mathbf{B} = \nabla\psi \times \nabla\theta_P + \nabla\phi \times \nabla\chi. \tag{1.1.3}$$

Here $\psi(p)$ is the toroidal magnetic flux divided by 2π,

$$\psi = \frac{1}{2\pi} \int \mathbf{B} \cdot \mathrm{d}\mathbf{S}_{\mathrm{pol}} = \frac{1}{2\pi} \int_0^{\psi} d\psi' \int_0^{2\pi} d\theta_P' \, (\mathbf{B} \cdot \nabla\phi) J, \tag{1.1.4}$$

---

[6] If the equilibrium is axisymmetric, i.e. a tokamak, all quantities, including $\mathbf{B}$ and $\mathbf{J}$, are independent of ϕ.

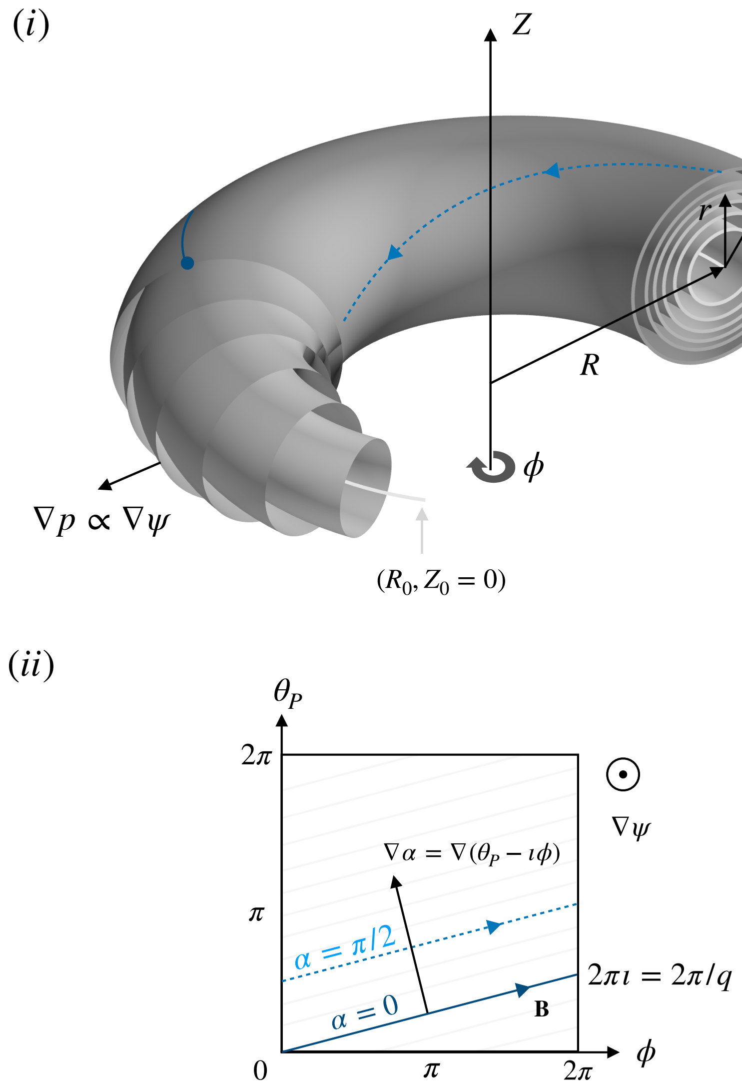

Figure 1.1.3. – Coordinate conventions for toroidally confined plasmas. See text for details. Note that while an axisymmetric configuration is shown, most of these conventions also apply to stellarators.

and $\chi(p)$ is the poloidal flux divided by 2π,

$$\chi = \frac{1}{2\pi}\int \mathbf{B}\cdot \mathrm{d}\mathbf{S}_{\mathrm{tor}} = \frac{1}{2\pi}\int_0^{\psi} d\psi' \int_0^{2\pi} d\phi\,(\mathbf{B}\cdot\nabla\theta_P)J. \tag{1.1.5}$$

A good illustration of the poloidal and toroidal areas which are used to measure these fluxes can be found for example in the two figures of chapter 5.5.1 of Ref. (Imbert-Gérard, Paul & Wright 2024). We therefore note that p, ψ, and χ are all equally suitable for labelling the nested surfaces. For this reason the nested surfaces are usually referred to as *flux surfaces*. The constancy of the magnetic flux on the nested surfaces is, of course, independent of their toroidal shaping (i.e. it works in both tokamaks and stellarators), whereas the geometric minor radius r which is shown in the left plot of figure 1.1.3 only labels the nested surfaces for axisymmetric circular shaping. This motivates the switch of the toroidal coordinate system from (r, θ, ϕ) to (X, θ, ϕ) where $X \in \{p, \psi, \chi\}$. In this Thesis we will use $X = \psi$.

In equation 1.1.3 we used a new poloidal angle θ_P which is related to the geometric poloidal angle via a general transformation function $\lambda(\psi, \theta, \phi)$,

$$\theta_P = \theta + \lambda. \tag{1.1.6}$$

This variable change turns the general coordinate system (ψ, θ, ϕ) into a so-called *magnetic coordinate system* (ψ, θ_P, ϕ) in which field lines are straight in the plane spanned by ϕ and θ_P. It is worth noting that for each choice of toroidal angle a special choice for the poloidal angle can be made such that field lines are straight in the plane of both angles, leading to different magnetic coordinate systems which are convenient in different settings (Helander 2014). For the geometric toroidal angle ϕ the corresponding straight-field-line angle is the above-mentioned θ_P, the so-called PEST poloidal angle. In the above expressions 1.1.4 and 1.1.5 for the magnetic fluxes, J is the Jacobian of the coordinate system (ψ, θ_P, ϕ), which when expressed with contravariant basis vectors reads

$$J^{-1} = (\nabla\theta_P \times \nabla\phi) \cdot \nabla\psi. \tag{1.1.7}$$

The *rotational transform* ι is defined as the change in poloidal flux with toroidal flux,

$$\iota(\psi) = \frac{\mathrm{d}\chi}{\mathrm{d}\psi}, \tag{1.1.8}$$

and is therefore constant on each nested surface. With this definition the magnetic field $\mathbf{B}$ from equation 1.1.3 can now be written even simpler,

$$\mathbf{B} = \nabla\psi \times \nabla\alpha, \tag{1.1.9}$$

where $\alpha = \theta_P - \iota\phi$ is a new variable that is constant along the field lines. The field lines therefore have a slope of

$$\frac{\mathrm{d}\theta_P}{\mathrm{d}\phi} = \iota. \tag{1.1.10}$$

This is indicated in the bottom plot of figure 1.1.3. Thus the rotational transform ι describes how many poloidal turns a field line completes in one toroidal turn (Helander 2014). Following convention, we will specify the slope of the magnetic field by the rotational transform whenever we deal with stellarators and by the so-called safety factor $q = 1/\iota$ whenever we deal with tokamaks.

1.2 Micro-turbulence

Just like fluids, plasmas also exhibit *turbulence* (Frisch 1995), the great unsolved problem of classical physics. Its existence in fusion plasmas first became apparent in the 1970s and was found to be due to *drift wave instabilities* which are excited by the large gradients in density and temperature that are required to make terrestrial fusion work. Its discovery in tokamaks followed the observation that the radial transport of heat is much larger than that expected from collisional processes. Turbulence in tokamaks and stellarators is mostly undesirable as it decreases the efficiency of obtaining high core temperatures and densities, i.e. it lowers the energy confinement time τ_E, making it more difficult to achieve ignition. However, turbulent transport can also be beneficial for flushing out impurities and alpha particles after they heated the plasma (so-called helium ash). The goal of the transport community is therefore to understand turbulence and to find means to control and optimise it.

The instabilities which excite the turbulence, the *drift waves*, exist on the small scale of the Larmor radius. The resulting turbulence was therefore coined *micro-turbulence*. In contrast, quantities that have spatial scales which are comparable to the device size, as measured for example by the minor radius a, are referred to as *macro*, whereas the region in between both is often referred to as the *meso scale*. Zonal flows, which play a central role in this work and will be introduced shortly, exist on this meso scale (Diamond et al. 2005).

Drift waves are named as such since their phase velocity is on the order of the diamagnetic (drift) velocity[7]

$$\mathbf{V}_s^* = \frac{\mathbf{B} \times \nabla p_s}{e_s B^2}, \tag{1.2.1}$$

where s is the species label, ∇p_s its pressure gradient, and $e_s = Z_s e$ its charge with charge number Z_s and elementary charge e. Over the years many such drift wave instabilities have been identified, broadly distinguished by their responsible particle type (ions vs. electrons), class (trapped vs. passing) and their regimes (electrostatic vs. electromagnetic). The focus of this work is one particular drift-wave instability which is believed to be responsible for most of the transport in today's fusion experiments, the *ion temperature gradient instability* (see Ref. (Brower et al. 1987) for the first experimental observation in a tokamak, and Ref.

---

[7] Note that the word drift is actually very misleading here since the diamagnetic velocity is only one of the plasma fluid of a given species and not of the individual particles of that species. This flow is a result of the single handedness of the gyro-motion with respect to the toroidal angle and the radially varying speed of this motion due to the pressure gradient. Being a result of gyro-motion it does not feature any net drift of the particles themselves.

(Bähner et al. 2021) for those from the first neoclassically optimised stellarator experiment, Wendelstein 7-X. We will introduce this instability next.

1.2.1 Ion temperature gradient driven turbulence

The focus of this work is turbulence driven by the *ion temperature gradient mode*, a micro-instability that is, as the name suggests, excited by the ion temperature gradient. The instability and its resulting turbulence is electrostatic, i.e. it only requires fluctuations in the electrostatic potential φ, and not in the magnetic vector potential $\mathbf{A}$. Its injection scale is on the order of the ion Larmor radius ρ_i, making it an 'ion-scale instability'. The ITG instability (often just called 'ITG') does not require the kinetic response of the electrons to be resolved, allowing us to model this species with a simple Boltzmann response, i.e. prescribing the log of the density to be proportional to φ.

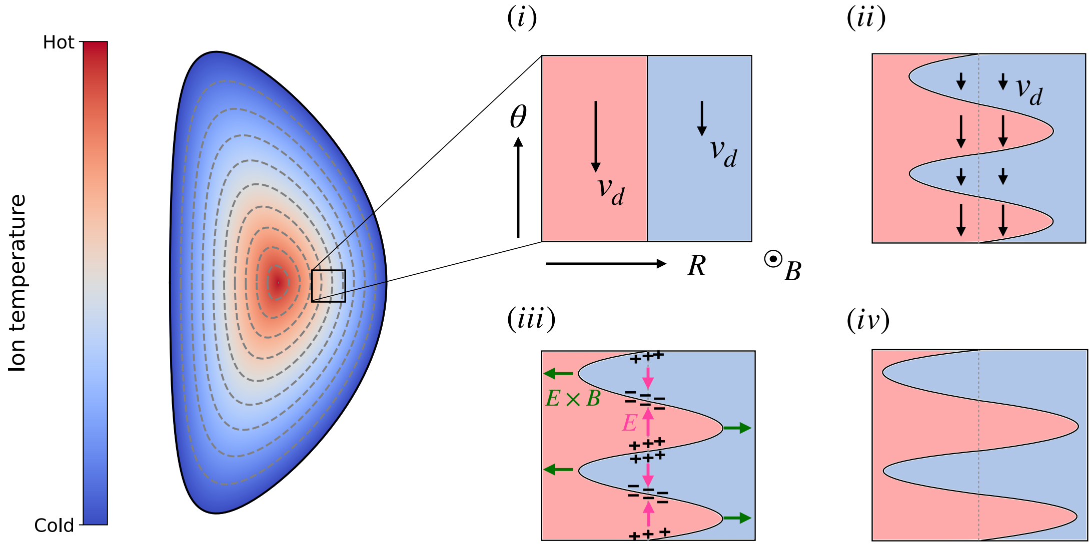

Figure 1.2.1. – A sketch of the ion temperature gradient instability. The four phases of its growth are shown, see text for more details.

ITG turbulence exhibits a toroidal branch and a slab branch. The slab branch does not require toroidal topology but is usually subdominant to the toroidal branch in the toroidal setting that we are interested in (Jenko & Dorland 2002). The basic dynamics of the toroidal ITG are depicted in figure 1.2.1 and features the following steps:

(i) The ion diamagnetic drift velocity is larger in regions of higher temperature.

(ii) An initial perturbation causes alternating regions of larger and smaller $\mathbf{v}_d$ along θ within the perturbed region.

(iii) The alternating regions of $\mathbf{v}_d$ in θ cause ions to accumulate or be depleted at the boundaries between the alternating regions, setting up an electric field.

(iv) The electric field combined with the equilibrium magnetic field causes the initial perturbation to grow more, due to the resulting $\mathbf{E} \times \mathbf{B}$ drift.

The same perturbation on the inboard side, however, causes the charge accumulation to be reversed, thereby reversing the direction of the resulting $\mathbf{E} \times \mathbf{B}$ drift

$$V_{\mathbf{E}\times\mathbf{B}} = \frac{\mathbf{E} \times \mathbf{B}}{B^2}, \tag{1.2.2}$$

i.e. the initial perturbation is brought back to equilibrium with the same dynamics and the instability is stable in that region. The toroidal ITG instability, like many other pressure driven instabilities, is therefore only driven on the so-called 'bad curvature region' on the outside of the torus where the curvature vector of the magnetic field line κ and the pressure gradient ∇p point in approximately the same direction. On the inboard side they point in approximately opposite directions, since κ is typically dominated by its toroidal component. While the ITG instability is driven by the ion temperature gradient, it is stabilised by the density gradient; this can be seen analytically by considering the mode in its two-fluid and slab limit (Dickinson 2012, chapter 3.2.2).

1.3 Zonal flows and their interaction with turbulence

Zonal flows (ZFs) are flows within a radial layer of a system that result from the self-organisation of turbulence. In toroidally confined plasmas these are directed in the binormal direction, i.e. the direction perpendicular to the field line but within the flux surface, such that the flow is mostly in the poloidal direction. Zonal flows also exist in spherical systems, such as on Earth and Jupiter, where they are directed in the azimuthal direction.

Unlike the ITG instability from figure 1.2.1, ZFs are unable to feed off the free energy supplied by the equilibrium gradients, i.e. they are linearly stable modes. Instead, the nonlinear interaction of different turbulent fluctuations feeds energy into the ZFs via their nonlinear interaction (Diamond et al. 2005). It is this nonlinear interaction which distinguishes the pure dynamics of a single unstable ITG mode from fully developed turbulence, that is, fully developed turbulence is inherently nonlinear featuring non-local coupling and local coupling (cascading).

In toroidally confined plasmas there are two main branches of ZFs, one with zero mean frequency (mean zonal flow or MZF) which enjoys an undamped residual (Rosenbluth & Hinton 1998), and a finite frequency mode called the geodesic acoustic mode (GAM) (Winsor, Johnson & Dawson 1968). While we do not make the distinction between both in this work, it is worth pointing out that we will also neglect neither. Our analysis includes the total zonal flow with all its frequencies, i.e. the MZF and any other finite frequency components such as the GAM.

While zonal flows can only be driven nonlinearly, their damping is determined through both linear and nonlinear processes. See Ref. (Diamond et al. 2005) for examples of both. In this work, however, we focus solely on the (nonlinear) *drive*.

1.3.1 Zonal flows for confinement

Both ZFs and equilibrium flow $\mathbf{E} \times \mathbf{B}$ can drastically improve confinement through their shearing action, i.e. their action on turbulence through their radial variation in strength or even direction. As depicted in figure 1.3.1, such flow *shear* breaks up turbulent eddies, especially those which are radially elongated, thereby reducing their radial transport ability as smaller eddy structures tend to cause less transport. However, due to their nonlinear interaction with turbulence, ZFs reduce turbulence in another way which the equilibrium $\mathbf{E} \times \mathbf{B}$ does not: when turbulent modes couple to feed the ZF they lose energy at the same time, i.e. energy is *transferred*, reducing the fluctuation strength and thus turbulent transport. This second effect is also depicted in figure 1.3.1.

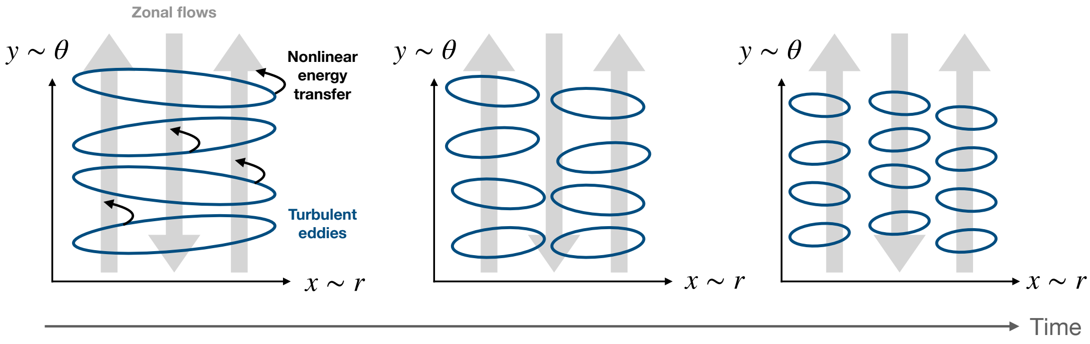

Figure 1.3.1. – Zonal flows reduce turbulent transport in two distinct ways, (1) by shearing eddies apart and (2) by taking energy out of them.

Zonal flows are not only of interest for general confinement, but also for confinement *transitions* such as the one from the low confinement mode to the high confinement mode of tokamaks, the so-called L-H transition (more on this later in section 2.1). While zonal flows are determined by turbulent processes, the equilibrium $\mathbf{E} \times \mathbf{B}$ shearing rate is determined by the mean radial electric field profile which in turn is determined by the radial component of the force balance equation (or momentum equation),

$$E_r = \frac{1}{n_s e_s} \frac{\mathrm{d}p_s}{\mathrm{d}r} + u_{\phi,s} B_\theta - u_{\theta,s} B_\phi, \tag{1.3.1}$$

and must hold for each species ‘s’. Here n_s in the equilibrium density, $u_{\phi,s}$ and $u_{\theta,s}$ are the toroidal and poloidal mean flow, respectively, and B_ϕ and B_θ are the toroidal and poloidal magnetic field components.

It is generally agreed that the mean shearing is responsible for ‘locking in’ confinement states *after* the L-H transition has happened (Burrell 1997). The large pressure gradient which develops *during* the transition results in a strong radial variation of E_r and thus mean shearing which reduces turbulence (Biglari, Diamond & Terry 1990; Terry 2000). The real question is what happens *before* the transition, i.e. what *triggers* it.

Since zonal flows arise from turbulence in a self-organised way they provide a compelling explanation for this initial trigger, and experimental evidence supports this theory (more on this later in section 2.1). More specifically it is the radial gradient of the Reynolds stress, $\partial_r \langle \tilde{v}_r \tilde{v}_\theta \rangle$, which drives poloidal flow (Diamond & Kim 1991), as we will show next. Since this mechanism drives the mean radial electric field through the poloidal velocity channel in equation 1.3.1, it falls into the category of ‘poloidal spin-up’ theories. Other theories for poloidal spin-up and other trigger mechanisms exist but have received less attention due to less experimental evidence pointing in their direction. A comprehensive overview is given in (McDermott 2009, Chapter 2).

1.3.2 The Reynolds stress

In Cartesian coordinates

We analyse the Reynolds stress in Cartesian perpendicular coordinates (x, y), where x is a radial-like coordinate and y a binormal-like coordinate. We start by writing the (ion) momentum equation in the general form

$$\frac{\partial \mathbf{v}}{\partial t_{\mathrm{fast}}} = -\mathbf{v} \cdot \nabla \mathbf{v} + \mathcal{L}[\mathbf{v}], \tag{1.3.2}$$

where we have grouped all linear terms into the general linear operator $\mathcal{L}$. We also introduced t_{fast} which is assumed as the fastest time scale of interest for our purposes here and is taken to be the turbulence time scale. We proceed by using the *Reynolds decomposition*,

$$\mathbf{v}(x,y,t_{\text{fast}}) = \langle \mathbf{v} \rangle (t,x) + \tilde{\mathbf{v}}(x,y,t_{\text{fast}}). \tag{1.3.3}$$

Here we used the *Reynolds average* or *ensemble average*

$$\langle \dots \rangle \equiv \frac{1}{TL_y} \int_0^{L_y} \int_t^{t+T} \dots \mathrm{d}y \, \mathrm{d}t'_{\text{fast}} \tag{1.3.4}$$

to separate mean from fluctuating quantities. The mean quantity now depends on a slower so-called *meso* time scale t where $t_{\text{fast}} \ll t \ll t_{\text{eq}}$, i.e. this time scale is much larger than the turbulent time scale but smaller than equilibrium time scales. Note that equation 1.3.3 and 1.3.4 define

$$\tilde{\mathbf{v}} \equiv \mathbf{v} - \langle \mathbf{v} \rangle \qquad \text{such that} \qquad \langle \tilde{\mathbf{v}} \rangle = 0 \tag{1.3.5}$$

by definition. We proceed by Reynolds-averaging the y-component of the momentum equation 1.3.2 and focusing on the nonlinear contributions

$$\left. \frac{\partial \langle v_y \rangle}{\partial t} \right|_{\text{NL}} = -\langle \tilde{v}_x \partial_x \tilde{v}_y + \tilde{v}_y \partial_y \tilde{v}_y \rangle - \langle \mathbf{v} \rangle \cdot \nabla \langle v_y \rangle \tag{1.3.6}$$

where we used 1.3.5 and we identify the last term as the *mean* convection term which does not feature any turbulent quantities. We proceed by making use of the incompressibility condition $\nabla \cdot \mathbf{v} = 0$ which gives

$$\begin{aligned} \nabla \cdot \mathbf{v} &= \partial_x(\langle v_x \rangle + \tilde{v}_x) + \partial_y(\langle v_y \rangle + \tilde{v}_y) = 0 \\ \langle \nabla \cdot \mathbf{v} \rangle &= \partial_x \langle v_x \rangle + \partial_y \langle v_y \rangle = 0 \end{aligned} \tag{1.3.7}$$

such that it follows $\partial_x \tilde{v}_x = -\partial_y \tilde{v}_y$. Note that incompressibility follows directly if one assumes the turbulent motion to be purely due to $\mathbf{E} \times \mathbf{B}$ motion (a common assumption which we will use throughout this work) since for perpendicular turbulence $\nabla \cdot \mathbf{b} = 0$, and $\nabla \times \nabla f = 0$ for any scalar function f, where here $f = \varphi$. We can now write

$$\langle \tilde{v}_x \partial_x \tilde{v}_y + \tilde{v}_y \partial_y \tilde{v}_y \rangle = \Big\langle \partial_x(\tilde{v}_x \tilde{v}_y) + \partial_y(\tilde{v}_y^2) \Big\rangle = \partial_x \langle \tilde{v}_x \tilde{v}_y \rangle. \tag{1.3.8}$$

In the last step we used $\langle \partial_y \dots \rangle = 0$ since the turbulence is expected to be *statistically periodic* at the boundaries of the y-domain as long as the domain is significantly larger than the turbulent correlation length in this dimension.

We therefore arrive at our final expression

$$\left.\frac{\partial \langle v_y \rangle}{\partial t}\right|_{\mathrm{NL,turb}} = -\partial_x \langle \tilde{v}_x \tilde{v}_y \rangle \tag{1.3.9}$$

which shows that the mean poloidal flow receives a non-zero contribution from turbulent fluctuations as long as the x-y component of the Reynolds stress tensor $\langle \tilde{v}_i \tilde{v}_j \rangle$ has a finite radial gradient, as first shown by (Diamond & Kim 1991). Physically this corresponds to a radial variation of the dominant eddy tilt angle as depicted in figure 1.3.2.

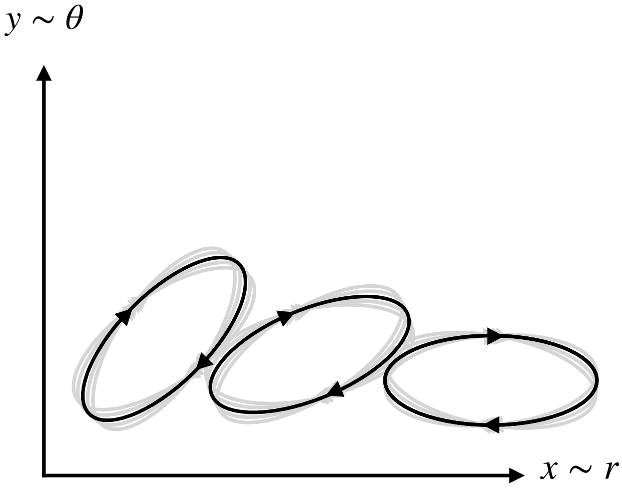

Figure 1.3.2. – Sketch of an eddy structure with radial gradient of Reynolds stress.

In cylindrical coordinates

The original reference (Diamond & Kim 1991) stated the Reynolds stress in cylindrical coordinates, motivating us to show that the analogous result to 1.3.9 indeed emerges for (v_r, v_θ). Using the standard form of the gradient operator in cylindrical coordinates, equation 1.3.6 becomes

$$\left.\frac{\partial \langle v_\theta \rangle}{\partial t}\right|_{\mathrm{NL}} = -\langle \tilde{v}_r \partial_r \tilde{v}_\theta + \frac{\tilde{v}_\theta}{r} \partial_\theta \tilde{v}_\theta \rangle - \langle \mathbf{v} \rangle \cdot \nabla \langle v_\theta \rangle \tag{1.3.10}$$

where we can write the first ensemble average on the right-hand side as

$$-\langle \partial_r (\tilde{v}_r \tilde{v}_\theta) - \tilde{v}_\theta \partial_r \tilde{v}_r + \frac{1}{2r} \partial_\theta \tilde{v}_\theta^2 \rangle \tag{1.3.11}$$

and notice that the last term in this expression vanishes due to $\langle \partial_\theta \ldots \rangle = 0$. The second term also vanishes (Diamond & Kim 1991). We thus obtain

$$\left.\frac{\partial \langle v_\theta \rangle}{\partial t}\right|_{\text{NL,turb}} = -\partial_r \langle \tilde{v}_r \tilde{v}_\theta \rangle, \tag{1.3.12}$$

the cylindrical coordinate analogue of equation 1.3.9.

1.4 Nonlinear turbulence-flow coupling and transfer functions

In this section we introduce the key analysis method for this work, namely the quantification of the nonlinear interaction of turbulence and ZFs. We start by introducing higher order spectra to sketch out the basic idea behind quantifying nonlinear mode interaction, before we introduce the energy transfer functions which will be at the centre of this work.

1.4.1 Higher order spectra

Higher order spectra give additional insight into the contributions to a signal compared to the well-known power spectrum, i.e. the (squared) magnitude of the Fourier transformed signal. Consider a simple nonlinearity of the form (Cziegler 2011)

$$\begin{aligned} f(x) &= [\cos(k_a x + \theta_a) + \cos(k_b x + \theta_b)]^2 \\ &= 1 + \cos[(k_a + k_b)x + \theta_a + \theta_b] + \cos[(k_a - k_b)x + \theta_a - \theta_b] \\ &+ \frac{1}{2}\{\cos[2k_a x + 2\theta_a] + \cos[2k_b x + 2\theta_b]\}. \end{aligned} \tag{1.4.1}$$

Further, consider the (related) linear signal

$$\begin{aligned} g(x) &= 1 + \cos[(k_a + k_b)x + \theta_A] + \cos[(k_a - k_b)x + \theta_B] \\ &+ \frac{1}{2}\{\cos[2k_a x + \theta_C] + \cos[2k_b x + \theta_D]\}. \end{aligned} \tag{1.4.2}$$

We define the first order and second order (auto-)correlation functions for some signal $y(x)$:

$$C_1(x') = \overline{y(x)y(x+x')}, \tag{1.4.3}$$

$$C_2(x', x'') = \overline{y(x)y(x+x')y(x+x'')}. \tag{1.4.4}$$

Here the overline denotes an averaging over the domain x. A quick calculation shows that taking the Fourier-transform of C_1 yields the familiar power spectrum:

$$S(k) \equiv \mathcal{F}(C_1(x')) = \int_{-\infty}^{+\infty} \left[\int_{-\infty}^{+\infty} y(x)y(x+x')dx\right] e^{-ikx'}dx' = \cdots = |\hat{y}(k)|^2. \tag{1.4.5}$$

The equivalent operation on C_2 yields the so-called bispectrum[8]

$$\begin{aligned} B(k_1, k_2) &\equiv \langle \mathcal{F}(C_2(x', x''))\rangle \\ &= \left\langle \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \left[\int_{-\infty}^{+\infty} y(x) y(x+x') y(x+x'') dx \right] e^{-ikx'} e^{-ikx''} dx' dx'' \right\rangle \\ &= \cdots = \langle \hat{y}^*(k_1) \hat{y}(k_2) \hat{y}(k_1 - k_2)\rangle \\ &= |\hat{y}^*(k_1)||\hat{y}(k_2)||\hat{y}(k_1 - k_2)| \left\langle e^{i(-\theta_1+\theta_2+\theta_3)} \right\rangle , \end{aligned} \tag{1.4.6}$$

where $\langle . \rangle$ denotes an ensemble average over realisations, i.e. an average over samples from distributions of phases θ_1, θ_2, and θ_3. Here θ_3 is the phase of the third frequency component $k_3 \equiv k_1 - k_2$. Applying such an ensemble average to the power spectrum in equation 1.4.5 does not make any difference since the power spectrum is not sensitive to the phases, evident by the absolute value operation that is performed on the complex-valued $\hat{y}(k)$.

Both signals, f and g, would yield the same power spectrum and would be indistinguishable, despite one being a nonlinear signal and the other being a linear one.

Due to the ensemble average the bispectrum B will only yield non-zero values when the three frequencies k_1, k_2, and $k_3 \equiv k_1 - k_2$ are phase-coupled, i.e. when the phase factor $\exp\left[i(-\theta_1 + \theta_2 + \theta_3)\right]$ is not randomly distributed due to an underlying nonlinearity. For the linear signal g the three phases would each be drawn from their own distribution, giving a randomly distributed sum, and thus $B = 0$. For f, however, this is not the case. For example, let us take $k_1 = k_a + k_b$ and $k_2 = k_a - k_b$, such that $k_3 = k_1 - k_2 = 2k_b$. Looking at the phases of the corresponding terms in f we see that $-\theta_1 = -(\theta_a + \theta_b)$, $\theta_2 = \theta_a - \theta_b$, and $\theta_3 = 2\theta_b$, such that their sum vanishes due to their coupled nature and the phase factor becomes unity.

Different to the power spectrum, the bispectrum is complex-valued, making it of more limited use for human intuition. A related quantity is the so-called (squared) bicoherence, defined in its normalised form as

$$b^2(k_1, k_2) = \frac{|B(k_1, k_2)|^2}{\langle |\hat{y}(k_1)|^2\rangle \, \langle |\hat{y}(k_2)\hat{y}(k_1 - k_2)|^2\rangle} \tag{1.4.7}$$

which takes on values between zero and unity, with zero indicating no coupling and unity indicating pure coupling. (A case in between would for example be the signal $f + g$.)

---

[8] Note that the symbol B refers to the bispectrum only in this section. In all other sections of this Thesis it refers to the magnetic field strength.

Evaluating the bicoherence on the entire valid part of the $k_1 - k_2$ plane of a given signal (i.e. the part where all three frequencies k_1, k_2, and k_3 are within bounds) is therefore a useful analysis to see all the nonlinear interactions within a complex signal, such as for example a fluctuating quantity in a turbulent fusion plasma. Such an analysis of the three-wave-coupling in a fusion plasma was first done by (Ritz, Powers & Bengtson 1989).

While the bicoherence tells us the strength of the coupling between two frequencies, it does not provide information on the directionality of the coupling, i.e. which mode tends to lose energy in the process and which tends to gain it. This shortcoming is addressed by the related energy transfer functions which will be introduced in the next sub-section.

1.4.2 Energy transfer functions

Energy transfer functions are a useful extension to the bicoherence. Their benefit over the bicoherence is additional information, and we will see how it allows us to specify zonal flow *drive* as opposed to general turbulence-flow *coupling*. The analysis here roughly follows the one in Ref. (Xu et al. 2009). However, we will take the coupling to take place in wavenumber space rather than temporal frequency space. The reason being that we focus on the analysis of gyrokinetic simulations which offer great wavenumber resolution and are naturally decomposed as such.

General formulation Consider the following set-up which is present in most fluid equations,

$$\frac{\partial \mathbf{X}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{X} = \mathcal{L}[\mathbf{X}], \tag{1.4.8}$$

i.e. the advection of some fluid quantity $\mathbf{X}$ by the fluid velocity (or flow) $\mathbf{V}$. In 1.4.8 $\mathcal{L}$ is some linear operator which stands for all other (linear) terms in the specific equation of interest. Here nonlinear means that it involves the product of two or more fluctuating quantities. This nonlinearity in the Navier-Stokes equation is the source of fluid turbulence, and in plasmas it is this nonlinearity in the continuity equation, momentum equation, or the evolution equation of higher moment fluid quantities which causes plasma turbulence.

We Fourier transform the equation, turning the product into a convolution, and bringing out the mode coupling implied by the nonlinearity:

$$\frac{\partial \hat{\mathbf{X}}_{\mathbf{k}}}{\partial t} + \sum_{\mathbf{k}_1} \hat{\mathbf{V}}_{\mathbf{k}-\mathbf{k}_1} \cdot (i\mathbf{k}_1) \hat{\mathbf{X}}_{\mathbf{k}_1} = \hat{\mathcal{L}}[\mathbf{X}]_{\mathbf{k}}. \tag{1.4.9}$$

Here we used $\nabla \rightarrow i\mathbf{k}$ under Fourier transformation. Upon multiplication by $\hat{\mathbf{X}}^*_{\mathbf{k}}$ and adding the equation to its complex conjugate we obtain an evolution equation for the spectral energy of the quantity $\mathbf{X}$:

$$\frac{\partial |\hat{\mathbf{X}}_{\mathbf{k}}|^2}{\partial t} = -2\,\mathrm{Re}\left\{\sum_{\mathbf{k}_1} \hat{\mathbf{X}}^*_{\mathbf{k}} \hat{\mathbf{V}}_{\mathbf{k}-\mathbf{k}_1} \cdot (i\mathbf{k}_1) \hat{\mathbf{X}}_{\mathbf{k}_1}\right\} + 2\,\mathrm{Re}\left\{\hat{\mathbf{X}}^*_{\mathbf{k}} \hat{\mathcal{L}}[\mathbf{X}]_{\mathbf{k}}\right\}. \tag{1.4.10}$$

In the final step we take the ensemble average of the equation to obtain the steady-state energy transfer[9], yielding

$$\left\langle \frac{\partial |\hat{\mathbf{X}}_{\mathbf{k}}|^2}{\partial t} \right\rangle = -2\,\mathrm{Re}\left\langle \sum_{\mathbf{k}_1} \hat{\mathbf{X}}^*_{\mathbf{k}} \hat{\mathbf{V}}_{\mathbf{k}-\mathbf{k}_1} \cdot (i\mathbf{k}_1) \hat{\mathbf{X}}_{\mathbf{k}_1} \right\rangle + 2\,\mathrm{Re}\left\langle \hat{\mathbf{X}}^*_{\mathbf{k}} \hat{\mathcal{L}}[\mathbf{X}]_{\mathbf{k}} \right\rangle. \tag{1.4.11}$$

The so-called three wave coupling inside the sum resembles the bispectrum of equation 1.4.6. However, here one mode is clearly distinguished as the receiver or source of energy, namely the mode $\mathbf{k}$. Thus equation 1.4.11 tells us how modes $\mathbf{k}_1$ and $\mathbf{k} - \mathbf{k}_1$ couple to drive the mode $\mathbf{k}$. Whether mode $\mathbf{k}$ is the receiver or the source of spectral energy in one particular three-wave interaction is determined by the sign of

$$-\,\mathrm{Re}\left\langle \hat{\mathbf{X}}^*_{\mathbf{k}} \hat{\mathbf{V}}_{\mathbf{k}-\mathbf{k}_1} \cdot (i\mathbf{k}_1) \hat{\mathbf{X}}_{\mathbf{k}_1} \right\rangle. \tag{1.4.12}$$

Positive values translate to mode $\mathbf{k}$ receiving spectral energy, whereas negative values imply that spectral energy is lost.

Note that since $\mathbf{X}$ is expected to represent a real-valued quantity, its Fourier transform satisfies the so-called reality condition $\hat{\mathbf{X}}_{\mathbf{k}} = \hat{\mathbf{X}}^*_{-\mathbf{k}}$, a relation which could be used to re-write equation 1.4.12 in a way such that it features no complex conjugate.[10] For sake of clarity, however, we keep it in its current form.

Note that this distinction of target, source, and mediator compared to general coupling as measured by the bicoherence is important since previous work showed that situations exist where the GAM strongly couples to turbulence but it does not actually *receive* substantial net kinetic energy, while the mean ZF does (see figure 6 in (Cziegler et al. 2015)). The first use of such transfer functions in the fusion plasma setting was for the simplified Hasegawa-Wakatani system (Hasegawa & Wakatani 1983; Wakatani & Hasegawa 1984) by (Camargo,

---

[9] Note that here the ensemble average $\langle \dots \rangle$ represents purely an average over time and not over the poloidal domain as it did in equation 1.3.4, which might seem inconsistent. However, it is not since this average is now enacted by choosing $k_y = 0$ to obtain the evolution of the mean poloidal component of $\mathbf{X}$.

[10] The reality condition is a welcome relation for solving the GKE numerically, since it allows one to simulate only the positive k_y range and automatically get those for negative values of k_y as well.

Biskamp & Scott 1995) and later applied to the general plasma continuity and ion momentum equation by (Xu et al. 2009).

Kinetic energy transfer function In this Thesis we are primarily interested in the nonlinear drive of zonal *flows*. We therefore choose the central fluid quantity of interest, $\mathbf{X}$, to be the flow itself, i.e. we replace it with $\mathbf{V}$. Equation 1.4.8 then becomes

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V} = \mathcal{L}[\mathbf{V}], \tag{1.4.13}$$

a way of casting the momentum equation of the plasma to focus on its nonlinearity.[11] *This equation is now the same as equation 1.3.2 and thus describes the Reynolds stress from section 1.3.2.* The operator $\mathcal{L}$ includes the pressure gradient force, the Lorentz force, the inter-species friction force, and the viscosity tensor. Equation 1.4.11 then gives us

$$\left\langle \frac{\partial |\hat{\mathbf{V}}_\mathbf{k}|^2}{\partial t} \right\rangle = -2\,\mathrm{Re} \left\langle \sum_{\mathbf{k_1}} \hat{\mathbf{V}}^*_\mathbf{k} \hat{\mathbf{V}}_{\mathbf{k}-\mathbf{k_1}} \cdot (i\mathbf{k_1}) \hat{\mathbf{V}}_{\mathbf{k_1}} \right\rangle + 2\,\mathrm{Re} \left\langle \hat{\mathbf{V}}^*_\mathbf{k} \hat{\mathcal{L}}[\mathbf{V}]_\mathbf{k} \right\rangle, \tag{1.4.14}$$

the evolution of the spectral kinetic energy.

Further progress can be made to bring equation 1.4.14 into more explicit form. We use the local coordinate system from figure 1.1.3, where x represents the radial coordinate, y the coordinate within the flux surfaces that is perpendicular to the magnetic field $\mathbf{B}$, and z the coordinate along said magnetic field; the x-y-plane is referred to as the 'perpendicular plane' and the direction z as the 'parallel' direction.

The movement of charged particles in a plasma is unrestricted in the parallel direction, making parallel wavelengths much longer than perpendicular ones, i.e.

$$|\mathbf{k}_\perp| \gg k_\parallel. \tag{1.4.15}$$

Thus the dominant mode coupling takes place in the perpendicular plane and from now on we take $\mathbf{k} = k_x \nabla x + k_y \nabla y$ and $\nabla = \nabla_\perp$. We make further progress by approximating the fluid motion by $\mathbf{E} \times \mathbf{B}$ motion, i.e. the fluctuating drifts that arise due a fluctuating electrostatic potential,

$$\mathbf{V} = B^{-1} \mathbf{b} \times \nabla \delta\varphi \quad \text{and thus} \quad \hat{\mathbf{V}}_\mathbf{k} = B^{-1} \mathbf{b} \times (i\mathbf{k}\delta\hat{\varphi}_\mathbf{k}). \tag{1.4.16}$$

---

[11] For ITG it is the momentum equation of ions which is of interest. To avoid clutter we omit the species label everywhere in this section.

Substituting this expression into equation 1.4.14 we obtain

$$\left\langle \frac{\partial |\hat{\mathbf{V}}_{\mathbf{k}}|^2}{\partial t} \right\rangle = \mathcal{T} + 2\,\mathrm{Re}\left\langle \hat{\mathbf{V}}^*_{\mathbf{k}} \hat{\mathcal{L}}[\mathbf{V}]_{\mathbf{k}} \right\rangle, \tag{1.4.17}$$

where we defined the kinetic energy transfer function

$$\mathcal{T} = 2B^{-3}\,\mathrm{Re}\left\langle \sum_{\mathbf{k_1}} [\mathbf{b}\times\mathbf{k}\cdot\mathbf{k}_1](\mathbf{b}\times\mathbf{k})\cdot(\mathbf{b}\times\mathbf{k}_1)\delta\hat{\varphi}^*_{\mathbf{k}}\delta\hat{\varphi}_{\mathbf{k}-\mathbf{k_1}}\delta\hat{\varphi}_{\mathbf{k_1}} \right\rangle. \tag{1.4.18}$$

Note that $\mathcal{T} = \mathcal{T}(k_x, k_y, k_{1x}, k_{1y}, z)$, with the z-dependence being introduced by $\delta\hat{\varphi}_{\mathbf{k}} = \delta\hat{\varphi}(k_x, k_y, z)$. With the expression 1.4.18 we can obtain the steady state drive of any individual mode (k_x, k_y). As the interest here is in the drive of ZFs we divide all modes into two groups: the 'zonal' modes where $k_y = 0$ and the 'turbulent' modes with $k_y \neq 0$. Modes with $k_y = 0$ are those where φ is constant in the direction within the flux surface, marking distinct radial zones determined by their radial wavenumber k_x. To get the total drive of ZFs, the central quantity which we are interested in in this work, we set $k_y = 0$ and sum over all k_x as well as all possible source modes $\mathbf{k}_1$:

$$\mathcal{T}_{ZF}(\theta) \equiv \sum_{k_x, k_{1x}, k_{1y}} \mathcal{T}(k_x, k_y = 0, k_{1x}, k_{1y}, z). \tag{1.4.19}$$

Note that when we evaluate $\mathcal{T}_{ZF}$ in gyrokinetic simulations we do not realise it as an average over N simulations as one might think. Instead we realise it as a temporal average. Provided that (1) the turbulence is stationary, i.e. its statistical properties do not change over time, and (2) the time over which the average is taken is substantially longer than the turbulent decorrelation time, the temporal average and ensemble average are equivalent due to the ergodic theorem (Frisch 1995). The reason for choosing the temporal average is a practical one: each (nonlinear) simulation includes a linear phase at the beginning until turbulent fluctuations saturate due to nonlinear coupling and fully developed turbulence is obtained. Running N nonlinear simulations would require calculating N such linear phases which have no use for the calculation of $\mathcal{T}_{ZF}$, thus greatly increasing the computational expense.

Symmetries of energy transfer functions The full transfer functions, such as $\mathcal{T}$ from equation 1.4.18, contain symmetries (Maeyama et al. 2021). To understand these symmetries it proves helpful to write $\mathcal{T} = \mathcal{T}(\mathbf{k}, \mathbf{k}_1, \mathbf{k}_2)$, even though $\mathbf{k}_2$ is not an independent parameter but is determined through the coupling condition $\mathbf{k}_2 = \mathbf{k} - \mathbf{k}_1$ by the other two modes.

The first symmetry appears under exchange of the source and target mode, that is one finds that

$$\mathcal{T}(\mathbf{k}=\mathbf{p},\mathbf{k}_1=\mathbf{q},\mathbf{k}_2=\mathbf{j}) = -\mathcal{T}(\mathbf{k}=\mathbf{q},\mathbf{k}_1=\mathbf{p},\mathbf{k}_2=-\mathbf{j}). \tag{1.4.20}$$

The second symmetry is with respect to parity,

$$\mathcal{T}(\mathbf{k}=\mathbf{p},\mathbf{k}_1=\mathbf{q},\mathbf{k}_2=\mathbf{j}) = \mathcal{T}(\mathbf{k}=-\mathbf{p},\mathbf{k}_1=-\mathbf{q},\mathbf{k}_2=-\mathbf{j}). \tag{1.4.21}$$

Note that we have adopted a different convention for our coupling condition than (Maeyama et al. 2021). We use the coupling condition $\mathbf{k}_2 = \mathbf{k}-\mathbf{k}_1$, whereas (Maeyama et al. 2021) uses $\mathbf{k}+\mathbf{k}_1+\mathbf{k}_2 = 0$.

Symmetries can become important when interpreting the transfer results and we will provide an example here with particular relevance for our context. Consider again the momentum equation 1.4.13. One might be inclined to consider the different contributions to the nonlinearity by writing it in the form

$$\partial_t \mathbf{V} + (V_x\partial_x + V_y\partial_y)V_x\nabla x + (V_x\partial_x + V_y\partial_y)V_y\nabla y = \mathcal{L}[\mathbf{V}],$$

and defining transfer functions $\mathcal{T}_{ij}$ for the nonlinearity $V_i\partial_i V_j$.[12] Here we have introduced the notation $\partial_x f = \partial f/\partial x$. Taking the same steps as before, namely Fourier transformation, multiplication by $\mathbf{V}^*_{\mathbf{k}}$, and adding of the complex conjugate equation, we obtain

$$\partial_t|\hat{\mathbf{V}}_{\mathbf{k}}|^2 + (\mathcal{T}_{xx}+\mathcal{T}_{yx}+\mathcal{T}_{xy}+\mathcal{T}_{yy}) = 2\hat{\mathcal{L}}[\mathbf{V}]_{\mathbf{k}},$$

where with the use of equation 1.4.16 we write

$$\begin{aligned}
\mathcal{T}_{xx} &= 2\,\mathrm{Re}\left(\hat{V}^*_{x,\mathbf{k}}\sum_{\mathbf{k}_1}\hat{V}_{x,\mathbf{k}-\mathbf{k}_1}(ik_{1x})\hat{V}_{x,\mathbf{k}_1}\right) = -Ck_y(k_y-k_{1y})k_{1x}k_{1y},\\
\mathcal{T}_{yx} &= 2\,\mathrm{Re}\left(\hat{V}^*_{x,\mathbf{k}}\sum_{\mathbf{k}_1}\hat{V}_{y,\mathbf{k}-\mathbf{k}_1}(ik_{1y})\hat{V}_{x,\mathbf{k}_1}\right) = +Ck_y(k_x-k_{1x})k_{1y}k_{1y},\\
\mathcal{T}_{xy} &= 2\,\mathrm{Re}\left(\hat{V}^*_{y,\mathbf{k}}\sum_{\mathbf{k}_1}\hat{V}_{x,\mathbf{k}-\mathbf{k}_1}(ik_{1x})\hat{V}_{y,\mathbf{k}_1}\right) = -Ck_x(k_y-k_{1y})k_{1x}k_{1x},\\
\mathcal{T}_{yy} &= 2\,\mathrm{Re}\left(\hat{V}^*_{y,\mathbf{k}}\sum_{\mathbf{k}_1}\hat{V}_{y,\mathbf{k}-\mathbf{k}_1}(ik_{1y})\hat{V}_{y,\mathbf{k}_1}\right) = +Ck_x(k_x-k_{1x})k_{1y}k_{1x},
\end{aligned}$$

---

[12] we are not using Einstein summation convention here; the indices are kept general only to shorten notation.

where $C \equiv 2B^{-3}\,\mathrm{Re}\,(\delta\hat{\varphi}^*_{\mathbf{k}}\delta\hat{\varphi}_{\mathbf{k}-\mathbf{k}_1}\delta\hat{\varphi}_{\mathbf{k}_1})$. For the drive of ZFs, for which we set $k_y = 0$, one might then conclude that the only contribution stems from the $\mathcal{T}_{yy}$ transfer term. ($\mathcal{T}_{xx}$, $\mathcal{T}_{yx}$, and the first term of $\mathcal{T}_{xy}$ vanish directly; the second term of $\mathcal{T}_{xy}$ cancels with the second term of $\mathcal{T}_{yy}$.)

Note that for the drive of ZFs, as considered here in the final step, it is enough to consider only the y-component of the momentum equation, i.e.

$$\partial_t V_y\Big|_{\mathrm{NL}} = -V_x\partial_x V_y - V_y\partial_y V_y, \tag{1.4.22}$$

giving the same result as before but now without the $\mathcal{T}_{ix}$ terms,

$$\begin{aligned}\partial_t|V_{y,\mathbf{k}}|^2\Big|_{\mathrm{NL}} &= 2B^{-3}\sum_{\mathbf{k}_1}\left[k_x(k_y-k_{1y})k_{1x}k_{1x} - k_x(k_x-k_{1x})k_{1y}k_{1x}\right] \\ &\quad\times \mathrm{Re}\,(\delta\varphi^*_{\mathbf{k}}\delta\varphi_{\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}_1}).\end{aligned} \tag{1.4.23}$$

However, by using the incompressibility condition,

$$\nabla\cdot\mathbf{V} = 0 \quad \text{implying} \quad \partial_x V_x = -\partial_y V_y,$$

we can rewrite equation 1.4.22 as

$$\partial_t V_y\Big|_{\mathrm{NL}} = -\partial_x(V_xV_y) - \partial_y V_y^2.$$

We then obtain

$$\begin{aligned}\partial_t|V_{y,\mathbf{k}}|^2\Big|_{\mathrm{NL}} &= 2B^{-3}\sum_{\mathbf{k}_1}\left[k_xk_x(k_y-k_{1y})k_{1x} - k_xk_y(k_x-k_{1x})k_{1x}\right] \\ &\quad\times \mathrm{Re}\,(\delta\varphi^*_{\mathbf{k}}\delta\varphi_{\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}_1})\end{aligned} \tag{1.4.24}$$

such that now the only contribution for $k_y = 0$ stems from the term $\partial_x(V_xV_y)$, which we identify as the Reynolds stress. Defining the two transfer functions from equations 1.4.23 and 1.4.24 as

$$\begin{aligned}\mathcal{T}_1 &= 2B^{-3}\left[k_xk_{2y}k_{1x}k_{1x} - k_xk_{2x}k_{1y}k_{1x}\right]\mathrm{Re}\,(\delta\varphi^*_{\mathbf{k}}\delta\varphi_{\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}_1}),\\ \mathcal{T}_2 &= 2B^{-3}\left[k_xk_xk_{2y}k_{1x} - k_xk_yk_{2x}k_{1x}\right]\mathrm{Re}\,(\delta\varphi^*_{\mathbf{k}}\delta\varphi_{\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}_1}),\end{aligned}$$

respectively, we observe that

$$\mathcal{T}_1(\mathbf{k}=\mathbf{p},\mathbf{k}_1=\mathbf{q},\mathbf{k}_2=\mathbf{j}) = 2B^{-3}\left[p_x j_y q_x q_x - p_x j_x q_y q_x\right]\mathrm{Re}\left(\delta\varphi^*_{\mathbf{p}}\delta\varphi_{\mathbf{j}}\delta\varphi_{\mathbf{q}}\right)$$

is in fact equal to $-\mathcal{T}_2(\mathbf{k}=-\mathbf{q},\mathbf{k}_1=-\mathbf{p},\mathbf{k}_2=\mathbf{j})$. With the second symmetry relation 1.4.21 we note that

$$-\mathcal{T}_2(\mathbf{k}=-\mathbf{q},\mathbf{k}_1=-\mathbf{p},\mathbf{k}_2=\mathbf{j}) = -\mathcal{T}_2(\mathbf{k}=\mathbf{q},\mathbf{k}_1=\mathbf{p},\mathbf{k}_2=-\mathbf{j}),$$

and due to the first symmetry relation 1.4.20

$$-\mathcal{T}_2(\mathbf{k}=\mathbf{q},\mathbf{k}_1=\mathbf{p},\mathbf{k}_2=-\mathbf{j}) = \mathcal{T}_2(\mathbf{k}=\mathbf{p},\mathbf{k}_1=\mathbf{q},\mathbf{k}_2=\mathbf{j}),$$

this finally implies

$$\mathcal{T}_1 = \mathcal{T}_2.$$

This shows that the symmetries reach across the individual terms in the nonlinearity and that one should therefore not write out and consider individual transfer functions.

1.5 Gyrokinetics

In this chapter gyrokinetics (GK) is summarised, the formalism which is the workhorse of turbulence theory and simulation in tokamaks and stellarators. Many extensive introductions to gyrokinetics have been written, and we will therefore not attempt to give a complete derivation with all its details here. For these purposes we refer the reader to the review by (Abel et al. 2013)[13] and for a slower-paced introduction we refer to the following Ph.D. Thesis (Dudding 2022). Here, we will only highlight the key ideas to establish the setting and terminology in which much of this work is placed.

1.5.1 Governing equations

Gyrokinetics is derived from the Vlasov-Maxwell system of equations and makes use of the time-scale separation between the fast gyro-motion of particles and the comparatively slow time-scale on which the drift-wave turbulence evolves. The fast gyro-motion itself is not resolved and averaged out in the formalism, providing dimensional reduction from six to five and thus making the problem computationally tractable.

In general, the gyrokinetic system of equations is fully electromagnetic and can include (statically imposed) equilibrium flow and flow shear. Here we make the assumptions of electrostatic turbulence, i.e. we assume magnetic fluctuations to be small. In addition, we assume that equilibrium flow and flow shear are negligible. Both of these effects constitute possible extensions to the work presented here.

The gyrokinetic ordering is given by

$$\epsilon \sim \frac{\omega_{ds}}{\Omega_s} \sim \frac{k_\parallel}{k_\perp} \sim \frac{\rho_s}{a} \sim \frac{\delta f_s}{f_s} \sim \frac{\delta\varphi}{\varphi} \tag{1.5.1}$$

and is imposed on the Vlasov-Maxwell system to order the terms. The gyrokinetic equation (GKE) is obtained at second order. Here $\omega_{ds} = \mathbf{v}_{ds} \cdot \mathbf{k}$ is the drift frequency, $\Omega_s = e_s B/m_s$ the gyrofrequency with mass m_s, ρ_s the Larmor radius, f_s the full distribution function with δf_s its perturbation due to turbulence, and the same for the electrostatic potential φ.

The zeroth and first order give valuable information which will be employed to make further progress. At zeroth order we learn that the Poisson equation reduces to imposing quasi-neutrality of the plasma. At first order we learn

---

[13] Note that the derivation in this review allows strong equilibrium flows, an assumption which we do not make in this work. However by dropping terms containing the flow velocity $\mathbf{u}$, this derivation corresponds exactly to our assumptions made here.

that the lowest order mean distribution function is a Maxwellian, and that the perturbed distribution function consists of a term proportional to the electrostatic potential and another which is already independent of the gyro-angle (Abel et al. 2013). That is, for each species (denoted by subscript *s*),

$$f_s = F_s + \delta f_s, \tag{1.5.2}$$

with the zeroth-order Maxwellian

$$F_s(\psi, \varepsilon_s) = n_s(\psi)\left(\frac{m_s}{2\pi T_s(\psi)}\right)^{3/2} \exp\left(-\frac{\varepsilon_s}{T_s(\psi)}\right), \tag{1.5.3}$$

and the first order perturbation

$$\delta f_s = h_s - \frac{e_s \delta\varphi}{T_s} F_s. \tag{1.5.4}$$

Here $\varepsilon_s = m_s v^2/2$ is the energy of the particle with velocity v, T_s the temperature and n_s the density. h_s is the 'non-adiabatic' distribution function which does not depend on the gyro-angle α_g, i.e. it depends only on the gyro-centre position $\mathbf{R}$ and not on the particle position $\mathbf{r} = \mathbf{R} + \boldsymbol{\rho}$, where $\boldsymbol{\rho}(\alpha_g)$ is the gyro-radius vector, see figure 1.5.1.

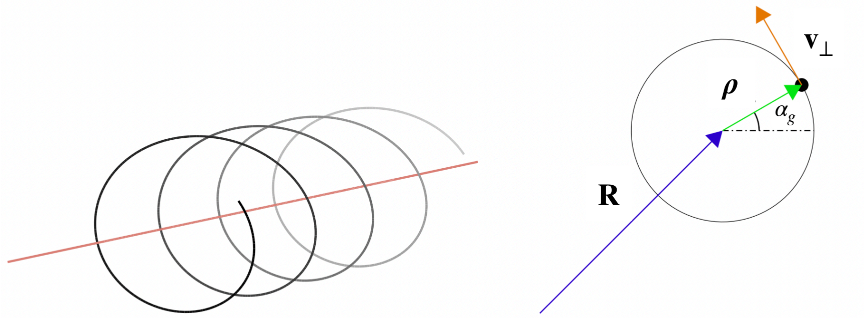

Figure 1.5.1. – The variables describing the gyro-motion. Taken from (Dickinson 2012), with variable names adapted to match our convention.

The nonlinear gyrokinetic equation is then given by (Frieman & Chen 1982)

$$\frac{\partial g_s}{\partial t} + (v_\parallel \mathbf{b} + \mathbf{V}_{Ds}) \cdot \nabla h_s + \mathcal{N}_s + \frac{\partial F_s}{\partial \psi} \langle \mathbf{V}_{\mathbf{E}\times\mathbf{B}} \rangle_s^{\mathbf{R}} \cdot \nabla\psi = 0, \tag{1.5.5}$$

which describes the temporal evolution of the gyro-averaged perturbed distribution function

$$g_s = \langle \delta f_s \rangle_s^{\mathbf{R}} = h_s - (e_s \langle \delta\varphi \rangle_s^{\mathbf{R}} / T_s) F_s, \tag{1.5.6}$$

where $\langle \ldots \rangle_s^{\mathbf{R}}$ denotes the gyro-averaging operation for species s at fixed guiding centre position $\mathbf{R}$. The gradient operator ∇ is also understood to be taken at constant $\mathbf{R}$. Note that after Fourier transforming the gyro-averaging operation $\langle X\rangle_s^{\mathbf{R}}$ turns into multiplication by a Bessel function, $J_0(k_\perp v_\perp/\Omega_s)X$.

The naming convention for the different distribution functions is quite inconsistent in the literature. The same symbols are usually used but what they are assigned to varies, increasing the possibility of confusion. We have therefore clarified our naming convention again in table 1.1.

Symbol	**Dependence**	**Explanation**
δf_s	$(t, \mathbf{r}, \mathbf{v})$	The full perturbed distribution function.
h_s	$(t, \mathbf{R}, \mathbf{v})$	The 'non-adiabatic' part of δf_s, i.e. the part which depends only on the gyro-centre $\mathbf{R}$ and not $\mathbf{r}$.
g_s	$(t, \mathbf{R}, \mathbf{v})$	The gyro-average of δf_s.

Table 1.1. – The symbols of distribution functions used in this Thesis.

Other quantities appearing in equation 1.5.5 are the parallel velocity $v_\parallel$ and the magnetic drift velocity

$$\mathbf{V}_{Ds} = \frac{1}{2\Omega_s}\mathbf{b} \times [2v_\parallel^2 \mathbf{b}\cdot\nabla\mathbf{b} + v_\perp^2 \nabla \ln \mathbf{B}], \qquad (1.5.7)$$

where $v_\perp$ is the perpendicular velocity, the gyrokinetic nonlinearity

$$\mathcal{N}_s = \langle \mathbf{V}_{\mathbf{E}\times\mathbf{B}}\rangle_s^{\mathbf{R}} \cdot \nabla h_s \qquad (1.5.8)$$

with

$$\mathbf{V}_s^{\mathbf{E}\times\mathbf{B}} = -\nabla(\langle\delta\phi\rangle_s^{\mathbf{R}}) \times \mathbf{B}/B^2. \qquad (1.5.9)$$

Equation 1.5.5 dictates the evolution of each species in the system which is treated gyro-kinetically. All are coupled together through Poisson's equation which describes how the electrostatic potential reacts to the densities and prevents inconsistent density accumulation in the system, i.e. it enforces quasi-neutrality. For the ITG instability the gyrokinetic dynamics of electrons is not crucial and as often done *we assume throughout this thesis that electrons follow a Boltzmann distribution* (also often referred to as the adiabatic electron assumption),

$$\delta n_e \equiv \int \langle g_e\rangle_{\mathbf{r}}\, d\mathbf{v} = \frac{e n_e}{T_e}(\delta\varphi - \overline{\delta\varphi}). \qquad (1.5.10)$$

Here the overline denotes the field line average, and with equation 1.5.10 the quasi-neutrality equation

$$\sum_s e_s \int \langle g_s \rangle_{\mathbf{r}} \, d\mathbf{v} = 0 \quad \Leftrightarrow \quad \sum_s \frac{n_s e_s^2}{T_s} \delta\varphi_{\mathbf{k}} = \sum_s e_s \int J_{0s\mathbf{k}} h_{s\mathbf{k}} d\mathbf{v} \tag{1.5.11}$$

reduces to

$$\int d\mathbf{v} \langle g_i \rangle_{\mathbf{r}} = \frac{e n_e}{Z_i T_e} (\delta\varphi - \overline{\delta\varphi}). \tag{1.5.12}$$

Now the gyro-average is taken at constant particle position $\langle . \rangle_{\mathbf{r}}$ since the fields, unlike the distribution functions, are not transformed to guiding centre coordinates (Abel et al. 2013).

1.5.2 Numerical treatment

In this work we use local gyrokinetics which takes advantage of the anisotropic length scales of the turbulence with a strong guide field such as toroidal plasmas. In such cases turbulence dynamics are much faster parallel to the guide field than perpendicular to it, resulting in turbulent eddies to be elongated along the field and of relatively short wavelength in the perpendicular direction, as indicated in the ordering in equation 1.5.1.

As a result, perpendicular correlation lengths are much shorter than parallel ones, motivating a radially local computational domain which is larger than the correlation lengths in each dimension but not significantly larger than it needs to be. The result is known as a flux-tube domain which extends for $\mathcal{O}(10\rho_i)$ in the perpendicular direction and for one poloidal turn in the parallel direction.

The radial local approach offers great computational speed-up compared to a radially global approach and allows one to understand how turbulence and zonal flows interact locally in realistic geometry. However, it does not allow one to dynamically adjust the gradients based on the particle and heat fluxes caused due to turbulence generated by previous gradient values, i.e. gradients and fluxes are inherently not self-consistent in a local approach. The fixed gradients are specified in the code through

$$\frac{a}{L_T} \equiv \frac{a}{T_0} \frac{dT}{dr} \quad \text{and} \quad \frac{a}{L_n} \equiv \frac{a}{n_0} \frac{dn}{dr} \tag{1.5.13}$$

which defines the temperature and density gradient scale lengths L_T and L_n. Here a is the devices minor radius as depicted in figure 1.1.3.

The flux tube is centred around a field line at position ψ_0 and α_0 and local radial and perpendicular coordinates are defined as (Beer, Cowley & Hammett 1995)

$$x = \frac{q_0}{B_0 r_0}(\psi - \psi_0) \quad \text{and} \quad y = -\frac{r_0}{q_0}(\alpha - \alpha_0) \tag{1.5.14}$$

such that they are proportional to ψ and α but have units of length. Here $q_0 = q(\psi_0)$ is the safety factor at the centre of the domain, B_0 is the field at the magnetic axis, and r_0 is the minor radius of the centre of the domain (Beer, Cowley & Hammett 1995).

The length of the computational domain is chosen such that it is longer than the turbulent correlation length in each dimension. Not only does this allow one to assume that the turbulent behaviour is statistically identical at both ends of the domain in both the binormal and parallel directions, but the same can be assumed in the radial direction since the domain is local around field lines and the energy-supplying equilibrium gradients can thus be assumed to be constant across the domain, i.e.

$$\left.\frac{a}{L_T}\right|_{x=0} = \left.\frac{a}{L_T}\right|_{x=L_x} \quad \text{and} \quad \left.\frac{a}{L_n}\right|_{x=0} = \left.\frac{a}{L_n}\right|_{x=L_x}. \tag{1.5.15}$$

The evolved quantities can therefore be expressed as Fourier series in the perpendicular domain,

$$Q(x, y, z, t) = \sum_{k_x, k_y} \hat{Q}(z, t)_{\mathbf{k}}\, e^{i(k_x x + k_y y)}. \tag{1.5.16}$$

Here Q stands for either the perturbed electrostatic potential $\delta\varphi$ or the perturbed distribution function δf, and z is the parallel coordinate.

For axisymmetric devices the parallel length of the flux-tube is chosen such that it begins and ends at the same poloidal location, in which case the turbulence is expected to be statistically identical in these places, since the geometry is exactly the same in both. In the case of exactly zero global magnetic shear $\hat{s}$ this permits periodic boundary conditions. With finite magnetic shear, different k_x modes must be coupled to each other, the so-called *twist-and-shift* boundary condition (Beer, Cowley & Hammett 1995).

In stellarators, after traversing the field line for one poloidal turn one is generally not at a location of the flux surface for which the geometry is the same as at the starting point. Instead, one can take advantage of field periodicity and stellarator symmetry to find locations of identical geometry and choose to let the flux tube end at those locations, provided that it is still longer than the parallel correlation length. Note that for a configuration with N_p number of field periods and stellarator symmetry the geometry repeats $2N_p$ times throughout the whole device.

The same subtleties as in axisymmetric devices apply for the coupling of modes at the ends of the parallel domain in stellarators, although the coupling can be generalised to take advantage of the integrated local shear rather than just the net global shear (Martin et al. 2018). Whenever we deal with non-axisymmetric configurations in this work we use this *generalised twist-and-shift* boundary condition by (Martin et al. 2018).

1.6 Motivation for plasma shaping

Once one has decided for toroidally nested flux surfaces as the topology for confining the plasma, one is left to choose the shaping of those flux surfaces. Since the shaping is predominantly controlled by external currents, i.e. by the coils, it is a direct design choice and less a product of the self-consistent evolution of the plasma, like many other processes within it. Thus the shaping of the coils and the currents within them mostly control the equilibrium and reflect a direct experimental control mechanism that can be exploited and optimised for desirable physics characteristics.

1.6.1 Axisymmetric shaping

In axisymmetry (tokamak) the shaping freedom reduces to the shape of the poloidal cross-section of each flux surface. The shaping components are specified by their poloidal Fourier-mode number m and named after their effect, giving elongation ($m = 1$), triangularity ($m = 2$), and squareness ($m = 3$)[14]. Higher Fourier-mode shaping is usually subdominant in axisymmetric equilibria and therefore not studied in isolation.

Over the past three decades the beneficial effect of axisymmetric shaping, i.e. flux surfaces that deviate from circularly-shaped poloidal cross-sections, has become increasingly clear. Increased elongation and positive triangularity (see figure 1.7.1 for their definition) increase the plasma beta threshold for Mercier stability (Freidberg 2014)[15], and also increase the stability of the plasma against micro-instabilities such as the ITG instability (Belli, Hammett & Dorland 2008).

However, these beneficial effects were not always known and were the primary reason that the *Tokamak Fusion Test Reactor (TFTR)* operated in Princeton during the ’80s and ’90s was – despite setting fusion records – limited in its performance, and lead to the next-generation large-scale and only other tritium-fueled device, the *Joint European Torus (JET)*, to be designed with elongated and (outward) D-shaped plasma shaping (Freidberg 2014).

---

[14] See figure 1.7.1 for a visualisation of elongation and triangularity.

[15] The Mercier criterion is a sufficient condition for interchange stability and a necessary condition for ballooning stability, i.e. iff the Mercier criterion is satisfied ballooning stability can be tested. This stability criterion will be addressed further in chapter 3.10.

1.6.2 Symmetry-preserving three-dimensional shaping

As alluded to earlier, breaking axisymmetry of the toroidally nested flux-surfaces is the key idea behind *stellarators*. Whereas in tokamaks a helical field can only be generated with a plasma current, in stellarators the three-dimensional (3D) shaping can generate a helical field on its own. This allows stellarators to be more easily designed for steady state operation, since external current drive is not necessarily required and coils can be designed to create flux-surfaces in the absence of any plasma such that no solenoid is required in the centre of the torus. Further, stellarators generally do not suffer from disruptions and are able to operate at densities that exceed the Greenwald limit (Greenwald 2002)[16]. It was shown in the compact toroidal hybrid (Hartwell et al. 2017) that densities much beyond the Greenwald limit could be obtained once the edge external rotational transform obtained values of $0.1 - 0.15$ or larger (Maurer 2015).

While finding non-axisymmetric ideal MHD equilibria with nested flux surfaces is relatively straightforward, finding ones which confine collisionless guiding centre trajectories is challenging and requires careful numerical optimisation of both the plasma equilibrium and coils to support it. The reasons for this will be further described in section 3.2.

Three-dimensional shaping of small magnitude is also being used in tokamaks. The high confinement operating regime (H-regime or H-mode), discovered in 1982 (Wagner et al. 1982), has long been seen as the most desirable operating mode for tokamak plasmas. Its steep edge gradients of temperature and density increase their core respective values, thus increasing performance. However, the H-mode comes with a list of drawbacks, one of the key ones being instabilities localised to the edge region which cause explosive transport events that damage plasma facing components. Therefore, these edge localised modes (ELMs) must be either avoided or mitigated. One of the key mitigation strategies is to increase their frequency, thereby reducing the energy expelled during each ELM event. One key idea which has been used on nearly all tokamaks around the world are so-called resonant magnetic perturbations (RMPs), small 3D fields with distinct toroidal mode numbers n which typically increase the transport in the plasma edge, thereby lowering the local pressure gradient, and avoiding the unstable ELM regime (Evans et al. 2004; Kirk et al. 2013).

---

[16] The density in stellarators only has a *soft* limit set by radiative losses from impurities (Helander et al. 2012), also known as the Sudo-limit (Sudo et al. 1990), and can therefore be increased through experimental scenario development (Fuchert et al. 2020).

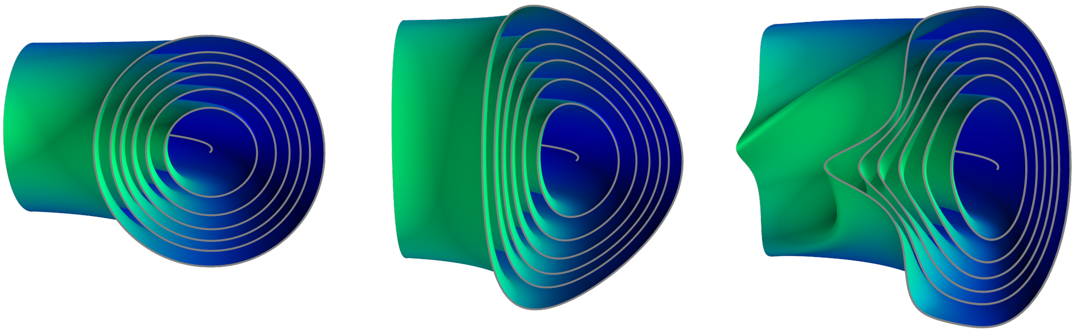

Figure 1.6.1. – Different degrees of plasma shaping shown as one quarter of the full torus of nested flux surfaces. **Left:** circularly shaped axisymmetric equilibrium. **Center:** axisymmetric equilibrium with elongation and (positive) triangularity. **Right:** non-axisymmetric equilibrium with axisymmetric and quasi-axisymmetric shaping. The field strength $|B|$ is overlaid on the flux surfaces with brighter colours corresponding to stronger field.

Since the successful use of RMPs for ELM suppression, numerous studies have explored how the currents of the saddle coils which generate the RMPs can be optimised to improve other plasma properties as well (Park et al. 2021). However, using the currents in those very simple coils as the degrees of freedom to optimise plasma performance is overly restrictive since their location, orientation, and shape is fixed and non-optimised. Further, the 3D shaping which these coils generate is relatively small and therefore unable to provide any of the key benefits of stellarators.

In this work we focus on optimised perturbations strong enough to provide stellarator benefits, and chosen in shape to preserve quasi-symmetry as much as possible to keep guiding centre trajectories confined. Figure 1.6.1 illustrates the change in the flux surfaces as they go from being circular to featuring axisymmetric shaping, and finally to feature both axisymmetric shaping and quasi-axisymmetric (QA) perturbations[17]. This figure already hints at a feature which we will elucidate further later on in section 3.8, namely QA perturbations tend to accumulate on the inboard side, making the use of RMPs for this purpose even less compelling. RMPs are usually located on the outboard side of tokamaks.

---

[17] The concept of quasi-symmetry will be introduced in detail in chapter 3.

1.7 Magnetohydrodynamic equilibrium

In section 1.1.2 we introduced the fact that the plasma must satisfy the ideal MHD equilibrium condition

$$\mathbf{J} \times \mathbf{B} = \nabla p \tag{1.7.1}$$

throughout the plasma volume $\mathcal{V}$. There are different methods of specifying and calculating such an equilibrium. Some are more convenient to handle analytically, thus allowing one to study isolated effects, others are closer to reality but must be performed numerically. In this work we make use of two methods. The first satisfies equation 1.7.1 locally around one flux surface of tokamaks, is analytic, and is often used to study axisymmetric shaping effects. Since tokamaks automatically confine all particles, calculating their global equilibrium is less interesting. The second method we introduce calculates stellarator equilibria globally and ensures that the 3D tokamak perturbations which we will study are indeed the most desirable ones we could consider.

1.7.1 Local tokamak equilibria with Miller formalism

When analysing changes to turbulence-flow interactions in tokamaks due to changes in the equilibrium we use the Miller formalism (Miller et al. 1998). This is the most commonly found approach in the literature, for the following reason. While equation 1.7.1 is only satisfied radially locally (i.e. around the flux surface of interest) within this framework, it comes with the great benefit that it has analytical expressions and that the parameters which describe the equilibrium can be varied individually, allowing all other parameters to stay fixed. With this feature the effect of individual equilibrium parameters can be studied in isolation, a feature which we take advantage of in this work. Moreover, since our simulations are local, we do not require equilibrium information on other flux surfaces. For the derivation we refer to the original paper (Miller et al. 1998).

In table 1.2 we list the nine independent equilibrium parameters which can be changed and what they correspond to physically. Importantly for this work, the plasma shape of the central flux surface is given by the following analytic expressions within the Miller formalism:

$$R(\theta) = R_0 + r\cos\left[\theta + \delta\ \sin\theta\right], \tag{1.7.2}$$

$$Z(\theta) = \kappa r \sin\theta. \tag{1.7.3}$$

Here κ is the elongation and δ the triangularity, see table 1.2 and figure 1.7.1.

Name	Symbol	Explanation
Aspect ratio	A	Ratio of major radius R_0 and minor radius a.
Elongation	κ	Vertical stretching of flux surface, $\kappa = 1$ for circular.
Radial variation of elongation	$\partial_r \kappa$	Change of κ in the neighbourhood of central flux surface.
Triangularity	δ	D-shaping of central flux surface, $\delta = 0$ for circular.
Radial variation of triangularity	$\partial_r \delta$	Change of δ in the neighbourhood of central flux surface.
Safety factor	q	Net twistedness of magnetic field.
Global magnetic shear	$\hat{s}$	Variation of q from one flux surface to another.
Radial variation of major radius	$\partial_r R_0$	Related to Shafranov shift.
Re-scaled pressure gradient	α_{MHD}	Related to pressure gradient of the equilibrium.

Table 1.2. – The nine independent parameters which describe a Miller equilibrium.

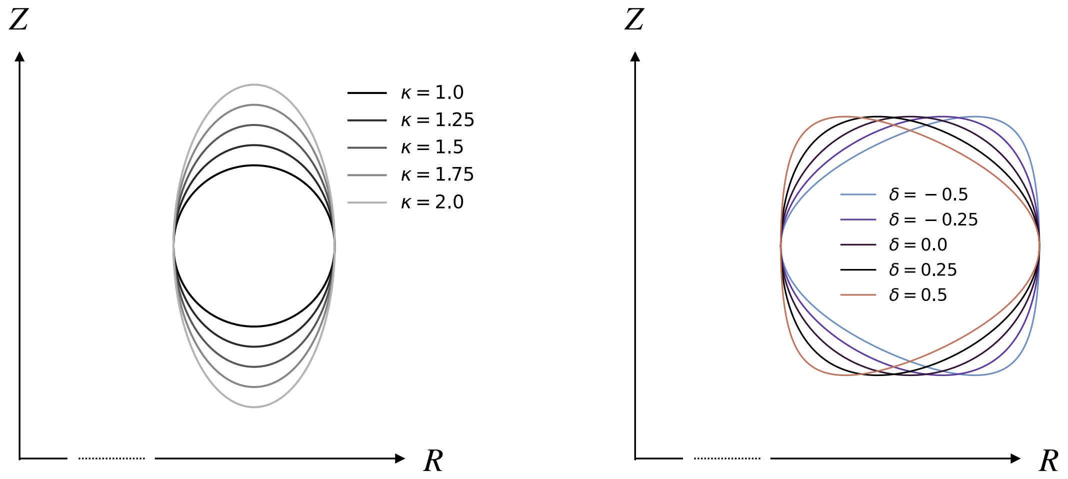

Figure 1.7.1. – Definitions of elongation κ (left) and triangularity δ (right).

1.7.2 Global three-dimensional equilibria with variational method

For the study of non-axisymmetric equilibria we compute global equilibria numerically with the Variational Moment Equilibrium Code (`VMEC`). No analytical local model for general 3D equilibria exists, making this the common approach in the literature. `VMEC` makes use of the result by Kruskal and Kulsrud (Kruskal

& Kulsrud 1958) which states that any 3D MHD equilibrium is determined by the shape of the plasma boundary $\delta\mathcal{V}$ and two radial functions. In `VMEC` these are the pressure profile $p(s)$, and either the toroidal current density profile $dI(s)/d(s)$ or the rotational transform profile $\iota(s)$. The field $\mathbf{B}$ inside the volume and therefore the shape of the flux surfaces are not known and are the output of the calculation. The code uses a gradient-descent method to minimise the MHD energy functional

$$W[\mathbf{B}, p] = \int_{\mathcal{V}} \left(\frac{|\mathbf{B}(\mathbf{r})|^2}{2\mu_0} - p(\mathbf{r}) \right) d^3r. \tag{1.7.4}$$

Finding stationary points of W corresponds to satisfying equation 1.7.1. $W = 0$ indicates that the plasma pressure is exactly balanced by the magnetic pressure everywhere inside the plasma. Other constraints are enforced such as enforcing nested flux surfaces, and requiring the magnetic field to be both divergence free, $\nabla \cdot \mathbf{B} = 0$, and tangential to the plasma boundary everywhere, i.e. $\mathbf{B} \cdot \mathbf{n} = 0$ on $\delta\mathcal{V}$ (Imbert-Gérard, Paul & Wright 2024).

2

Nonlinear energy transfer in tokamaks

In this chapter we will explore the poloidal distribution of zonal flow drive in tokamaks with strong shaping. We will use turbulence modelled with gyrokinetics in a flux tube to obtain a theoretical estimate for the drive strength at each poloidal location. This approach allows us to show that there is a clear distinction between the poloidal distribution of turbulent fluctuations and the nonlinear drive of zonal flows.

This chapter is largely based on the content in reference

> T M Schuett, I Cziegler, and D Dickinson (Nov. 2025a). The poloidal distribution of electrostatic zonal flow drive in strongly shaped tokamaks. In: *Plasma Physics and Controlled Fusion* 67.11, p. 115022

2.1 Motivation

Tokamaks exhibit confinement *phases or modes*, where turbulent fluctuations and thus confinement exhibits step-like differences across the *phase transitions*. The most prominent is the high confinement mode, discovered in 1982 on ASDEX in Garching (Wagner et al. 1982). The conditions under which a confinement transition occurs in a tokamak plasma is affected by many different physical processes and no clear theoretical picture exists to date. Scaling laws (Takizuka 2004; Martin et al. 2008) capture some trends but miss many crucial effects (Gohil et al. 2008; Righi et al. 1999; Howlett et al. 2023).

One dominant theory describes the nonlinear interaction between turbulence and zonal flows to set off a feedback loop and trigger such a transition. Indeed, the nonlinear drive of zonal flows has been experimentally linked to the transitions between these confinement regimes (Moyer et al. 2001; Diamond et al. 2005; Fujisawa 2008). Numerous studies have measured energy transfer in detail and showed that the nonlinear energy transfer to zonal flows increases in the period leading up to the transition (Manz et al. 2012; Cziegler et al.

2013; Yan et al. 2014; Cziegler et al. 2015; Schmitz 2017). Even a previously unexplained asymmetry in the heating power required to reach H-mode, depending on the direction of the grad-B drift, can be explained by a corresponding asymmetry in the nonlinear energy transfer rate to zonal flows (Cziegler et al. 2017); see reprinted in figure 2.1.1 for convenience. Therefore, to further our physics understanding of confinement transitions we must start by advancing our understanding of nonlinear turbulence-flow interaction. It should be noted that even in the absence of confinement transitions, an increased drive of zonal flows improves confinement and motivates understanding of its physical processes by itself.

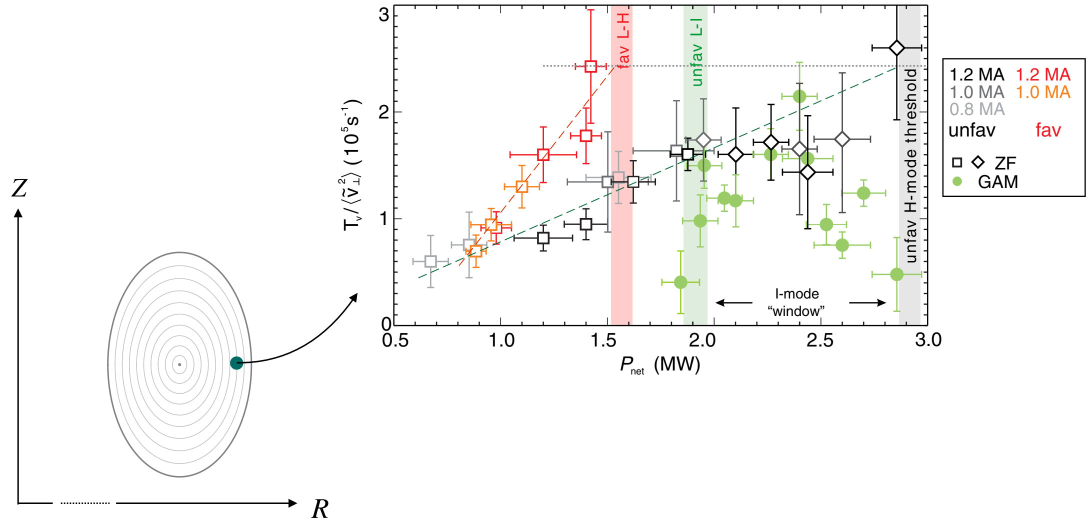

Figure 2.1.1. – The nonlinear drive of (mean) zonal flows reaches the *same* critical values at the L-H transition. Note that these measurements are taken at a single poloidal location, the outboard midplane. Image adapted from (Cziegler et al. 2017).

In the experimental studies cited above the turbulence diagnostics are typically either beam emission spectroscopy (BES), gas puff imaging (GPI), or probes. These cover the perpendicular plane of a flux tube at only *one poloidal location* on most tokamaks. This location is usually the outboard midplane where the drift curvature is the 'worst' and the turbulence intensity is therefore the strongest. This location is marked by a green dot in figure 2.1.1.

While the turbulent potential fluctuations typically balloon out on the outboard midplane, the zonal potential is poloidally symmetric[1]. One might therefore wonder about their interaction poloidally. *How strong is the drive of ZFs at poloidal locations which are not the outboard midplane? Is the nonlinear driving strength simply proportional to the turbulent intensity?* The lack of answers to these questions has meant in the past that macroscopic plasma parameters

---

[1] Note that the zonal *flow* has an in-out asymmetry due to the variation in the field strength B.

such as shaping and rotational transform profiles have been kept as constant as possible within a given experimental study, in which case differences in local drive measurements *are* still expected to be indicative of changes to the total ZF drive. With a theoretical estimate on whether the ZF drive distribution varies with macroscopic plasma parameters, and if so how, one could meaningfully compare experimental results with different macroscopic parameters.

The most natural tool to answer this question is local gyrokinetics, introduced in section 1.5. Local gyrokinetic codes solve the gyrokinetic equation along a flux tube for a full poloidal turn. In tokamaks, due to axisymmetry, this gives information on how the turbulence behaves on the full flux surface. The perpendicular plane resembles those covered by the turbulence diagnostics, i.e. it covers a few correlation lengths in both the radial and binormal directions. The crucial point, however, is that in gyrokinetics this plane is covered at *every* parallel point along the flux tube. By performing a nonlinear drive analysis at each of these points and projecting the results onto the poloidal plane we can obtain an estimate for the poloidal distribution of the ZF drive. Figure 2.1.2 illustrates how the poloidal projection of the dynamics parallel to the flux tube give information on the poloidal distribution.

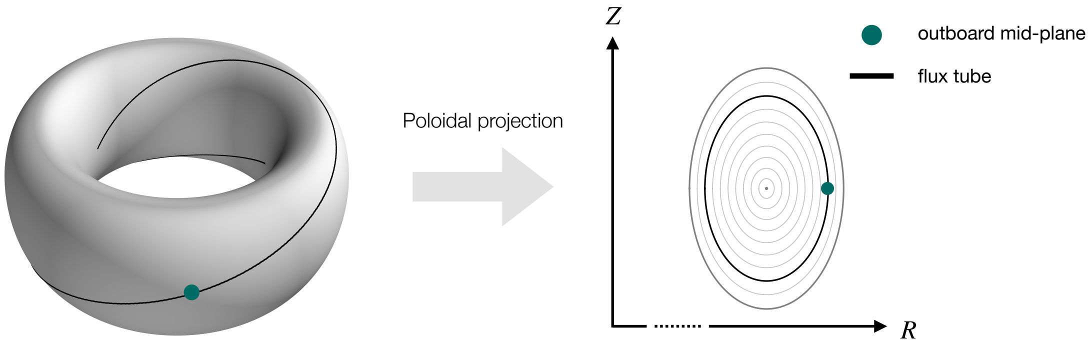

Figure 2.1.2. – Dynamics along a field line give direct information on the poloidal distribution due to axisymmetry. **Left:** a flux tube with $q = 1.4$ on a flux surface. **Right:** the poloidal projection, showing that the field line covers one poloidal turn. The green dot shows the outboard midplane in all three plots, the poloidal location which has been mostly used for turbulence-flow interaction studies in the past.

Local gyrokinetics has been previously used to validate nonlinear transfer dynamics, and it was shown that the forward cascade of internal energy transfer mediated by ZFs agrees qualitatively at the outboard midplane (Holland et al. 2007). Nonlinear transfer functions from gyrokinetics (referred to as free energy transfer or entropy transfer in the literature, more on this topic in chapter 6) as well as transfer of temperature fluctuations have been used to highlight the importance of multi-scale coupling in gyrokinetic simulations to accurately match experimental heat flux measurements (Maeyama et al. 2015; Maeyama et al. 2017; Howard et al. 2015). This study is therefore not the first time gyrokinetics has been used to evaluate nonlinear energy transfer. However, it is the first time, along with the previous work for circular shaping by (Biggs-Fox 2022), that the full poloidal coverage is taken advantage of, and that the analysis is specifically aimed at the drive of zonal *flows*, since it is the flow which shears apart turbulent eddies and appears to trigger confinement transitions in experiments.

2.2 Cyclone base case

The cyclone base case (CBC) (Miller et al. 1998) is the most used tokamak benchmark equilibrium. It is based on an equilibrium from the DIII-D tokamak at General Atomics, features a circularly shaped plasma cross-section, an aspect ratio $A = 2.72$, safety factor $q = 1.4$ and global magnetic shear $\hat{s} = 0.78$. Its gradients are $a/L_T = 2.49$ and $a/L_n = 0.81$.

The poloidal distribution of ZF drive $\mathcal{T}_{ZF}$ was investigated for the CBC by Biggs-Fox (Biggs-Fox 2022). It was found that while the ZF drive is not proportional it is correlated to the envelope of the turbulent fluctuations. A slight dependency of $\mathcal{T}_{ZF}$ on the temperature gradient was observed for fixed density gradient.

In current and future tokamaks, plasmas typically have a non-circular cross section, and in this chapter we investigate the dependence of $\mathcal{T}_{ZF}$ on non-circular plasma shaping. We retain the CBC as the reference case, i.e. we do not change A, q, or $\hat{s}$ unless otherwise noted. We also retain the assumption of adiabatic electrons (see equation 1.5.10).

2.3 A note on computational methods

The calculation of the sum over triple products which is at the heart of the calculation of the nonlinear transfer functions is computationally challenging from a memory perspective when done as a post-processing script in `Python`. Such a calculation requires the handling of the full complex-valued $\delta\hat{\varphi}(t, k_x, k_y, \theta)$ to calculate equation 1.4.19, stated here again for convenience:

$$\mathcal{T}_{ZF} = 2B^{-3}\,\mathrm{Re}\left\langle \sum_{k_x, k_{1x}, k_{1y}} c_{\mathbf{k},\mathbf{k_1}} (\delta\hat{\varphi}_{\mathbf{k}} \delta\hat{\varphi}_{\mathbf{k}-\mathbf{k_1}} \delta\hat{\varphi}_{\mathbf{k_1}}) \right\rangle. \tag{2.3.1}$$

Typical values in gyrokinetic simulations are $N_x = 128$, $N_y = 96$, and $N_\theta = 64$. After de-aliasing with the $2/3$-rule the wavenumber grid resolutions are $N_{k_x} = 85$ and $N_{k_y} = 32$ (the binormal wavenumber gets halved again due to the previously mentioned reality condition). A typical simulation time is $1000\,a/v_{th}$ with the first $\sim 250\,a/v_{th}$ being dominated by linear physics, not saturated, and thus excluded from our analysis. Typical time steps by the solver are $\Delta t = 0.01\,a/v_{th}$, controlled by the Courant–Friedrichs–Lewy condition. Thus a total of 75k time steps are taken in the code. Diagnostics are typically written out every $\texttt{nwrite} = 20$ time steps or similar, such that $3,750$ snap shots of $\hat{\varphi}$ are typically written out for one simulation. For typical double precision of real and imaginary part the total size of $\hat{\varphi}(t, k_x, k_y, \theta)$ is

$$3750 \,*\, 85 \,*\, 32 \,*\, 64 \,*\, 2 \,*\, 8\,\text{bytes} \sim 10\,\text{GB} \tag{2.3.2}$$

To make use of optimised tensor multiplication libraries in `Python` such as `Numpy`, the mediator mode tensor in equation 2.3.1 is typically expanded into a tensor of dimensions $[N_t, N_{k_x}, N_{k_{1x}}, N_{k_{1y}}, N_\theta]$, making the memory requirements of the post-processing script quickly approach $\mathcal{O}(100\,\text{GB})$. Simulations with finer grid resolution, such as those needed to check convergence, quickly make the memory requirements prohibitively large. Further, writing the full $\delta\hat{\varphi}$ to disk during the simulation can take up significant time, increasing the total wall clock time a simulation takes significantly. Should one be interested in the full transfer tensor of the form $\mathcal{T} = \mathcal{T}(k_x, k_y, k_{1x}, k_{1y}, \theta)$ to analyse coupling dynamics beyond those that feed the ZF, memory requirements quickly exceed those available on one compute node, requiring MPI-based post-processing tools.

For all those reasons it is desirable to calculate the transfer tensor during the simulation as part of the standard diagnostics routine of the GK code. Parallelisation can be fully leveraged, no large files need to be written to disk, speed-ups from the more efficient programming language used for these codes (Fortran in the case of `GS2`) come for free, and diagnostic capabilities can be expanded relatively straightforwardly without hitting memory limits.

The computation of $\mathcal{T}_v$ was implemented as a python post-processing tool for `GS2` as part of prior work by Biggs-Fox (Biggs-Fox 2022) and has been implemented as a diagnostic to `GS2` as part of this Thesis, allowing for memory efficient MPI parallelised computation at runtime. The diagnostic has been first made available in `GS2` version 8.2.0 (Barnes et al. 2024).

2.4 Convergence

2.4.1 Routine convergence tests for grid resolutions

For each shaping scan convergence was checked for the most extreme cases, e.g. for triangularity we checked convergence for $\delta = \pm 0.75$. It was ensured that simulations are sufficiently resolved such that both the heat flux and the fluctuations spectra of the electrostatic potential do not change meaningfully when the resolution on one of the grids is increased by at least 50%. We also checked that the poloidal distribution of the ZF drive does not change meaningfully. Upon finding converged resolutions at the extrema of a given shaping scan, the higher resolution of either extremum are selected for all cases in between which are then expected to be converged also.

2.4.2 Nonlinear coupling convergence

Beyond the checks that our observations of the previous sections are not sensitive to the grid resolutions used, we also check that the simulation time is sufficient for convergence of the transfer functions.

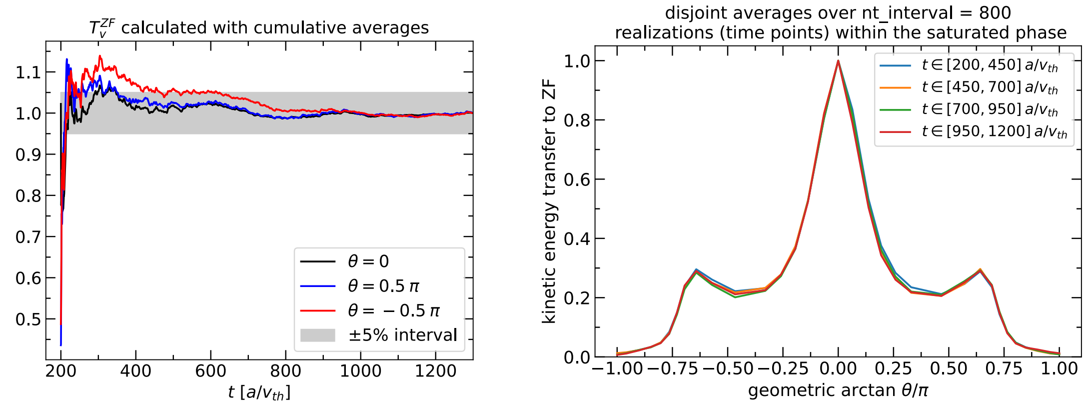

Figure 2.4.1. – Convergence tests for the nonlinear coupling behind the ZF drive. The case with $A = 2.72$, $\kappa = 1.0$, $\delta = +0.5$, and $a/L_T = 6.0$ (which will be discussed in subsection 2.5.2) is used as an example here. **Left:** Cumulative averages within the saturated phase to test convergence with respect to the ensemble size. Each data is normalised to its final value as the relative differences between these locations were established in the previous sections. **Right:** Poloidal distribution of the ZF drive stemming from averages over disjoint intervals within the saturated regime.

This is shown in the left plot of figure 2.4.1 which shows cumulative averages for three somewhat arbitrarily selected poloidal locations. We see that after $t = 600\ a/v_{th}$, the changes are $< 5\%$. Note that the time stepping did not change during this phase such that spacing in t is representative of spacing in realisations.

Moreover, the right plot of figure 2.4.1 shows that the poloidal distributions of the ZF drive are stationary and not dominated by a large-scale event. To this end we calculated the ZF drive over four disjoint intervals of length $\Delta t = 250\, a/v_{th}$. Almost no variation is seen in the distributions, showing that they are not dominated by singular events.

2.5 Strongly shaped tokamaks

2.5.1 Elongation

To motivate the parameter range which we choose for the elongation study we first map out the dependence of the linear ITG growth rate on both the temperature gradient and elongation. The results in the left plot of figure 2.5.1 show the known trend, namely the ITG mode is stabilised by both a lower temperature gradient and higher elongation, the latter occurring when ∇T is specified at the outboard midplane (Angelino et al. 2009), as customary in tokamak gyrokinetics.

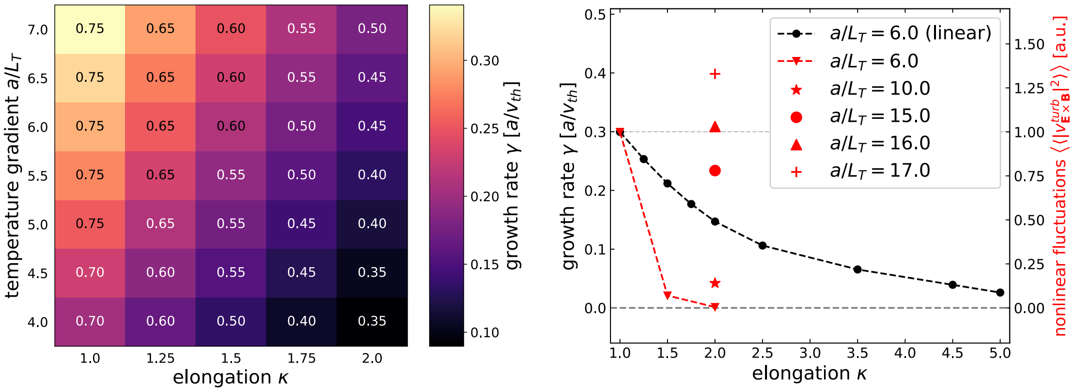

Figure 2.5.1. – Trends of ITG turbulence with elongation and temperature gradient when specified at the outboard midplane. **Left:** The linear growth rate of the most unstable ITG mode for a range of temperature gradients a/L_T and values of elongation κ. The binormal wavenumber of the mode is overlaid. **Right:** The slice of $a/L_T = 6.0$ with the nonlinear turbulent activity added.

These results motivate the choice $a/L_T = 6.0$ for our nonlinear study. This slice is shown again in the right plot in figure 2.5.1, where the turbulent activity from the nonlinear simulations are added. Here we chose a three-point-scan with $\kappa \in \{1.0, 1.5, 2.0\}$. We define the turbulent activity to be the fluctuations of the $\mathbf{E} \times \mathbf{B}$ velocity, defined as

$$\left\langle\left\langle |\delta v^{turb}_{\mathbf{E}\times\mathbf{B}}|^2 \right\rangle\right\rangle = \frac{1}{B^2} \sum_{k_x, k_y \neq 0} (k_x^2 + k_y^2) \left\langle\left\langle |\delta\phi(\mathbf{k})|^2 \right\rangle\right\rangle . \tag{2.5.1}$$

Here $\langle\langle . \rangle\rangle$ denotes temporal averaging over the saturated period of the simulations as well as poloidal averaging. Previous results by Biggs-Fox (Biggs-Fox 2022) showed hints that the envelope of the ZF drive might have a dependency on the temperature gradient itself, i.e. on the degree of non-marginality of the simulations. We thus attempt to find the temperature gradient at increased elongation ($\kappa = 2.0$) for which the fluctuation level approximately matches those of the CBC ($\kappa = 1.0$). We find this gradient to be $a/L_T = 16$. While this is a very large value compared to typical experimental value, it serves as a useful tool to isolate the effect of elongation itself, and to prevent the mixing of two separate effects from affecting our study.

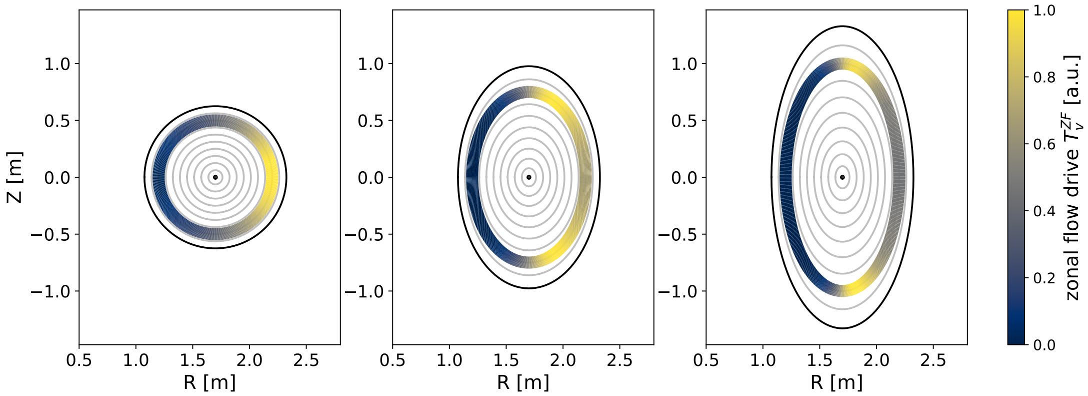

Figure 2.5.2. – The poloidal distribution of zonal flow drive for varying elongation κ, mapped onto the simulated flux surface $\rho = 0.8$. The temperature gradient is fixed to $a/L_T = 6.0$. Note that the radial thickness of the coloured region is shown for illustrative purposed only. **Left:** CBC, **Middle:** $\kappa = 1.5$, **Right:** $\kappa = 2.0$.

The poloidal distribution of the ZF drive T_v^{ZF} are shown for varying elongation in figure 2.5.2. As the flux surfaces depart from being circular the single maximum of the drive at outboard midplane bifurcates, and two maxima develop which are aligned with the locations of maximum poloidal curvature near the top and bottom of the equilibrium. Here, and in all future plots, we normalise the zonal flow drive to its maximum value, such that we gain information on its relative distribution. We often observe the zonal flow level to be only quasi-stationary throughout the simulation in these strongly shaped equilibria, making it difficult to compare absolute values between different equilibria. We note, however, that with respect to experimental validation this ability could be restored through a single localised measurement at e.g. the outboard midplane.

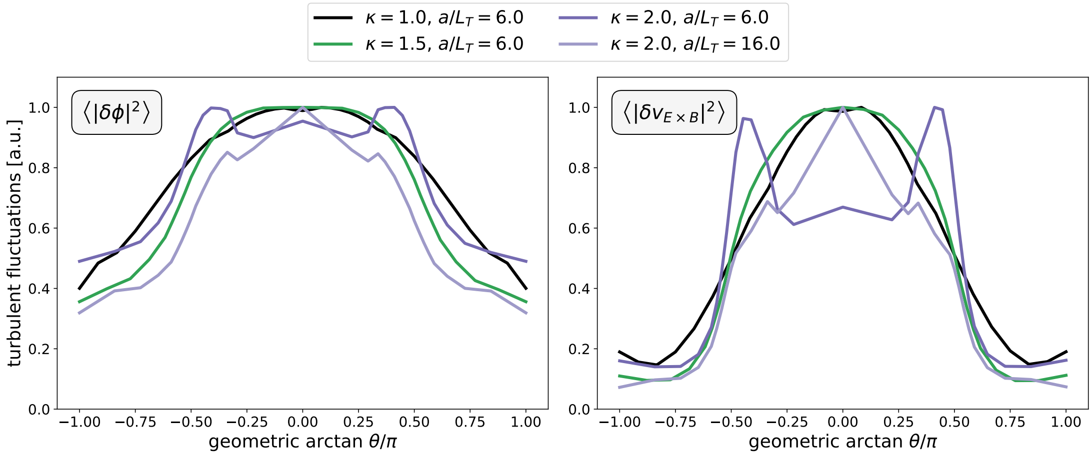

Figure 2.5.3. – The poloidal distributions of the turbulent activity, time-averaged over the saturated phase of the simulations. **Left:** Electrostatic potential. **Right:** $\mathbf{E}\times\mathbf{B}$ velocity.

The poloidal distributions of the turbulent activity are shown in figure 2.5.3. In addition to the total fluctuations of the $\mathbf{E}\times\mathbf{B}$ velocity, defined in equation 2.5.1, we also show the poloidal distribution of electrostatic potential fluctuations. Since the difference in total fluctuation levels was shown in figure 2.5.1, we focus on the normalised poloidal envelopes here to bring out their structure.

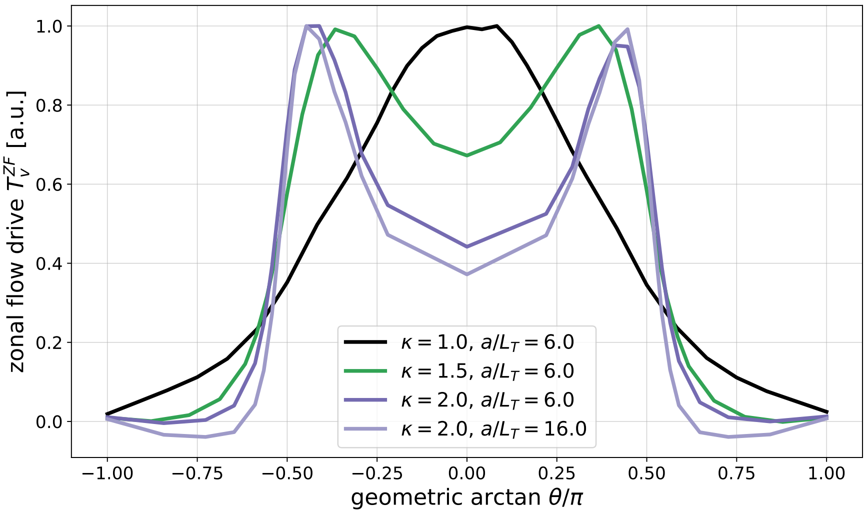

Figure 2.5.4. – The poloidal distribution of the ZF drive for varying elongation.

For the strongly elongated equilibrium, results are shown both for $a/L_T = 6.0$ and for the temperature gradient of matched fluctuation levels with the CBC, $a/L_T = 16$. While the poloidal distributions of the potential fluctuations remain broadly peaked across the outboard side, those of the $\mathbf{E}\times\mathbf{B}$ velocity develop local maxima for the marginal case ($\kappa = 2.0$, $a/L_T = 6.0$). However, this effect disappears for increased temperature gradient ($\kappa = 2.0$, $a/L_T = 16$), for which the envelopes of both fluctuation metrics are strongly peaked at the outboard midplane.

However, figure 2.5.4 shows that the distribution of T_v^{ZF} retains a similar character for both temperature gradients. The transfer is strongly reduced on the outboard midplane. We note in passing that the $\mathbf{E} \times \mathbf{B}$ fluctuation envelope shows local maxima on the inboard midplane for the cases with $a/L_T = 6.0$, while no such peaking is seen in the ZF drive envelope. This could be explained by the extra factor of B^{-1} that features in $\mathcal{T}_{ZF}$ compared to $\left\langle |\delta v^{turb}_{\mathbf{E}\times\mathbf{B}}|^2 \right\rangle$.

2.5.2 Triangularity

The plasma shaping on most tokamaks today features not only elongation but also triangularity. While positive triangularity (PT) has been traditionally used due to its beneficial effects for MHD stability (Freidberg 2014), more recently negative triangularity (NT) has been increasingly studied due to its high performing L-mode characteristics which avoids the ELM-suppression problem of H-mode or H-mode-like regimes (Marinoni, Sauter & Coda 2021). Therefore, both PT and NT are included here.

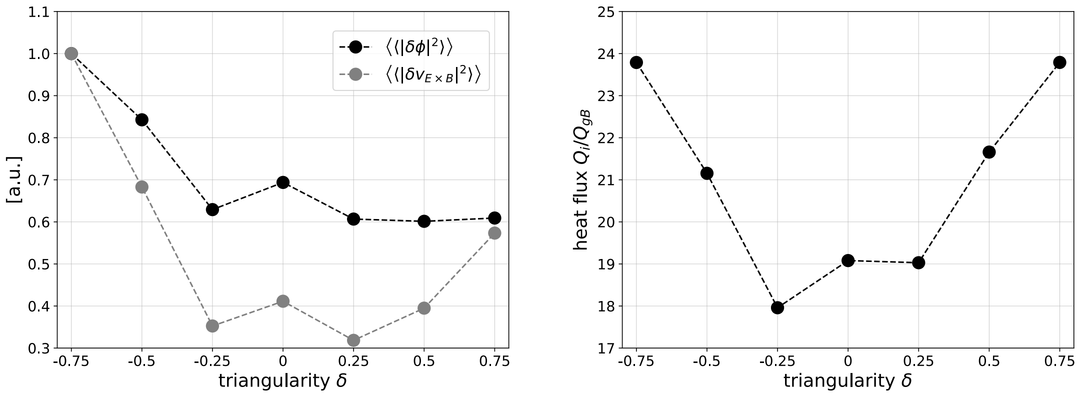

Figure 2.5.5. – fluctuation metrics (left) and time-averaged heat flux (right) against triangularity for no elongation ($\kappa = 1$).

In figure 2.5.5 we show that the fluctuation levels are relatively insensitive to triangularity alone, compared to the dependency on elongation. The results for elongation of the previous section (figure 2.5.4) showed that the effects from the proximity to marginality appear to be subdominant to shaping effects on the ZF drive envelope. Motivated by these two facts we keep the temperature gradient constant at $a/L_T = 6.0$ throughout this part of the study.

We vary the triangularity in the range $\delta \in [-0.75, 0.75]$ and set its radial derivative to $\partial_r \delta = \delta/\sqrt{1-\delta^2}$ (Miller et al. 1998). Figure 2.5.6 shows that the

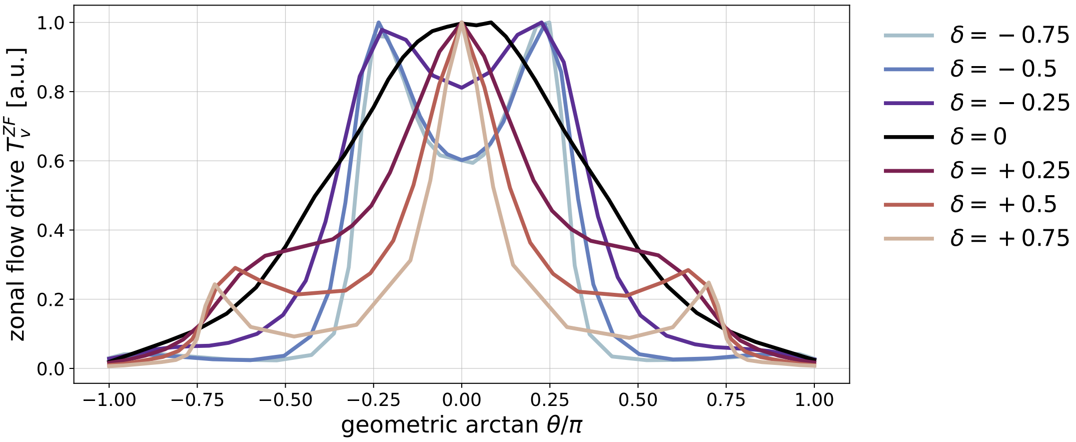

Figure 2.5.6. – The poloidal distribution of the ZF drive for varying values of triangularity.

trend observed for elongation continues, namely as the flux surface departs from being circular, local maxima of the ZF drive emerge near locations of strong poloidal curvature. We note that the fall-off near the maxima is more pronounced for higher absolute values of δ.

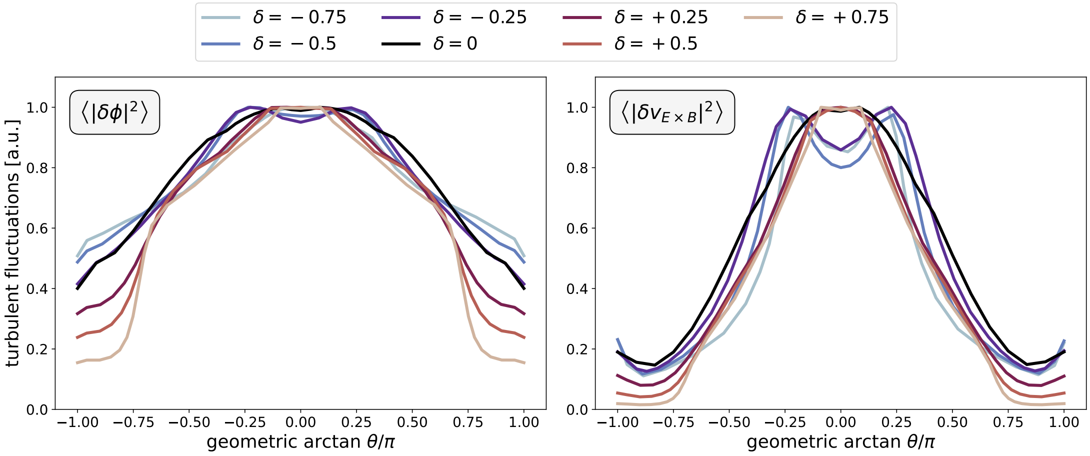

Figure 2.5.7. – The poloidal distribution of fluctuations for varying values of triangularity. **Left:** Electrostatic potential fluctuations. **Right:** $\mathbf{E} \times \mathbf{B}$ velocity fluctuations.

The distributions of the fluctuation metric are shown in figure 2.5.7. Unlike the ZF drive, these are broadly peaked on the outboard side for all levels of triangularity.

For easier visual interpretation figure 2.5.8 shows the ZF drive again, but mapped onto the cross-section for $\delta = \pm 0.5$.

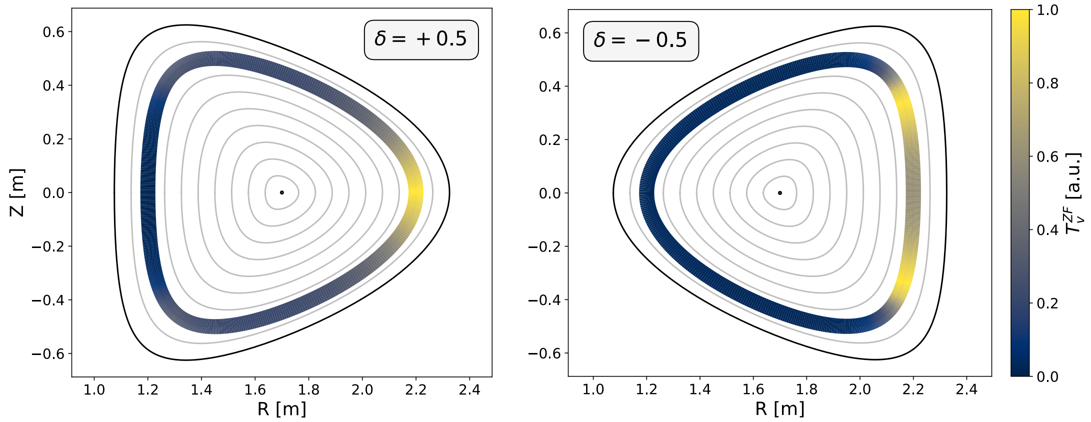

Figure 2.5.8. – ZF drive mapped onto relevant flux surface $\rho = 0.8$. **Left:** Positive triangularity $\delta = +0.5$, **Right:** Negative triangularity $\delta = -0.5$.

2.5.3 Aspect ratio

Aspect ratio alone shows to have very little effect on the ZF drive envelope for circularly shaped cross-sections. Thus, the results are not shown here. The general shape remains the same as for the CBC, but a slight widening is observed towards small aspect ratio when comparing the results for the CBC at $A = 2.72$ to the ones at spherical tokamak aspect ratio of $A = 1.4$. This remains the case even when the stabilising effect of aspect ratio is taken out by selecting a higher temperature gradient to match the fluctuation levels. The poloidal distribution of the fluctuation metrics also show to be insensitive to aspect ratio.

2.5.4 Spherical tokamak equilibrium

After having considered the isolated effects of elongation, triangularity and aspect ratio in the previous subsections, we turn to the geometry of a typical spherical tokamak equilibrium which features all of these effects combined, as seen in the photo 2.5.9 of a MAST plasma. Typical values for MAST-Upgrade equilibria are $\kappa = 2.0$ and $\delta = +0.5$ at an aspect ratio of $A = 1.4$. All three of these effects combined are very stabilising, so we increase the temperature gradient to a very high level of $a/L_T = 18.0$ to match the fluctuation levels of the CBC.

Previous linear gyrokinetic studies of spherical tokamaks with negative triangularity have shown that detrimental transport is to be expected from electromagnetic turbulence excited by kinetic ballooning modes (Davies, Dickinson & Wilson 2022). Nonetheless this parameter regime is interesting for validation purposes and therefore we also include an equilibrium with flipped triangularity of $\delta = -0.5$. That such equilibria are possible in STs and produce ELM-free

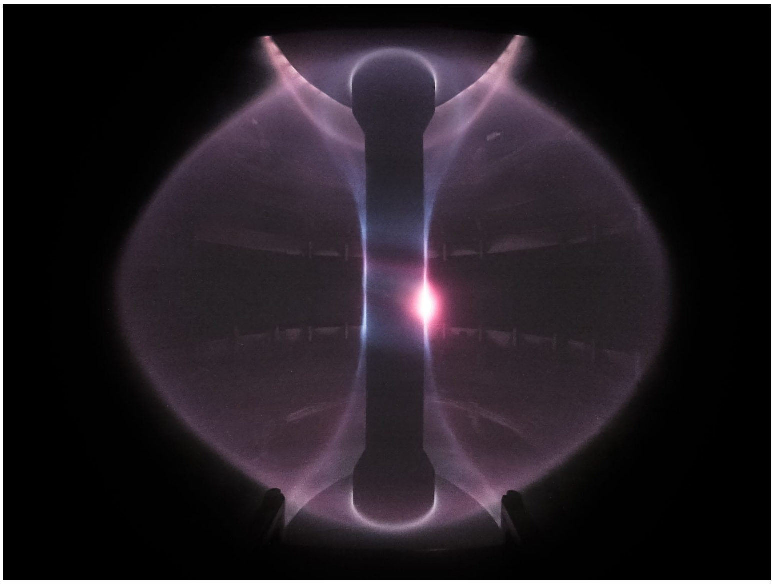

Figure 2.5.9. – Plasma in the MAST spherical tokamak in Culham, UK. Image from UKAEA.

regimes as on TCV has been recently demonstrated on MAST-Upgrade (Nelson et al. 2024).

Equilibria in MAST-Upgrade also feature significant squareness. However, including squareness requires a reformulation of the Miller equilibria (Joiner & Dorland 2010). We therefore do not include it here for simplicity, but note that it is one of the possible extensions of this study.

Figure 2.5.10 shows the poloidal distributions of the ZF drive and the fluctuation metrics, for positive and negative triangularity, respectively. It appears that the effect of elongation dominates over the one from triangularity for the ZF drive, as now a local minimum is observed at the outboard midplane for PT, despite the local maximum of poloidal curvature at this location. *Is this caused only by the combination of elongation and triangularity, or does the compact aspect ratio play a critical role as well?* Figure 2.5.11 shows the result for combined elongation and PT at conventional aspect ratio $A = 2.72$ (same as CBC). While the ZF drive at the outboard midplane is now raised compared to the compact equilibrium there is no local maximum at this location. We thus conclude that the effect of elongation is stronger than the one of triangularity. This is the case independent of aspect ratio, consistent with the results from the previous subsection 2.5.3 that the ZF drive is relatively insensitive to aspect ratio.

The question still remains as to what physical mechanism is responsible for these results, namely the non-trivial poloidal distribution of the ZF drive, and we will return to this in subsection 2.7.

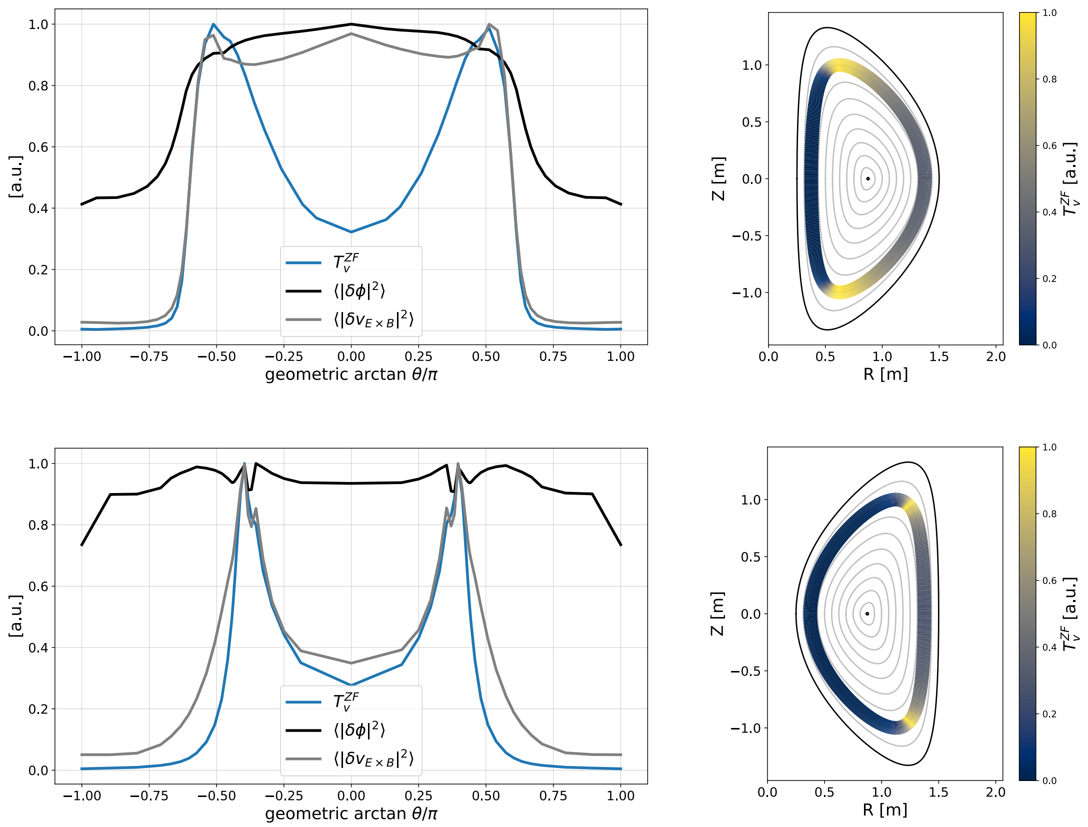

Figure 2.5.10. – Zonal flow drive distribution in spherical tokamak with positive and negative triangularity. **Top:** Spherical tokamak equilibrium with positive triangularity. **Bottom:** Negative triangularity counterpart. **Left:** The ZF drive envelope alongside those of the fluctuation metrics. **Right:** ZF drive mapped onto the cross-sections. Note that the left axis also corresponds to the symmetry axis.

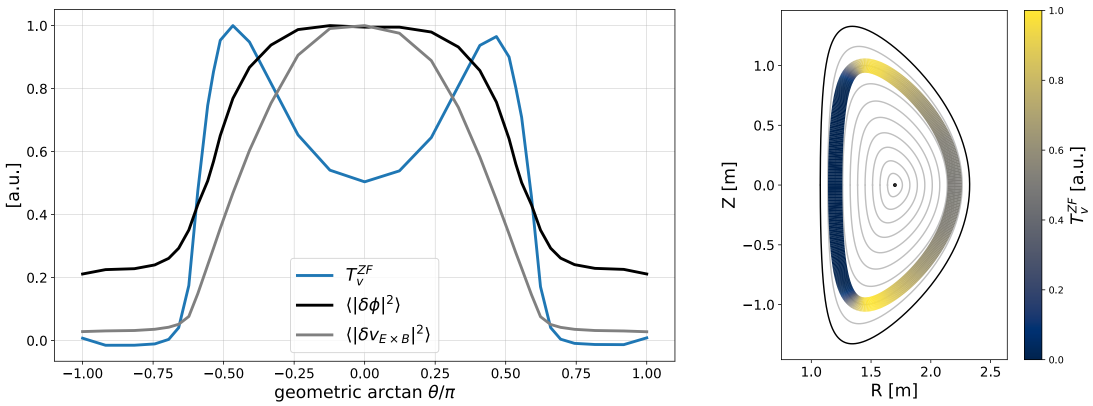

Figure 2.5.11. – Elongation and PT at conventional aspect ratio. **Left:** ZF drive and fluctuation metrics envelopes. **Right:** ZF drive mapped onto the cross-sections.

2.6 Up-down asymmetric equilibria

2.6.1 Implications of breaking up-down symmetry

The shaping cases considered thus far all retain up-down symmetry which describes the mirror symmetry of the flux surfaces across the midplane. Breaking up-down symmetry has strong implications, especially in the case considered here where the equilibrium toroidal flow and flow shear are assumed to vanish. In this case turbulent fluctuations are unable to transport toroidal angular momentum radially for up-down symmetric equilibria due to a symmetry of the gyrokinetic equation (Parra, Barnes & Peeters 2011). The radial momentum flux from above the midplane cancels with the one from below. Hence, rotation and rotation shear cannot be intrinsically generated from an initial stationary equilibrium. However, strongly up-down asymmetric equilibria, as considered in this section, are able to transport toroidal angular momentum even in the absence of externally driven flow and flow shear (Parra, Barnes & Peeters 2011).

The toroidal angular momentum can be split into a parallel and perpendicular (to the field) component within the flux surface. Zonal flows contribute to the perpendicular component. For zero equilibrium flow, the general expression for the radial flux of toroidal angular momentum is given by (Parra, Barnes & Peeters 2011, eq. 32)

$$\Pi_s^\phi = m_s \left\langle \left\langle R \int d^3v (\mathbf{v}\cdot\hat{\zeta})(\mathbf{v}\cdot\nabla\psi) f_s \right\rangle_\psi \right\rangle_{\Delta\psi}, \tag{2.6.1}$$

where f_s is the full distribution function, ζ is the toroidal angle and $\hat{\zeta}$ is the unit vector in the toroidal direction.

$$\langle \dots \rangle_{\Delta\psi} \equiv \frac{1}{\Delta\psi} \int_{\Delta\psi} d\psi (\dots) \tag{2.6.2}$$

is an average over the radial domain (the simulation box) and

$$\langle \dots \rangle_\psi \equiv \frac{1}{V'} \oint d\theta \oint d\zeta \mathcal{J} (\dots) \tag{2.6.3}$$

is a flux surface average with the gradient of the volume element $V' \equiv dV/d\psi = 2\pi \oint d\theta \mathcal{J}$ and the Jacobian $\mathcal{J} \equiv 1/|(\nabla\psi \times \nabla\theta)\cdot\nabla\zeta|$ (Parra, Barnes & Peeters 2011).

It can be shown (Parra, Barnes & Peeters 2011, eq. 36) that the turbulent contribution to Π_s^ϕ, in the following denoted by the superscript 'tb', reduces to

$$\Pi_s^{\phi,\mathrm{tb}} = m_s \left\langle \left\langle R^2 (\nabla_\perp \delta\phi \cdot \hat{\zeta}) \int d^3(\mathbf{v} \cdot \hat{\zeta}) h_s \right\rangle_\psi \right\rangle_{\Delta\psi} \tag{2.6.4}$$

in the absence of magnetic perturbations and flow shear. Upon using the relation

$$R\hat{\zeta} = \frac{I(\psi)\mathbf{b}}{B} - \frac{\mathbf{b} \times \nabla\psi}{B}, \tag{2.6.5}$$

with $I(\psi) = RB_\zeta$, the expression simplifies to

$$\Pi_s^{\phi,\mathrm{tb}} = m_s \left\langle \left\langle R(\delta v_{\mathbf{E}\times\mathbf{B}} \cdot \nabla\psi) \int d^3(\mathbf{v} \cdot \hat{\zeta}) h_s \right\rangle_\psi \right\rangle_{\Delta\psi}, \tag{2.6.6}$$

where again $\delta v_{\mathbf{E}\times\mathbf{B}} = (\mathbf{b} \times \nabla\delta\phi)/B$ (Parra, Barnes & Peeters 2011). This highlights that the turbulent radial flux of momentum is entirely due to turbulent $\mathbf{E} \times \mathbf{B}$ motion. Its decomposition into perpendicular and parallel components are then given by

$$\Pi_{s,\perp}^{\phi,\mathrm{tb}} = m_s \left\langle \left\langle R(\delta v_{\mathbf{E}\times\mathbf{B}} \cdot \nabla\psi) \int d^3(\mathbf{v}_\perp \cdot \hat{\zeta}) h_s \right\rangle_\psi \right\rangle_{\Delta\psi} \tag{2.6.7}$$

$$\Pi_{s,\parallel}^{\phi,\mathrm{tb}} = m_s \left\langle \left\langle R(\delta v_{\mathbf{E}\times\mathbf{B}} \cdot \nabla\psi) \int d^3(\mathbf{v}_\parallel \cdot \hat{\zeta}) h_s \right\rangle_\psi \right\rangle_{\Delta\psi} \tag{2.6.8}$$

$$\Pi_s^{\phi,\mathrm{tb}} = \Pi_{s,\perp}^{\phi,\mathrm{tb}} + \Pi_{s,\parallel}^{\phi,\mathrm{tb}} \tag{2.6.9}$$

The poloidal distribution of $\Pi_{s,\perp}^{\phi,\mathrm{tb}}$ is given by $\pi_{s,\perp}^{\phi,\mathrm{tb}}(\theta)$ (`es_mom_flux_perp_dist` in `GS2`) such that

$$\Pi_{s,\perp}^{\phi,\mathrm{tb}} = \oint d\theta J_\psi \pi_{s,\perp}^{\phi,\mathrm{tb}}. \tag{2.6.10}$$

2.6.2 Up-down asymmetric equilibria through poloidal tilting

A straightforward method of introducing up-down asymmetry is by tilting the Miller equilibria in the poloidal plane. This allows the equilibria to be expressed with generalised Miller equations which are available in `GS2` (Ball et al. 2014; Ball 2016). For convenience the equations for the resulting flux surface shapes are stated in Appendix A. The two additional equilibrium parameters are θ_κ and θ_δ which describe the poloidal tilt angle of elongation and triangularity, respectively.

We note that up-down asymmetric equilibria which can be described by a single tilt angle, and thus retain a mirror symmetry in the poloidal plane, are a special case of up-down asymmetry since for small scale shaping (i.e. high Fourier mode shaping as opposed to the large scale, small Fourier mode shaping used here) the gyrokinetic equation possesses another symmetry which results in only small changes to the turbulent transport of momentum and energy (Ball, Parra & Barnes 2016).

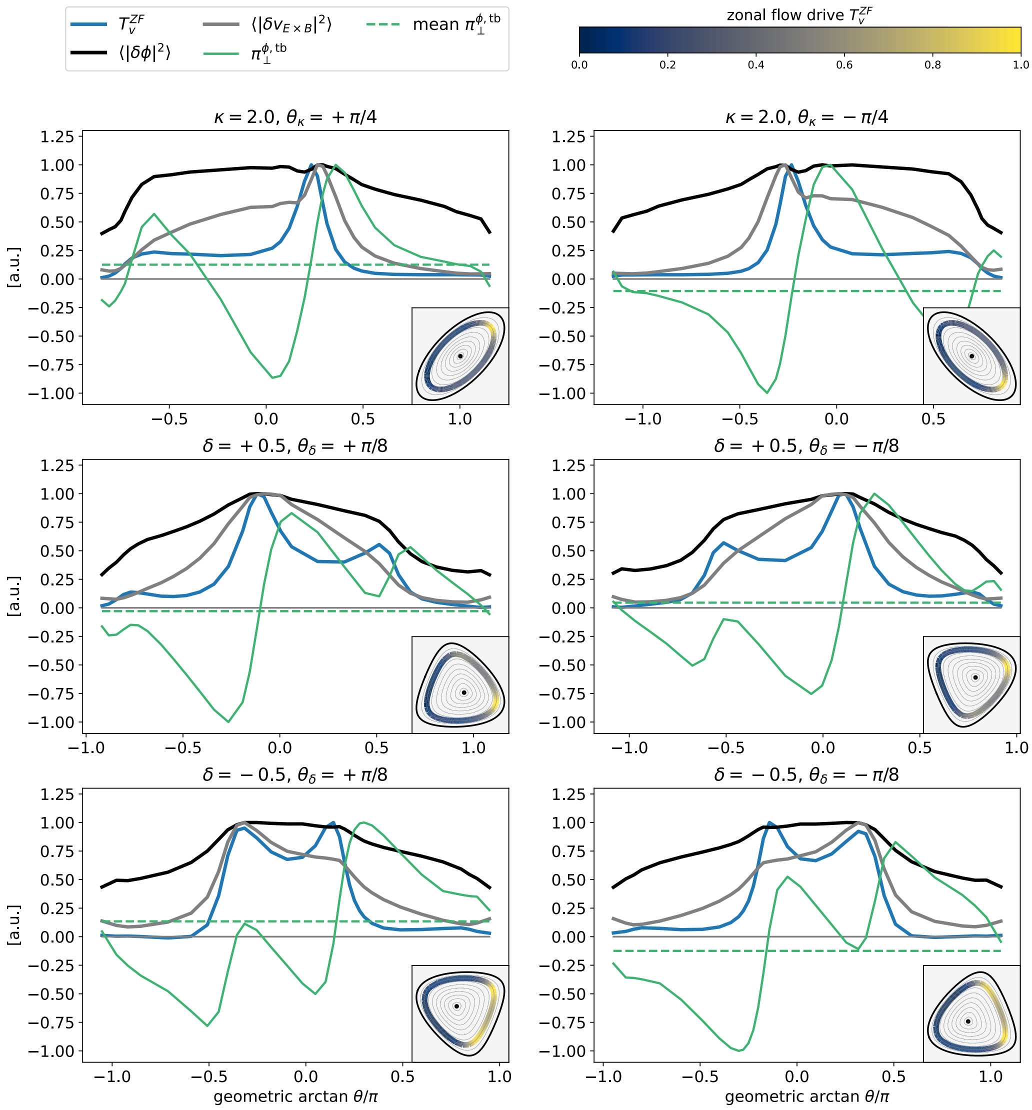

Figure 2.6.1. – Nonlinear metrics for up-down asymmetric equilibria through tilting of Miller equilibria. Shown are the ZF drive, the turbulent fluctuation metrics, and the envelopes of the perpendicular component of the angular momentum transfer. Gradients are chosen to match fluctuation levels of the CBC. Results are shown for positive and negative tilt angles to highlight the expected symmetry. **Top row:** tilted elongation. **Middle row:** tilted positive triangularity. **Bottom row:** tilted negative triangularity. (Left column: counter-clockwise tilting, Right column: clockwise tilting).

In figure 2.6.1 we show the results for tilted elongation with tilt angles $\theta_\kappa = \pm\pi/4$, as well as those for tilted positive and negative triangularity each with tilt angles $\theta_\delta = \pm\pi/8$. We first note that the breaking of up-down symmetry in the flux-surfaces also breaks such a symmetry in the fluctuations and ZF drive envelope. We also note that this symmetry breaking is symmetric with respect to the sign of the tilt angle. This would not be expected to occur with equilibrium flow or flow shear which would introduce additional symmetry breaking.

For elongation the tilting effectively removes one of two local maxima in the ZF drive that were present for pure elongation without tilting. Most of the drive now occurs near the maximum of poloidal curvature which is at a larger major radius, even though the turbulent activity is broadly distributed. A similar effect occurs for titled positive triangularity (see middle row of figure 2.6.1) where the local maximum of poloidal curvature is favoured which has moved towards larger major radius rather than the one which has been moved further inward. For negative triangularity (bottom row of figure 2.6.1), the tilting merely introduced a slight asymmetry in the ZF drive.

In addition to the fluctuation metric and the ZF drive, figure 2.6.1 also shows the poloidal distributions of the perpendicular component of toroidal angular momentum transfer, $\pi_\perp^{\phi,\mathrm{tb}}(\theta)$. The resulting net radial flux is also shown, and highlights that most of the transfer is still cancelled out. Importantly we see that there is no clear correlation between the ZF drive envelope and $\pi_\perp^{\phi,\mathrm{tb}}(\theta)$. Instead, maxima in ZF drive often occur at locations where very little momentum transfer takes place.

This leads us to conclude that the locations where perpendicular momentum is driven on a flux surface do not correspond to the locations where it is most strongly transferred between flux surfaces.

2.7 What role does geometry and curvature play?

The overarching and striking result from the previous sections is the influence of curvature on the zonal flow drive. Visually in the mapping of the drive on the poloidal cross-section, drive maxima appear near those of poloidal curvature with a tendency to be shifted towards the bad curvature side, see figures 2.5.2, 2.5.8, 2.5.10, 2.5.11, and 2.6.1. In a tokamak, the toroidal and poloidal curvatures are equal to the principal curvatures κ_1 and κ_2, respectively.

The other important curvatures are the components of the field line curvature vector $\boldsymbol{\kappa} = \mathbf{b} \cdot \nabla \mathbf{b}$. This vector is orthogonal to the field line everywhere due to

$\mathbf{b} \cdot \nabla \mathbf{b} \cdot \mathbf{b} = 0$. A natural way to think about $\boldsymbol{\kappa}$ is to decompose it into the parts that are normal to the flux surface and within the surface, i.e. into normal and *geodesic* curvature. These are obtained through

$$\kappa_n = \mathbf{n} \cdot \boldsymbol{\kappa} \quad \text{and} \quad \kappa_g = (\mathbf{b} \times \mathbf{n}) \cdot \boldsymbol{\kappa}, \tag{2.7.1}$$

where $\mathbf{n} = \nabla\psi / |\nabla\psi|$ is the flux surface normal (Helander 2014). Field line curvature impacts turbulence in gyrokinetics by featuring in the magnetic drift frequency $\omega_d = \mathbf{v}_d \cdot \mathbf{k}$. Here

$$\mathbf{v}_d = \frac{m_s}{2e_s B}(v_\perp^2 \mathbf{b} \times \nabla \mathbf{b} + 2v_\parallel^2 \mathbf{b} \times \boldsymbol{\kappa}) \tag{2.7.2}$$

is the magnetic drift velocity. We note that

$$\boldsymbol{\kappa} = \frac{\mu_0 \nabla p}{B^2} + \frac{\nabla_\perp B}{B} = \frac{\beta}{2p}\nabla p + \frac{\nabla_\perp B}{B} \tag{2.7.3}$$

where $\beta = 2\mu_0 p / B^2$ is the plasma beta.

Since in gyrokinetics we write the wave vector in the natural flux tube coordinates (ψ, α),

$$\mathbf{k} = k_\psi \nabla\psi + k_\alpha \nabla\alpha, \tag{2.7.4}$$

the parts of the magnetic drift that feature are $\mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\psi$, $\mathbf{b} \times \nabla\mathbf{b} \cdot \nabla\alpha$, and $\mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\alpha$. Note that since $\nabla p = (\mathrm{d}p/\mathrm{d}\psi)\nabla\psi$ the radial component of the grad-B and curvature drift are the same up to constants and thus have the same poloidal variation,

$$\mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\psi = \mathbf{b} \times \nabla\mathbf{b} \cdot \nabla\psi, \tag{2.7.5}$$

regardless of the assumption on the plasma beta. If, however, the plasma beta is indeed assumed to be small, then also

$$\mathbf{b} \times \nabla\mathbf{b} \cdot \nabla\alpha \approx \mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\alpha. \tag{2.7.6}$$

Note that by permutation rules of the triple vector product the geodesic curvature is therefore equal to the radial component of the magnetic drift

$$\kappa_g \propto \mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\psi, \tag{2.7.7}$$

whereas the $\nabla\alpha$-component of the magnetic drift is not completely in the same direction as the normal curvature, since generally $\mathbf{b} \times \nabla\alpha$ has a component in a direction other than just $\mathbf{n}$ due to the non-orthonormality of the coordinate system. However, the fastest growing modes generally have zero radial wavenumber such that it is solely $\mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\alpha$ that determines the parallel profile of ω_d in the small beta limit and is therefore often called the *drift curvature* κ_d. The sign of ω_d and thus of $\kappa_d \equiv \mathbf{b} \times \boldsymbol{\kappa} \cdot \nabla\alpha$ determines whether curvature driven ITG modes are locally unstable or not, and thus separates the so-called good from bad curvature regions. The geodesic curvature is known to be important for linear zonal flow physics; stronger geodesic curvature causes stronger linear damping of zonal flows (Xanthopoulos et al. 2011; Nakata & Matsuoka 2022). We will thus consider the pair (κ_d, κ_g) in the following instead of (κ_n, κ_g).

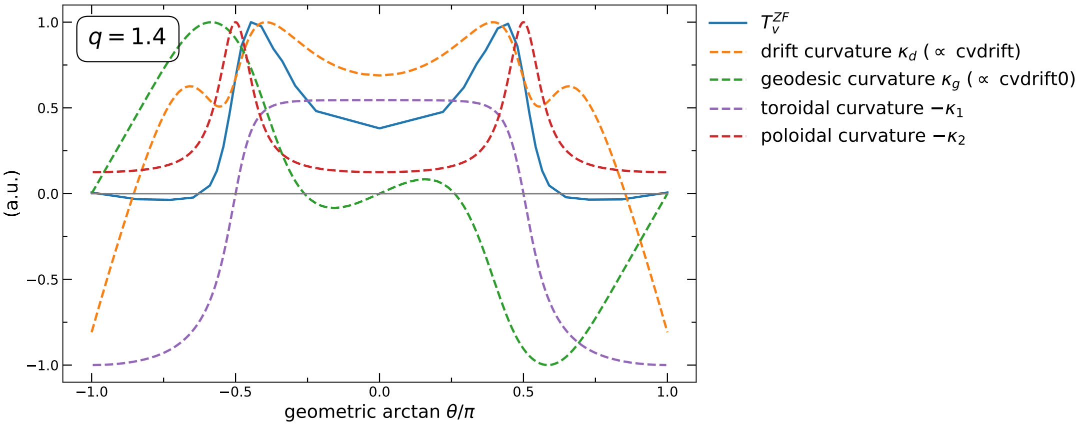

Figure 2.7.1. – The poloidal distribution of zonal flow drive alongside those of curvatures.

All four curvatures are plotted alongside the zonal flow drive envelope in figure 2.7.1 for the case with elongation $\kappa = 2.0$[2]. The drift curvature κ_d and geodesic curvature κ_g are obtained from the geometry output of `GS2`[3], and the principle curvatures were calculated analytically from the Miller equations, see Appendix B.

In figure 2.7.1, the two curvatures which have local maxima most closely aligning with those of zonal flow drive are the drift curvature and the poloidal curvature. The drift curvature can be varied without changing the poloidal curvature by modifying the safety factor. To understand why we note that (1) if $\nabla\psi$ and $\nabla\alpha$ were orthogonal, the drift curvature would correspond to the normal curvature, and (2) that the normal curvature is related to the principle curvatures via (Kreyszig 1959, theorem 42.1)

$$\kappa_n = \kappa_1 (\cos\xi)^2 + \kappa_2 (\sin\xi)^2, \tag{2.7.8}$$

---

[2] The symbol κ has unfortunately many uses in the literature. Its meaning here should be clear from the context. Expressed as a vector or with subscript it refers to curvatures, as a scalar without subscript it refers to elongation.

[3] from the standard output `cvdrift` and `cvdrift0`, respectively.

where $\xi = \arctan(\iota)$ is the pitch angle, i.e. the angle between the field line and the toroidal direction. From this we conclude that the coupling between the drift curvature and the poloidal curvature can be reduced by decreasing the pitch angle ξ, i.e. by increasing the safety factor q.

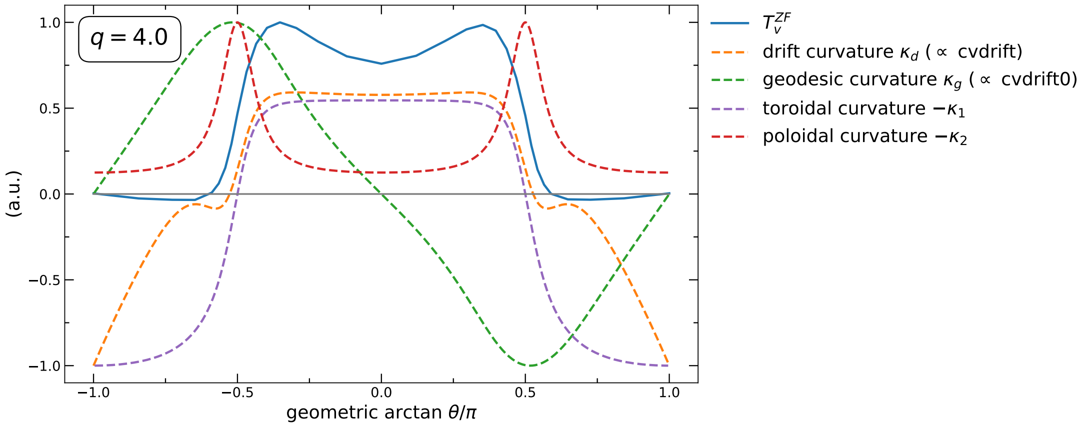

Figure 2.7.2. – The poloidal distribution of zonal flow drive alongside those of curvatures for increased safety factor q to vary the contribution from toroidal and poloidal curvatures to the drift curvature (see text for more details).

The results for $q = 4.0$ are shown in figure 2.7.2. Indeed, compared to the results with $q = 1.4$ from figure 2.7.1, the maxima in the drift curvature are significantly reduced, and the profile now follows more closely the one of toroidal curvature which does not have such maxima. The zonal flow drive still features its previous local maxima, though less pronounced. *We are therefore at least able to conclude that the drift curvature profile is unable to explain the poloidal distribution of the zonal flow drive.*

In terms of geometry, aside from the drift curvature profiles, the magnetic field magnitude B and the (contravariant) metric terms $g^{\psi\psi} = |\nabla\psi|^2$, $g^{\psi\alpha} = \nabla\psi \cdot \nabla\alpha$, and $g^{\alpha\alpha} = |\nabla\alpha|^2$ directly enter the gyrokinetic equation. For completeness, figure 2.7.3 shows the zonal flow drive distribution next to all six unique geometric quantities that enter the equations[4]. None exhibit a poloidal structure with similar maxima as the zonal flow drive, aside from the curvature drift and the geodesic curvature which we showed in the previous plots.

We therefore conclude by noticing that it is far from obvious how the magnetic geometry, through the nonlinearity of the turbulent dynamics, affects the zonal flow drive distribution. Formally such understanding is a mapping of the form

$$\mathcal{M} : \left(\mathbb{R}^{[-\pi,\pi]}\right)^6 \to \mathbb{R}^{[-\pi,\pi]}, \quad \text{such that} \quad \mathcal{T}^{ZF}(\theta) = \mathcal{M}(g_1, \ldots, g_6)(\theta) \tag{2.7.9}$$

---

[4] Note that the ZF drive distribution originates from a nonlinear simulation where additional parallel points add significant computational expense, whereas the geometric quantities were generated with a computationally cheap linear simulation with very high parallel resolution.

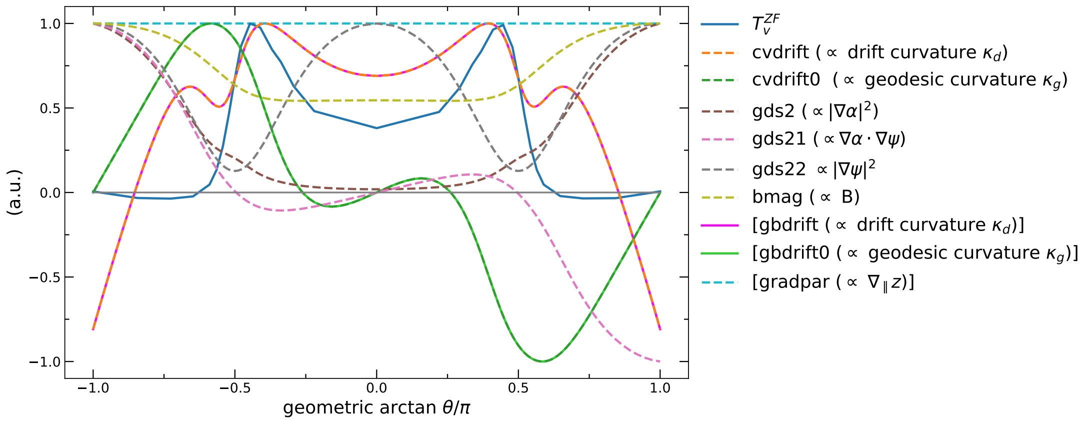

Figure 2.7.3. – The zonal flow drive distribution alongside the six independent geometric functions that enter the gyrokinetic equation for the case with $q = 1.4$. For completeness three other geometric quantities are plotted which for us are either equal to one of the other six, or just constants, but with the exception of `gbdrift0` need not be trivial in general.

where the $g_i(\theta)$ are the geometric quantities. Note that only recently a significantly simpler mapping of the form

$$\mathcal{M}' : \left(\mathbb{R}^{[-\pi,\pi]}\right)^6 \to \mathbb{R}, \quad \text{such that} \quad Q = \mathcal{M}'(g_1, \ldots, g_6) \tag{2.7.10}$$

was found to map the geometric quantities to the single-valued (time-averaged) heat flux Q and required sophisticated machine learning techniques, both for the mapping and for its interpretability (Landreman et al. 2025).

The goal of this chapter was to investigate if the zonal flow drive distribution changes with plasma shaping. We uncovered a strong dependence on plasma shaping with the possibility of multiple local maxima. However, despite the computational nature of this work which makes understanding easier than in experiments, understanding what the origins are here likely requires significantly more work, and is therefore outside the scope of this Thesis.

However, one does wonder whether simple proxies are able to capture the trends, a question which we will address next.

2.7.1 Transfer proxies

The common trend we have observed in all of the results so far is that the ZF drive envelope aligns with maxima of poloidal curvature but tends to be shifted towards the outboard side where the turbulence is stronger. As argued in the previous section, understanding why this is the case is non-trivial. In general, getting an analytical handle on the complex nonlinear interaction of turbulence has occupied researchers for a long time, and has proven incredibly difficult.

This inherent difficulty to make analytical progress on this topic has in fact been one of the motivations of approaching this problem in the way we have done here, i.e. via direct measurement of three-wave-coupling.

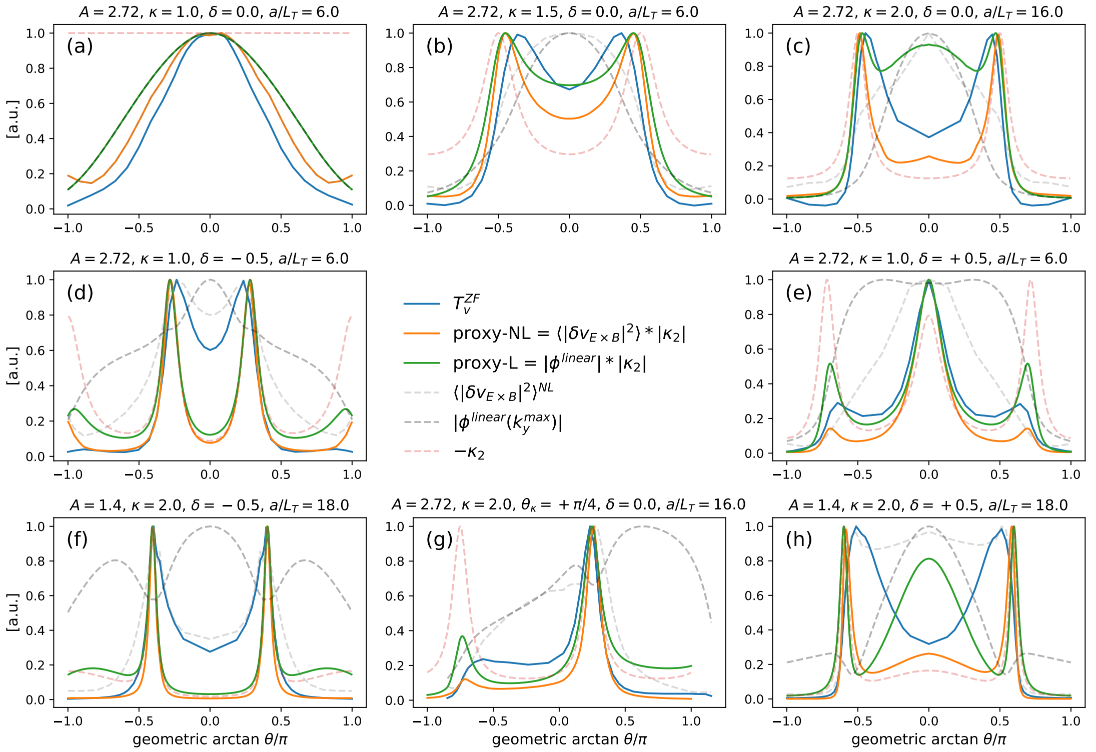

Figure 2.7.4. – Proxies for the zonal flow drive for a range of equilibria from the previous sections. Both the linear and nonlinear proxy are shown (see text). The contributing quantities are shown in light dashed lines. The configurations are **(a)** 'CBC', **(b)** 'slightly-elongated-CBC', **(c)** 'elongated-CBC', **(d)** 'NT-CBC', **(e)** 'PT-CBC', **(f)** 'elongated-NT-ST', **(g)** 'tilted-elongation', **(h)** 'elongated-PT-ST'.

One useful approach we can take, however, is to find suitable proxies *a posteriori*. Here we test simple proxies for the drive envelope of the form

$$\begin{aligned} \text{proxy-NL} &= \langle|\delta v_{\mathbf{E}\times\mathbf{B}}|^2\rangle * |\kappa_2|, \\ \text{proxy-L} &= |\phi^{linear}_{t=t_{end}}| * |\kappa_2|, \end{aligned} \tag{2.7.11}$$

where, again, κ_2 is the poloidal curvature. It should be noted that there is no specific theoretical motivation for the use of κ_2 here, but we showed in the previous sections that its poloidal dependence is a good predictor for the number and approximate locations of the local minima in the ZF drive.

Aside from the interest of how well such simple proxies can reproduce the observed trends, one of the main motivations for a proxy is computational speed-up, since for experimental analysis one would like to be able to evaluate the ZF drive envelope quickly for varying equilibria. Evaluating the ZF drive is expensive as it requires a dedicated converged nonlinear simulation for each equilibrium. The nonlinear proxy, 'proxy-NL', relies on the nonlinear turbulent activity envelope and would thus not give any speed-up. We therefore also introduce a linear proxy, 'proxy-L', which uses the linear mode envelope of the fastest growing mode (in the pure linear picture) instead.

Figure 2.7.4 shows both proxies evaluated for many of the equilibria of the previous sections. Both the linear and nonlinear proxies reproduce the qualitative shape, and number and location of poloidal maxima remarkably well. The influence from the poloidal curvature is often too strong. This could be adjusted for example by introducing exponents on the term in equation 2.7.11 and fitting exponents for a desired parameter space.

2.8 Transfer contributions

Additional insight into the details behind the observed poloidal distributions can be obtained by omitting part of the summation in the transfer functions as defined in equation 1.4.19. We define three auxiliary functions

$$\begin{aligned} \mathcal{T}_{x}^{ZF}(\theta, k_x) &\equiv \sum_{k_{1x},k_{1y}} \mathcal{T}(k_y = 0), \\ \mathcal{T}_{1x}^{ZF}(\theta, k_{1x}) &\equiv \sum_{k_x,k_{1y}} \mathcal{T}(k_y = 0), \\ \mathcal{T}_{1y}^{ZF}(\theta, k_{1y}) &\equiv \sum_{k_x,k_{1x}} \mathcal{T}(k_y = 0), \end{aligned} \tag{2.8.1}$$

which now resolve the drive poloidally as well as along a specific scale. As an example we use the case of large elongation here with $\kappa = 2.0$ (and otherwise CBC values). Figure 2.8.1 shows the poloidal distribution for each scale separately.

The top plot in figure 2.8.1 shows the binormal size of eddies which drive ZFs. Near the maxima of the total drive the origin is concentrated in lobes within $0.1 \leq k_y\rho_i \leq 0.5$. Towards the midplane the binormal source range for ZFs is broad and extends across $k_y\rho_i > 0.25$.

The middle plot in figure 2.8.1 shows the radial size of eddies that are driving ZFs at each poloidal location. Near the local maxima of the total drive a significantly wider range of structures are depositing their energy into the ZFs, while the spectral range is narrower at the outboard midplane.

Finally, the bottom plot in figure 2.8.1 shows the poloidal distribution of the drive of each radial structure of ZFs. A broad range of ZFs are driven with $0.25 \leq k_x\rho_i \leq 1.0$, the majority of which follow the same poloidal distribution as the total drive. The drive appears slightly more broadly at $k_x\rho_i \sim 0.75$ compared

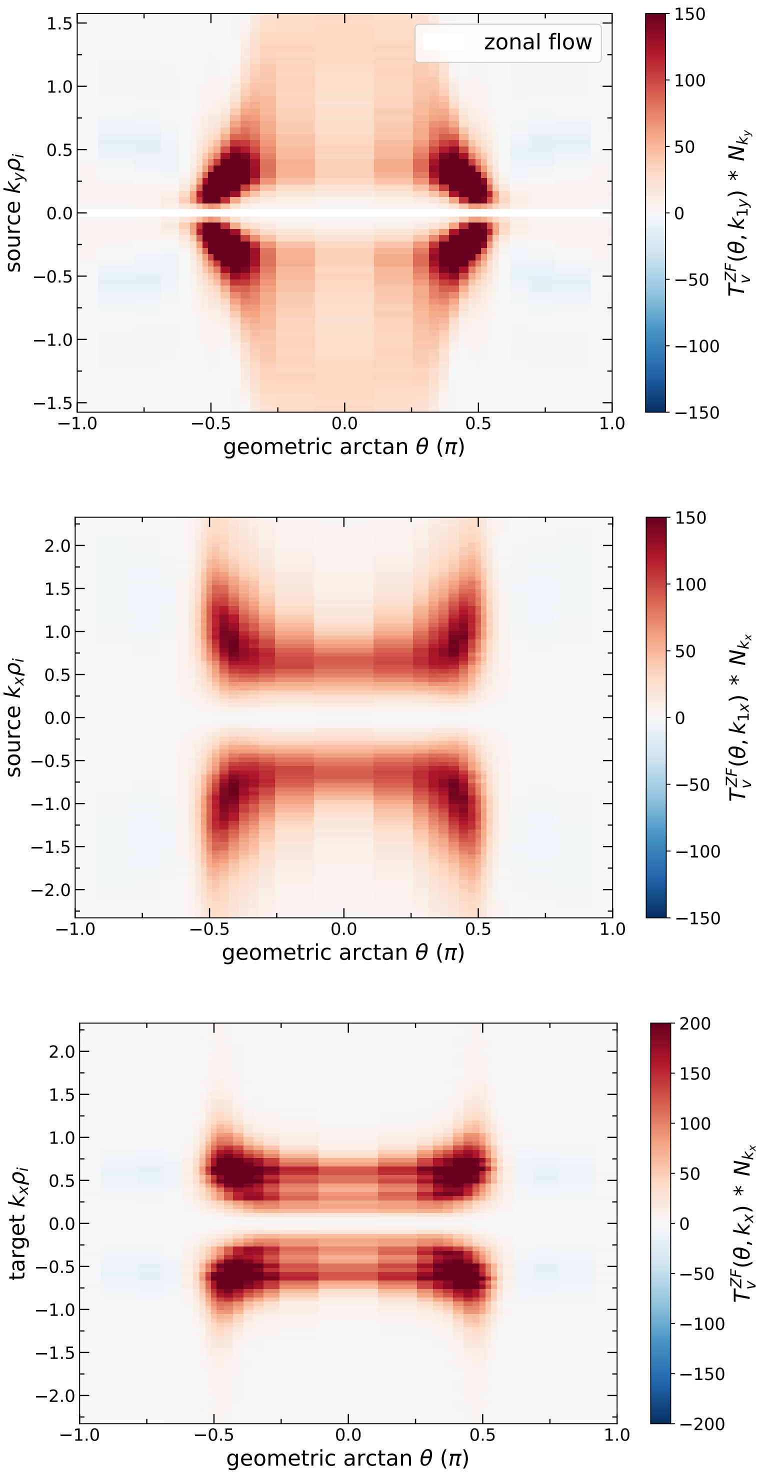

Figure 2.8.1. – Transfer contributions from different scales at every poloidal location for pure elongation ($A = 2.72$; $\kappa = 2.0$; $\delta = 0.0$; $a/L_T = 6.0$). **Top:** Drive from each driving binormal mode. **Middle:** Drive from each driving radial mode. **Bottom:** Drive to each radial structure of ZFs.

to the other scales. Note that the drive over the binormal target scale does not exist since we selected the zonal flow as the single binormal target mode.

2.9 Towards experimental validation

The presented zonal flow drive analysis was performed in a way that can directly be validated (Holland 2016) with experiments. Velocimetry algorithms such as cross-correlation time delay estimation (CCTD) (Holland et al. 2004) or dynamic time warping (DTW) (Quénot, Pakleza & Kowalewski 1998; Kriete et al. 2018) can be used to obtain the velocity fluctuations from spatially resolved measurements of density fluctuations (Enters et al. 2023), as it was used in for example (Cziegler et al. 2017). While both beam emission spectroscopy (BES) and gas puff imaging (GPI) can provide such measurements, data from GPI usually features much better signal-to-noise ratio which is absolutely crucial for the measurement of the nonlinear coupling which is fundamental to the ZF drive.

For such validation one could make use of the opposite trends for the ZF drive for positive and negative triangularity around the outboard midplane, presented in section 2.5.2. For PT the outboard midplane is predicted to be the global drive maximum, while for NT it is a local minimum. Thus, as one moves away from the midplane, the drive is expected to decrease (increase) for PT (NT). This effect should be measurable in experiments. Higher absolute values of triangularity would be desirable here since we showed that maxima and minima are sharper in this case.

Even if the imaging position of the turbulence diagnostic is fixed within the vessel, a different poloidal location could be observed through vertical shifting and, if needed, some change in aspect ratio to select the same flux surface. Such a small aspect ratio correction should not spoil the experiment, since we found pure aspect ratio effects to be negligible for the ZF drive.

While our analysis was already guided by validation with experiment, i.e. we chose a Reynolds stress analysis that has been performed in experiments before, one step further into this direction could be taken through the use of synthetic turbulence diagnostics in our gyrokinetic codes (Holland 2016).

We hope that these ideas presented here will help future validation efforts.

2.10 Discussion and conclusion

We have shown that the zonal flow drive develops a non-trivial poloidal distribution for realistic plasma shaping. Once elongation is introduced to the circular plasma shaping of the Cyclone Base Case, the drive maximum splits into two, which appear near the top and bottom of the tokamak; The outboard midplane now becomes a local minimum.

The study of triangularity showed clearly an asymmetry of the drive distribution with respect to the sign of triangularity, highlighting that the position of poloidal curvature relative to the distribution of turbulence is important: positive triangularity obtains its global drive maximum at the outboard midplane, while negative triangularity does not, as expected, obtain such a maximum at the inboard midplane where the turbulent fluctuations are most suppressed. Instead two local maxima appear near the outer corners, shifted towards the outboard midplane.

By combining compact shaping with significant elongation and triangularity we examined a typical spherical tokamak equilibrium. The outboard midplane becomes a local minimum of the drive due to the dominating effect of elongation over triangularity. This hierarchy between elongation and triangularity in setting the drive maxima was also observed for at standard aspect ratio[5]. *The important consequence is that in typical plasmas in spherical tokamaks such as MAST-U and NSTX-U, as well as in conventional tokamaks such as ASDEX-U and TCV, the zonal flow drive will likely not peak at the outboard midplane.*

Note that it was shown recently that the *mean* $\mathbf{E} \times \mathbf{B}$ shearing rate exhibits similar shaping effects for triangularity as we observe for the Reynolds stress (Singh, Diamond & Nelson 2023). The authors of that study argue that the Reynolds stress could have a similar structure in experiments, as the poloidally varying mean shearing rate tilts eddies differently at different poloidal locations. It should be pointed out that the work in this Thesis found a non-trivial structure for the Reynolds stress *without a background* $\mathbf{E} \times \mathbf{B}$ *shearing rate and the results are therefore to be viewed as independent.* Our results show that the Reynolds stress exhibits a nonlinear poloidal distribution without any eddy tilting from background $\mathbf{E} \times \mathbf{B}$ motion since gyrokinetics only evolves fluctuating quantities. This is important since a sufficiently strong background radial electric field might not always exist.

Further, we showed that the ZF drive envelope responds to up-down asymmetry of the gyrokinetic equation. Since this symmetry is broken by shaping but also non-zero flow and flow shear, it is essentially always broken to some degree in experiments. *The zonal flow drive envelope therefore is not only expected to peak away from the midplane in experiments, but asymmetrically so.*

In the final sections we highlighted the fundamental difficulty of understanding the root cause of how geometry affects the drive envelope. We highlighted that

---

[5] From this one would expect that higher order shaping, i.e. squareness, has a subdominant effect. This would be worth investigating in future studies.

none of the gyrokinetic geometric quantities seem directly correlated but that proxies formed by combining turbulence distributions with those of poloidal curvature perform reasonably well. Finally, we showed how the transfer functions can give additional insight, and we showed how the poloidal drive dynamics unfold over each of the involved scales.

2.10.1 Future work

These non-trivial results open the door to many avenues of future work. One aspect concerns validation with experimental data as well as applying these results to the interpretation of existing data. Note that for validation it is not required to measure the full poloidal distribution in experiments. Such measurements are not possible in reality due to access restrictions to the vacuum vessel. Instead, a two-point comparison near the midplane for example could be enough to validate the predicted trends. By comparing measurements at the midplane and slightly away from the midplane for negative and positive triangularity, one would expect opposite trends from the results presented in this chapter.

It would be desirable to gain more understanding on why the zonal flow level is often observed to be only quasi-stationary for some of the zonal modes, and to find ways to avoid this problem without introducing unrealistic levels of hyper-viscosity.

It would also be instructive to include more physics into the analysis. Including equilibrium flow and flow shear would teach us important lessons on how exactly the drive envelope would develop an up-down asymmetry in experiments. Including electromagnetic fluctuations with the associated Maxwell stress as well as kinetic electrons might modify the results. Note that usually the drive from the Reynolds stress is believed to be stronger than the one from the Maxwell stress (Diamond et al. 2005), though plasmas with a high plasma beta such as the designs of STEP (Meyer & STEP Plasma Team 2024) might have a non-trivial electromagnetic contribution. Moreover, including the effect of squareness in the analysis of spherical tokamaks in particular would be a natural extension towards even more realism.

Even though in this Thesis the potential application to validation studies motivated the use of a turbulence model that is as realistic as possible while computationally feasible – gyrokinetics – it would also be interesting to explore how much the turbulence model can be reduced while retaining the essential trends in the ZF drive distribution. One such simplification would be the use of gyrofluid instead of gyrokinetic simulations. Gyrofluid simulations have the benefit of being computationally less expensive while retaining much of the same physics. Going even further, one could test whether 'quasilinear' estimates for the Reynolds stress structure can be constructed by suitably combining the linear mode envelopes of a group (or groups) of three modes (so-called triads). A suitable choice of triad(s) could be those whose coupling is known from gyrokinetic simulations to contribute strongly to the drive of ZFs. Mode envelopes could be taken from either linear gyrofluid or linear gyrokinetic calculations, though the computational speed-up is less important here due to the linear nature of the calculation. An important caveat, however, is that such a quasilinear estimate would assume that a few triads contain the relevant poloidal dependency of the ZF drive in their linear phase. If instead, as one might expect, the ZF drive fundamentally depends on the envelope of the nonlinear fluctuations, larger collective effects, or both, such an estimate would likely not work well.

Finally, a deeper understanding of the exact interplay of geometry and the ZF drive distribution would be desirable. Given how many resources the community has spent on understanding the simpler nonlinear interplay of geometry and the nonlinear heat flux, this is not a trivial task. Approaches similar to the ones from Ref. (Landreman et al. 2025) appear like a promising starting point.

2.10.2 Key findings and implications

The key points of this chapter and its implications are:

- The poloidal distribution of the zonal flow drive is generally not correlated to the distribution of the turbulent fluctuations and responds sensitively to the flux surface shaping, specifically its curvature.
- Due to the strong flux surface shaping of today's tokamaks and future reactor designs, the ZF drive is not expected to be strongest at the outboard midplane where the turbulence is usually most active.
- Global ZF drive estimates can now be obtained from an experimental measurement at a single poloidal location, making a comparison of global ZF drive between differently shaped equilibria possible.
- Validation studies should consider these results to avoid measurements for a ZF drive analysis to be taken in locations of particularly weak drive.
- The results contribute to the development of comprehensive physics-based models of confinement and its transitions.

3

Bridging tokamaks and stellarators: optimising perturbed tokamaks for quasi-axisymmetry

In this chapter we will investigate how perturbations to tokamaks can be numerically optimised such that they obey so-called quasi-axisymmetry to very high degree, with a quasi-symmetry error $< 0.1\%$ globally at an aspect ratio of $A \sim 2.4$. We will develop a new optimisation target function which allows us to design equilibria that feature perturbations which are strong enough that stellarator benefits (see section 1.6.2) can be expected while having negligible fast particle losses, as well as incorporating additional physics targets into the designs. In the context of this Thesis, this chapter will serve as a stepping stone to expanding our zonal flow drive analysis to stellarators in the most natural way.

This chapter is largely based on the content in references

> T. M. Schuett and S. A. Henneberg (Dec. 2024b). Exploring novel compact quasi-axisymmetric stellarators. In: *Phys. Rev. Res.* 6, p. L042052
>
> Tobias M Schuett and Sophia A Henneberg (June 2025). Optimization of compact quasi-axisymmetric stellarators. In: *Plasma Physics and Controlled Fusion* 67.6, p. 065024

3.1 Motivation – the potential of compact QA

So far we have investigated the effect of axisymmetric (tokamak) shaping on the zonal flow drive. While simple to build due to their axisymmetry, tokamaks also have their drawbacks. Due to the large requirement on the plasma current, steady state operation becomes difficult to achieve economically. Further, tokamaks can suffer from disruptions (Helander et al. 2012), which release the stored plasma energy in a very short time and must therefore be avoided in any future reactor. Stellarators on the other hand have a reduced need for a plasma current, and, with sufficiently high amount of twist from shaping, they do not suffer from disruptions. They are, however, more difficult to build due to their often complex coils, and typically require larger aspect ratio, increasing the cost substantially.

In this chapter we describe a method to obtain stellarator equilibria via numerical optimisation which are compact, quasi-axisymmetric (what this means and why it is required will be explained shortly), and retain an axisymmetric volume that is large enough such that they can be realised as a stellarator-tokamak hybrid as envisioned in the work by Henneberg and Plunk (Henneberg & Plunk 2024), depicted in figure 3.1.1.

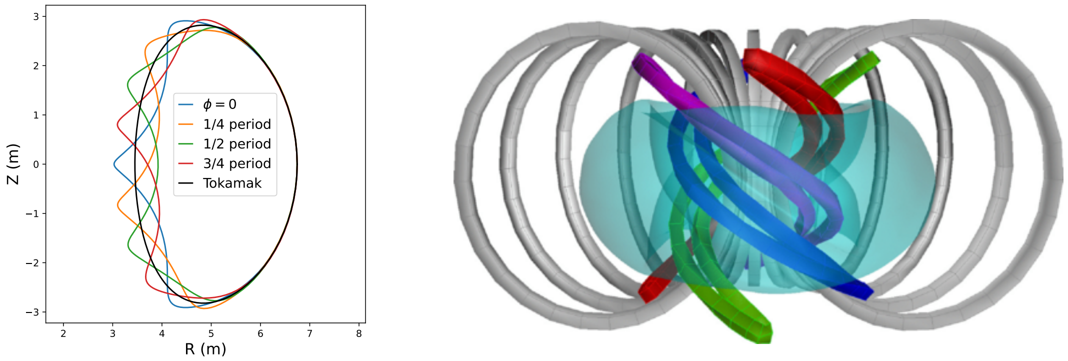

Figure 3.1.1. – The stellarator-tokamak hybrid design by Henneberg and Plunk (Henneberg & Plunk 2024). **Left:** Poloidal cross-sections at four toroidal locations within one of the four field periods. The axisymmetric boundary is also shown. **Right:** 3D rendering. The plasma is shown in turquoise, the QA-coils in dark colours, and the standard tokamak coils in gray. Image on the right from Ref. (Henneberg & Plunk 2024).

These compact QA plasma equilibria which retain a large axisymmetric part were first found analytically, both under the vacuum approximation (Plunk & Helander 2018) and for finite plasma pressure (Plunk 2020). The design from Ref. (Henneberg & Plunk 2024) was based on an analytic equilibrium and showed they are well suited to obtain a first-of-its-kind *optimised* hybrid design. This hybrid design can operate continuously between tokamak and stellarator based on the choice of the currents in the QA-coils. When no current is run in the QA-coils one obtains a tokamak, whereas with full currents in the QA-coils a stellarator is obtained (Henneberg & Plunk 2024). Different coil sets are able to generate these equilibria (Henneberg & Plunk 2025). Neoclassical changes are minimal due to QA. Not only is such a device a promising solution to merge the advantages of tokamaks and stellarators, but it is also ideal for the exploration of effects from confinement-preserving 3D shaping, such as those on the zonal flow drive.

It is desirable to have a method to optimise this class of equilibria numerically as this allows one to mould them into different shapes to adhere to various requirements. In this chapter we present such a method, and present various equilibria with improved characteristics and which are suitable for a hybrid. These include a self-consistent current and improved MHD stability properties. We also corroborate our understanding of where 3D perturbations naturally appear poloidally to preserve QA as much as possible.

3.2 Why do stellarators need to be optimised?

The optimisation of the plasma equilibrium of stellarators is required for the following reason. While nested magnetic surfaces, produced through twist in the field, confine particles that are unrestricted in their motion along magnetic field lines (called parallel motion), they generally do not confine particles that *are* restricted in their parallel motion, and as we will elucidate now some particles do face such a restriction. In the single-particle picture of the plasma, all particles conserve their kinetic energy[1] $\varepsilon = mv^2/2$ as well as their magnetic moment[2] $\mu = mv_\perp^2/(2B)$ along their trajectories (Helander 2014). Due to the toroidal nature of the magnetic field, the field strength depends inversely on the major radius R, i.e. $B \sim 1/R$. As particles follow the twisted field lines they therefore travel through regions of varying magnetic field strength.

As particles move away from a magnetic field minimum B_{min} (it can just be a local minimum) the magnetic field strength increases and $v_\perp$ must therefore increase to keep μ conserved. This must happen at the expense of $v_\parallel$ since the total velocity $v^2 = v_\perp^2 + v_\parallel^2$ is conserved through ε. As a consequence the parallel motion slows down as the particles enter regions of stronger magnetic field.

Particles which already had a small parallel velocity relative to their perpendicular velocity at B_{min} will eventually get repelled and reverse their motion. These particles are referred to as *trapped particles* since they are only allowed on a specific section of the field line depending on where they entered. The two locations along the field line to either side of B_{min} where a given class of

---

[1] The kinetic energy is conserved since the electrostatic potential is a flux function to leading order.

[2] Strictly speaking μ is an adiabatic invariant but can be regarded as constant for the purposes here.

particles reverse their motion are called the *bounce points*, and their motion is referred to as *bounce motion*.

Particles which start with a large parallel velocity relative to their perpendicular velocity are able to pass through all points along the field line and are thus called *passing particles*.

In a tokamak there is only one minimum B_{min} which typically sits at the outboard midplane, such that passing and trapped particles alike spend the same amount of time on the top and bottom half of the torus and the radial magnetic drift motion induced by the toroidicity still averages to zero, keeping them confined.

In a stellarator this need not be the case as multiple magnetic field minima are possible and they might be located somewhere where the magnetic geometry is not symmetric to either side of them, causing an asymmetry in the particle drifts. This leaves particles which are trapped around this minimum with a net radial drift, i.e. no confinement. It is for this reason that stellarators must be optimised in their magnetic geometry such that the average radial drift does vanish for all trapped particles and confinement is restored.

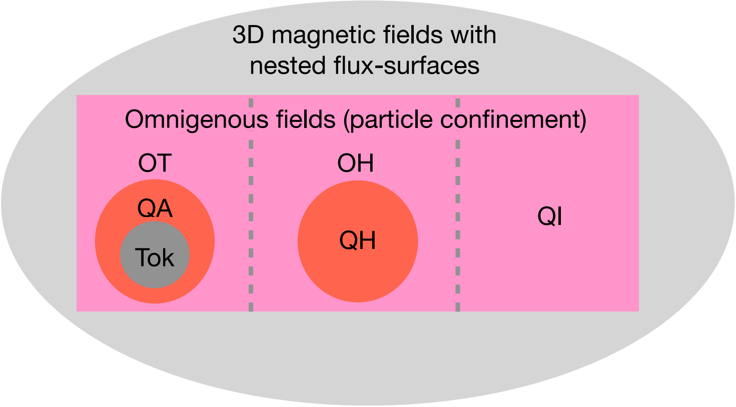

Figure 3.2.1. – Another distinction of toroidal confinement devices is by the topology of the magnetic field strength $B = |\mathbf{B}|$ on the nested surfaces. From left to right within the pink region: toroidal, helical, and poloidal topology of B. Here, the tokamak labelled as 'Tok', is just a special limit of confinement devices with toroidal B-topology. The other regions are toroidal omnigeneity (TO or OT), quasi-axisymmetry (QA), helical omnigeneity (HO or OH), quasi-helical symmetry (QH), and poloidal omnigeneity more commonly referred to as quasi-isodynamicity (QI).

What is the minimum requirement on the magnetic geometry to confine all particles, passing and trapped? For this we must first note that trapped particles do not only drift radially during their bounce motion but also *precess* from one field line to another as part of their drift motion. Their confinement therefore does not only depend on how the geometry and B changes along one particular field line, but instead on how B behaves on the entire flux surface. One requirement turns out to be that the maximum and minimum of B on a flux surface are lines that circle around the plasma, rather than being a point on the surface. From this it follows that all contours of B must encircle the plasma, which they can do either toroidally, helically, or poloidally. (The plasma itself, of course, still has toroidal topology.) Another requirement is that the distance between two bounce points along the field line in straight-field-line angles must be the same on all field lines. This must be the case for all possible bounce points, i.e. all values of B (Landreman & Catto 2012; Helander 2014).

Equilibria which fulfil this requirement are called *omnigenous*. Equilibria in which all contours of B not only encircle the plasma but do so as straight lines for some special poloidal and toroidal angle definitions (e.g. so-called Boozer angles) are said to satisfy *quasisymmetry*, since then B only depends on one angle χ in these coordinates, i.e. $B = B(\psi, \chi)$ with $\chi = M\theta_B - N\phi_B$ where M and N specify the topology of the symmetry. These equilibria are only quasi-symmetric and not symmetric like a tokamak since the magnetic field vector still depends on both angles, i.e. $\mathbf{B} = \mathbf{B}(\psi, \theta_B, \phi_B)$.

The fact that omnigeneity can be obtained according to these three topological variants on how B-contours encircle the plasma offers a more fundamental perspective on the classification of confining toroidal magnetic fields (in contrast to the well-known tokamak versus stellarator classification). This classification is depicted in figure 3.2.1, featuring {*toroidal/axi, helical, poloidal*} *omnigeneity* and *quasi-*{*toroidal/axi, helical, (poloidal)*} *symmetry*[3]. In this picture the tokamak is just a special case of omnigenous equilibria, one with a symmetry much beyond what is required for confinement. Omnigeneity is difficult to achieve analytically beyond asymptotic expansions, motivating numerical optimisation of 3D equilibria to design stellarators.

3.3 A short introduction to stellarator optimisation

The details of the three-dimensional shaping of a stellarator dictates, to a large extent, their physical properties. The 3D shaping represents a large number of degrees of freedom. Boozer estimated these to be ~ 50 compared to typically ~ 4

---

[3] Note that quasisymmetry with B-contours that are straight lines which close poloidally is not possible near the magnetic axis (Helander 2014), and thus quasi-poloidal symmetry is omitted from this figure.

for tokamaks (aspect ratio, elongation, triangularity, and squareness) (Boozer 2005).

Fundamental to stellarator optimisation is the result by Kruskal and Kulsrud (Kruskal & Kulsrud 1958) which states that the global MHD equilibrium is specified by the shape of the plasma boundary and two flux functions. These are typically the pressure profile and either the toroidal current density profile or the rotational transform profile. A choice must be made for the representation of the plasma boundary, and this choice matters as it affects the details of the optimisation landscape. The most common choice, and the one used for this work, is a double Fourier-series for both cylindrical coordinates[4]

$$\begin{aligned} R(\vartheta,\varphi) &= \sum_{m=1}^{m_{\max}} \sum_{n=-n_{\max}}^{n_{\max}} R_{m,n}\cos(m\vartheta - N_p n\varphi) + \sum_{n=0}^{n_{\max}} R_{0,n}\cos(N_p n\varphi), \\ Z(\vartheta,\varphi) &= \sum_{m=1}^{m_{\max}} \sum_{n=-n_{\max}}^{n_{\max}} Z_{m,n}\sin(m\vartheta - N_p n\varphi) - \sum_{n=1}^{n_{\max}} Z_{0,n}\sin(N_p n\varphi). \end{aligned} \tag{3.3.1}$$

Here, we make the common choice to assume stellarator symmetry, which eliminates half of the modes. Stellarator symmetry, also colloquially referred to as 'pancake symmetry', requires $R(\vartheta,\varphi) = R(-\vartheta,-\varphi)$ and $Z(\vartheta,\varphi) = -Z(-\vartheta,-\varphi)$.

The optimisation space, which we shall call $\mathcal{M}$, then consists of the Fourier modes $R_{m,n}$ and $Z_{m,n}$ and sometimes other parameters such as the current profile or the field strength. The optimisation problem is then commonly cast as a sum of squares

$$\min_{x\in\mathcal{M}} \quad f(\mathbf{x}) = \sum_i w_i\,(f_i(\mathbf{x}) - f_i^*)^2, \tag{3.3.2}$$

where constraints are included by penalising constraint violations. $f_i(\mathbf{x})$ are the different objectives, f_i^* their target values, and w_i are user-specified weights.

Casting the problem as a nonlinear least-squares problem has the benefit that part of the Hessian is known from the Jacobian (Nocedal & Wright 1999, Chapter 10). While varying the user-specified weights is not optimal for exploring the Pareto frontier[5], this approach has been the workhorse for many years and has generally shown to give the best results so far in the literature (Henneberg et al. 2019; Landreman & Paul 2022; Landreman, Buller & Drevlak 2022; Goodman et al. 2023), though other approaches to include constraints have been developed as well (Conlin et al. 2024; Gil et al. 2025).

---

[4] In this representation ϑ is only a parameter and the representation is non-unique. Other boundary representations exist which remove this non-uniqueness such as the one by Henneberg *et al.* (Henneberg, Helander & Drevlak 2021).

[5] The Pareto frontier is the set of all points that are Pareto optimal. A design point is said to be Pareto optimal with respect to a multi-objective optimisation (MMO) problem when none of the objectives can be improved without worsening another. The Pareto frontier thus captures trade-offs between the objectives in the MMO problem.

The stellarator optimisation space is typically non-convex and plagued by a plethora of local minima. To avoid these local minima, a range of different hyper-parameter settings for the finite-difference step sizes are scanned for each optimisation problem. Another common strategy to avoid local minima which we also adopt is to stage the optimisation in a series of steps where the lowest-order Fourier-coefficients are optimised first and additional modes are added to the optimisation space once the minimum of the reduced space was found.

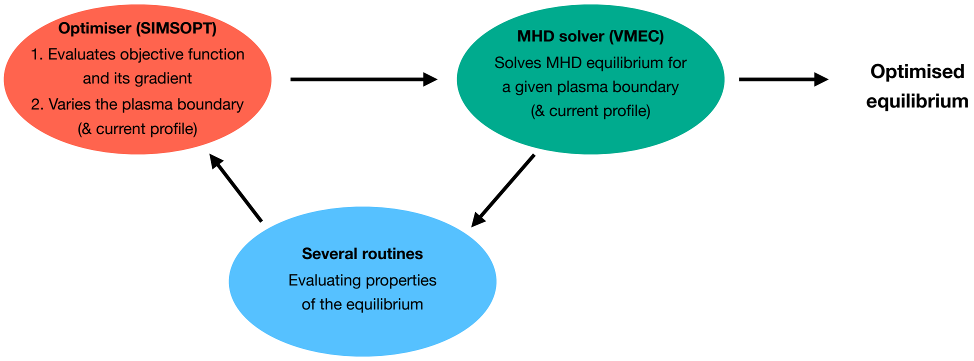

Figure 3.3.1. – The stellarator optimisation workflow used in this Thesis featuring the code `VMEC` and the optimisation framework `SIMSOPT`. Diagram based on lecture slides from Sophia Henneberg.

In this work we use the `SIMSOPT` framework (Landreman et al. 2021; Medasani et al. 2023) to perform the optimisation where we use `SCIPY`'s nonlinear least-squares method with the default algorithm 'trust region reflective'. Inside the optimisation loop `SIMSOPT` calls `VMEC` to evaluate MHD equilibria. While this is a local optimisation method, one could consider the hyper-parameter scans to be a crude form of globalisation. A schematic of the optimisation workflow is shown in figure 3.3.1.

3.4 Quasi-axisymmetry optimisation

The Lagrangian for particle motion in an electromagnetic field is given by

$$\mathcal{L} = \frac{mv^2}{2} + Ze\mathbf{A}\cdot\mathbf{v} - Ze\varphi. \tag{3.4.1}$$

The Lagrangian for the motion of the guiding centers of particles is obtained by gyro-averaging $\mathcal{L}$, and after transforming to Boozer coordinates is given by

$$\mathcal{L}_{\mathbf{gc}} = \frac{m}{2B^2}(I\dot{\theta}_B + G\dot{\phi}_B)^2 + Ze(\psi\dot{\theta}_B - \chi\dot{\phi}_B) - \mu B - Ze\varphi. \tag{3.4.2}$$

Since the electrostatic potential is usually a flux function $\varphi = \varphi(\psi)$, the guiding centre Lagrangian can only depend on the toroidal angle ϕ through the magnetic field strength $B = |\mathbf{B}|$. In a general stellarator field $B = B(\psi, \theta_B, \phi_B)$, but if B is independent of ϕ_B so is the guiding centre Lagrangian, and the canonical momentum is a conserved quantity:

$$\dot{p}_\phi = \frac{\partial L}{\partial \phi_B} = 0. \tag{3.4.3}$$

This guarantees that guiding centres remain confined to flux surfaces (Helander 2014)[6]. It is therefore sufficient for the magnitude of the magnetic field to have this property, not for all components of the magnetic field vector. The latter is the case in axisymmetry (i.e. a tokamak) and is a condition beyond what is required to confine collisionless guiding centre orbits.

While it has been shown that obtaining exact quasi-axisymmetry globally (on all flux surfaces) is not possible to first order in the perturbation (Plunk & Helander 2018), it has been shown numerically that it can be obtained to such a high degree that the error is truly negligible (Landreman & Paul 2022). Optimisation tools are thus used to tailor the magnetic field to minimise some measure of the symmetry error E_{QA}. A very intuitive measure is obtained by first transforming the magnetic field from `VMEC`'s PEST coordinates to Boozer coordinates, then Fourier-expanding the magnetic field strength in Boozer coordinates on each flux surface[7]

$$B(s, \theta_B, \phi_B) = \sum_{m,n} B_{m,n}(s) \cos\left(m\theta_B - N_p n \phi_B\right), \tag{3.4.4}$$

and finally penalising those Fourier-modes which break toroidal symmetry in Boozer coordinates through a measure like

$$E_{QA}(s) = \frac{\max\{|B_{m,n\neq 0}(s)|\}}{B_{0,0}(s)}. \tag{3.4.5}$$

This method was used for many years in stellarator optimisation and led to the designs of the HSX (Anderson et al. 1995) and NCSX (Zarnstorff et al. 2001) experiments, as well as other candidate designs, e.g. (Drevlak et al. 2013; Henneberg et al. 2019).

The downside of a target of this form is that it requires the coordinate transformation into Boozer coordinates. This transformation is usually performed with

---

[6] This is also true when B only depends on a linear combination of both angles $B = B(\psi, M\theta_B - N\theta_B)$, i.e. for quasi-symmetry as discussed in section 3.2.

[7] Only cosine modes appear due to the assumption of stellarator symmetry. See discussion around equation 3.3.1.

a code like `BOOZXFORM` (Sanchez et al. 2000) and introduces numerical noise as well as non-differentiability. In this work we use the target which was first used by Landreman and Paul (Landreman & Paul 2022) and yielded equilibria with much lower quasi-symmetry error than what was obtained before. This target reads

$$f_{\mathrm{QS}} = \sum_{s_j} w_j \left\langle \left(\frac{1}{B^3} [(N - \iota M)\mathbf{B} \times \nabla B \cdot \nabla \psi - (G + NI)\mathbf{B} \cdot \nabla B] \right)^2 \right\rangle, \qquad (3.4.6)$$

and is based on the idea that for perfect QA the two terms in the square brackets are equal. Here ι is the rotational transform, $2\pi G/\mu_0$ is the poloidal current, $2\pi I/\mu_0$ is the toroidal current, and M and N specify the type of quasi-symmetry. For QA, as considered here, $N = 0$ and $M = 1$. The extra factor B^{-3} is used to make the target scale-independent and w_j are weights which can be set by the user to prioritise some flux surfaces over others in the optimisation. Here we do not explore this freedom and set all weights equal to one, $w_j = 1$.

3.5 Optimising finite-beta QA equilibria close to axisymmetry

Optimising plasma equilibria near axisymmetry that are realistic for a future reactor, i.e. those with finite plasma pressure, is not a trivial task. The results in Ref. (Landreman & Paul 2022) were obtained by making the vacuum assumption, i.e. by assuming that the plasma pressure and net toroidal current are negligible. Under these assumptions the departure from axisymmetry is easily obtained by requiring some finite amount of rotational transform. Since no net toroidal current is present which can provide external rotational transform, the optimiser is forced to achieve this constraint through three-dimensional (3D) shaping.

Recently an efficient method has been introduced for calculating the plasma-driven bootstrap current in quasi-symmetric stellarators by making use of the isomorphism with axisymmetry (Landreman, Buller & Drevlak 2022; Landreman 2022). The quasi-axisymmetric results which now assume finite pressure and net toroidal current, were obtained by initialising the optimiser with the optimised vacuum result of Ref. (Landreman & Paul 2022) to prevent the optimiser from falling back or staying at axisymmetry. Axisymmetry is the only global optimum for quasi-axisymmetry and thus a very attractive solution for the optimiser.

The optimisation of compact QA equilibria for hybrid designs faces the same problem. Initialising optimisations far from axisymmetry is not possible here as it is precisely this parameter space in which we want to find solutions. Figure

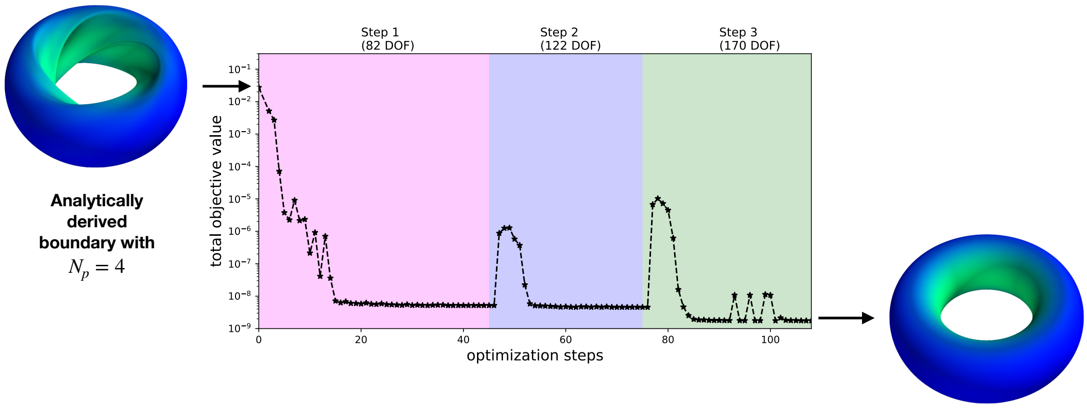

Figure 3.5.1. – A typical optimisation path without the use of $f_{\iota,\mathrm{ext}}$ when initialised with the $N_p = 4$ configuration of Refs. (Henneberg & Plunk 2024; Plunk 2020).

3.5.1 shows the results of an optimisation initialised with the analytically derived configuration of Refs. (Henneberg & Plunk 2024; Plunk 2020) for the target

$$f = f_{\mathrm{QA}} + (A - A^*)^2, \tag{3.5.1}$$

where A is the aspect ratio and was set to stay at the value of the analytic solution, that is we set the target value $A^* = 2.42$. The optimiser reduces the objective value by seven orders of magnitude within < 20 optimisation steps and the final result is essentially axisymmetric. Both from the final objective value and through inspection of the boundary modes of the final result it becomes apparent that the optimiser does not actually find axisymmetry to machine precision. This represents yet another manifestation of the various local minima present in the stellarator optimisation space.

3.6 A new target for optimisation in the vicinity of axisymmetry

To prevent the optimiser from falling back to axisymmetry a constraint must be added to the optimisation problem. Since previous stellarator optimisation for quasi-axisymmetry was done either in vacuum (Landreman & Paul 2022) or initialised with significant non-axisymmetry when allowing for finite beta (Landreman, Buller & Drevlak 2022), this problem of finite-beta optimisation near axisymmetry was not solved before. A method which we found to work well in practice to solve this problem is to estimate the external rotational transform with the help of the underlying axisymmetric equilibrium. For a general stellarator equilibrium the rotational transform can be formally decomposed as

$$\iota(s) = \iota_{\text{current}}(s) + \iota_{\text{ext}}(s), \tag{3.6.1}$$

where ι_{current} is the contribution from the plasma current and ι_{ext} the contribution from external currents, i.e. the coils. For a tokamak, of course, $\iota_{\text{ext}} = 0$. We estimate the external rotational transform of the stellarator by computing the `VMEC` equilibrium of the underlying tokamak at each step during the optimisation. We define the underlying tokamak to be the one with the same current profile, pressure profile, and axisymmetric modes, but with all non-axisymmetric boundary modes set to zero. We then estimate the external rotational transform through

$$\iota_{\text{ext,proxy}}(s) = \iota_{\text{stellarator}}(s) - \iota_{\text{tokamak}}(s). \tag{3.6.2}$$

To give the profile of $\iota_{\text{ext,proxy}}(s)$ the freedom to evolve to its naturally preferred evolution we constrain only its radial average

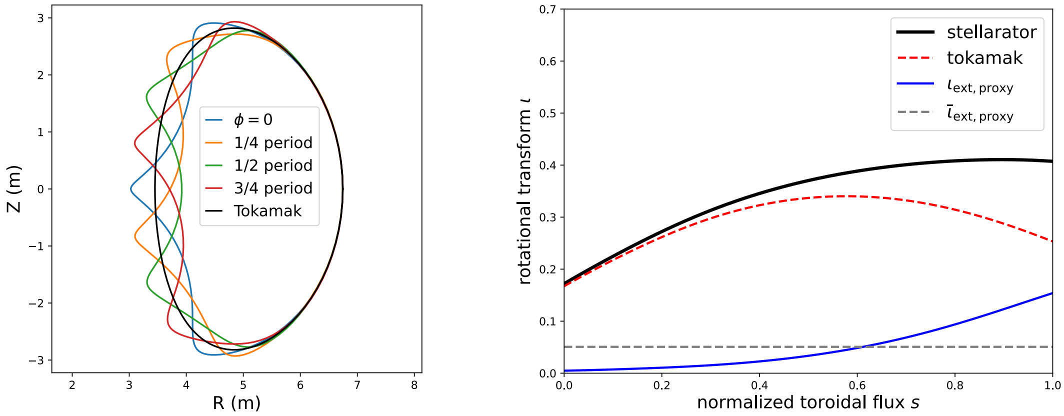

Figure 3.6.1. – The construction of $\bar{\iota}_{\text{ext,proxy}}$ for the $N_p = 4$ configuration of Refs. (Henneberg & Plunk 2024; Plunk 2020) which was already shown in figure 3.1.1. Left: The cross-sections of the stellarator equilibrium and the underlying tokamak. Right: The resulting rotational transform profiles.

$$\bar{\iota}_{\text{ext,proxy}} = \int_0^1 \iota_{\text{ext,proxy}}(s)ds. \tag{3.6.3}$$

The construction of this proxy is illustrated in figure 3.6.1. By formulating the new target as

$$f = f_{\text{QA}} + (A - A^*)^2 + f_{\iota,\text{ext}} \tag{3.6.4}$$

with

$$f_{\iota,\text{ext}} = w_\iota * \max[\bar{\iota}^*_{\text{ext,proxy}} - \bar{\iota}_{\text{ext,proxy}}, 0]^2 \tag{3.6.5}$$

we ensure that $\bar{\iota}_{\text{ext,proxy}} \geq \bar{\iota}^*_{\text{ext,proxy}}$, where $\bar{\iota}^*_{\text{ext,proxy}}$ is a user-specified lower bound. Choosing the same value as for the analytically derived configuration yields the optimisation path shown in figure 3.6.2. The improvement in quasi-axisymmetry which is achieved by the numerical optimisation is substantial. This is shown in figure 3.6.3 which shows the contours of B in Boozer coordinates at four radial locations $s = 0.05, 0.25, 0.5, 1.0$. For perfect QA these contours would be horizontally straight lines.

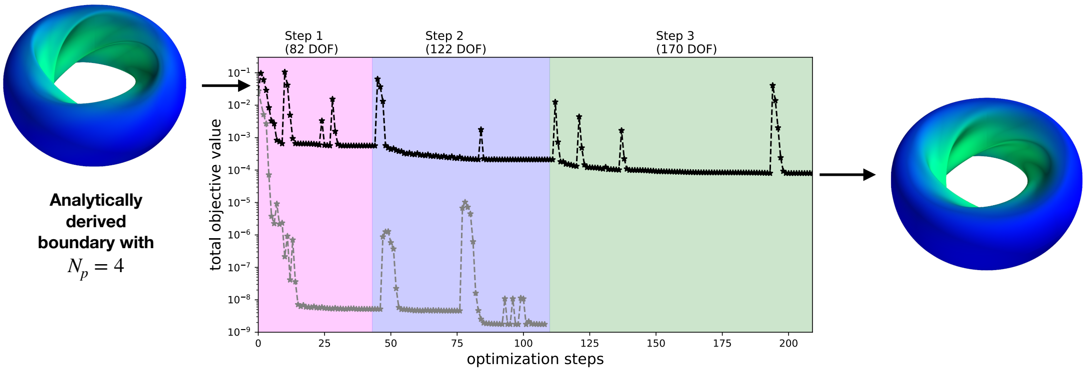

Figure 3.6.2. – The optimisation path with the use of $f_{\iota,\text{ext}}$. The optimisation without this target, which was shown before in figure 3.5.1, is shown again here in gray.

3.6.1 Optimised equilibria of varying perturbation

Configurations of varying non-axisymmetry can be obtained by setting different values for $\bar{\iota}^*_{\text{ext,proxy}}$. Figure 3.6.4 shows additional four field-period configurations ($N_p = 4$) where $\bar{\iota}^*_{\text{ext,proxy}}$ has been set to twice and four times the value of the analytically derived configuration of Refs. (Henneberg & Plunk 2024; Plunk 2020). The middle row, which shows the poloidal cross-sections at different toroidal locations within one field period, shows that the effect of setting higher values of the minimum external rotational transform increases the strength of the inboard perturbations, as opposed to for example introducing perturbations on the outboard side. The outboard side remains axisymmetric even at strong perturbations. This observation will be further addressed in section 3.8.

The bottom row in figure 3.6.4 shows the radial profiles of the rotational transform of the stellarator and the underlying tokamak, highlighting the details of how the optimiser chose to increase the mean external rotational transform. We observe that most of the external ι is accumulated in the outer half of the volume. This can be understood intuitively from the fact that the perturbations of the plasma boundary are created by high poloidal Fourier mode numbers. The magnetic field amplitude falls off faster than lower mode numbers and thus a

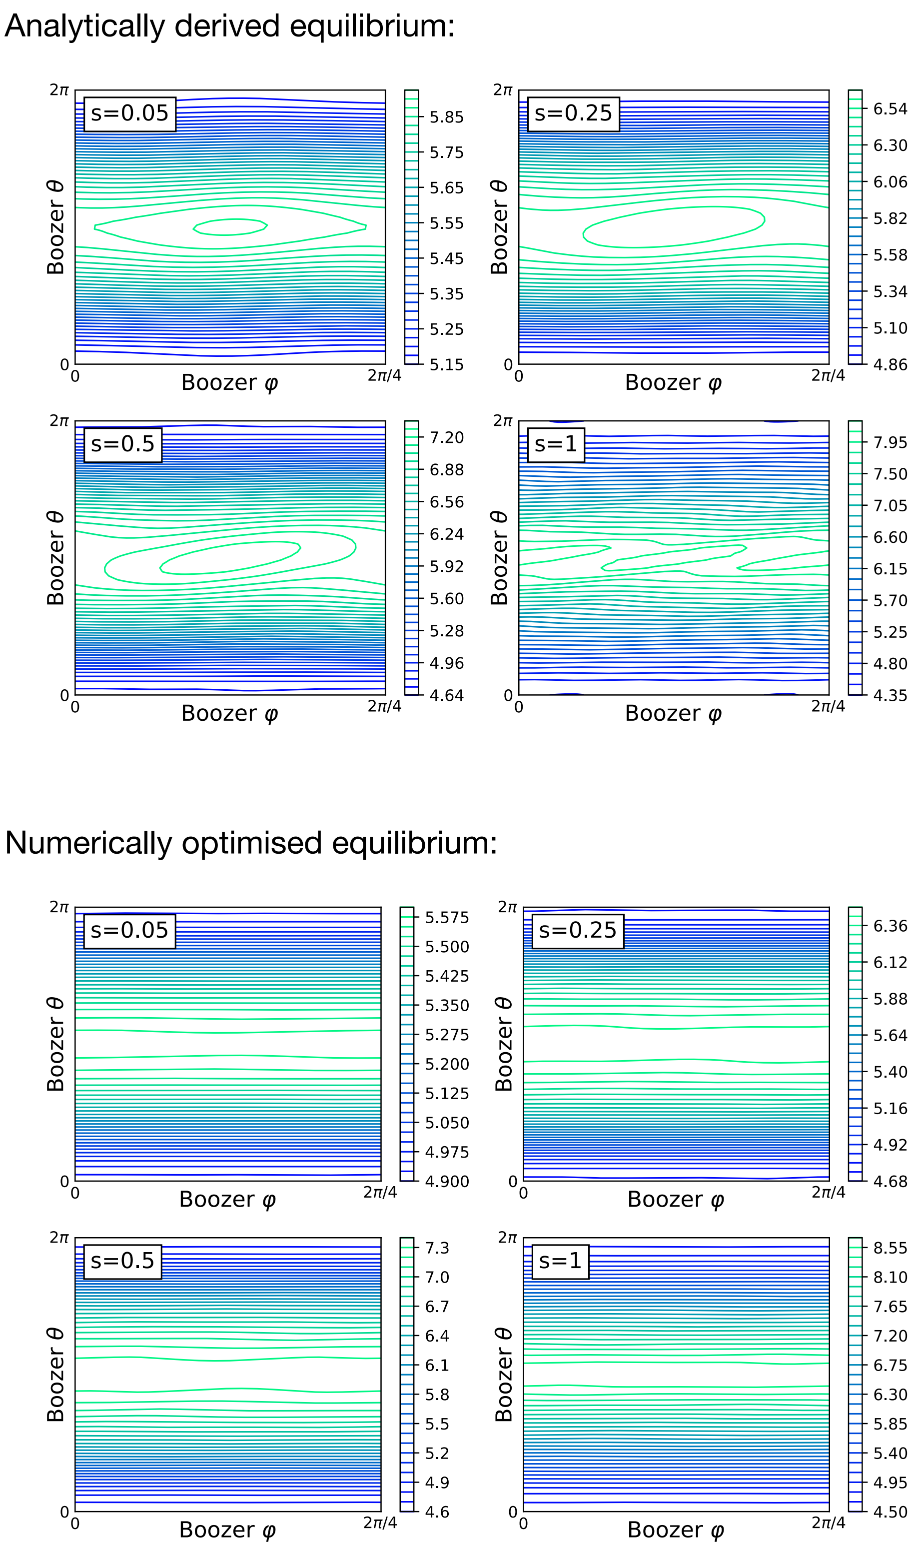

Figure 3.6.3. – Contours of the magnetic field strength in Boozer coordinates at four radial locations. **Top:** the analytically derived $N_p = 4$ solution of Refs. (Henneberg & Plunk 2024; Plunk 2020). **Bottom:** the optimised boundary with the same estimated external rotational transform.

configuration with four field periods will have a much more axisymmetric core than a more conventional two-field-period QA stellarator.

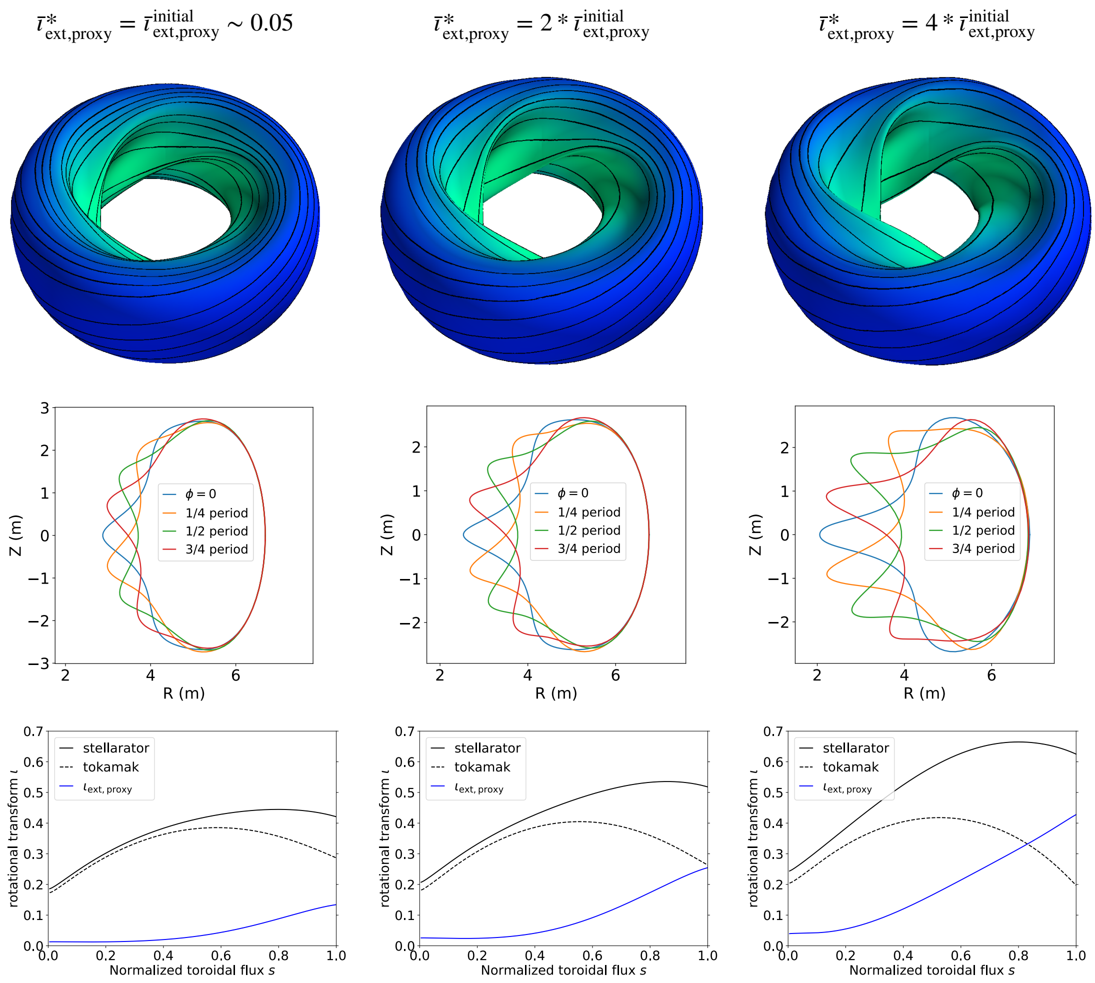

$\bar{\iota}^{*}_{\text{ext,proxy}} = \bar{\iota}^{\text{initial}}_{\text{ext,proxy}} \sim 0.05$

$\bar{\iota}^{*}_{\text{ext,proxy}} = 2 * \bar{\iota}^{\text{initial}}_{\text{ext,proxy}}$

$\bar{\iota}^{*}_{\text{ext,proxy}} = 4 * \bar{\iota}^{\text{initial}}_{\text{ext,proxy}}$

Figure 3.6.4. – Optimised QA equilibria of varying external rotational transform achieved by setting different values of $\bar{\iota}^{*}_{\text{ext,proxy}}$. **Top:** the final plasma boundary with the magnetic field strength overlaid where brighter colours correspond to stronger field. Black lines are field lines. **Middle:** the poloidal cross-sections at four toroidal locations within one of the four field periods. **Bottom:** the rotational transform profiles.

While not the optimal target during optimisation, the QA error metric defined in equation 3.4.5 can be used to assess the QA quality of the final optimised solution. It is worth noting that neither this Boozer-mode metric nor the so-called two-term metric from equation 3.4.6 are more QA than the other, see Ref. (Rodríguez, Paul & Bhattacharjee 2022).

The top plot in figure 3.6.5 shows the radial profiles of E_{QA} for the configurations of varying perturbations. For the same perturbation level, the QA error is reduced by about an order of magnitude across the radius. This reduction was already seen in the much straightened contours shown in figure 3.6.3. Even for much stronger perturbations with four times the external rotational transform, the QA error is still lower than the analytic initial condition.

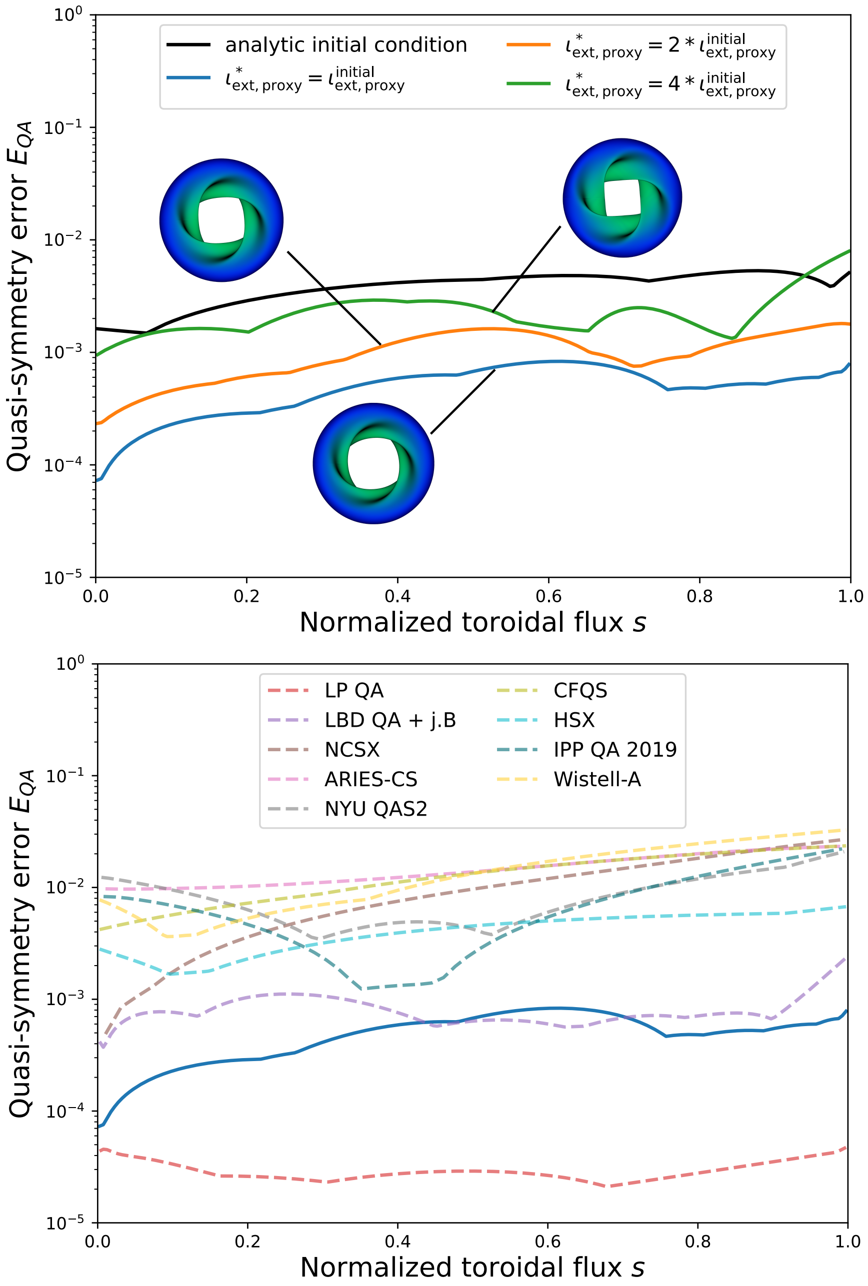

Figure 3.6.5. – The radial profiles of the QA error for the new configurations with four field periods. **Top:** Configurations with varying perturbation compared to the analytic solution, showing the substantial improvements achieved by the optimisation. **Bottom:** The configuration with $\iota^*_{\text{ext,proxy}} = \iota^{\text{initial}}_{\text{ext,proxy}}$ from the top plot in context with legacy configurations. The configuration *LBD QA + j.B* is the one from (Landreman, Buller & Drevlak 2022), all other configurations are taken from Ref. (Landreman & Paul 2022; Landreman 2021).

The bottom plot in figure 3.6.5 shows a comparison to legacy configurations and the new 'precise' QA configuration from Ref. (Landreman & Paul 2022) and the one with self-consistent bootstrap current from (Landreman, Buller & Drevlak 2022). While our configuration is not as QA as the 'precise' configuration, it has a substantially lower error than the legacy configurations. We will see later in section 3.11 that our QA error is sufficiently low for negligible alpha particle losses.

3.7 Self-consistent bootstrap current optimisation

While stellarators do not rely on the bootstrap current for creating rotational transform, it still arises naturally. In quasi-isodynamic stellarators, like Wendelstein 7-X, it is vanishingly small, while it is of substantial magnitude in quasi-helical stellarators and even more so in quasi-axisymmetric stellarators, which behave most like tokamaks.

When optimising finite-beta MHD equilibria some assumption for the toroidal current profile $I(s)$ must be made. So far we have chosen a simple quadratic profile which has some finite value on-axis, peaks off-axis, and goes to zero at the plasma boundary. Note that $I(s)$ is not the net toroidal current profile. However, the most natural choice would be the one which causes the implicitly assumed net parallel current profile to be that which would be expected from the bootstrap current for the chosen density and temperature profiles. Since any feasible tokamak reactor design should produce $\geq 75\%$ of its current from the bootstrap current (Freidberg 2008), creating an equilibrium which is fully bootstrap driven is a natural first design choice. Different to tokamaks, the current profile can impact the confinement properties, and must therefore be treated with care.

Since the parallel current profile is an output of the MHD equilibrium calculation and not an input, optimisation must be used to achieve self-consistency. This problem was solved in Ref. (Landreman, Buller & Drevlak 2022) where analytic tokamak formulae by Redl *et al.* (Redl et al. 2021) were adapted for quasi-symmetry based on the isomorphism between axisymmetry and quasi-symmetry. Here we use that target to integrate a self-consistent bootstrap current into our compact QA designs, i.e. we use the target

$$f_{\text{boot}} = w_{\text{boot}} * \frac{\int_0^1 ds[\langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{vmec}} - \langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{Redl}}]^2}{\int_0^1 ds[\langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{vmec}} + \langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{Redl}}]^2}, \tag{3.7.1}$$

where $\langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{vmec}}$ is the parallel current profile given as an output by `VMEC` and $\langle \mathbf{j}\cdot\mathbf{B}\rangle_{\text{Redl}}$ the expected bootstrap current profile for the current choice of plasma boundary and toroidal current profile.

In the optimisation the degrees of freedom are the values I_i' which represents the values of $dI(s)/ds$ at the equally spaced spline points s_i. The interpolation is done with `SCIPY`'s default cubic spline interpolation method. The optimisation space is thus expanded to be

$$\mathcal{M} = \{R_{m,n}, Z_{m,n}, \psi_a, I_i'\} \tag{3.7.2}$$

and the optimisation problem reads

$$\begin{aligned} \min_{x\in\mathcal{M}} f &= f_{\text{QS}} + (A - A^*)^2 + f_{\iota,\text{ext}} + f_{\text{boot}} + f_\beta + f_V + f_{\text{elongation}}, \\ &\text{where} \\ f_\beta &= w_\beta * \{\max\left[\langle\beta\rangle_{\min} - \langle\beta\rangle, 0\right] + \max\left[\langle\beta\rangle - \langle\beta\rangle_{\max}, 0\right]\}^2, \\ f_V &= w_V * \{\max\left[V_{\min} - V, 0\right] + \max\left[V - V_{\max}, 0\right]\}^2, \\ f_{\text{elongation}} &= w_\kappa * \max\left[\kappa - \kappa_{\max}, 0\right]^2 \end{aligned} \tag{3.7.3}$$

restrict the volume-averaged plasma beta to the range $4.5\% \leq \langle\beta\rangle \leq 5\%$[8] and the plasma volume to $\pm 20\%$ around the value $V = 444\ \text{m}^3$ of the ARIES-CS reactor study (Najmabadi et al. 2008). The elongation constraint restricts $\kappa \leq 2.0$. For cold-start optimizations (those initialised from axisymmetry) and optimizations with $N_p = 2$ we define the elongation to be the one of the underlying tokamak; for optimizations starting from the analytically derived four-field-period configuration we define the elongation as the ratio of the midplane width and the maximum height when taken at the half period cross-section. This constraint is only used to prevent pathological results, but is often not active in the final solutions and the details of it showed to have negligible impact on the optimization results.

The temperature and density profiles are chosen somewhat arbitrarily and represent typical modelling parameters that were also used in Ref. (Landreman, Buller & Drevlak 2022):

$$\begin{aligned} T_i(s) = T_e(s) &= T_0(1 - s), \\ n_i(s) = n_e(s) &= n_0(1 - s^5). \end{aligned} \tag{3.7.4}$$

It is worth noting that recent work by Buller *et al.* has shown that changes in the pressure profile which maintain the same plasma beta do not affect the QA quality substantially (Buller 2024).

Figure 3.7.1 (left) shows the parallel current profiles before and after the optimisation for a configuration optimised for $\iota^*_{\text{ext,proxy}} = 0.1$. For the optimised

---

[8] The plasma beta $\langle\beta\rangle$ typically increases during the optimizations such that the exact value of the lower bound is unimportant in practice.

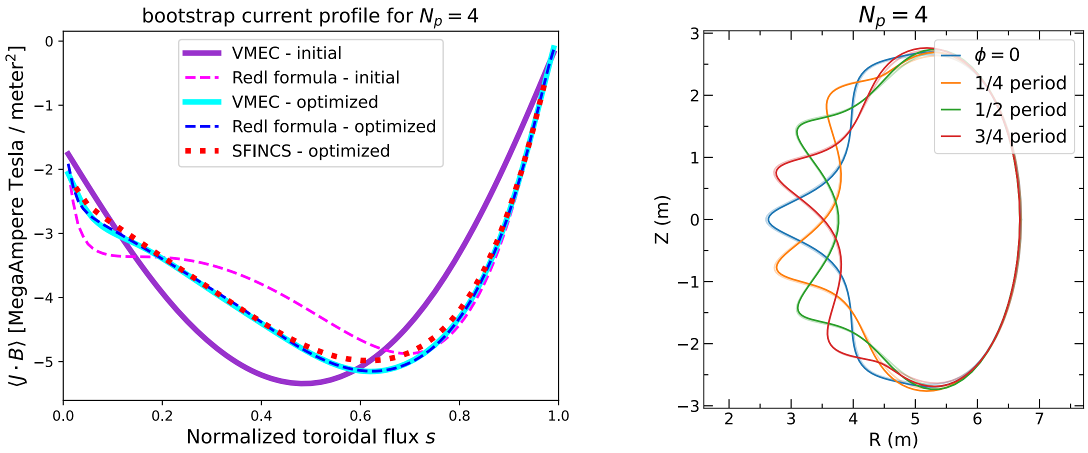

Figure 3.7.1. – A four-field-period configuration optimised for QA and self-consistent bootstrap current. **Left:** The parallel current profile before and after the optimisation. **Right:** The final plasma boundary in solid lines. In shaded lines is the plasma boundary of the purely QA optimised configuration, which served as the initial condition for this optimisation. Both are very similar and are almost on top of each other.

configuration the profiles from `VMEC` and the Redl-formulae are identical to the naked eye. The profile of the final result was also checked with the fully 3D code `SFINCS`, which accounts for imperfections in quasi-symmetry. This profile, shown in red in figure 3.7.1 (left), agrees well with the other two, corroborating the self-consistency.

Figure 3.7.1 (right) shows cross-sections of the optimised plasma boundary in solid lines and those of the purely QA optimised case in shaded lines, which served as the initial condition for this optimisation. Both are very similar, indicating that the self-consistency is mostly achieved through changes in the current profile alone.

3.8 Localisation of quasi-axisymmetric perturbations

All compact QA configurations shown thus far showed that perturbations accumulate on the inboard side and the outboard side is left mostly axisymmetric. This has some benefits, e.g. imperfections in quasi-symmetry also accumulate on the inboard side and therefore only affect passing and shallowly-trapped particles, and not those which are more deeply trapped (Henneberg & Plunk 2025). Nonetheless, one might wonder whether it would be possible to move perturbations to the outboard side instead. There would be some benefits, e.g. 3D coils might be easier to place on the outboard side.

This tendency for perturbations to accumulate on the inboard side was also noticed in the analytical studies by Plunk and Helander (Plunk & Helander 2018, Appendix E), where it was traced back to the local magnetostatics constraint

$$\oint d\theta \frac{\partial Z_0}{\partial \theta} R_0^{N_p^2-1} \hat{Z}_1 = 0, \tag{3.8.1}$$

where $R_0(\theta)$ and $Z_0(\theta)$ are the axisymmetric shaping and $\hat{Z}_1(\theta, \phi)$ the first order perturbation. This implies that for low aspect ratio, i.e. where the relative difference in R_0 between outboard and inboard side is large, or for higher numbers of field periods, perturbations $\hat{Z}_1$ must be much stronger on the inboard side than outboard side. Put more precisely, it implies that the difference in perturbation between outboard side and inboard side scales as (Henneberg & Plunk 2025)

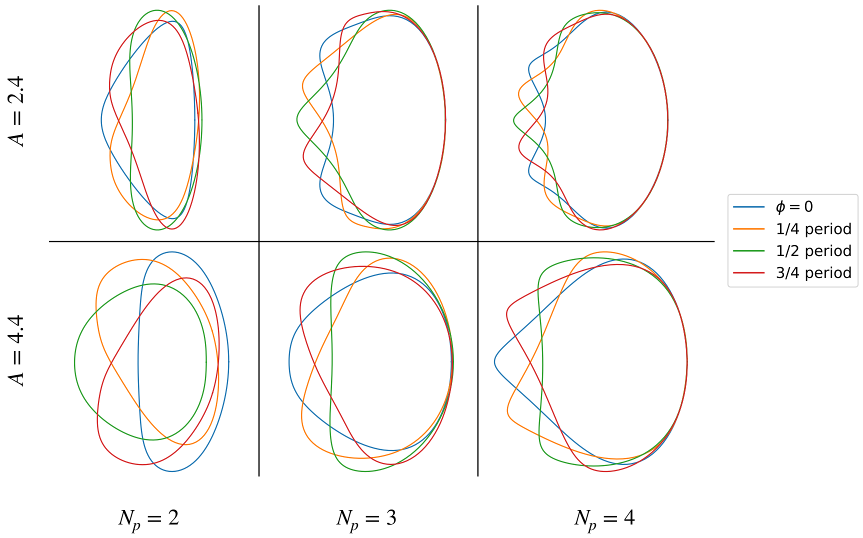

Figure 3.8.1. – The localisation of quasi-axisymmetry preserving perturbations. Shown are the cross-sections within one field period for optimised cases of varying N_p and A. The initial condition for all cases was a finite-beta axisymmetric torus with circular cross-section. Cylindrical coordinate axes are omitted to avoid clutter.

$$\left(\frac{1 + 1/A}{1 - 1/A} \right)^{N^2-1} . \tag{3.8.2}$$

Some indications of such a trend were also seen in results from adjoint optimisation of vacuum fields (Nies et al. 2024).

Here we want to explore whether this observation holds in optimisation as well, unconstrained by the analytical method (asymptotic expansion in axisymmetry perturbation) or the vacuum field approximation. We optimise configurations of $N_p = 2, 3, 4$ with aspect ratios $A = 2.4, 4.4$, and each initialised from axisymmetry and with the same objective function. This arguably most neutral initial condition allows the optimiser to go into any ‘perturbation-direction’ compared to a warm-start approach. The final optimised plasma boundaries are shown in figure 3.8.1. Towards lower aspect ratio and higher numbers of field periods, perturbations are more localised to the inboard side, inline with the expectations from analytical theory discussed above around equations 3.8.1 and 3.8.2. This corroborates that this is a fundamental feature of QA fields, and that configurations with perturbations on the outboard side are disadvantageous from a neoclassical confinement perspective.

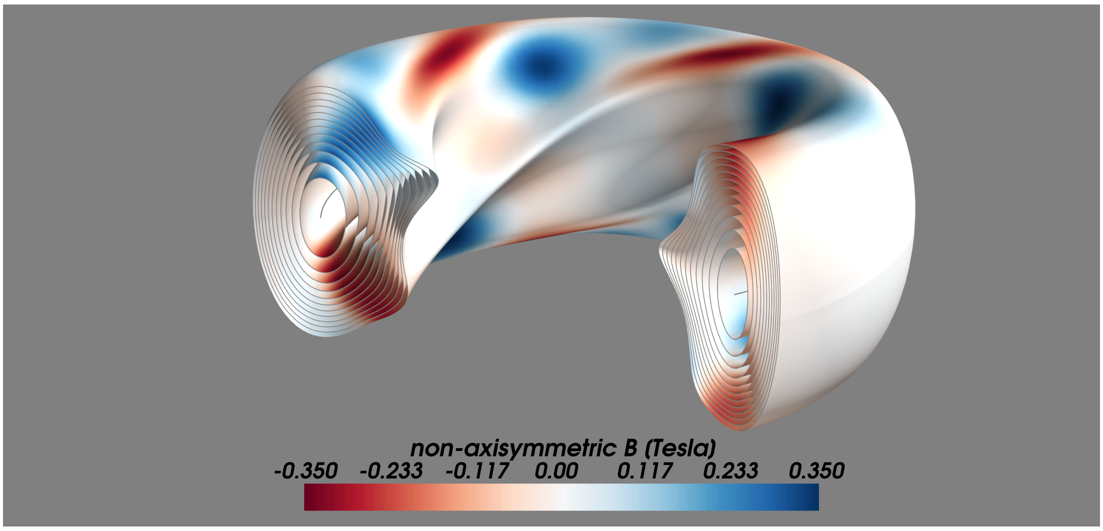

Figure 3.8.2. – The non-axisymmetric magnetic field component on the flux surfaces for $A = 2.4$ and $N_p = 4$. Like the 3D flux surface shaping, the non-axisymmetric field is confined to the inboard side.

Furthermore, figure 3.8.2 shows the non-axisymmetric magnetic field strength for the compact four-field-period configuration, highlighting that the outboard side is also mostly axisymmetric in the magnetic field structure.

3.9 Optimised equilibria of varying field periods

Most stellarators possess a unique discrete toroidal symmetry specified by the number of field periods N_p. The number of field periods states the integer amount of identical toroidal segments within the device such that for example $B(\theta, \phi) = B(\theta, \phi + 2\pi J/N_p)$ for any integer J. Wendelstein 7-X (Klinger et al. 2016) has five field periods and HSX (Anderson et al. 1995) has four. When a tokamak is perturbed to create a compact QA device this choice is not obvious and implies trade-offs. Further, due to its discrete nature it must be decided *a priori* for the optimisation. While we have so far shown examples limited to four field periods, our method is able to find similar configurations at other field periodicities as well.

Figure 3.9.1 shows examples of $N_p = 2, 3, 5$. The configuration with $N_p = 5$ is optimised for QA, while those with $N_p = 2, 3$ are optimised for QA, self-consistent bootstrap current, and Mercier stability. The details of how Mercier stability is included are discussed in the next section 3.10. The optimisations for three and five field periods were initialised with the analytically derived boundary for four field periods by simply taking out or adding one field period. The value of $\bar{\iota}^*_{\text{ext,proxy}}$ was targeted to be the same as the optimisation starting point. The design with two field periods was initialised with a circularly shaped tokamak and due to the lack of a reference point the external rotational transform was set to $\bar{\iota}^*_{\text{ext,proxy}} = 0.1$.

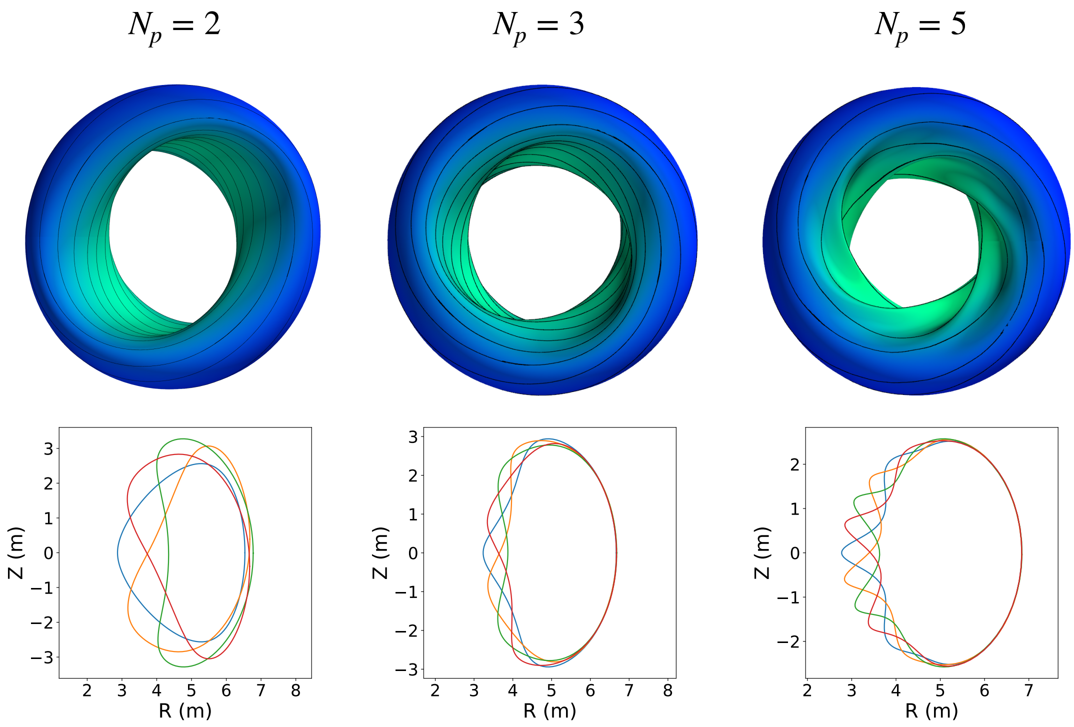

Figure 3.9.1. – optimised configurations with varying field periods. The top row shows the optimised plasma boundaries with the field strength overlaid, and black lines are field lines. The bottom row shows the cross-sections within one field period.

3.9.1 Trade-offs between field periods

The cross-sections in figure 3.9.1 already suggest that the two-field-period design is unlikely to be realisable as a stellarator-tokamak hybrid in the form designed by Henneberg and Plunk (Henneberg & Plunk 2024) due to the relatively global perturbations and the much reduced axisymmetric volume. In figure 3.9.2 the evolution of the axisymmetric volume fraction for the different number of field periods and varying perturbation is shown. Here we follow Henneberg and Plunk (Henneberg & Plunk 2025) and define the axisymmetric volume fraction as the volume of the stellarator compared to the largest axisymmetric volume that can fit inside the cross-sections, where we restrict ourselves to elliptical shaping for the latter. This estimate is somewhat conservative as it assumes that coils and therefore the vacuum vessel must be close to the plasma boundary everywhere.

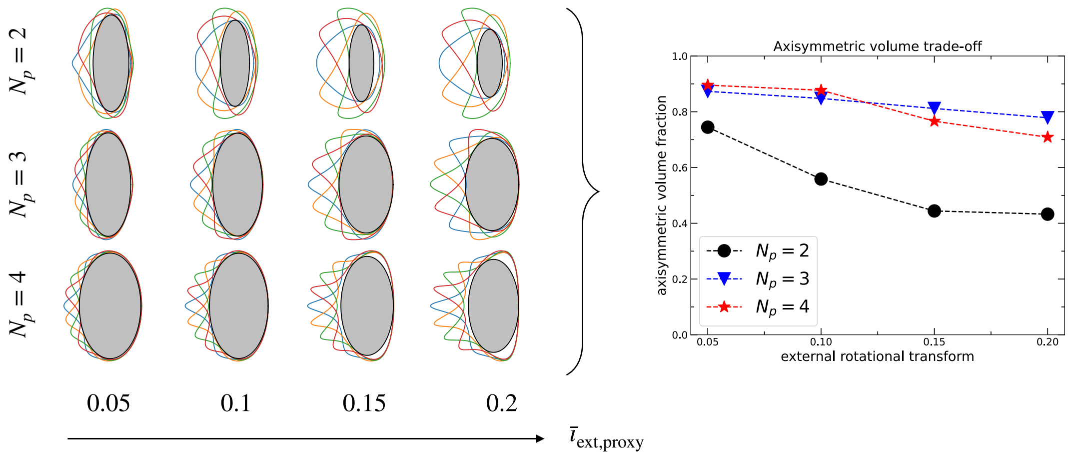

Figure 3.9.2. – The reduction in the axisymmetric volume fraction with increasing external rotational transform for different field periods. Left: The cross-section and the estimated axisymmetric volume as an elliptical torus shown in gray. Right: Summary of the fractions of the axisymmetric volume to the full stellarator volume.

Figure 3.9.2 shows that the two-field-period case stands out with its low axisymmetric volume, dropping below 50% for an external rotational transform of 0.15. The cases with three and four field periods perform much better, with $> 70\%$ axisymmetric volume up to the largest external rotational transform tested.

3.10 Mercier-stability optimisation

Another desirable feature which is not guaranteed is Mercier stability (Mercier 1962; Mercier 1964; Greene & Johnson 1962b; Greene & Johnson 1962a; Landreman & Jorge 2020; Bauer, Betancourt & Garabedian 2012). Not only is it important to find a good equilibrium, but further an equilibrium that is as resilient as possible to perturbations, i.e. one that is as *stable* as possible. Mercier stability is a sufficient criterion for stability against Mercier interchange modes, and a necessary criterion for stability against ballooning modes (Freidberg 2014). In tokamaks, Mercier stability is not usually the primary concern because current-driven instabilities can set a more stringent limit (Ramasamy et al. 2025; Zohm 2015). In stellarators the necessity of Mercier-stability is not clear. Experiments are sometimes able to operate above the critical plasma beta of the Mercier criterion (Freidberg 2014; Weller et al. 2006; Geiger et al. 2004). However, due to its unknown character in stellarators, Mercier-stability is typically included in more developed candidate designs and its trade-off with other desirable physics properties is of interest.

Since the QA perturbations tend to localise on the inboard side, i.e. the good curvature side, as corroborated in section 3.8, one might suspect that Mercier-stability could be hindered in such designs. Looking at the evolution of the Mercier coefficient for QA-optimised solutions of increasing perturbation, figure 3.10.1, one indeed observes an increasing unstable region for stronger perturbations.

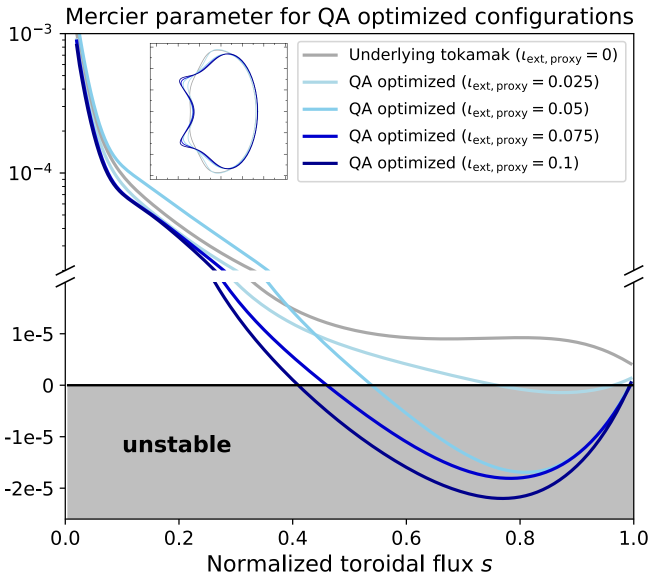

Figure 3.10.1. – The Mercier coefficient for (purely) QA-optimised cases with four field periods and of increasing perturbation.

To design compact QA configurations which are Mercier stable we add the target

$$f_{\mathrm{Merc}} = w_{\mathrm{Merc}} * \int_0^1 \mathrm{d}s \max\left[D^*_{\mathrm{Merc}} - D_{\mathrm{Merc}}(s), 0\right]^2, \tag{3.10.1}$$

where D_{Merc} is the Mercier coefficient as output by `VMEC` and D^*_{Merc} is some suitably chosen lower bound. The expression for D_{Merc} is lengthy and since we will not discuss its details we point to Ref. (Landreman & Jorge 2020) where it is stated and discussed around equation A.16. `VMEC`'s Mercier coefficient has previously been shown to be sensitive to the radial resolution (Landreman & Jorge 2020; Panici et al. 2023). However, as usually done, our optimisations are performed with a relatively low radial resolution of $N_s = 49$ flux surfaces. This

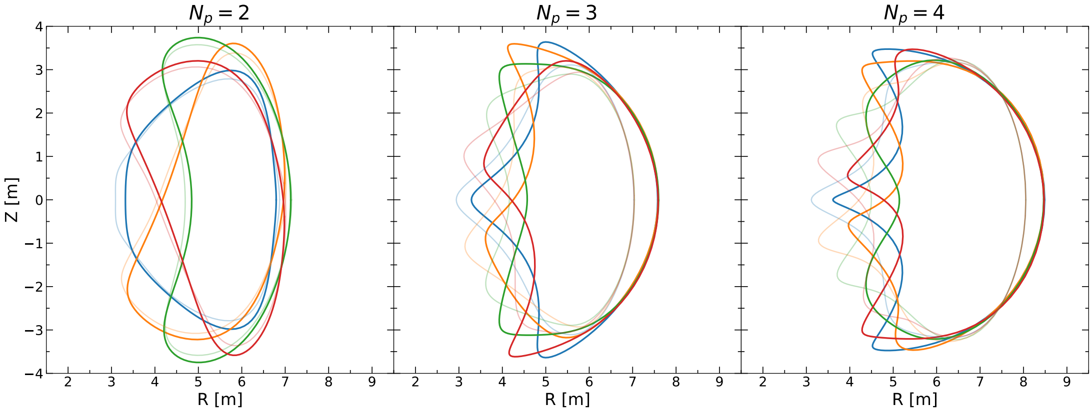

Figure 3.10.2. – The cross-sections within one field period. Solid lines are those of configurations optimised for QA, self-consistent bootstrap current, and Mercier stability. Shaded lines show those optimised purely for QA and with the same external rotational transform. **Left:** two field periods. **Center:** three field periods. **Right:** four field periods.

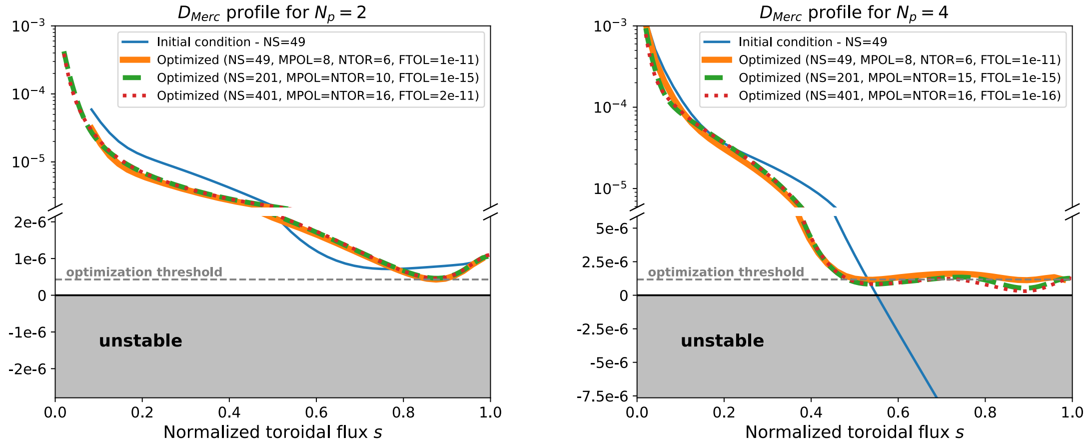

Figure 3.10.3. – The profile of the Mercier coefficient, showing stability. The optimised solution is Mercier-stable up to high radial resolution. **Left:** two field periods. **Right:** four field periods.

is motivated by computational cost and by the observation that low resolution equilibria are very good proxies of higher resolution ones for most metrics. To make sure that our equilibria are Mercier-stable up to high radial resolutions the necessary value of D^*_{Merc} was identified beforehand by doing optimisations with one set of hyper-parameters for a range of values of D^*_{Merc} and checking whether the resulting configurations are stable up to higher resolutions. The identified value was then used for an optimisation of a higher number of hyper-parameters and the best results are shown in figure 3.10.2. The cross-sections change visibly for $N_p > 2$ by developing a stronger bean-like cross-section and increasing perturbations towards the top and bottom of the configurations. The fact that there is some significant trade-off with QA quality will be reflected in the alpha particle confinement, which we discuss next.

3.11 Collisionless confinement of alpha particles

Quasi-symmetry serves as a proxy for the confinement of thermal trapped particles and energetic (non-thermal) particles, namely alpha particles produced by the fusion reactions. Perfect quasi-symmetry implies perfect alpha particle confinement, and stronger deviations increase the alpha particles losses. However, the exact details of the connection between QS and alpha particle confinement remain somewhat elusive. A study by Henneberg *et al.* highlighted that the details of the deviation matter and can be chosen favourably (Henneberg, Drevlak & Helander 2019).

Here we test the confinement of collisionless alpha particle trajectories using the code `SIMPLE` (Albert, Kasilov & Kernbichler 2020a; Albert, Kasilov & Kernbichler 2020b) for our optimised configurations to see whether the QA error is sufficiently low for these designs to be considered for serious applications. For this we launch 5000 alpha particles on the $s = 0.25$ flux surface and follow them for $250\,\mathrm{ms}$. After this time they are either thermalised or lost. We record the final fraction of lost particles. While this choice of flux surface, $s = 0.25$, is somewhat arbitrary, it is commonly used in the literature (Landreman & Paul 2022; Goodman et al. 2023; Bindel, Landreman & Padidar 2023a), and therefore a useful choice to aid comparison. A radial distribution of alpha particles that is consistent with the fusion reaction rate for the assumed temperature and density profiles would be a possible extension of this analysis, but the results are not expected to change significantly [see figure 6 in (Bindel, Landreman & Padidar 2023a) which shows a comparison of both approaches]. The confine-

N_p	ID	Configuration	$\iota_{\mathrm{ext}}^{\mathrm{proxy}}$	Init	a-sc [%]	V-sc [%]	Fig.	Ref.
2	2.1	QA	0.1	Tok	0.22	0.10	3.10.2	[SH25]
2	2.2	QA-j.B	0.1	Tr. 2.1	0.10	No loss	–	[SH25]
2	2.3	QA-j.B-Merc	0.1	Tr. 2.1	0.62	No loss	3.9.1	[SH24]
3	3.1	QA	0.1	Tok	0.44	No loss	3.10.2	[SH25]
3	3.2	QA-j.B-Merc	0.05	Tr. 3.1	0.34	No loss	3.9.1	[SH24]
3	3.2	QA-j.B-Merc	0.1	Tr. 3.1	3.86	1.14	3.10.2	[SH25]
4	4.0	Analytical	0.05	–	0.52	0.08	3.1.1	[HP24]
4	4.1	QA	0.05	Tr. 4.0	No loss	No loss	3.6.4	[SH24]
4	4.2	QA	0.1	Tr. 4.0	0.46	No loss	3.6.4	[SH24]
4	4.2	QA	0.2	Tr. 4.0	5.06	0.80	3.6.4	[SH24]
4	4.3	QA-j.B	0.1	Tr. 4.0	3.46	No loss	3.7.1	[SH25]
4	4.4	QA-j.B-Merc	0.1	Tr. 4.0	9.06	3.04	3.10.2	[SH25]
5	5.0	QA	0.05	Tr. 4.0	0.04	No loss	3.9.1	[SH24]

Table 3.1. – Alpha losses for compact QA configurations of this chapter. All configurations have aspect ratio $A \sim 2.4$. a*-sc* refers to scaling to the same minor radius as ARIES-CS, and V*-sc* refers to scaling to the volume (see text for more details). In the last column [SH24] refers to (Schuett & Henneberg 2024b), [SH25] to (Schuett & Henneberg 2025), and [HP24] to (Henneberg & Plunk 2024). In the third column *j.B* refers to self-consistent bootstrap current, and *Merc* to Mercier stability. In the fifth column *Tr.* refers to Fourier-mode truncated initialisation.

ment of trapped thermal particles is quantified with the so-called effective ripple ϵ_{eff}, which encapsulates the geometry dependent contribution to neoclassical transport in the stellarator-specific $1/\nu$ regime (Nemov et al. 1999; Helander et al. 2012). However, once alpha particle losses are sufficiently low, neoclassical transport is typically in the range where it is subdominant to turbulent transport. For this reason we only check the alpha particle confinement for our configurations and not ϵ_{eff}. We will see that we can expect ϵ_{eff} to be sufficiently low.

In order to make a fair comparison of alpha particle confinement at reactor scale between configurations, it is important that they be scaled to the same length scale and magnetic field, since the confinement of any given configuration can always be improved by simply scaling it up in size. A common reference point for a reactor scale in the literature is the one of the ARIES-CS design (Najmabadi et al. 2008), with a minor radius of $a = 1.7\mathrm{m}$, volume of $V = 444\mathrm{m}^3$, and volume-averaged magnetic field of $\langle B \rangle = 5.86\ \mathrm{T}$. We therefore scale the compact QA configurations to this magnetic field and length scale to test their confinement of alpha particles. While scaling configurations to the same minor radius is motivated by a better physics comparison, it is the volume which sets the length scale for a fusion reactor, since both the fusion power and cost scale with the volume. We provide results for both length scalings to ease comparison with existing literature (Landreman, Buller & Drevlak 2022; Bader et al. 2021), and refer to them as 'a-sc' and 'V-sc'.

Table 3.1 shows the final loss fractions for the configurations discussed in this chapter. For most of the configurations the alpha particle losses are negligible and significantly lower than those of any of the current stellarator experiments, see for example figure 17 in Ref. (Landreman, Buller & Drevlak 2022). Note that Wendelstein 7-X, the largest stellarator experiment, loses $> 5\%$ of alpha particles (a-scaling) with an aspect ratio $A = 10$. *This excellent performance of the compact QA configurations, despite the added physics criteria and the compact aspect ratio, further corroborates the potential of this design space.* Note that performance can usually be improved by increasing the aspect ratio[9], where configurations can be kept suitable for a stellarator-tokamak hybrid by increasing the number of field periods, using the insight from section 3.8 and Ref. (Plunk & Helander 2018).

3.12 Coil solutions

Here we briefly show that our optimised plasma boundaries are indeed realisable with 3D electromagnetic coils.

3.12.1 Henneberg modular and banana coil designs for three field period configuration

Figure 3.12.1 shows two optimised coil sets designed by Sophia Henneberg (Henneberg 2024) for a three-field-period configuration with QA, self-consistent bootstrap current, and Mercier stability – similar to the configuration from figure 3.9.1. Both a modular coil set and a 'banana' coil set as in figure 3.1.1 is shown.

The modular coil set consists of only three unique coils, resulting in a total of 18 coils due to the field periodicity (factor 3) and stellarator symmetry (factor 2). The banana coil set only features one type of QA coil, as well as standard tokamak toroidal field (TF) and poloidal field (PF) coils which were omitted in the figure for visual clarity.

For the modular coil set the field error, $\mathbf{B}\cdot\mathbf{n}$ on the plasma boundary, takes on a mean of 0.2% and a maximum value of 0.7%. At a reactor volume of $1900\,\mathrm{m}^3$ the

---

[9] Quasisymmetry is generally easier to achieve the larger the aspect ratio, see for example figure 4 in Ref. (Bindel, Landreman & Padidar 2023b).

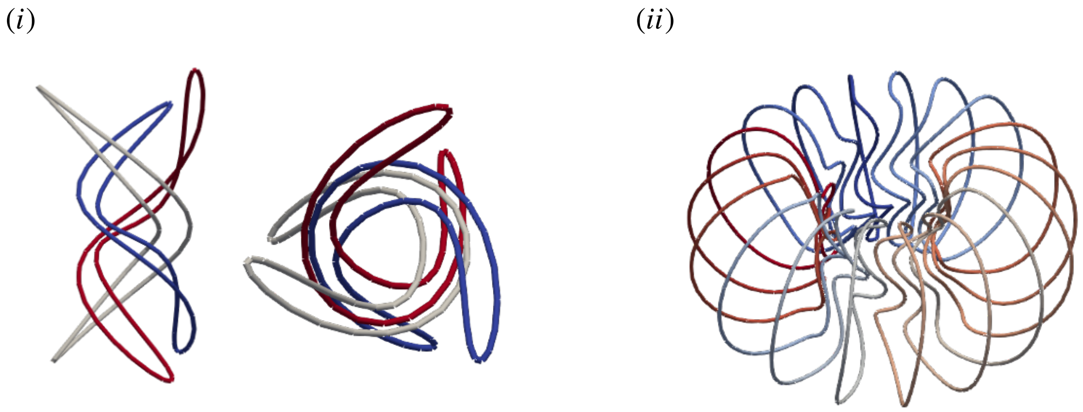

Figure 3.12.1. – Coil designs for three field period configuration designed by Sophia Henneberg. **Left:** banana coil set from two perspectives, with TF and PF coils omitted from figure. **Right:** Modular coil set, with PF coils omitted.

plasma-coil distance is $1.21\,\mathrm{m}$ and the coils have a minimum radius of curvature of $73\,\mathrm{cm}$.

For the banana coil set the field error takes on a mean of 0.7% and maximum of 3%, with a plasma-coil distance of $1.75\,\mathrm{m}$ and a minimum radius of curvature of $1\,\mathrm{m}$.

The values for the field error are better than those of Wendelstein 7-X, the plasma-coil distance at this reactor scale is sufficient to provide space for a breeding blanket, and from an engineering perspective these coil sets appear much easier to build than those of other similarly compact QA configurations published in the past.

3.12.2 Kaptanoglu coil designs for two field period configuration

While the $N_p = 2$ configuration from figure 3.9.1 is unlikely to work as a stellarator-tokamak hybrid, it is still an interesting 'stand-alone' stellarator equilibrium and has been used in the works by Kaptanoglu *et al.* for the design of a coil set consisting of modular coils and dipole coils which together support this equilibrium well. This has been done for active dipole coils with driven current as well as for passive coils with induced currents. See references (Kaptanoglu et al. 2025) and (Kaptanoglu, Landreman & Zarnstorff 2025), respectively. The coil solutions from both papers are also shown here for convenience in figure 3.12.2.

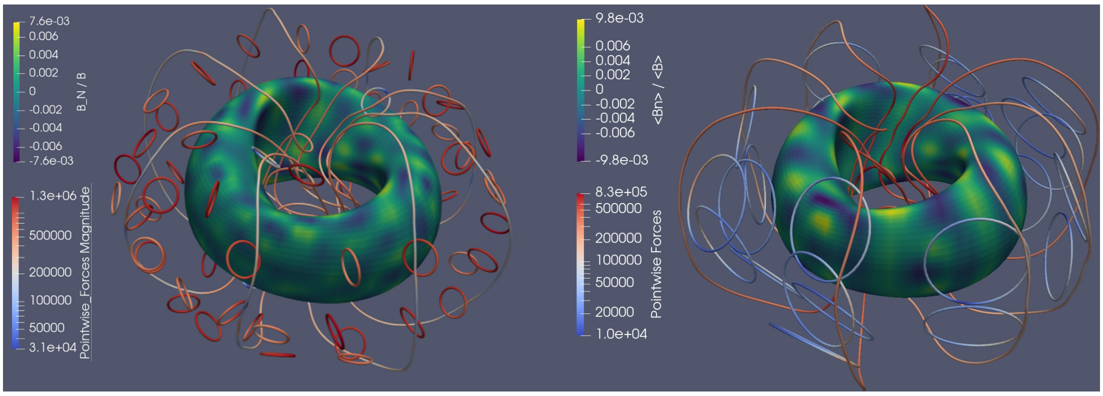

Figure 3.12.2. – The coil designs by Kaptanoglu *et al.* for the two-field-period configuration from figure 3.9.1. **Left:** Coil set solution with active dipoles, reprinted from (Kaptanoglu et al. 2025). **Right:** Coil set solution with passive dipoles, reprinted from (Kaptanoglu, Landreman & Zarnstorff 2025).

3.13 Discussion and conclusion

We have shown, via optimisation, that tokamaks can be perturbed three-dimensionally to a degree that offers stellarator benefits, i.e. significant external rotational transform (see section 1.6.2), while preserving quasi-symmetry surprisingly well. A critical piece here was the new optimisation target for the external rotational transform to prevent the optimiser from falling back to a tokamak. Further, we demonstrated that additional features such as a bootstrap-current profile and enhanced MHD stability can be included in these designs. These results corroborate the stellarator-tokamak hybrid concept (Henneberg & Plunk 2024) as a promising fusion device.

Moreover, this new parameter space is perfectly suited to investigate stellarator effects on tokamak dynamics *step-by-step*. Given that turbulence will dictate confinement in this device, one wonders what the turbulence properties are and how they fit in compared to tokamaks and well-known stellarators. It has long been difficult to compare turbulence dynamics between tokamaks and stellarators and this parameter space offers a promising pathway towards answering questions in this space. This, as well as the turbulence-flow interaction, will be the topic of the next two chapters.

3.13.1 Future work

Given the coil concepts from Refs. (Henneberg & Plunk 2024; Henneberg & Plunk 2025) and the results presented here, the primary next step is to investigate whether existing tokamaks can be upgraded towards such a stellarator-tokamak hybrid configuration. The methods presented in this chapter play a central role in such an investigation which is currently underway.

There are additional fundamental questions such as investigation and optimisation of global MHD stability. However, since the configurations feature significant axisymmetric shaping one might suspect that this shaping dominates their global MHD stability also. Moreover, the possible divertor solutions should be investigated. Collaborations aiming to explore these avenues have been started. Regarding the optimisation methodology, it might prove useful to move to combined plasma-coil optimisation (single-stage optimisation) to allow for an integrated approach of plasma optimisation and engineering constraints.

Finally, it would be good to extend the methods to the optimisation of equilibria that break stellarator symmetry. In the axisymmetric limit, stellarator symmetry corresponds to up-down symmetry,

$$\left.\begin{aligned} R(\theta,\phi) &= R(-\theta,-\phi) \\ Z(\theta,\phi) &= -Z(-\theta,-\phi) \end{aligned}\right\} \xrightarrow{\text{axisymmetry}} \begin{cases} R(\theta) = R(-\theta) \\ Z(\theta) = -Z(-\theta) \end{cases} \tag{3.13.1}$$

which is usually broken in tokamaks at least to some degree. Such an extension requires some extra work since most stellarator tools do not yet support non-stellarator symmetric equilibria. Initial work has been started for `SIMSOPT` support. Another reason this extension might be interesting, is that non-stellarator symmetry is required for intrinsic momentum transport, i.e. flow generation, and such support would allow one to investigate these possibilities in a quasi-symmetric stellarator.

3.13.2 Key findings and implications

The key points of this chapter and its implications are:

- Using optimisation, strong three-dimensional perturbations to tokamaks can be tailored to obey quasi-symmetry to a sufficiently high degree to keep alpha particles confined, while satisfying other desirable physics properties.
- Novel stellarator-tokamak hybrid equilibria have been created, enabling development of improved hybrid devices.
- The presented methods enable computational studies regarding the effect of confinement-preserving 3D perturbations on the physical behaviour of the plasma, such as e.g. turbulence and turbulence-flow interaction.

4

Ion temperature gradient turbulence in compact quasi-axisymmetric stellarators

In this chapter we will explore how the transport from ion temperature gradient (ITG) turbulence behaves generally in the new compact quasi-axisymmetric equilibria which we developed in the last chapter. The unique continuous deformation from axisymmetry of this configuration space allows us to study the effect of the three dimensional shaping in isolation. After having shed some light on the transport properties, we will then move onto the zonal flow drive in non-axisymmetric configurations in the next chapter.

Much of the content in this chapter was presented as:

> T. M. Schuett, S. A. Henneberg, D. Dickinson, and I. Cziegler (2025). Ion temperature gradient turbulence in compact quasi-axisymmetric stellarators [Poster]. *Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting, New York, USA*

Acknowledgements We are grateful for helpful conversations with Per Helander, Gabriel Plunk, Alan Goodman, Gareth Roberg-Clark, Paul Costello and Joaquim Loizu.

4.1 Motivation

Before we embark to gain insight into the zonal flow drive distribution as tokamaks are perturbed towards compact QA configurations, it is worth understanding how ITG turbulence generally behaves in this parameter space. Since this stellarator parameter space was previously unexplored, their turbulent properties are not known either. Now that we have gained some insight into which shaping tends to arise at varying number of field periods, we are able to perform a targeted study of the ITG stability and nonlinear heat flux and how it and its implied confinement compares to other well-known stellarators. This shall provide us with the necessary context before moving onto the ZF drive details.

As previously mentioned, the MHD equilibrium is described by the plasma boundary shape, the pressure profile, and by either the toroidal current profile or rotational transform profile. These MHD input parameters determine the magnetic geometry in the entire volume. For a chosen flux tube there are seven[1] specific geometric quantities which enter the gyrokinetic equation and thus determine the resulting heat flux. These depend only on the field-line-following coordinate l and are given by

$$\mathcal{G} = \{B, |\nabla\alpha|^2, |\nabla\psi|^2, \nabla\alpha\cdot\nabla\psi, \mathbf{b}\times\boldsymbol{\kappa}\cdot\nabla\alpha, \mathbf{b}\times\nabla\mathbf{b}\cdot\nabla\alpha, \mathbf{b}\times\boldsymbol{\kappa}\cdot\nabla\psi\}. \quad (4.1.1)$$

The MHD input parameters therefore determine, albeit indirectly through $\mathcal{G}$, the nonlinear turbulent heat flux. In this study we keep the pressure profile for the equilibrium construction fixed. The remaining MHD inputs can be grouped into the following three classes:

1. The axisymmetric shaping, specified through $\{R_{m,n=0}, Z_{m,n=0}\}$, and its known influence on axisymmetric gyrokinetic dependencies (Belli, Hammett & Dorland 2008; Angelino et al. 2009).

2. The current profile or rotational transform profile, which can directly set the (global) magnetic shear, and indirectly influences the Shafranov shift and thus all other geometric inputs to the gyrokinetic equation.

3. The non-axisymmetric shaping specified through $\{R_{m,n\neq 0}, Z_{m,n\neq 0}\}$.

For the work in this section we use the code `stella` (Barnes, Parra & Landreman 2019) for the linear calculations and `GX` (Mandell et al. 2024) for the nonlinear ones. The choice of `GX` is motivated by computational speed-up. Each simulation takes $< 30\,\mathrm{min}$ on one of the NVIDIA A40s on the Viking2 cluster at York. This allows us to perform more simulations. *The simulations are electrostatic and use adiabatic electrons.*

---

[1] It can be only six or eight depending on whether the equilibrium is electrostatic and on the choice of the parallel coordinate.

4.2 A comment on normalisations and scaling of configurations

For the nonlinear simulations we are particularly interested in the time-averaged ion heat flux which is given in units of gyro-Bohm

$$Q_{i,\mathrm{gB,ref}} = \rho_i^2 n_i T_i v_{\mathrm{th},i}/a^2 \sim 1/(a_{\mathrm{ref}} B_{\mathrm{ref}})^2 \tag{4.2.1}$$

with the ion Larmor radius ρ_i and thermal speed $v_{\mathrm{th},i}$. The ion density n_i, temperature T_i, and mass m_i are all set to unity for each simulation and define the reference value. The values of a_{ref} and B_{ref}, however, can change from one simulation to another. In `GX` the minor radius output from `VMEC` is used for a_{ref}, and B_{ref} is set based on the enclosed toroidal flux ϕ_{edge}, which is an input to `VMEC`, through

$$B_{\mathrm{ref}} = |\phi_{\mathrm{edge}}|/(\pi a_{\mathrm{ref}}^2). \tag{4.2.2}$$

One could compare different stellarator configurations without applying any scaling such that each configuration comes with their own gyro-Bohm units. When comparing configurations, one would then compare how each configuration is performing with respect to their own gyro-Bohm reference value. However, to draw the conclusion that this would translate to the same differences at a given reactor scale one would have to rely on the gyro-Bohm scaling, which can be broken (Dudding et al. 2022).

To avoid this, we scale configurations to the same a_{ref} and ϕ_{edge} before simulating and comparing the turbulent transport in them. For comparisons within one configuration family, as we will do in section 4.3, we choose the reference a_{ref} and B_{ref} to be those of the originally optimised configuration. For comparisons between configuration families, i.e. between different number of field periods, as we will do in section 4.5 and different types of stellarators as in section 4.6, we choose values based on the ARIES-CS minor radius of $a_{\mathrm{ARIES}} = 1.7\,\mathrm{m}$ and volume-averaged magnetic field of $\langle B\rangle_{\mathrm{ARIES}} = 5.86\,\mathrm{T}$. This is the same reference scale that was selected for the alpha particle confinement analysis in section 3.11. We then define our reference toroidal flux as

$$\phi_{\mathrm{edge,ref}} = \pi a_{\mathrm{ARIES}}^2 \langle B\rangle_{\mathrm{ARIES}}. \tag{4.2.3}$$

4.3 The connection to axisymmetry

One can think of different approaches to making the connection between the compact QA configurations and axisymmetry. These generally trade off realism for cleaner isolated treatment of QA-shaping. One realistic approach would be to study the free-boundary equilibria which result from varying currents in specific 3D coils, such as the banana coils by (Henneberg & Plunk 2024), which support the equilibrium alongside separate axisymmetric coils. If instead one wants to stay within the framework of *fixed-boundary* MHD problems one could alternatively optimise equilibria of varying values of external rotational transform. However, in our experience, this requires the axisymmetric boundary modes to be free parameters in the optimisation as well, to avoid the optimiser from getting stuck in local minima. Thus both of these methods would vary the axisymmetric shaping which can strongly affect the gyrokinetic stability and nonlinear heat flux (Belli, Hammett & Dorland 2008; Angelino et al. 2009).

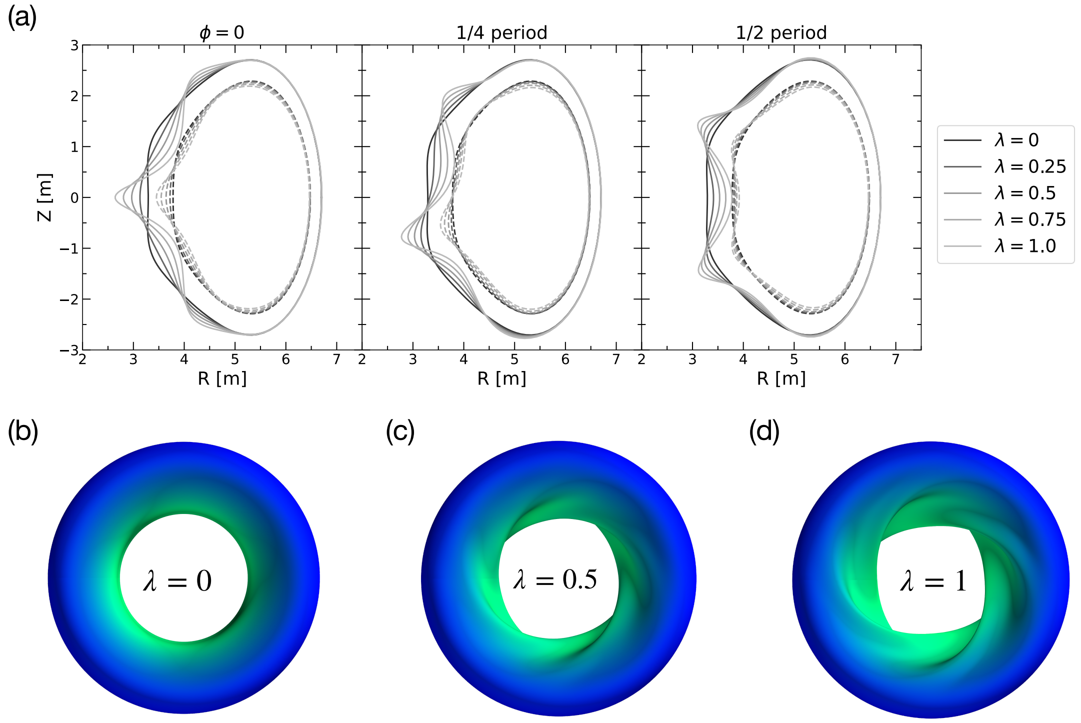

Figure 4.3.1. – Interpolated plasma boundaries to make the connection to axisymmetry. **(a):** The cross-sections at three different unique toroidal sections within one field period. Solid lines represent the plasma boundary $s = 1$, whereas dashed lines represent $s = 0.64$, the flux surface selected for the gyrokinetic analysis. **(b)-(d):** Top view of the plasma axisymmetric plasma boundary for $\lambda = 0$, $\lambda = 0.5$, and the original optimised configuration with $\lambda = 1.0$.

To avoid this variation in the axisymmetric shaping we opt for a third approach and create interpolated equilibria between the optimised QA configuration and its underlying tokamak by scaling only the non-axisymmetric boundary modes. To create the interpolated equilibria of varying external rotational transform we first recall the boundary representation from equation 3.3.1, and then scale the non-axisymmetric Fourier modes according to

$$\tilde{R}_{m,n\neq0} = \lambda R_{m,n\neq0} \quad \text{and} \quad \tilde{Z}_{m,n\neq0} = \lambda Z_{m,n\neq0}, \tag{4.3.1}$$

where $\lambda \in [0, 1]$ is the scaling factor. $\lambda = 1$ is the original configuration and we refer to the configuration with $\lambda = 0$ as the underlying tokamak.

The boundaries of the resulting equilibria are shown in figure 4.3.1, which shows that this approach keeps the axisymmetric shaping the same between all equilibria.

4.3.1 Maintaining a constant rotational transform profile

One way in which the choice of the current or rotational transform profile can directly influence gyrokinetic results is through the (global) magnetic shear $\hat{s}$ on the flux surface under consideration. Larger absolute values of $\hat{s}(s) = -2(s/\iota)(d\iota/ds)$[2] cause faster secular growth of

$$k_\perp = \sqrt{k_\psi^2|\nabla\psi|^2 + 2k_\alpha k_\psi \nabla\psi\cdot\nabla\alpha + k_\alpha^2|\nabla\alpha|^2}$$

through the magnitude of $\nabla\alpha = \nabla\theta_P - \iota\nabla\zeta - \iota'\zeta\nabla\psi$ where $\iota' = d\iota/ds$, and therefore provides linear stabilisation through the Bessel function J_0 in the Fourier-transformed gyrokinetic equation (Plunk et al. 2014). Increased *negative* shear can have an additional stabilising effect as it can cause $\mathbf{k}_\perp$ to become misaligned with $\mathbf{v}_M$ more quickly, narrowing the drift well $\omega_M = \mathbf{v}_M \cdot \mathbf{k}_\perp$ (Choi & Horton 1980; Antonsen et al. 1996; Plunk et al. 2014). This effect was also observed experimentally on TFTR (Levinton et al. 1995).

To keep the magnetic shear the same between the original optimised configuration and all interpolated equilibria we impose the rotational transform profile of the optimised configuration as an input to the MHD equilibrium problem of the interpolated equilibria. Here we use the iota profile of the stellarator configuration ($\lambda = 1$). Note that *constant ι profile* here thus refers to *constant*

---

[2] This definition is motivated by the tokamak definition of $\hat{s}(s) = (r/q)(dq/dr)$ where $q = 1/\iota$.

between equilibria and not *constant across the volume*, an assumption sometimes made in other theoretical studies.

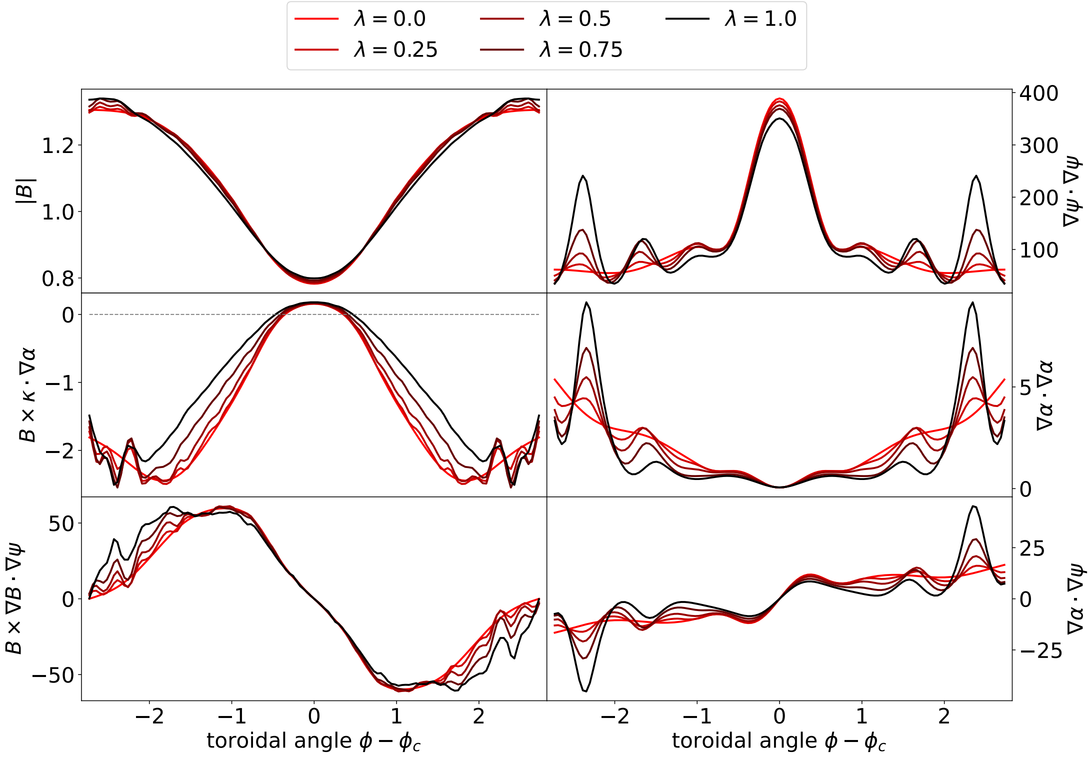

Figure 4.3.2. – Gyrokinetic geometric quantities (from equation 4.1.1) of equilibria with varying non-axisymmetric shaping and constant rotational transform profiles at $s = 0.64$ over one poloidal turn. Here $\phi_c = \pi/N_p$ is the toroidal location of the bean flux tube, and sets the toroidal centre of the flux tube and the location at which the flux tube is at $\theta_P = 0$.

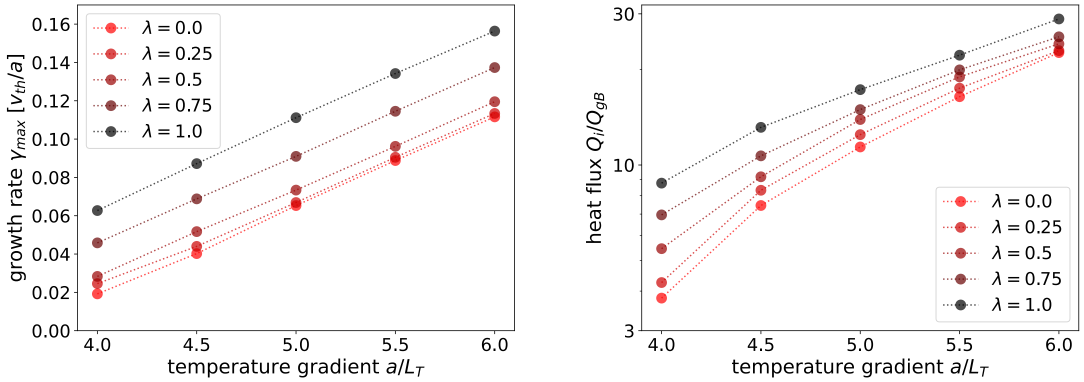

Figure 4.3.3. – ITG properties of equilibria with varying non-axisymmetric shaping and constant rotational transform profiles at $s = 0.64$. **Left:** the linear growth rate of the ITG instability. **Right:** the nonlinear heat flux.

The rotational transform profile ranges from ~ 0.15 on the axis to ~ 0.45 on the plasma boundary, with stellarator-typical negative shear throughout the majority of the volume. At $s = 0.64$, the location chosen for the analysis here, the values are $\iota(0.64) \sim 0.37$ and $\hat{s}(0.64) \sim -1.43$. We choose to analyse the flux surface at a normalised minor radius of $\rho = \sqrt{s} = 0.8$, as it is close enough to the edge that the flux-surfaces are still strongly non-axisymmetric (see figure 4.3.1), while far enough from the edge that the standard gyrokinetic framework is clearly applicable.

The geometric inputs to the gyrokinetic equation which result from the MHD equilibria are shown in figure 4.3.2. Configurations with stronger 3D shaping exhibit slightly wider drift wells, indicated by the positive region of $\mathbf{B} \times \boldsymbol{\kappa} \cdot \nabla\alpha$. It is also noteworthy that the flux surface compression, indicated by $|\nabla\psi|^2$, is slightly increased towards axisymmetry at the outboard midplane ($\phi - \phi_c = 0$), thought to be destabilising. However, the differences are small, and the 3D equilibria instead feature significant spikes in flux surface compression away from the midplane. This might have some effect. In a recent study by Landreman *et al.* (Landreman et al. 2025) the geometric feature which best predicted the heat flux in stellarators was one which took into account both the flux surface compression in regions of bad curvature as well as good curvature[3]. Lastly, we note that the locations away from the outboard midplane where $|\nabla\psi|^2$ peaks locally have similar peaks in $|\nabla\alpha|^2$ which is expected to be stabilising through $k_\perp$; Note that in the case of $k_x = 0$, which usually features the most unstable mode, $k_\perp \propto |\nabla\alpha|$.

The linear growth rate, calculated with `stella`, and the saturated nonlinear heat flux, calculated with `GX`, are shown in figure 4.3.3. *The numerical details, such as resolutions and convergence tests, are detailed in Appendix C.* The results in figure 4.3.3 show that both the growth rate and the heat flux are decreased across varying temperature gradients towards axisymmetry for the assumptions made here, though the differences in the heat flux are not that significant.

4.3.2 Maintaining a self-consistent rotational transform profile

The constant rotational transform profiles of the previous section served as a useful tool to isolate the effect of the 3D shaping with constant magnetic shear, but it also implies somewhat unrealistic assumptions for the current profile. While the chosen rotational transform is the one which yields a configuration with self-consistent bootstrap current for the original configuration ($\lambda = 1$), it assumes a large current drive at larger radii for the other configurations. This can be seen from figure 4.3.4 which shows the assumed profile of the parallel

---

[3] In this study this meant shifting the result of $\mathcal{H}(\mathbf{B} \times \boldsymbol{\kappa} \cdot \nabla\alpha)$, where $\mathcal{H}$ is the Heaviside function, to give some non-zero weight to all regions.

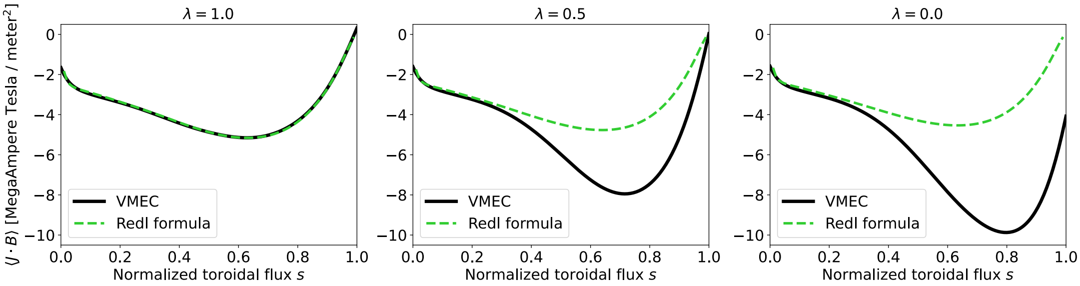

Figure 4.3.4. – The parallel current profiles for the interpolated equilibria from the previous section which share the same rotational transform profile. Solid black lines show the parallel current assumed in the `VMEC` equilibrium, and dashed green lines show the bootstrap prediction from the Redl-formulae.

current in these `VMEC` equilibria and the expectation from the bootstrap current using the Redl-formulae.

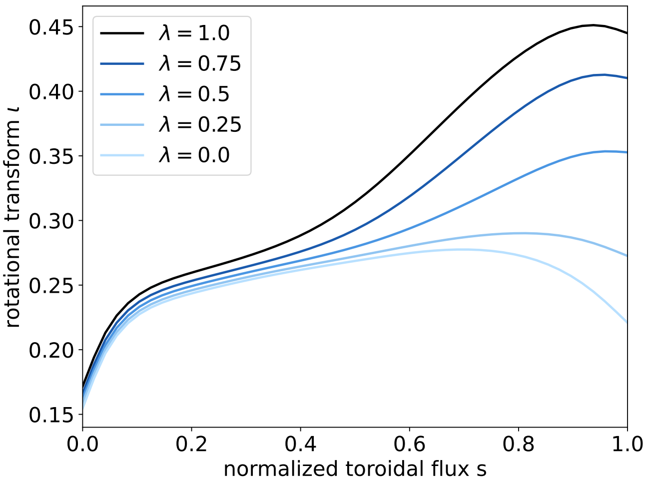

Figure 4.3.5. – The rotational transform profiles for configurations with interpolated boundaries and optimised self-consistent bootstrap current profiles. Simulations are performed at $s = 0.64$.

To obtain interpolated configurations which are *each* fully driven by the bootstrap current we re-optimise only the current profile by minimising the bootstrap-current-objective from equation 3.7.1, stated here again for convenience:

$$f_{boot} = \frac{\int_0^1 ds[\langle \mathbf{j} \cdot \mathbf{B} \rangle_{\mathrm{vmec}} - \langle \mathbf{j} \cdot \mathbf{B} \rangle_{\mathrm{Redl}}]^2}{\int_0^1 ds[\langle \mathbf{j} \cdot \mathbf{B} \rangle_{\mathrm{vmec}} + \langle \mathbf{j} \cdot \mathbf{B} \rangle_{\mathrm{Redl}}]^2}. \tag{4.3.2}$$

The resulting rotational transform profiles are shown in figure 4.3.5. Towards axisymmetry the rotational transform in the outer volume is reduced due to the reducing contribution from 3D shaping.

Figure 4.3.6 shows the geometric quantities which enter the gyrokinetic equations over one poloidal turn at $s = 0.64$. Due to the different values of ι

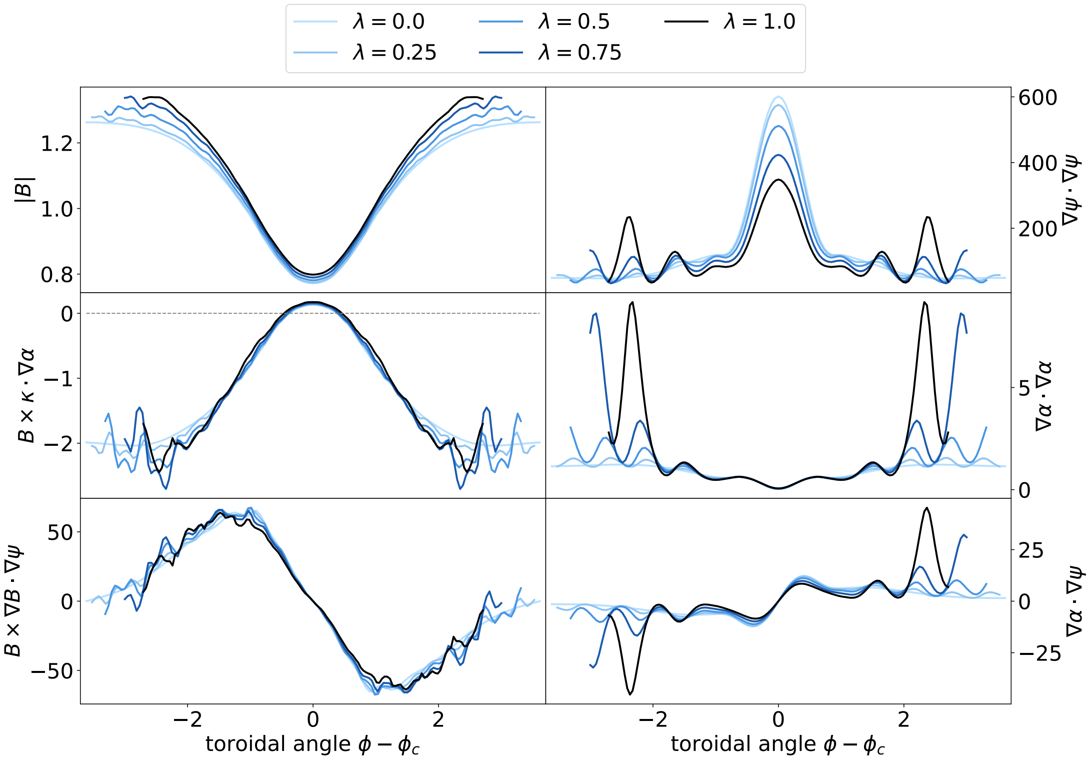

Figure 4.3.6. – Gyrokinetic geometric quantities (from equation 4.1.1) of equilibria with varying non-axisymmetric shaping and self-consistent bootstrap current at $s = 0.64$ over one poloidal turn. Here $\phi_c = \pi/N_p$ is the toroidal location of the bean flux tube, and sets the toroidal centre of the flux tube and the location at which the flux tube is at $\theta_P = 0$.

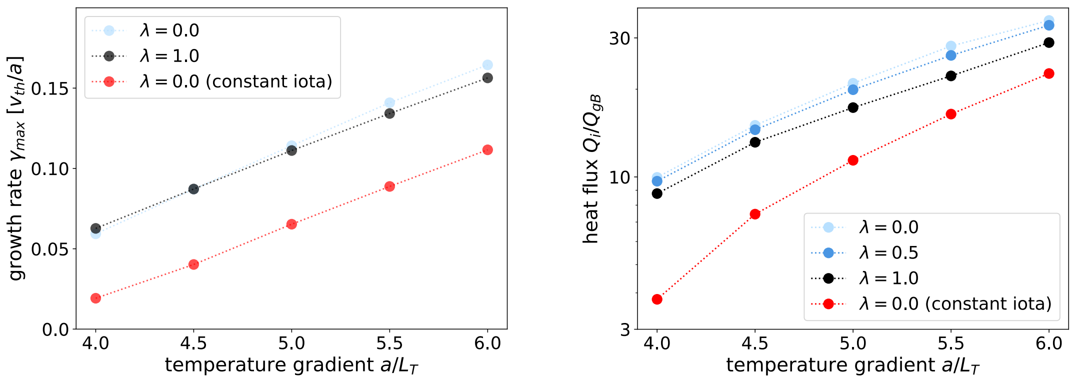

Figure 4.3.7. – ITG properties of equilibria with varying non-axisymmetric shaping and self-consistent current profile at $s = 0.64$. The axisymmetric result with constant iota profile from figure 4.3.3 is shown for comparison. **Left:** Linear growth rate. **Right:** Nonlinear heat flux.

the toroidal extent of flux-tubes which complete one poloidal turn varies. The variations in the flux surface compression $|\nabla\psi|^2$ on the outboard midplane, i.e. around $\phi - \phi_c = 0$ are now much larger compared to figure 4.3.2, with more compressions towards axisymmetry.

Figure 4.3.7 shows that both the linear growth rate and the heat flux are now similar between the QA configuration and the underlying tokamak, since the artificial stabilisation from the strong negative shear is now not present. The results from the previous section are shown in light red to highlight the differences arising from the self-consistent treatment of the rotational transform profile.

4.4 Impact of Mercier-stability

In section 3.10 it was shown that the Mercier-stability criterion strongly affects the optimised plasma boundary for $N_p = 4$, introducing a bean-like cross-section at the half period ($\phi = \pi/N_p$).

We also found there to be a significant trade-off between Mercier-stability and QA quality. *Does such a trade-off also exist between Mercier-stability and ITG turbulent transport?* A previous study (Roberg-Clark, Xanthopoulos & Plunk 2024) concerned with the turbulence optimisation of quasi-helically (QH) symmetric configurations showed such a trade-off to be present. It was reported there that the turbulence-optimised configuration does not have a bean-like cross-section, the absence of which allowed for smaller values of $|\nabla\psi|$ on the outboard midplane. It was therefore Mercier-unstable, i.e. it possessed a vacuum magnetic hill.

The top plots of figure 4.4.1 show the flux surfaces of the four-field-period compact QA configuration from section 3.10 which includes a self-consistent bootstrap current and Mercier-stability, comparing it to the flux surfaces of the same configuration without Mercier-stability, highlighting again the visible differences that the inclusion of Mercier-stability caused here.

The panels in the bottom section of figure 4.4.1 show the gyrokinetic geometric quantities in both configurations. Despite the differences in the flux surfaces, the differences in the geometric quantities are not significant. Even though the Mercier-stable configuration has a bean-like cross-section, the surface compression, quantified by $|\nabla\psi|^2$ is only slightly larger on the outboard midplane (see top right panel in the bottom plot of figure 4.4.1). The differences in the other geometric quantities are similarly small, e.g. the Mercier-stable configuration has only a slightly narrower drift well.

Figure 4.4.2 shows the growth rates and heat fluxes for both configurations across a range of temperature gradients on the flux surface $s = 0.64$. Even though the inboard shaping is visibly very different, the ITG characteristics are barely affected. The growth rate is slightly larger in the Mercier-stable

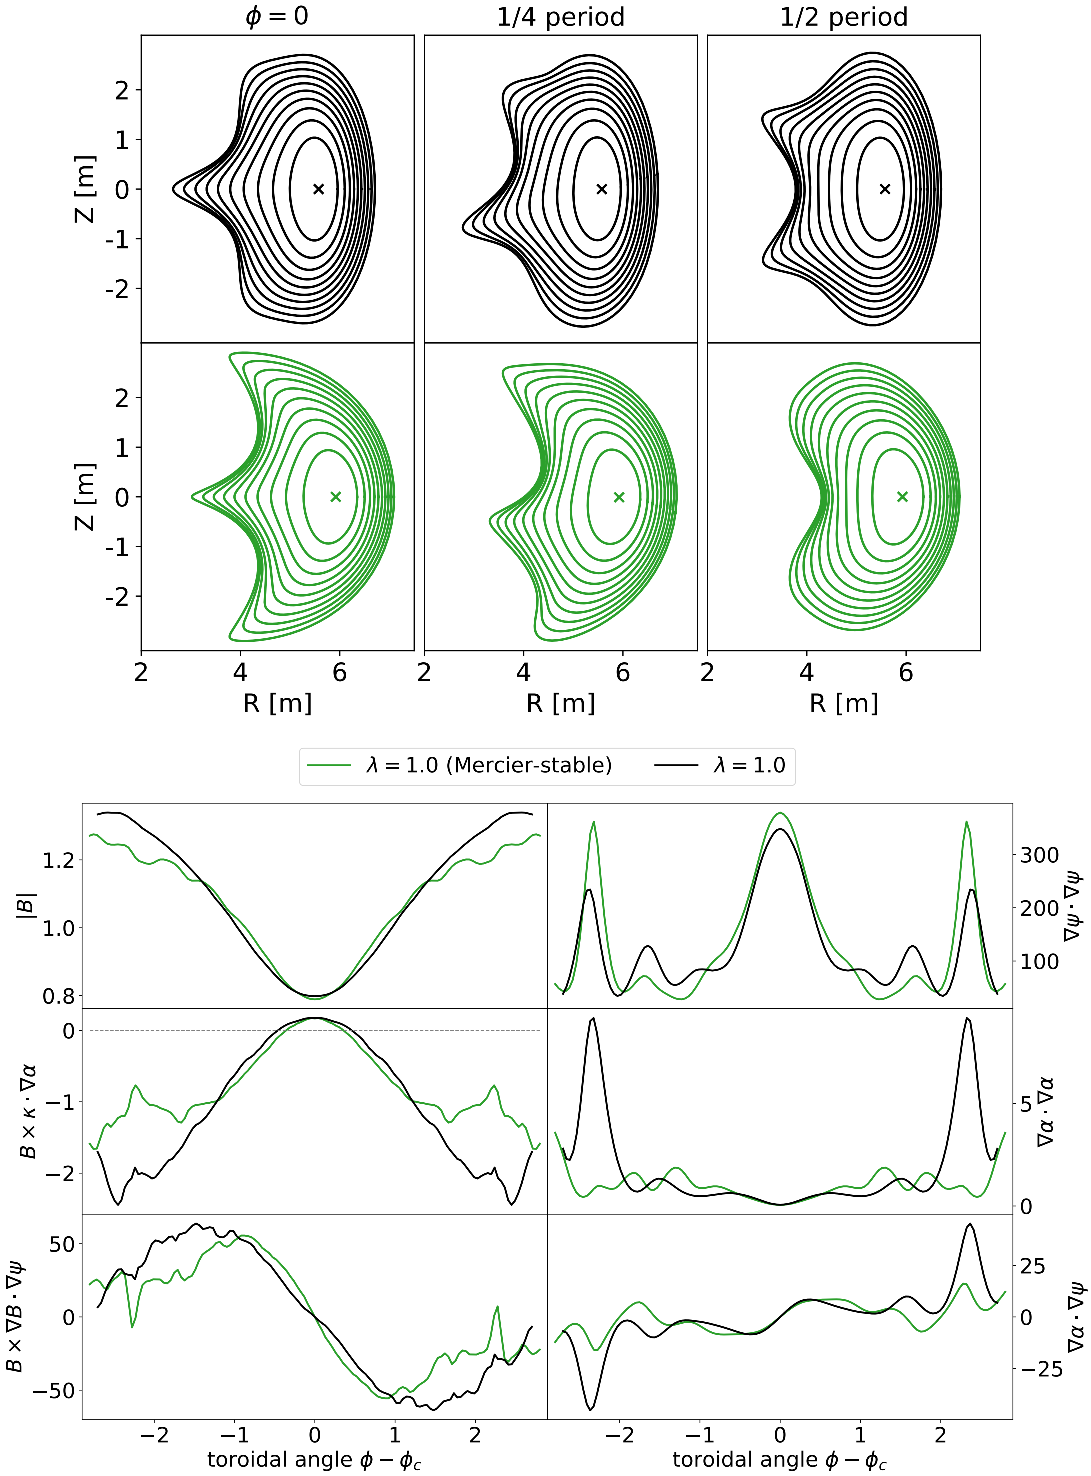

Figure 4.4.1. – Equilibria with (green) and without (black) Mercier-stability. **Top:** Flux surfaces at three different toroidal cuts within one field period. **Bottom:** Gyrokinetic geometric quantities from equation 4.1.1 at $s = 0.64$ along bean-like flux tube for one poloidal turn.

configuration but virtually no differences are observed in the heat flux across all temperature gradients. This is consistent with the linear picture of toroidal ITG

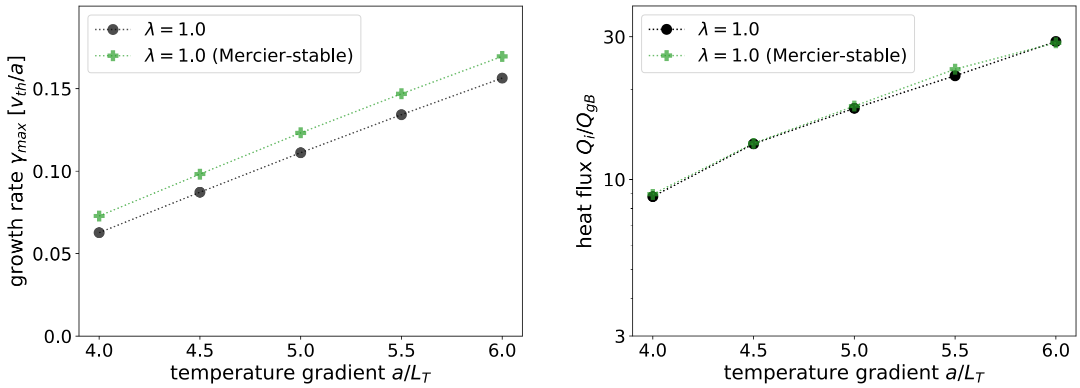

Figure 4.4.2. – ITG properties of equilibria with and without Mercier-stability at $s = 0.64$. **Left:** Linear growth rate. **Right:** Nonlinear heat flux.

modes which are driven only on the bad curvature side, i.e. the axisymmetric outboard side.

4.5 The impact of field-periodicity

In section 3.9 we optimised configurations of different field periods. Some trade-offs were identified:

1. Configurations with $N_p > 2$, even in the compact design space, are much more suitable for a stellarator-tokamak hybrid (Henneberg & Plunk 2024).
2. Two-field-period configurations are more naturally Mercier-stable, facing less of a trade-off with QA than those with $N_p > 2$.
3. The predicted coil complexity, as measured by the magnetic gradient scale length (Kappel, Landreman & Malhotra 2024), is increased towards a larger number of field periods[4].

Here we ask the question *how does the choice of the number of field periods impact ITG turbulent confinement?* We choose the configurations with most estimated external rotational transform of 0.1, optimised for QA, self-consistent bootstrap current, as well as Mercier-stability.

We first note that the chosen field periodicity along with the choice of self-consistent bootstrap current causes significant differences in the total rotational transform profile as well as magnetic shear profile, see figure 4.5.1. Three locations are marked in the figure with varying trends:

- $s = 0.09$ *(*$\rho = 0.3$*):* The magnetic shear is similar for three and four field periods, while the magnetic shear of the two-field-period configuration

---

[4] This is not discussed in this Thesis but is discussed in reference (Schuett & Henneberg 2025).

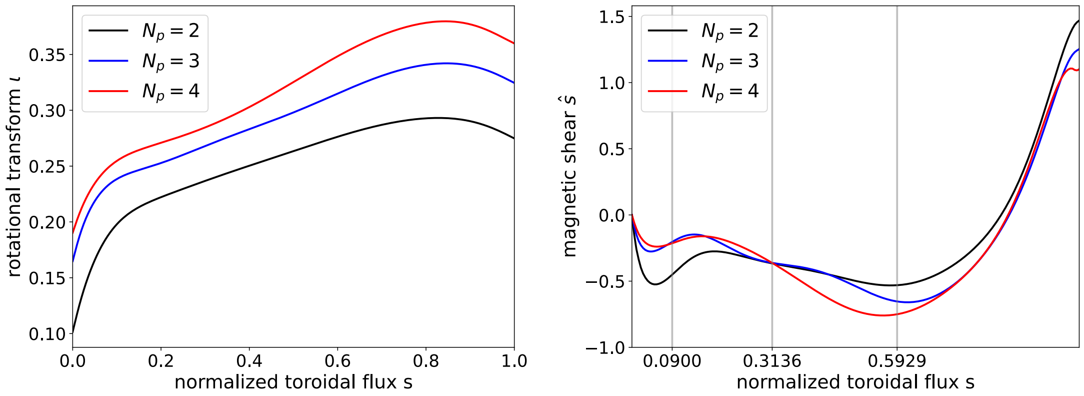

Figure 4.5.1. – Properties of configurations with varying field periods. **Left:** The rotational transform. **Right:** The magnetic shear. Three locations of varying trends are marked by vertical gray lines.

is more negative by a factor of ~ 2. Here the flux surfaces of the configurations with $N_p > 2$ are mostly axisymmetric, such that we can expect tokamak-like behaviour.

- *$s = 0.3136$ ($\rho = 0.56$):* All three configurations have approximately the same value of $\hat{s}$.
- *$s = 0.5929$ ($\rho = 0.77$):* The magnetic shear is more negative the higher the number of field periods.

The impact of magnetic shear on ITG turbulence is usually significant due to the reasons alluded to in section 4.3.1, so these choices of flux-surfaces allow us to investigate its effect as well as other differences between these configurations.

Figure 4.5.2 shows the time-averaged heat flux from `GX` between the configurations at the three radial locations marked in figure 4.5.1. For this comparison the three configurations are scaled to the ARIES-CS scale (minor radius and reference toroidal flux). Results are shown for a range of temperature gradients, to highlight the consistency of the trends within the non-marginal regime.

Despite the different trends of the magnetic shear at the three radial locations the trend of the heat flux is very similar at all three, *namely the heat flux is decreasing with increasing number of field periods. Only for the innermost surface $s = 0.09$ there are no differences observed between the configurations with three and four field periods.* A possible explanation for this observed monotonic trend, that appears insensitive to the magnetic shear, will be discussed in section 4.7, after these three configurations are compared to a range of other well-known stellarators in section 4.6.

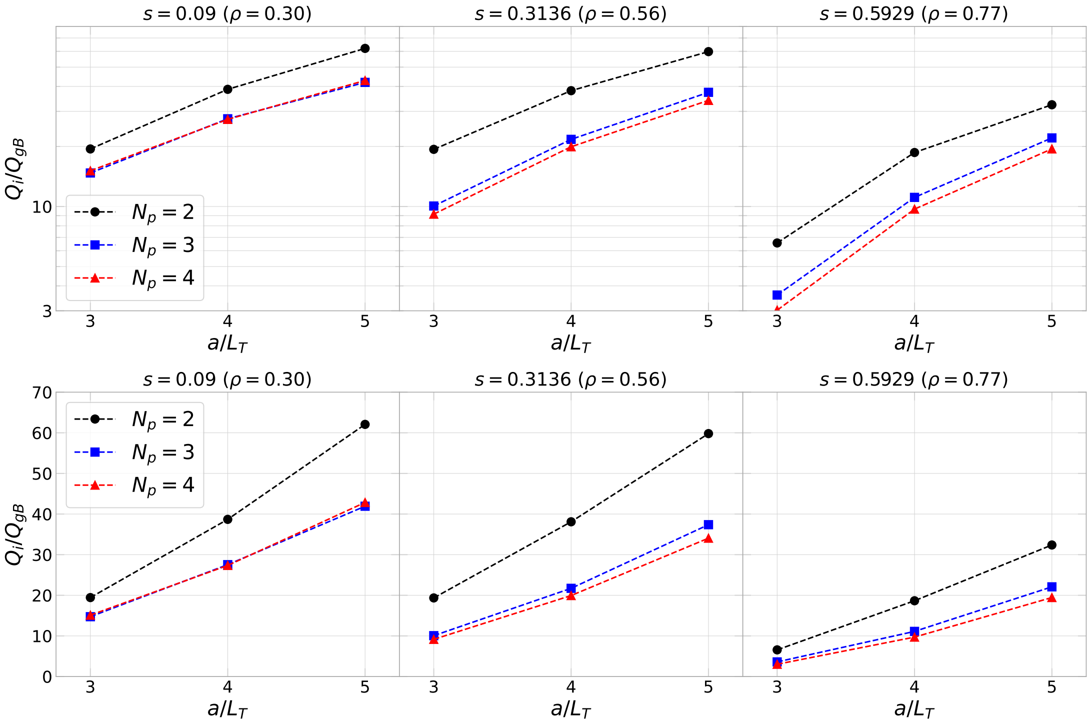

Figure 4.5.2. – The time-averaged heat flux from `GX` for configurations with two, three, and four field periods at three radial locations. The chosen locations are the ones which are marked in figure 4.5.1. **Top row:** Log scale. **Bottom row:** Linear scale.

4.6 In comparison with other legacy stellarators

Here we seek to answer the question *how does the turbulent transport from ITG of these configurations fit into the broader stellarator field?* To answer this question, we run GK simulations with `GX` in the bean-like flux-tube of a range of other well-known configurations. These include NCSX (Zarnstorff et al. 2001), HSX(Anderson et al. 1995), W-7X (Klinger et al. 2016), LHD (Iiyoshi et al. 1999), CFQS (Liu et al. 2018), the design by (Landreman & Paul 2022), the design by (Landreman, Buller & Drevlak 2022), the design by (Henneberg et al. 2019), and a previous compact design by (Garabedian & McFadden 2009). The corresponding `VMEC` equilibria were taken from the dataset Ref. (Landreman 2022).

4.6.1 Constant minor radius scaling

We first compare the saturated heat flux directly between configurations when scaled to the same minor radius and reference magnetic field, for which we choose the values from the ARIES-CS design. It follows that the gyro-Bohm units are identical between configurations with this choice, see section 4.2, but, of course, their volumes differ. The latter will be the focus of the next sub-section.

Figure 4.6.1 shows that the saturated heat flux in this set of configurations varies widely, with the two-field-period compact QA performing similar to many other configurations. The compact QA configurations with three and four field periods, however, are among the configurations with the lowest heat flux.

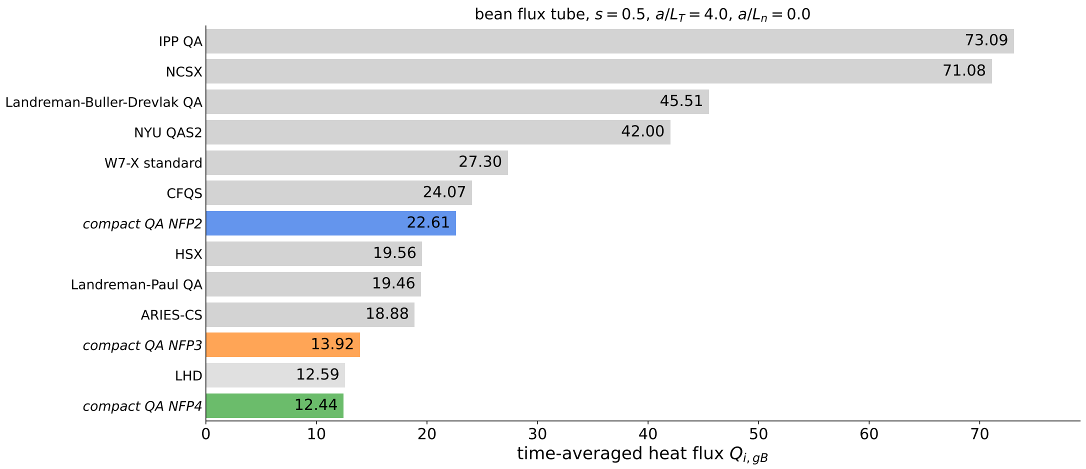

Figure 4.6.1. – Comparison of the time-averaged heat flux across a range of devices. Flux tubes are centred on the bean-like cross section at $s = 0.5$, and a fixed temperature gradient of $a/L_T = 4.0$ and no density gradient are set.

4.6.2 Constant volume scaling

The last section kept a constant minor radius between configurations with varying volumes. When considering the implications on the confinement time, however, it makes most sense to compare configurations of the same volume. This is motivated by the fact that both the cost and the fusion power scale with the volume.

Following a simple random walk estimate the confinement time is given by

$$\tau_E \sim a^2/\chi, \tag{4.6.1}$$

where the heat diffusivity χ is related to the heat flux via

$$Q = -\chi n \frac{dT}{dr}. \tag{4.6.2}$$

Thus for constant values of $a/L_T = (a/T)(dT/dr)$, as is the case in our simulations, the heat flux scales as

$$Q \sim \chi/a. \tag{4.6.3}$$

Using this relation and equation 4.2.1 to write $Q = Q_{\mathrm{gB}} Q_{\mathrm{gB,ref}} \sim Q_{\mathrm{gB}} (a_{\mathrm{ref}} B_{\mathrm{ref}})^{-2}$, where Q_{gB} refers to the heat flux expressed in gyro-Bohm values, we arrive at

$$\tau_E \sim a/Q \sim (a_{\mathrm{ref}}^3 B_{\mathrm{ref}}^2)/Q_{\mathrm{gB}} \sim \frac{V B_{\mathrm{ref}}^2}{A} \frac{1}{Q_{\mathrm{gB}}}. \tag{4.6.4}$$

In the last step we used the fact that the volume scales as $V \sim a^2 R \sim a^3 A$.

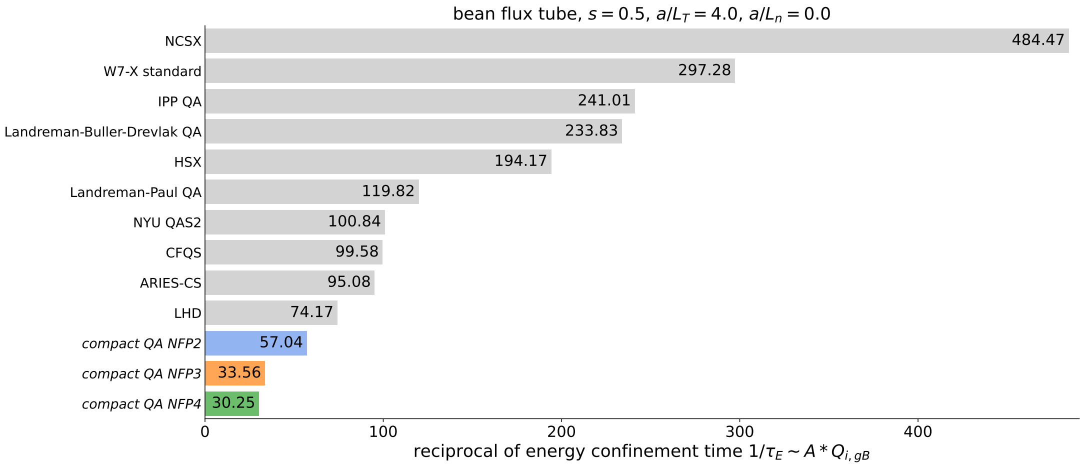

Figure 4.6.2. – Comparison of inverse confinement time estimates from flux tube simulations across a range of devices. Flux tubes are centred on the bean-like cross section at $s = 0.5$, and a fixed temperature gradient of $a/L_T = 4.0$ and no density gradient are set.

We re-scale our configurations to have the same reference field and volume as ARIES-CS, in which case the confinement time scales as $\tau_E \sim (AQ_{\mathrm{gB}})^{-1}$, with the rest up to a constant which can be used to express the estimates in units of e.g. seconds if one desires to do so. The gyro-Bohm normalisations now differ between configurations since ϕ_{edge} is constant but a_{ref} is not. However, we note that this difference is taken into account, since the gyro-Bohm units themselves enter in equation 4.6.4, and we are not using these units for the normalisation of the confinement time.

In figure 4.6.2 the reciprocal of the confinement time, proportional to AQ_{gB}, is compared for the same set of configurations as in figure 4.6.1. The same flux-tube and the same equilibrium gradients are chosen as before. The three compact QA configurations perform even more favourably here due to their small aspect ratio.

4.7 Consistency with picture of flux surface compression

In figure 4.7.1 we show the metric of mean bad-curvature flux surface compression (Angelino et al. 2009; Xanthopoulos et al. 2014; Mynick et al. 2014; Stroteich et al. 2022; Goodman et al. 2024; Landreman et al. 2025) along the analysed flux tubes. Specifically, here we use the expression from Ref.

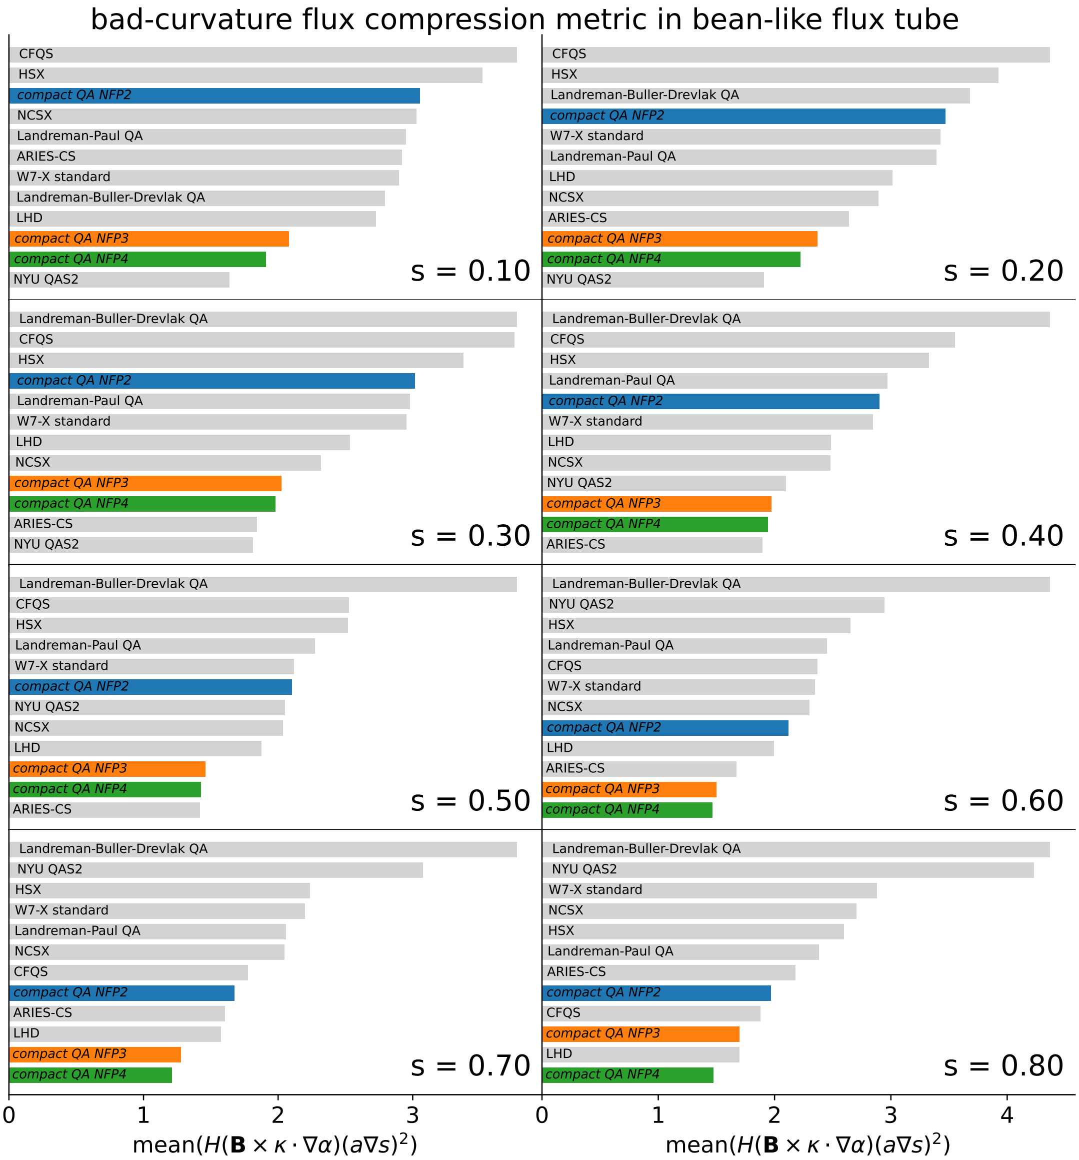

Figure 4.7.1. – The metric of flux surface compression between the compact QA configurations and a range of legacy configurations. Each panel corresponds to a different radial locations and the bean-like flux tube was selected.

(Landreman et al. 2025) for this metric:

$$\mathcal{X} = \text{mean}[\mathcal{H}(\mathbf{B} \times \kappa \cdot \nabla\alpha)(a\nabla s)^2], \tag{4.7.1}$$

where $\mathcal{H}$ is the Heaviside step function, the sign of $\mathbf{B} \times \kappa \cdot \nabla\alpha$ denotes whether the field line curvature is locally 'good' or 'bad' as analysed before in sections 4.3.1 and 4.3.2, and the mean is taken along the field line for one poloidal turn.

The figure features the compact QA configurations of different field-periodicity that were compared in section 4.5, as well as the comparison to the same set of legacy configurations as in section 4.6.

We see that the monotonic trend with field-periodicity observed in section 4.5 is also seen in this metric, with lower values of flux compression towards larger number of field periods. Similarly, the trend from section 4.6 is observed here as well, i.e. the compact QA configurations with $N_p > 2$ score among the lowest in the comparison with this set of legacy configurations.

4.8 Conclusion and discussion

In this chapter we have carefully explored the impact of quasi-axisymmetric perturbations on ion temperature gradient driven turbulence. When keeping the radial profile of rotational transform constant between configurations, some differences are observed. However, we showed that when the rotational transform profile is treated self-consistently, stabilisation is regained, resulting in negligible differences between the QA configuration and the underlying tokamak – despite the strong modification of the good curvature side. Even Mercier-stability, the inclusion of which can strongly change the details of the inboard side shaping, did not result in a noticeable difference in the turbulent transport properties. The results are consistent with the linear physical picture of toroidal ITG modes.

By comparing configurations across varying field periods at different radii we found that (1) the (negative) magnetic shear does not set the differences in their nonlinear heat flux as one might expect in a tokamak, but instead (2) the nonlinear heat flux increases towards lower number of field periods, in line with the flux surface compression at the outboard midplane. The latter has been shown to be the best geometric indicator for the general stellarator space (Landreman et al. 2025).

Lastly we provided some context by comparing the heat fluxes to some well-known stellarators. The compact QA configurations perform well, likely due to the modest elongation of the bean-like cross-section which strongly squeezes flux surfaces together at the outboard midplane. These configurations perform even more favourably in a comparison of the implied confinement time of the heat fluxes, due to their much lower aspect ratio than traditional stellarators.

All in all these results corroborate the notion that the parameter space of QA-perturbed tokamaks offers the potential of combining the strengths of tokamaks and stellarators. With respect to transport from ITG driven turbulence, one can likely expect tokamak-like confinement, while enjoying stabilisation effects against disruptions from the external rotational transform produced on the inboard side.

4.8.1 Future work

It would be instructive to run the calculations in more flux tubes of the configurations and at more radii to obtain an even better picture. It would also be instructive to run full flux surface calculations for these configurations. This is motivated by the fact that in stellarators different flux tubes generally experience different geometry.

Beyond this extension of the study, the natural next step is to use these findings for the purpose of optimisation, i.e. to find hybrid configurations that maximise confinement. Since we showed that the strong QA-shaping of the inboard side does not modify the turbulence properties substantially, one could likely optimise the underlying tokamak towards confinement and these gains would be reflected in QA-perturbed version of the configuration. Since we find in the QA-optimisation of these configurations that certain axisymmetric shaping allows for lower QA error and MHD stability, the best solution likely is to couple the confinement optimisation of the underlying tokamak into the QA-optimisation loop of the stellarator-tokamak hybrid.

Moreover, it would be valuable to investigate turbulence from instabilities that are traditionally extended, such as microtearing modes or other electron driven modes. With respect to trapped particle modes, it should be noted that many of the trapped particles only spend time on the nearly axisymmetric outboard side of the hybrid configurations and therefore essentially do not experience any 3D shaping. Further, as seen in figures 4.3.2 and 4.3.6, despite the strong 3D shaping, the magnitude of the magnetic field along a field line at most only experiences small variations on the inboard side. No substantial additional trapping wells are created, in contrast to more traditional stellarators. This can be explained by the strong $1/R$ dependence of the magnetic field, and suggests that the trapping well of the underlying tokamak remains the main one in these configurations, even for shallowly trapped particles which do sample a significant fraction of the inboard side. Therefore, similar to the results for ITG turbulence presented in this chapter, there is reason to expect transport from trapped electron modes (TEM) to behave similarly to the equivalent tokamak as well, but future work should verify this assumption.

4.8.2 Key findings and implications

The key points of this chapter and its implications are:

- Quasi-axisymmetrically perturbed tokamaks have similar ITG transport properties as their equivalent tokamak.
- Transport reduces with increasing number of field periods, consistent with the picture of flux surface compression from the literature.
- The compact nature of these designs is beneficial for the confinement time.
- Future optimisation of stellarator-tokamak hybrids to reduce transport from ion temperature gradient turbulence could likely utilise tokamak tools inside the optimisation loop.

5

Nonlinear energy transfer in stellarators

In this chapter we will explore how three-dimensional shaping affects the spatial distribution of zonal flow drive in stellarators. By porting our diagnostic from chapter 2 to `stella` we will first explore the effect of the continuous quasi-axisymmetric perturbations which we optimised and studied in part in the last two chapters, before applying the method to a flux tube of the Wendelstein 7-X device.

5.1 Introduction

For tokamaks, a single flux tube which traverses the device for one poloidal turn captures all the geometrical information of the flux surface. Hence the drive of ZFs along this flux tube can be mapped directly to its poloidal distribution, which is then the same at every toroidal angle.

As mentioned in the previous chapter, flux tubes which are centred at different toroidal angles will generally experience different geometry, and thus their turbulent properties and ZF drive properties will generally be different. We cannot draw a conclusion of some poloidal distribution of the ZF drive from its distribution along a field line. Due to broken axisymmetry no such poloidal distribution is expected to exist and a flux surface dependence of the form

$$\mathcal{T}_v^{ZF} = \mathcal{T}_v^{ZF}(\theta, \phi) \tag{5.1.1}$$

would be expected.

That being said, many insights can be gained from flux tube simulations in stellarators. Since turbulent eddies are, of course, still highly elongated along the magnetic field line and short in their perpendicular extent away from it, flux tube simulations let us understand how stellarator specific geometry affects the gyrokinetic equation and thus the turbulence (Landreman et al. 2025). Flux tube simulations have also been used to understand how geometry can be changed to increase the critical gradient (Roberg-Clark, Xanthopoulos & Plunk 2024;

Roberg-Clark et al. 2023). Understanding such dependencies can help one to understand how the dynamics in the other flux tubes might behave.

We found in chapter 2 that the ZF drive envelope in tokamaks is highly sensitive to changes in the geometry. Since stellarators feature even stronger and more localized variations in the geometry, an analysis of the ZF drive lets us further understand how these are connected. Since axisymmetry is broken our results represent the parallel distribution of the ZF drive along a chosen field line rather than a general poloidal distribution. We highlight this by adopting the coordinate change

$$\mathcal{T}_v^{ZF}(\theta) \to \mathcal{T}_v^{ZF}(z), \tag{5.1.2}$$

although z could of course still represent a poloidal angle for the parameterisation if one wanted to make that choice. Following conventions in `stella`, z will typically be parametrized by the toroidal angle ϕ. *As before all simulations in this chapter are electrostatic and assume adiabatic electrons.*

5.2 Benchmarking zonal flow drive between `GS2` and `stella`

Whereas for axisymmetric configurations we have used `GS2`, in this chapter we will use `stella` (Barnes, Parra & Landreman 2019) since it was built for non-axisymmetric fields in mind from the beginning.

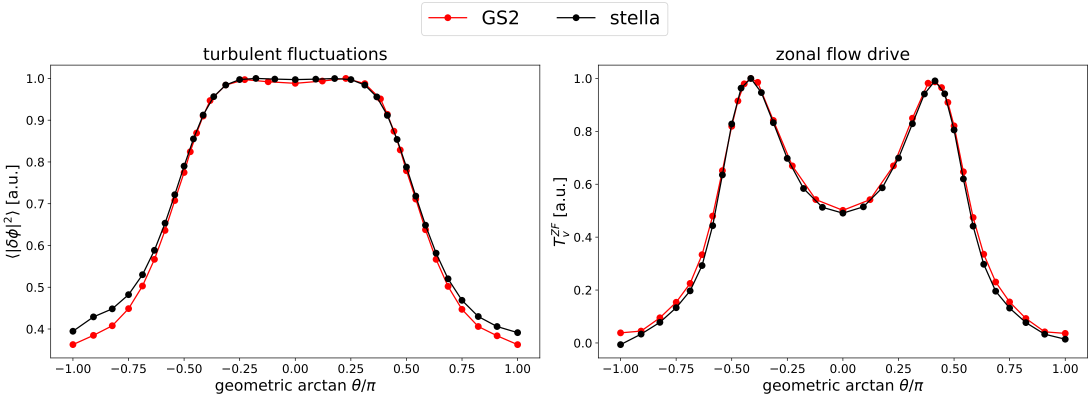

Figure 5.2.1. – Benchmark between `stella` and `GS2`. The CBC equilibrium was chosen, with modified elongation $\kappa = 1.5$, temperature gradient $a/L_T = 4.0$, and flux surface $\rho = 0.8$. **Left:** distribution of turbulent fluctuations. **Right:** distribution of ZF drive.

While `GS2` features a non-axisymmetric extension for the linear GKE, it does not for the nonlinear GKE. `stella` was successfully benchmarked against `GS2` in the saturated heat flux of nonlinear tokamak equilibria and for the linear growth rate spectra and envelopes of the electrostatic potential for a non-axisymmetric configuration (NCSX) (Barnes, Parra & Landreman 2019).

Before we explore the ZF drive in non-axisymmetric configurations we first benchmark it as well as the distribution of turbulent fluctuations between the two codes. Here we use the CBC equilibrium with modified elongation of $\kappa = 1.5$ (i.e. one of the cases from section 2.5.1). The result is shown in figure 5.2.1 and shows excellent agreement between both codes.

5.3 Zonal flow drive from tokamak to QA stellarator

In chapter 4 we showed that the compact QA stellarators which were created in chapter 3 feature geometric flux tube quantities that are mostly the same as in axisymmetric configurations, albeit with strong local variations on the inboard side. This, of course, is due to the localized QA shaping and the ridges which are the characteristic feature of this parameter space.

One wonders how this geometry affects the parallel distributions of turbulence and ZF drive. We take advantage of the perturbed character of the compact QA configurations by making the connection to the underlying tokamak following the same method which we used in section 4.3 of chapter 4. There we introduced the perturbation parameters λ which interpolates the fully perturbed configuration with the underlying tokamak without changing the axisymmetric shaping.

Figure 5.3.1 shows the parallel distributions of the zonal flow drive as the perturbation parameter λ is increased. A clear gradual response to the gradual change of QA-shaping can be seen as the four local maxima of zonal flow drive emerge on the inboard side. The two local maxima towards the outboard side remain relatively constant. We do note however, that their width also increases as the QA-shaping is introduced.

Note that while the tokamak configurations in chapter 2 all had positive global magnetic shear $\hat{s} = 0.78$ which they inherited from the CBC, the underlying tokamak here has significant negative global magnetic shear of $\hat{s} = -1.43$ at the chosen flux surface $\rho = 0.8$. Nonetheless the poloidal structure of the zonal flow drive in the underlying tokamak here is qualitatively very similar to the one for the elongated CBC from figure 2.5.4. Figure 4.3.1 shows that the poloidal cross-section of the underlying tokamak does feature mostly elongation. We thus conclude that the global magnetic shear does not seem to alter the results of the zonal flow drive envelope significantly.

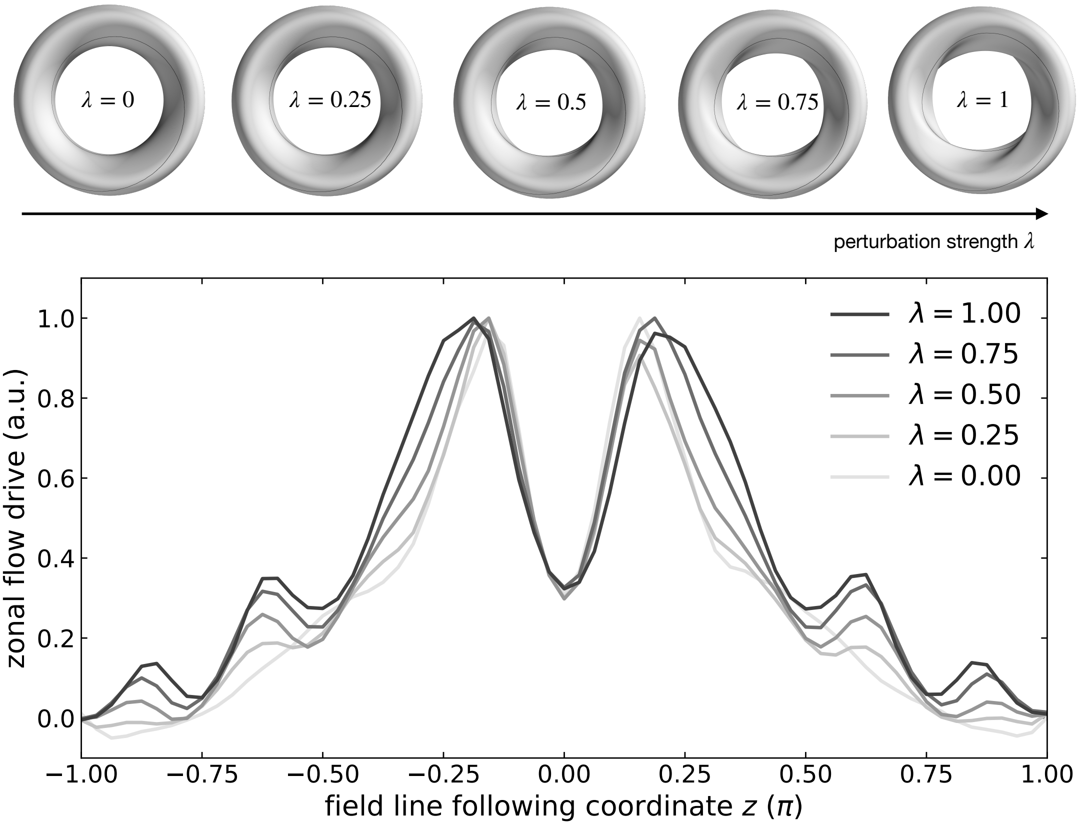

Figure 5.3.1. – The parallel distribution of zonal flow drive under continuous departure from axisymmetry towards a fully QA-perturbed configuration. The change of the plasma boundary as the perturbation parameter λ increases is shown on the top. The results show how the local maxima develop *step-by-step* as λ goes from zero to unity.

5.4 Zonal flow drive in compact QA stellarator

Figure 5.4.1 shows the ZF drive for the fully perturbed case from the previous subsection, i.e. the case with $\lambda = 1$ from figure 5.3.1, alongside our measures of turbulent fluctuations from chapter 2.

Our measures of turbulent activity show the usual ballooning towards the outboard midplane with essentially no signature of the localized shaping. However, the ZF drive does show a significant signature of the localized shaping. This is seen by four additional local maxima which have emerged on the inboard side. Figure 5.4.1 also shows the central field line of the flux tube which we considered on the flux surface which was chosen, $\rho = \sqrt{s} = 0.8$. Upon closer inspection it can be seen that the field line crosses the ridges on the inboard side four times within its one poloidal turn. The local maxima in the drive correspond to these four locations, whereas the two maxima at $z \sim 0.25$ appear to be the usual ones from the underlying tokamak shaping which features significant elongation, as confirmed in figure 5.3.1 for the underlying tokamak ($\lambda = 0$).

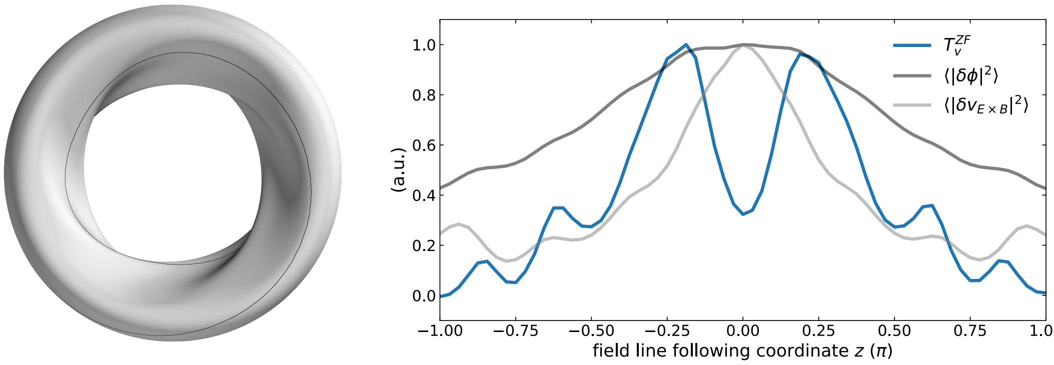

Figure 5.4.1. – The zonal flow drive in a QA-perturbed tokamak configuration. **Left:** the chosen flux surface ($\rho = 0.8$) with the central field line of the chosen flux tube. **Right:** the parallel distributions of turbulent fluctuations and zonal flow drive, showing a unique signature of the QA-shaping.

5.5 Zonal flow drive in Wendelstein 7-X

Wendelstein 7-X (W7-X) (Klinger et al. 2016), which was depicted in figure 1.1.2, is the newest large-scale stellarator experiment and is currently operational in Greifswald, Germany. With a diameter of $\sim 10\,\mathrm{m}$ and an average minor radius of $\sim 0.5\,\mathrm{m}$ it has a more typical, stellarator-like, larger aspect ratio $A \sim 10$ than any of the previous configurations considered in this work. The configuration has five field periods, $N_p = 5$, and is approximately quasi-isodynamic, see figure 3.2.1 for context.

W7-X, like most stellarators optimised for improved MHD stability, features a bean-like cross-section, which tends to drive most ITG turbulent transport due to the tighter flux surface spacing on the outboard side within this cross-section. The flux tube which is centred on the outboard midplane of this cross-section in W7-X has therefore been extensively studied in the past, e.g. (Xanthopoulos et al. 2007; González-Jerez et al. 2022; Stroteich et al. 2022; Podavini et al. 2024).

As a widely studied and experimentally relevant configuration this flux tube is therefore ideally suited to study the zonal flow drive envelope in strongly shaped geometry. The configuration and the central field line of the flux tube are shown in figure 5.5.1.

Figure 5.5.2 shows the parallel distributions of both zonal flow drive and turbulent fluctuations within this flux tube. Both measures of turbulent fluctuations peak strongly on the outboard midplane of the bean cross-section, as expected.

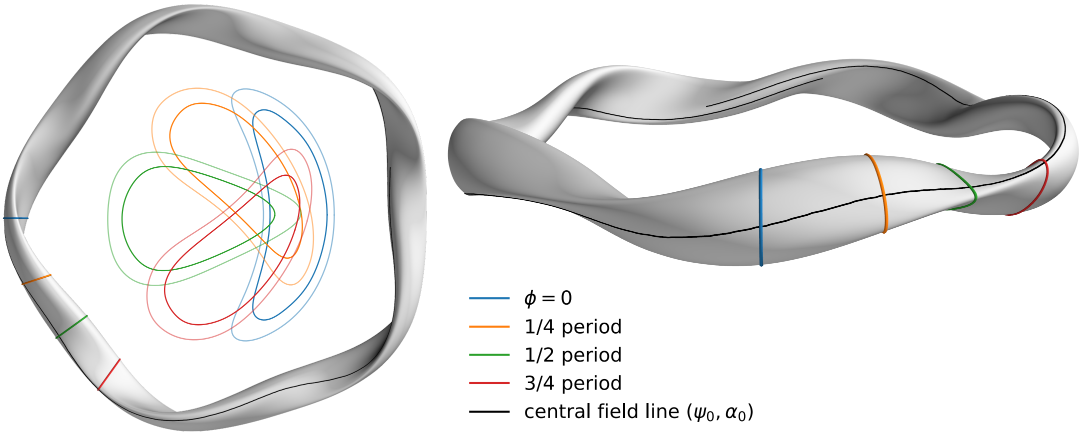

Figure 5.5.1. – The bean flux tube on the $s = 0.64$ flux surface of the Wendelstein 7-X standard configuration. **Left:** The analysed flux surface shown in gray from the top. Four poloidal cuts are indicated in colour, and are displayed in the centre in the usual poloidal R-Z-plane. The same cuts of the plasma boundary ($s = 1$) are also shown in shaded lines. The bean cross-section is shown in blue. The central field line of the flux tube, centred toroidally on the bean cross-section, is shown in black. **Right:** View from the side.

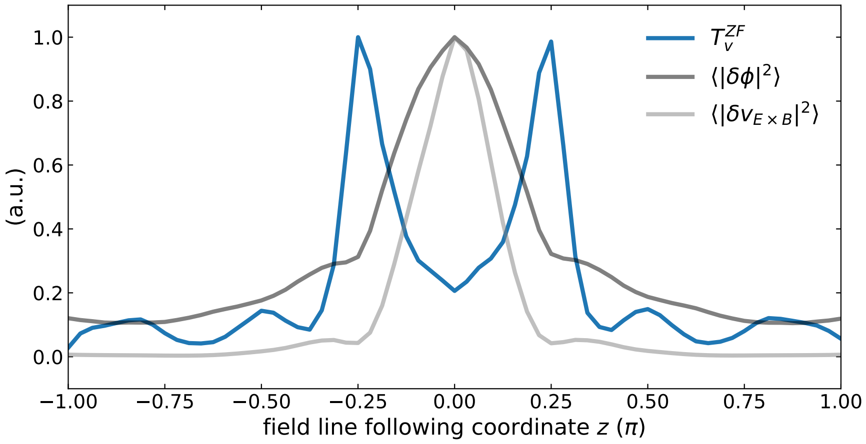

Figure 5.5.2. – The parallel distributions of zonal flow drive and turbulent fluctuations in the bean flux tube of the Wendelstein 7-X standard configuration, visualized in figure 5.5.1. The chosen flux surface is $s = 0.64$. The zonal flow drive does not peak where the turbulence does.

However, as for shaped tokamaks and the compact QA configurations, the zonal flow drive obtains a local minimum here and instead features two local maxima to either side of the centre of the flux tube.

Note that W-7X features near-zero global shear on this flux surface, $\hat{s}(s = 0.64) = -0.19$. Given that the CBC featured moderate positive shear, and the compact QA configuration strongly negative shear, this corroborates the observation that the bifurcation of the zonal flow drive is relatively insensitive to the values of global shear. Figure 5.5.2 also has two small additional local maxima on either side of the flux tube center and beyond the first dominant local maximum. These are likely due the locally varying geometry, similar to the local maxima seen in the compact QA configurations in figure 5.3.1.

5.6 Discussion and conclusion

By extending the analysis from tokamaks to stellarators, we have further shown that geometry affects the poloidal distribution of zonal flow drive in a non-trivial way. This results in local maxima of the drive envelope on the inboard side for quasi-axisymmetrically perturbed tokamaks. Further, a local minimum of the drive is expected at the location of maximum turbulence in Wendelstein 7-X, i.e. at the outboard midplane on the bean-like cross-section. Given the recent observation of zonal flows in W7-X (Carralero et al. 2025), and the installation of gas puff imaging (GPI) diagnostics on the device (Terry et al. 2024), these results can inform the interpretation of experimental measurements of turbulence-flow coupling in W7-X.

We have only scratched the surface of stellarator geometry in this chapter. The analysis can be extended to various other regions of the stellarator geometry space. This is a promising route to understanding how geometry affects the ZF drive distribution, since stellarators offer vastly more geometric possibilities than tokamaks. A promising approach might be one similar to the methodology in (Landreman et al. 2025), but rather than finding interpretable models for influence of the geometric quantities on the time-averaged heat flux, one could perhaps adapt the methods to find their influence on the zonal flow drive distribution. Note that the difference between the two mappings lies in the nature of the quantity one is interested in. While the time-averaged heat flux is a scalar, the ZF drive distribution is inherently a function of the parallel coordinate z. See the discussion in chapter 2 around equations 2.7.9 and 2.7.10. This likely makes the methodology more complicated than the one from (Landreman et al. 2025).

Moreover, due to the non-axisymmetry of stellarators, it is a debated topic whether or not a full flux surface (FFS) treatment is required to correctly calculate the zonal flow and heat flux. Some codes such as `GENE` (Jenko et al. 2000; Xanthopoulos et al. 2007; Xanthopoulos et al. 2014) as already available in this mode while others such as `stella` are under active development to support it. It would be interesting to adapt the nonlinear coupling analysis to the FFS treatment. This would allow one to obtain the distribution $\mathcal{T}_v^{ZF}(\theta,\phi)$ from equation 5.1.1. However, this will likely require a reformulation of the drive formalism and some additional computational challenges will have to be overcome to prevent this analysis from becoming computationally intractable. FFS simulations themselves, i.e. without the inclusion of the compute and memory intensive three-wave-couplings, require significant computational resources.

5.6.1 Key findings and implications

The key points of this chapter and its implications are:

- The localised 3D shaping of stellarator-tokamak hybrid equilibria creates additional local maxima in the zonal flow drive on the inboard side, in addition to the distribution expected from the underlying tokamak.
- The distribution of ZF drive is also expected to be non-trivial in the Wendelstein 7-X stellarator, where local maxima are expected away from the outboard midplane of the toroidal bean-like cross-section, the location of strongest turbulent activity.
- These results can aid the interpretation of measurements of nonlinear transfer in Wendelstein 7-X, as well as the interpretation and validation of nonlinear gyrokinetic simulations.

6

Beyond the fluid picture for energy transfer

In this chapter we will derive a kinetic extension to the electrostatic zonal flow drive directly from gyrokinetics. Borrowing methods from (Bañón Navarro et al. 2011a; Bañón Navarro et al. 2011b; Helander & Plunk 2022) we will focus the phase space coupling specifically on the drive of zonal modes, resulting in expressions for the nonlinear drive of zonal modes in the free energy, entropy, and fluctuation energy, with the latter describing the direct drive of zonal flows. We highlight the benefits of such an analysis as well as computational challenges.

6.1 Motivation

So far we have focused on the fluid picture for the zonal flow drive, in which it is the radial derivative of the Reynolds stress $\langle \delta v_y \delta v_x \rangle$ that is responsible for the drive. For the generation of the turbulence data for this analysis we made the case that gyrokinetics is the theoretical framework which captures tokamak and stellarator turbulence in the most realistic way while being computationally feasible. Gyrokinetics, of course, features its own nonlinearity. Unlike the momentum equation with the advection of the flow by the flow itself, the gyrokinetic nonlinearity features the advection of the perturbed distribution function by the gyrokinetic $\mathbf{E} \times \mathbf{B}$ drift, i.e. in gyrokinetics the general convective derivative 1.4.8 takes the form:

$$\frac{\partial g_s}{\partial t} + \langle \mathbf{V}_{\mathbf{E}\times\mathbf{B}} \rangle_s^{\mathbf{R}} \cdot \nabla h_s = \mathcal{L}[g_s], \tag{6.1.1}$$

where from equation 1.5.9 in the electrostatic limit

$$\langle \mathbf{V}_{\mathbf{E}\times\mathbf{B}} \rangle_s^{\mathbf{R}} = B^{-1} \mathbf{b} \times \nabla \langle \delta\phi \rangle_s^{\mathbf{R}}.$$

After Fourier transforming this becomes

$$\frac{\partial g_{s\mathbf{k}}}{\partial t} + \frac{1}{B^2} \sum_{\mathbf{k}_1} \mathbf{B} \cdot (\mathbf{k}_1 \times \mathbf{k}) J_{0s(\mathbf{k}-\mathbf{k}_1)} \delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} = \mathcal{L}_{\mathbf{k}}[g_s] \tag{6.1.2}$$

where from equation 1.5.6 $g_{s\mathbf{k}}$ is the gyro-average of the full perturbed distribution function:

$$g_{s\mathbf{k}} = h_{s\mathbf{k}} - (Z_s e J_{0s\mathbf{k}} \delta\phi_{\mathbf{k}} / T_{0s}) F_{0s}. \tag{6.1.3}$$

Here $J_{0s\mathbf{k}} = J_0(k_\perp v_\perp / \Omega_s)$ is a Bessel function which arises from the gyro-averaging operation $\langle \ldots \rangle_s^{\mathbf{R}}$ under Fourier transformation.

In this chapter we derive the transfer functions which arise from gyrokinetics and discuss their potential benefits and challenging numerical realisation. A similar derivation has been done before, see Refs. (Bañón Navarro et al. 2011a; Bañón Navarro et al. 2011b), but here we focus specifically on the drive of zonal *flows*.

The analysis is electrostatic throughout but does not make assumptions on the electron dynamics, i.e. we do not assume electrons to behave Boltzmann-like. A generalization of all the derivations for the electromagnetic case is not known at the time of writing this. A difficulty arises since the two components of the Helmholtz free energy, the fluctuation energy and the entropy, are not conserved separately any longer once magnetic fluctuations are allowed (Bañón Navarro et al. 2011a; Bañón Navarro et al. 2011b).

6.2 Energy transfer functions within gyrokinetics

6.2.1 Helmholtz free energy transfer

We first demonstrate, following some of the steps from (Helander & Plunk 2022), how we can obtain an evolution equation for the Helmholtz free energy from equation 6.1.1. We start by applying the operator

$$\mathrm{Re} \sum_s T_s \int (\ldots) \frac{h^*_{s\mathbf{k}}}{F_s} d\mathbf{v}$$

on equation 6.1.2. From the first term on the left-hand side we obtain

$$\mathrm{Re} \sum_s \int \frac{T_{0s}}{2F_{0s}} \frac{\partial |h_{s\mathbf{k}}|^2}{\partial t} d\mathbf{v} - \mathrm{Re} \sum_s Z_s e \left(\frac{\partial \delta\varphi_{\mathbf{k}}}{\partial t} \right) \int h^*_{s\mathbf{k}} J_{0s\mathbf{k}} \, d\mathbf{v}. \tag{6.2.1}$$

This expression can be further simplified by using the general quasi-neutrality condition (from which equation 1.5.12 followed for the case of adiabatic electrons)

$$\sum_s \frac{n_s e_s^2}{T_s} \delta\varphi_{\mathbf{k}} = \sum_s e_s \int h_{s\mathbf{k}} J_{0s\mathbf{k}} \, d\mathbf{v}, \tag{6.2.2}$$

where the factor $J_{0s\mathbf{k}}$ results from the dual gyro-average $\langle h_s \rangle_{\mathbf{r}}$ under Fourier transformation. The total evolution equation can thus be written as

$$\frac{\partial H_{\mathbf{k}}}{\partial t} = 2\left(N_{\mathbf{k}}^{H} + L_{\mathbf{k}}^{H}\right) \tag{6.2.3}$$

where

$$H_{\mathbf{k}} = \sum_s \left\{ T_s \int \frac{|h_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v} - \frac{n_s e_s^2}{T_s} |\delta\varphi_{\mathbf{k}}|^2 \right\}, \tag{6.2.4}$$

$$\boxed{N_{\mathbf{k}}^{H} = B^{-2}\,\mathrm{Re} \sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}\times\mathbf{k}_1)\, T_s \int \frac{h_{s\mathbf{k}}^*}{F_s} J_{0s\mathbf{k}-\mathbf{k}_1} \delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},} \tag{6.2.5}$$

$$L_{\mathbf{k}}^{H} = \mathrm{Re} \sum_s T_s \int \frac{h_{s\mathbf{k}}^*}{F_{0s}} \mathcal{L}_{\mathbf{k}}[g_s] d\mathbf{v}. \tag{6.2.6}$$

Here H is the Helmholtz free energy, $N_{\mathbf{k}}^{H}$ its nonlinear drive, and $L_{\mathbf{k}}^{H}$ the linear drive. Note that the nonlinear contribution vanishes when the summation over all $\mathbf{k}$ is taken due to symmetry cancellations (Helander & Plunk 2022). That however, does not mean that energy is not redistributed between modes, e.g. the transfer of free energy to zonal modes and local transfer cascading dynamics are still contained within it (Bañón Navarro et al. 2011b; Helander & Plunk 2022). The naming of H can be understood by first re-writing equation 6.2.4 in terms of the gyro-averaged full perturbed distribution function $g_{s\mathbf{k}}$ from equation 6.1.3,

$$H_{\mathbf{k}} = \sum_s \left\{ T_s \int \frac{|g_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v} + \frac{n_s e_s^2}{T_s} [1 - \Gamma_{0s\mathbf{k}}(b_s)] |\delta\varphi_{\mathbf{k}}|^2 \right\}. \tag{6.2.7}$$

Here $\Gamma_0(x) = I_0(x)e^{-x}$ where I_0 is the so-called modified Bessel function of the first kind, and we define $b_s = k_\perp^2 \rho_s^2 = k_\perp^2 T_s/(m_s \Omega_s^2)$ (Helander & Plunk 2022), and we moved the details of this calculation to appendix D.3. We can now write

$$H_{\mathbf{k}} = U_{\mathbf{k}} - \sum_s T_s S_{s\mathbf{k}} \tag{6.2.8}$$

where the first term is the total spectral fluctuation energy

$$U_{\mathbf{k}} = \sum_s \frac{n_s e_s^2}{T_s} [1 - \Gamma_{0s\mathbf{k}}(b_s)] |\delta\varphi_{\mathbf{k}}|^2, \tag{6.2.9}$$

and the second term features the spectral entropy of each species (see appendix D.1)

$$S_{s\mathbf{k}} = -\int \frac{|g_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v}. \tag{6.2.10}$$

The three-wave-coupling in the nonlinear drive thus describes the redistribution of both fluctuation energy and entropy, i.e. the redistribution of Helmholtz free energy. The derived nonlinear drive term can thus be used to analyse the drive of *zonal modes of the Helmholtz free energy*.

Note that in some parts of the literature a different naming has been adopted and this transfer of Helmholtz free energy has somewhat confusingly been referred to as 'entropy transfer' (Sugama, Watanabe & Nunami 2009; Nakata, Watanabe & Sugama 2012; Maeyama et al. 2017).

Even in electromagnetic gyrokinetics, where an extra term must be added to $U_{\mathbf{k}}$ describing magnetic fluctuation (Helander & Plunk 2022), the total Helmholtz free energy is conserved (Schekochihin et al. 2009).

For the case of electrostatic turbulence, as we are considering here, the fluctuation energy and entropy are conserved independently (Bañón Navarro et al. 2011b; Bañón Navarro et al. 2011a). This was argued by Navarro *et al.* to hold for the 3D gyrokinetic case in these papers where `GENE` simulations were used in standard tokamak geometry (i.e. with parallel dynamics). They found that the power laws which they observe are very similar to those expected from strictly two dimensional turbulence. However, Plunk and Tatsuno argue in their work from equation 1.5.6 that strictly speaking, the fluctuation energy and entropy are only conserved separately in 2D electrostatic gyrokinetics, where the parallel streaming term is neglected. Even if not conserved exactly, such a separation shall still prove useful in the current context, similar to Refs. (Bañón Navarro et al. 2011b; Bañón Navarro et al. 2011a).

In the following two sub-sections we show how these can be separated, following the methods of (Bañón Navarro et al. 2011b; Bañón Navarro et al. 2011a) and use this separation to describe the different contributions to the gyrokinetic drive of the zonal modes.

6.2.2 Fluctuation energy transfer

The main goal of this section is to seek to isolate the nonlinear drive of the fluctuation energy $U_{\mathbf{k}}$ itself. To this end we apply the operator

$$\mathrm{Re}\sum_s e_s \int (\ldots) J_{0s\mathbf{k}} \delta\varphi_{\mathbf{k}}^* d\mathbf{v}$$

on equation 6.1.2. Using similar steps as in the previous subsection for the left-hand side[1] we see that we now obtain

$$\frac{\partial U_{\mathbf{k}}}{\partial t} = 2\left(N_{\mathbf{k}}^{U} + L_{\mathbf{k}}^{U}\right) \tag{6.2.11}$$

where the nonlinear and linear contributions are, respectively,

$$\boxed{N_{\mathbf{k}}^{U} = B^{-2}\,\mathrm{Re}\sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}\times\mathbf{k}_1)\, e_s \int J_{0s\mathbf{k}}\delta\varphi_{\mathbf{k}}^{*} J_{0s\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},} \tag{6.2.12}$$

$$L_{\mathbf{k}}^{U} = \mathrm{Re}\sum_{s} e_s \int J_{0s\mathbf{k}}\delta\varphi_{\mathbf{k}}^{*}\mathcal{L}_{\mathbf{k}}[g_s] d\mathbf{v}. \tag{6.2.13}$$

Equation 6.2.11 thus describes the temporal evolution of the spectral fluctuation energy. U only features the electrostatic potential in the electrostatic limit. Thus, by calculating $N_{\mathbf{k}}^{U}$ over all zonal modes for $\mathbf{k}$ we obtain the direct nonlinear drive of the zonal flow within the gyrokinetic framework. While the coupling takes place in the perpendicular plane, i.e. between mode pairs (k_x, k_y), in the same way as it does for the Reynolds stress, the individual components depend on velocity space. This offers the great advantage that $N_{\mathbf{k}}^{U}$ allows us to determine which species and which parts of velocity space are particularly involved in driving zonal flows. To name just one example, one could split the velocity integral from equation 6.2.12 in two parts, trapped and passing, giving the ZF drive from each group of particles for a given species. It is this potential for insight which motivated this work.

6.2.3 Entropy transfer

Finally, we demonstrate how the entropy transfer can be isolated from the free energy transfer. Here we use the operator

$$-\,\mathrm{Re}\sum_{s} T_s \int (\ldots)\frac{g_{s\mathbf{k}}^{*}}{F_s} d\mathbf{v}$$

on equation 6.1.2. We then trivially obtain

$$\frac{\partial}{\partial t}\sum_{s} T_s S_{s\mathbf{k}} = 2\left(N_{\mathbf{k}}^{S} + L_{\mathbf{k}}^{S}\right) \tag{6.2.14}$$

---

[1] specifically we use the quasi-neutrality relation from equation 6.2.2 and the integral relation from equation D.2.5

with

$$\boxed{N^S_{\mathbf{k}} = -B^{-2}\,\mathrm{Re}\sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}\times\mathbf{k}_1)\,T_s \int \frac{g^*_{s\mathbf{k}}}{F_s} J_{0s\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},} \tag{6.2.15}$$

$$L^S_{\mathbf{k}} = -\,\mathrm{Re}\sum_s T_s \int \frac{g^*_{s\mathbf{k}}}{F_s}\mathcal{L}_{\mathbf{k}}[g_s] d\mathbf{v}. \tag{6.2.16}$$

describing the nonlinear and linear drive of spectral entropy. This concludes the separation of the evolution for the full free energy. One can also convince oneself that indeed

$$N^H_{\mathbf{k}} = N^U_{\mathbf{k}} - N^S_{\mathbf{k}}$$
$$L^H_{\mathbf{k}} = L^U_{\mathbf{k}} - L^S_{\mathbf{k}}$$

as required.

6.3 Numerical implementation & computational challenges

As described in the previous subsection, the nonlinear transfer functions within (electrostatic) gyrokinetics is given by (see equations 6.2.5, 6.2.12, and 6.2.15)

$$N^H_{\mathbf{k}} = B^{-2}\,\mathrm{Re}\sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}\times\mathbf{k}_1)\,T_s \int \frac{h^*_{s\mathbf{k}}}{F_s} J_{0s\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},$$
$$N^U_{\mathbf{k}} = B^{-2}\,\mathrm{Re}\sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}_1\times\mathbf{k})\,e_s \int J_{0s\mathbf{k}}\delta\varphi^*_{\mathbf{k}} J_{0s\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},$$
$$N^S_{\mathbf{k}} = -B^{-2}\,\mathrm{Re}\sum_{s,\mathbf{k}_1} \mathbf{B}\cdot(\mathbf{k}\times\mathbf{k}_1)\,T_s \int \frac{g^*_{s\mathbf{k}}}{F_s} J_{0s\mathbf{k}-\mathbf{k}_1}\delta\varphi_{\mathbf{k}-\mathbf{k}_1} h_{s\mathbf{k}_1} d\mathbf{v},$$

which involves the three-wave coupling of the fields

$$\delta\varphi_{\mathbf{k}} \equiv \delta\varphi(t,k_x,k_y,\theta),$$
$$h_{s\mathbf{k}} \equiv h_s(t,k_x,k_y,\theta,\mu,\varepsilon,\sigma).$$

As we pointed out in section 6.2.2, this allows us to get additional insight into which parts of velocity space are particularly active in driving zonal modes.

This phase-space-resolved analysis also comes with increased numerical challenges. We outlined in section 2.3 how the calculation of the Reynolds stress via post-processing can quickly cause the required file sizes to be prohibitively large. This hindrance becomes yet stronger when writing out the distribution function $h_{s\mathbf{k}}$ for a potential calculation of $N^H_{\mathbf{k}}$, $N^U_{\mathbf{k}}$ or $N^S_{\mathbf{k}}$ via post-processing. We argued in section 2.3 that a typical file size for $\delta\varphi$ when written out every `nwrite` $= 20$

time steps is $\sim 10\,\mathrm{GB}$. For typical velocity space resolutions $N_\mu = 47$, $N_\varepsilon = 16$, and with $N_\sigma \equiv 2$ the file size of the distribution function would take

$$10\,\mathrm{GB} \, * \, 47 \, * \, 16 \, * \, 2 \sim 15\,\mathrm{TB}$$

per simulation. Thus, any calculation of these transfer functions must be done during the simulation.

Another complication is that, unlike $\delta\varphi_{\mathbf{k}}$, the distribution function $h_{s\mathbf{k}}$ is split across processors in `GS2`. Therefore, an implementation of the gyrokinetic transfer functions requires increased use of MPI communication. Further, since it now involves coupling with a full phase-space-resolved quantity, the distribution function, this analysis features $\mathcal{O}\left(N_\mu \, * \, N_\varepsilon\right)$ more operations than the calculation of the Reynolds stress from section 1.4.2.

6.3.1 Initial implementation in `GS2`

An initial implementation can be found at (Schuett 2025). This implementation is split into six diagnostics:

- `free_energy_transfer_4D` and `free_energy_transfer_theta_lambda`
- `field_energy_transfer_4D` and `field_energy_transfer_theta_lambda`
- `entropy_transfer_4D` and `entropy_transfer_theta_lambda`

The `4D` diagnostics resolve the transfer functions $N^H_{\mathbf{k}}$ (6.2.5), $N^U_{\mathbf{k}}$ (6.2.12) and $N^S_{\mathbf{k}}$ (6.2.15) as a function of $(\mathbf{k}, \mathbf{k}_1) = (k_x, k_y, k_{1x}, k_{1y})$, with the velocity space integration having been carried out and the outboard midplane selected as the poloidal location. *These diagnostics therefore include the transfer to zonal modes as well as all other transfers* and are thus useful to examine the dynamics in the perpendicular wavenumber space.

The `theta_lambda` diagnostics isolate the total transfer to zonal modes, i.e. $k_y = 0$ has been set and the contributions over (k_x, k_{1x}, k_{1y}) have been summed up. However, the velocity space integral in the expressions 6.2.5, 6.2.12 and 6.2.15 has now been taken outside this summation and only been carried out partially. That is, the integration over the energy dimension and the summation over the velocity sign have been carried out, but the integration over the trapping parameter λ has been omitted. Also, here all poloidal locations are resolved.[2] *Therefore, this diagnostic gives the total drive of zonal modes at each*

---

[2] While both θ and λ are resolved, the integration weight `wl`(θ, λ) has been applied, such that for example $\mathcal{T}(\theta)$ can be obtained simply by summation of the data along the λ-dimension.

poloidal location and from each particle class, e.g. enabling a distinction between contributions from trapped particles versus those from passing ones.

Initial results & current limitations

Initial results from this implementation are shown in figures 6.3.1 and 6.3.2. *As before in this thesis, the assumption of adiabatic electrons has been made.* Figure 6.3.1 shows the free energy transfer dynamics in the perpendicular plane at the outboard midplane, produced with the `free_energy_transfer_4D` analysis.

The left plot shows the transfer in radial wavenumber space that is mediated by zonal modes, i.e. $k_{2y} = k_y - k_{1y} = 0$. This transfer shows a forward cascading behaviour, i.e. free energy is transferred from larger scales to smaller scales, i.e. from smaller wavenumbers to larger wavenumbers. This can be seen from the fact that the positive transfer (red) region in the upper right quadrant is in the region $k_x > k_{1x}$, whereas the region $k_x < k_{1x}$ features the corresponding negative transfer (blue). The transfer in the lower left quadrant contains the same information and repeats in this fashion due to the symmetries of the transfer function. Such a forward cascade for the free energy was also observed in Ref. (Bañón Navarro et al. 2011b) for a shell-to-shell transfer-grouping approach.

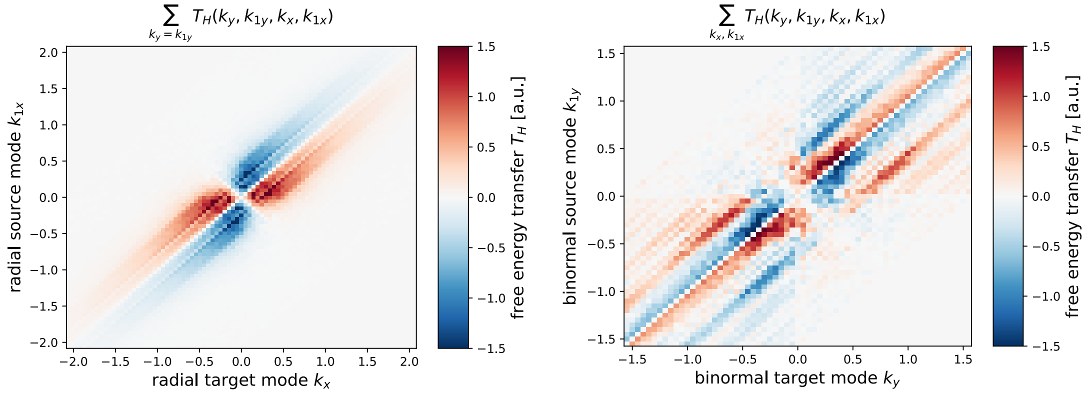

Figure 6.3.1. – Transfer of free energy at the outboard midplane. The chosen equilibrium is ($A = 2.72$; $\kappa = 1.5$; $\delta = 0.0$; $a/L_T = 6.0$). **Left:** the total transfer in radial wavenumber space mediated by zonal modes. **Right:** the total transfer in binormal wavenumber space.

The right plot in figure 6.3.1 shows the total transfer dynamics in the binormal wavenumber space. Transfer mediated by larger binormal structures, i.e. those with small $k_{2y} = k_y - k_{1y}$ and therefore close to the diagonal, shows a inverse cascade. Larger binormal mediators, i.e. transfer further away from the diagonal, produce a weaker forward cascade.

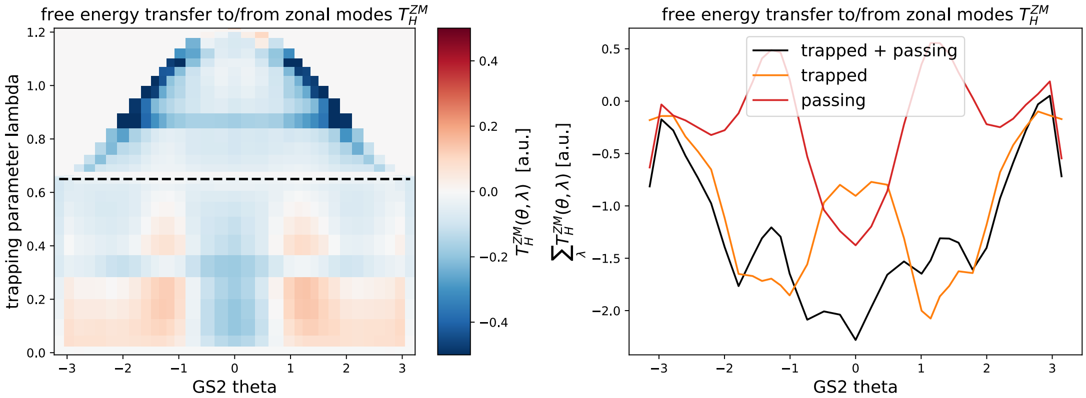

Figure 6.3.2. – The poloidal distribution of free energy transfer to and from zonal modes across the trapping parameter. The chosen equilibrium is ($A = 2.72$; $\kappa = 1.5$; $\delta = 0.0$; $a/L_T = 6.0$). **Left:** the fully resolved transfer over both the poloidal angle θ and the trapping parameter λ. The trapped-passing boundary is indicated. **Right:** The total poloidal distribution as well as the contributions from the passing and the trapped group.

Figure 6.3.2 shows the free energy transfer to and from zonal modes resolved over the poloidal angle and the trapping parameter. This result was created with the `free_energy_transfer_theta_lambda` analysis.

In the left plot, regions of both positive and negative transfer are seen, with virtually all positive transfer coming from the passing part of phase space. Particularly strong transfer appears to occur near the bounce points of the trapped particles, consistent with the physical picture that trapped particles spend a lot of their time near these points.

However, when summing the contributions over the trapping parameter, the poloidal distributions do not always appear smooth and symmetric with respect to the sign of the poloidal angle as one would expect. This is shown in the right plot of figure 6.3.2. Interestingly, the passing contribution appears both relatively smooth and symmetric, highlighting that in the case considered here the trapped contribution is responsible for the non-physical behaviour in the total distribution. We will discuss possible solutions to address this shortcoming next in the discussion.

6.4 Discussion and conclusion

With a particular focus on the drive of zonal modes we have derived the transfer functions of free energy, entropy, and fluctuation energy within gyrokinetics. Thereby, this chapter highlighted the theory of how a Reynolds-stress-based analysis of zonal flow drive can be expanded to include kinetic effects.

We have motivated why such an analysis must be implemented as a native diagnostic in a gyrokinetic code and presented an initial implantation in the gyrokinetic code `GS2`. Initial results from this diagnostic show clear cascading dynamics in the perpendicular domain, as previously reported in the literature (Bañón Navarro et al. 2011b). Further work is required to understand the origins of non-smooth and non-symmetric poloidal distributions of the total drive. Such smoothness and symmetry are a good test since they are physically expected but nothing is strictly enforcing this; all the sums over three-wave-couplings for the total drive are carried out independently at each poloidal location.

The summations of triple products across velocity space can introduce noise, and it might be the case that much higher resolutions are required to obtain more physical poloidal distributions. The poloidal distribution of the transfer might be more susceptible to such noise than e.g. the qualitative structure of the transfer in the perpendicular domain, which we observed to be physical.

6.4.1 Future work

Before a comparison between the spatial distributions of gyrokinetic transfer and fluid-based transfer – such as the Reynolds stress – can be carried out, future work first needs to identify the exact origins of some of the non-physical behaviour of the poloidal distributions and find solutions to address these. A natural first step would be to increase the resolutions substantially. However, since this analysis handles three-wave-couplings that are resolved in phase space, the memory requirements are already very large. The simulations that produced the current results used resolutions $(N_x \times N_y \times N_\theta \times N_\lambda \times N_\varepsilon = 128 \times 96 \times 33 \times 32 \times 16)$ which required most of the available memory of the compute nodes on the viking2 cluster. Thus, further memory optimisation of the diagnostics might be required first. Finally, since the gyrokinetic drive of the electrostatic potential is relatively weak, as alluded to in Ref. (Bañón Navarro et al. 2011b), it would be useful to gain understanding on the relative strength of the different channels through which energy ends up in the zonal flow within gyrokinetics.

We hope that the derived transfer functions and the initial implementation thereof in `GS2` guide future efforts on the analysis of the spatial distribution of gyrokinetic transfer dynamics.

6.4.2 Key findings and implications

The key points of this chapter and its implications are:

- Using suitable operators, nonlinear transfer functions of free energy, entropy, and fluctuation energy can be obtained directly from the nonlinearity in the gyrokinetic equation.
- This approach enables a velocity-space-resolved measurement of the drive of zonal structures within gyrokinetics.
- An initial diagnostic implementation of the transfer functions in `GS2` gives expected cascade dynamics in the perpendicular domain but unphysical poloidal distributions, thus requiring additional investigation before further conclusions can be drawn.

7

Discussion and conclusion

Zonal flows play a key role in determining the turbulent confinement in toroidal fusion plasmas. Moreover, experimental evidence points to ZFs as the trigger mechanism for confinement transitions in tokamaks. A missing piece of our understanding of zonal flows is how their nonlinear drive by turbulence is distributed across the flux surface. The main goal of this Thesis therefore was to gain insight into the spatial distribution of nonlinear zonal flow drive under varying toroidal geometry, including the geometry of newly designed hybrid equilibria, which enabled us to perform a step-by-step analysis of stellarator effects.

We highlighted the importance of the turbulent Reynolds stress as the key player in the convective nonlinearity of the momentum equation to transfer energy from turbulent modes to zonal modes. A numerical diagnostic that was previously developed as a post-processing script (Biggs-Fox 2022) was implemented as a native diagnostic in the gyrokinetic code `GS2` (Barnes et al. 2024) to explore these dependencies efficiently.

The contributions of this Thesis answer the questions:

1. How does the nonlinear drive of zonal flows vary spatially on the flux surface of realistic tokamaks?

   Unlike the turbulent fluctuations themselves, the drive of zonal flows depends strongly on the flux surface shaping, appearing to be strongest near locations of large poloidal curvature.

2. Can perturbed tokamaks be optimised to integrate stellarator benefits into a compact design?

   With a suitable combination of optimisation targets, significant stellarator shaping can be produced at a compact aspect ratio while maintaining excellent levels of confinement-preserving symmetry in the magnetic field as well as other desirable physics properties.

3. How does the highly localised shaping of these perturbed tokamaks affect the turbulent transport and zonal flow drive distribution?

   While the transport from ion temperature gradient turbulence remains broadly similar to that of the underlying tokamak, the localised three-dimensional shaping introduces additional local maxima in the zonal flow drive, adding to those appearing from the underlying axisymmetric shaping.

4. How can the fluid based Reynolds stress analysis of zonal flow drive be generalised within the gyrokinetic framework?

   Transfer functions for free energy, entropy, and fluctuation energy can be obtained directly from the nonlinearity of the gyrokinetic equation, enabling the analysis of the nonlinear transfer to their zonal component.

More in-depth answers to the questions are given in the same order below.

Poloidal distribution of zonal flow drive

We uncovered a non-trivial variation of the poloidal distribution of zonal flow drive with plasma shaping in chapter 2. Once the axisymmetric flux surfaces depart from circular, the global maximum of the drive at the outboard midplane is no longer guaranteed and often turns into a local minimum. We bridged the gap from circular shaping to realistic plasma shaping by first varying the dominant shaping contributions, elongation and triangularity, separately. For the variation of elongation we saw that instead local maxima emerge near locations of maximal poloidal curvature, with a consistent shift towards the bad curvature side. The combined apparent effect of both curvature and turbulence distributions was brought out further through the variation of triangularity. A clear asymmetry of the drive envelope with respect to the sign of triangularity was uncovered, as would be expected when both turbulence and curvature distributions are key components. For realistic plasma shaping, with combined elongation and triangularity, as observed on most tokamaks, the zonal flow drive distribution is more strongly affected by elongation than triangularity. This results in maxima near the top and bottom of the device with the outboard midplane being a local drive minimum. We saw that this is both the case in spherical and conventional tokamaks.

Not only do plasmas in experiment feature significant shaping, but also break the up-down asymmetry of the gyrokinetic equation through non-zero flow and flow shear as well as shaping. An asymmetric zonal flow drive distribution is seen in up-down asymmetric equilibria, leading us to the conclusion that the drive is not only expected to peak away from the midplane in experiments, but asymmetrically so.

We motivated the difficulty in understanding the exact mapping from the geometric quantities that enter the gyrokinetic equation to the zonal flow drive distribution, but showed that simple proxies perform reasonably well in predicting both the amount and poloidal position of the local maxima.

Optimising quasi-axisymmetrically perturbed tokamaks

We motivated why any three-dimensional (3D) shaping in toroidal plasmas must be optimised for collisionless guiding centre trajectories to be confined as in tokamaks. Tokamaks, while being simple, have many drawbacks which can be remedied by introducing 3D shaping. The introduction of optimised 3D shaping with significant external rotational transform has been a previously unexplored area, and was shown to be realisable as a stellarator-tokamak hybrid by (Henneberg & Plunk 2024).

A new optimisation target was developed which let us optimise this parameter space of compact ($A \sim 2.4$) global MHD equilibria numerically. Using this approach we achieved about an order of magnitude improvement in the quasi-axisymmetric quality in this parameter space, allowing us to introduce additional desirable physics properties such as a self-consistent current profile and improved magnetohydrodynamic stability. This approach expands our capabilities of optimising stellarators into the hybrid space and also generated new compact two-field period stellarators. The collisionless confinement of alpha particles showed to be negligible in most cases, performing better than traditional stellarators, despite the very compact aspect ratio.

Turbulent transport and zonal flow drive from perturbed tokamaks to stellarators

Since the equilibria we presented are realisable as a stellarator-tokamak hybrid, they allow us to study stellarator effects on tokamak dynamics step-by-step through continuous departure from axisymmetry. Before embarking on the impact of stellarator shaping on the zonal flow drive envelope, we discussed the impact on ion temperature gradient turbulence more generally. We showed that any small differences in the stability and transport which emerge when rotational transform profiles are held constant, are off-set by the additional magnetic shear that is introduced for self-consistent rotational transform profiles. We also highlighted a monotonic trend with the number of field periods, regardless of the values of magnetic shear, though consistent with the picture of flux surface compression in the bad curvature region.

Having found ways to optimise 3D-perturbed tokamaks and shedding some light on their general transport properties, we showed that the now parallel distribution of zonal flow drive is impacted in a non-trivial way by the localised inboard shaping of this parameter space. The analysis was applied to the gyrokinetic code `stella` and benchmarked successfully against `GS2` for the tokamak results, increasing our confidence in those results. Additional local maxima emerge, consistent with the amount of inboard-side ridges the flux tube crosses during one poloidal turn. Finally, we showed that the envelope of zonal flow drive is non-trivial in the most unstable flux tube of the largest stellarator experiment, Wendelstein 7-X. These results pave the way towards understanding total zonal flow drive dynamics in stellarators. A full flux surface analysis represents the natural, though computationally challenging, next step.

Zonal flow drive in the gyrokinetics framework

In the final chapter 6, we showed how the zonal flow drive analysis can be generalised from the Reynolds stress of the momentum equation to the nonlinearity of the gyrokinetic equation itself. We derived transfer functions for free energy, field energy, and entropy with the goal of isolating the quantity that describes the drive of zonal *flows*, which is the transfer of field energy. Though more work is required to obtain physical smooth poloidal distributions of the gyrokinetic transfer, the numerical challenges were laid out and an initial implementation in the gyrokinetic code `GS2` has been completed.

7.1 Closing comments

One of the key insights from this Thesis is the difference between the spatial distribution of the zonal flow drive by the turbulence and the distribution of the turbulence itself. This degeneracy is present in realistic tokamaks, in a traditional optimised stellarator, as well as in a hybrid of both. By uncovering this degeneracy, this work contributes to our understanding of zonal flows and turbulent confinement in tokamaks and stellarators. These results will aid the development of alternative and advanced confinement regimes for tokamaks, and inform validation efforts of gyrokinetic codes. We also hope that future work may be able to provide a simple physical picture for the observed behaviour. The derivations (and initial implementation thereof) presented in the final chapter contribute towards a purely gyrokinetics-based analysis of zonal flow drive, that holds the promise of additional insight.

Further, this work corroborated the attractiveness of the stellarator-tokamak hybrid design space by showing that quasisymmetry can be satisfied to sufficiently high levels for alpha particle confinement at compact, tokamak-like aspect ratio. The presented optimization method and resulting three-dimensional equilibria contribute towards further investigation of this design space for a fusion reactor concept that combines the advantages of tokamaks and stellarators. The unique feature of this design space to continuously connect tokamaks to optimised stellarators, and the ability to create equilibria within it through optimisation, enable further exploitation for theoretical studies. The specific turbulence and zonal flow drive analysis of this space presented in this Thesis represents a first such example.

A

Miller equations for up-down asymmetry

Miller equilibria with poloidal titling to introduce up-down asymmetry in a simple way was added to `GS2` as part of the work from references (Ball et al. 2014; Ball 2016). Then the flux surface shape is given by

$$R(\theta_{GS2}) = R_0 + R_\delta^{tilt}(\theta_{GS2}) \tag{A.0.1}$$

$$Z(\theta_{GS2}) = Z_\delta^{tilt}(\theta_{GS2}) \tag{A.0.2}$$

$$R_\delta^{tilt}(\theta_{GS2}) = R_\delta(\theta_{GS2}) \cos\theta_\delta + Z_\delta(\theta_{GS2}) \sin\theta_\delta \tag{A.0.3}$$

$$Z_\delta^{tilt}(\theta_{GS2}) = Z_\delta(\theta_{GS2}) \cos\theta_\delta - R_\delta(\theta_{GS2}) \sin\theta_\delta \tag{A.0.4}$$

$$R_\delta(\theta_{GS2}) = R_\kappa^{tilt}(\theta_{GS2}) + r[\cos(\theta_{GS2} + \delta\sin\theta_{GS2}) - \cos\theta_{GS2}] \tag{A.0.5}$$

$$Z_\delta(\theta_{GS2}) = Z_\kappa^{tilt}(\theta_{GS2}) \tag{A.0.6}$$

$$R_\kappa^{tilt}(\theta_{GS2}) = R_\kappa(\theta_{GS2}) \cos\theta_\kappa - \theta_\delta + Z_\kappa(\theta_{GS2}) \sin\theta_\kappa - \theta_\delta \tag{A.0.7}$$

$$Z_\kappa^{tilt}(\theta_{GS2}) = z_\kappa(\theta_{GS2}) \cos\theta_\kappa - \theta_\delta - R_\kappa(\theta_{GS2}) \sin\theta_\kappa - \theta_\delta \tag{A.0.8}$$

$$R_\kappa(\theta_{GS2}) = R_c(\theta_{GS2}) \tag{A.0.9}$$

$$Z_\kappa(\theta_{GS2}) = Z_c(\theta_{GS2}) + (\kappa - 1) r \sin\theta_{GS2} + \theta_\kappa - \theta_\delta \tag{A.0.10}$$

$$R_c(\theta_{GS2}) = r\cos(\theta_{GS2} + \theta_\kappa - \theta_\delta) \tag{A.0.11}$$

$$Z_c(\theta_{GS2}) = r\sin(\theta_{GS2} + \theta_\kappa - \theta_\delta) \tag{A.0.12}$$

where the two new parameters are θ_κ and θ_δ which describe the poloidal tilt angle of elongation and triangularity, respectively.

B

Poloidal and toroidal curvatures of Miller equilibria

If the coordinate curves on a surface are lines of curvatures then the principle curvatures are given by

$$\kappa_1 = \frac{b_{\phi\phi}}{g_{\phi\phi}} \quad \text{and} \quad \kappa_2 = \frac{b_{\theta\theta}}{g_{\theta\theta}}, \tag{B.0.1}$$

where

$$g_{\alpha\beta} = \mathbf{x}_\alpha \cdot \mathbf{x}_\beta \tag{B.0.2}$$

are the coefficients of the first fundamental form and

$$b_{\alpha\beta} = \mathbf{x}_{\alpha\beta} \cdot \mathbf{n} \tag{B.0.3}$$

are the coefficients of the second fundamental form (Kreyszig 1959). This is the case for an axisymmetric torus[1]. Here

$$\mathbf{n} = \frac{\mathbf{x}_\phi \times \mathbf{x}_\theta}{|\mathbf{x}_\phi \times \mathbf{x}_\theta|} \tag{B.0.4}$$

is the surface normal and the partial derivatives of the position vector $\mathbf{x}(\theta, \phi)$ are defined as (Kreyszig 1959)

$$\mathbf{x}_\alpha = \frac{\partial \mathbf{x}}{\partial \alpha} \quad \text{and} \quad \mathbf{x}_{\alpha\beta} = \frac{\partial^2 \mathbf{x}}{\partial \alpha \partial \beta}, \tag{B.0.5}$$

---

[1] This becomes apparent from the fact that $g_{\phi\theta} = b_{\phi\theta} = 0$ (Kreyszig 1959, Theorem 41.3, see also section 43)

where $\alpha, \beta \in \{\phi, \theta\}$. For up-down symmetric Miller equilibria the position vector is given by

$$\mathbf{x}(\theta, \phi) = \begin{bmatrix} R(\theta)\cos(\phi) \\ R(\theta)\sin(\phi) \\ Z(\theta) \end{bmatrix}, \tag{B.0.6}$$

$$R(\theta) = R_0 + r\cos(\theta + \delta\sin(\theta)), \tag{B.0.7}$$

$$Z(\theta) = \kappa r \sin(\theta). \tag{B.0.8}$$

While the calculation to obtain κ_1 and κ_2 is simple, it is also tedious and was therefore performed with a computer algebra system, `Mathematica 12`. The final result is given by

$$\kappa_1(\theta, \phi) = \frac{a(\theta, \phi)}{b(\theta, \phi)} \tag{B.0.9}$$

for the toroidal curvature, where

$$\begin{aligned} a(\theta, \phi) &= 2\kappa\cos(\theta), \\ b(\theta, \phi) &= c(\theta, \phi)(r\cos(\delta\sin(\theta) + \theta) + R_0), \\ c(\theta, \phi) &= [\delta^2 + 2\kappa^2 + \left(\delta^2 + 2\kappa^2\right)\cos(2\theta) - \frac{1}{2}\delta^2\cos(2\delta\sin(\theta)) \\ &\quad - \delta^2\cos(2(\delta\sin(\theta) + \theta)) - \frac{1}{2}\delta^2\cos(2\delta\sin(\theta) + 4\theta) \\ &\quad + 4\delta\cos(\theta) - 2\delta\cos(2\delta\sin(\theta) + \theta) - 2\delta\cos(2\delta\sin(\theta) + 3\theta) \\ &\quad - 2\cos(2(\delta\sin(\theta) + \theta)) + 2]^{1/2}, \end{aligned} \tag{B.0.10}$$

and the poloidal curvature is given by

$$\kappa_2(\theta, \phi) = \frac{d(\theta, \phi)}{e(\theta, \phi)}, \tag{B.0.11}$$

where

$$\begin{aligned}
d(\theta,\phi) &= \kappa f(\theta,\phi),\\
f(\theta,\phi) &= \left(3\delta^2+8\right)\cos(\delta\sin(\theta))+\delta g(\theta,\phi),\\
g(\theta,\phi) &= 4\cos(\theta-\delta\sin(\theta))\\
&\quad+\delta\cos(2\theta-\delta\sin(\theta))+8\cos(\delta\sin(\theta)+\theta)\\
&\quad+3\delta\cos(\delta\sin(\theta)+2\theta)+4\cos(\delta\sin(\theta)+3\theta)\\
&\quad+\delta\cos(\delta\sin(\theta)+4\theta),\\
e(\theta,\phi) &= 2rh(\theta,\phi)k(\theta,\phi),\\
h(\theta,\phi) &= [4\left(\delta^2+2\kappa^2\right)\cos(2\theta)-2l(\theta,\phi)+16\delta\cos(\theta)]^{1/2},\\
l(\theta,\phi) &= -2\delta^2-4\kappa^2+\delta^2\cos(2\delta\sin(\theta))\\
&\quad+\delta^2\cos(2\delta\sin(\theta)+4\theta)+2\left(\delta^2+2\right)\cos(2(\delta\sin(\theta)+\theta))\\
&\quad+4\delta\cos(2\delta\sin(\theta)+\theta)+4\delta\cos(2\delta\sin(\theta)+3\theta)-4,\\
k(\theta,\phi) &= \cos^2(\phi)\sin^2(\delta\sin(\theta)+\theta)(\delta\cos(\theta)+1)^2\\
&\quad+\sin^2(\phi)\sin^2(\delta\sin(\theta)+\theta)(\delta\cos(\theta)+1)^2+\kappa^2\cos^2(\theta).
\end{aligned} \tag{B.0.12}$$

These equations are given here for completeness, though they are cumbersome and do not give particular insight. However, it is worth pointing out that in the circular limit where $\kappa=1.0$ and $\delta=0$ we recover the known formulae for the circular torus (Kreyszig 1959, p. 135)

$$\kappa_1=\pm\frac{\cos(\theta)}{r\cos(\theta)+R_0} \quad \text{and} \quad \kappa_2=\pm\frac{1}{r}. \tag{B.0.13}$$

The surface curvatures for the up-down asymmetric cases were obtained analo-

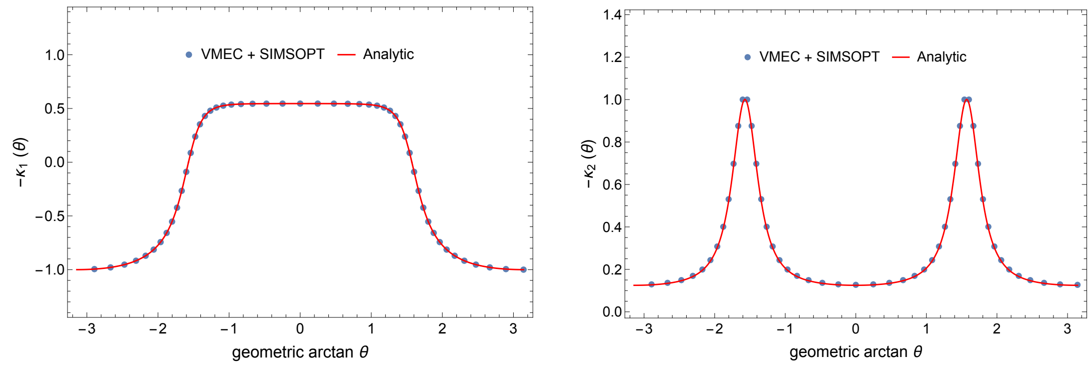

Figure B.0.1. – Comparison of the poloidal distributions of the principal curvatures from the analytic formulae (red) and the numerical calculation from the general three-dimensional codes `SIMSOPT` & `VMEC` (blue). Here $A=2.72$, $\kappa=2.0$, $\delta=0$, and $\rho=0.8$. **Left:** The toroidal curvature κ_1. **Right:** The poloidal curvature κ_2.

gously by replacing equations B.0.7 and B.0.8 with equations A.0.1 and A.0.2. Figure B.0.1 shows a comparison between the analytic result and the one computed numerically by the general three-dimensional codes `VMEC` (Hirshman & Whitson 1983) and `SIMSOPT` (Landreman et al. 2021). We observe excellent agreement between both for both κ_1 and κ_2. Since the procedure of finding `VMEC` equilibria with the exact same flux-surface shape as specified by the Miller equations is more tedious once the triangularity is finite, we use the analytic procedure for all equilibria in this paper.

C

Numerical details of stellarator turbulence simulations

Here we provide the details of the simulations from chapter 4. The simulations are for collisionless ITG turbulence, and are therefore run electrostatic with adiabatic electrons.

C.1 Linear simulations

The linear simulations are performed with the `stella` code v0.7 (commit `f64addc`). The simulations use the velocity space resolutions $N_{v_\parallel} = 24$ and $N_\mu = 12$ with domain bounds $v_{\parallel,c} = v_{\perp,c} = 3v_{\mathrm{th,i}}$. The parallel (ballooning) domain corresponds to two poloidal turns with $N_z = 256$, i.e. 128 points per poloidal turn [similar resolutions were used in the benchmark by (González-Jerez et al. 2022)]. Simulations are run until $t_{\mathrm{end}} = 500\, a/v_{\mathrm{th}}$ with a time-stepping of $\Delta t = 0.15$ and growth rates are averaged over the last quarter of the simulation. The simulations use $k_x = 0$.

Resolution checks for the growth rate spectra were performed at the highest and lowest temperature gradient for each configuration by doubling any of the velocity space grid resolutions, halving the time step, doubling the simulation time, doubling the parallel resolution, and doubling the parallel extent of the domain while increasing N_z correspondingly to keep at least 128 parallel grid points per poloidal turn. Moreover, it was checked that the mode envelopes, $\delta\varphi_{t=t_{\mathrm{end}}}(z, k_y)$, are zero at the parallel boundaries for each of the simulated binormal wavenumbers k_y, such that the ballooning formalism is valid.

C.2 Nonlinear simulations

The nonlinear simulations are performed with the `GX` code from commit `5ea4f95`. The simulations use the velocity space resolutions $N_{\mathrm{hermite}} = 8$ and $N_{\mathrm{laguerre}} = 4$.

The flux tube extends for one poloidal turn with $N_\theta = 96$ parallel points. The perpendicular domain is approximately square at the outboard midplane with $L_x \sim L_y \sim 82\rho_i$ (`y0` $= 13$) with perpendicular real-space grid resolutions $N_x = N_y = 64$. Simulations are run until $t_{\text{end}} = 800\, a/v_{\text{th}}$ and heat fluxes are averaged over the saturated period, which typically corresponds to $t > 250\, a/v_{\text{th}}$. Similar resolutions were used in the study by (Landreman et al. 2025).

Resolution checks for the time-averaged heat flux were performed by increasing any of the resolutions in the velocity space, the parallel coordinate, or in the perpendicular real-space domain by at least 50%.

D

Useful relations in gyrokinetics

D.1 Entropy of distribution function

To goal of this section is to show that

$$S_{s\mathbf{k}} = -\int \frac{|g_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v}. \tag{D.1.1}$$

from equation 6.2.10 corresponds to the entropy of the fluctuations in the distribution function. This has been shown before in less detail, see for example (Sugama & Horton 1995; Sugama, Watanabe & Nunami 2009; Schekochihin et al. 2009).

To this end we write the gyro-average of the full distribution function f_s from equation 1.5.2 as

$$\tilde{f} = \langle f_s \rangle_s^{\mathbf{R}} = F_s + \langle \delta f_s \rangle_s^{\mathbf{R}} = F_s + g_s$$

The Gibbs entropy for the full distribution of gyro-centres $\tilde{f}$ is given by its well-known general expression (up to a Boltzmann constant)

$$\tilde{S}_s = -\left\langle \int d\mathbf{v} \tilde{f} \ln \tilde{f} \right\rangle$$

where $\langle \dots \rangle$ represents ensemble-averaging. We then expand it in the gyrokinetic small parameter ϵ, giving

$$\begin{aligned}
\tilde{S}_s &= -\left\langle \int d\mathbf{v}(F_s + g_s) \left[\ln\left(1 + g_s/F_s\right) + \ln F_s\right] \right\rangle \\
&= -\left\langle \int d\mathbf{v}(F_s + g_s) \left[(g_s/F_s) - (1/2)(g_s/F_s)^2 + \ln F_s\right] + \mathcal{O}(\epsilon^3) \right\rangle \\
&= -\left\langle \int d\mathbf{v} g_s^2/2F_s + (1 + \ln F_s) g_s + F_s \ln F_s \right\rangle + \mathcal{O}(\epsilon^3) \\
&= -\int d\mathbf{v} F_s \ln F_s - \int d\mathbf{v} \left\langle g_s^2 \right\rangle /2F_s + \mathcal{O}(\epsilon^3)
\end{aligned}$$

where we used that

$$\langle g_s \rangle = 0.$$

Thus $S_{s\mathbf{k}}$ from equation D.1.1 is the only contribution from fluctuations to the full entropy, adding to the macroscopic contribution from the background Maxwellian F_s.

D.2 Re-writing quasi-neutrality

Here we show a convenient way of re-writing the quasi-neutrality condition. In its original form it reads

$$\sum_s \frac{n_s e_s^2}{T_s}\delta\varphi_{\mathbf{k}} = \sum_s e_s \int h_{s\mathbf{k}} J_{0s} d\mathbf{v}. \tag{D.2.1}$$

Upon replacing $h_{s\mathbf{k}}$ by $g_{s\mathbf{k}}$ with equation D.3.3 we obtain

$$\sum_s \frac{n_s e_s^2}{T_s}\delta\varphi_{\mathbf{k}} = \sum_s e_s \left[\int g_{s\mathbf{k}} J_{0s} d\mathbf{v} + \frac{e_s}{T_s}\delta\varphi_{\mathbf{k}} \int F_s J_{0s\mathbf{k}}^2 \, d\mathbf{v}\right]. \tag{D.2.2}$$

We focus on the second integral

$$\int F_s J_{0s\mathbf{k}}^2 \, d\mathbf{v}, \tag{D.2.3}$$

and specify the velocity coordinates as

$$d\mathbf{v} = v_\perp d\alpha \, dv_\perp dv_\parallel = 2\pi \, v_\perp dv_\perp \, dv_\parallel \tag{D.2.4}$$

where the gyro-angle $\alpha \in [0, 2\pi)$, $v_\perp \in [0, \infty)$, and $v_\parallel \in (-\infty, \infty)$, and the gyro-angle integration will always yield 2π since by definition all quantities in the gyrokinetic equation are gyro-averaged and thus do not depend on α.

The integral yields

$$\begin{aligned}\int F_s J_{0s\mathbf{k}}^2 \, d\mathbf{v} &= n_s \left(\frac{m_s}{2\pi T_s}\right)^{3/2} \int J_0^2(k_\perp v_\perp/\Omega_s) e^{-(mv^2/2T_s)} \, d\mathbf{v} \\ &= n_s \left(\frac{m_s}{2\pi T_s}\right)^{3/2} \left\{2\pi \int_{-\infty}^{\infty} e^{-mv_\parallel^2/(2T_s)} dv_\parallel \int_0^\infty J_0^2(k_\perp v_\perp/\Omega_s) e^{-mv_\perp^2/(2T_s)} v_\perp dv_\perp\right\} \\ &= n_s \left(\frac{m_s}{2\pi T_s}\right)^{3/2} \left\{2\pi \sqrt{\frac{2\pi T_s}{m_s}} \int_0^\infty J_0^2\left(x\sqrt{2b_s}\right) e^{-x^2} (2T_s/m_s) x dx\right\} \\ &= n_s \Gamma_0(b_s).\end{aligned} \tag{D.2.5}$$

Here we used the substitution $x = \sqrt{m_s v_\perp^2/(2T_s)}$ and introduced $b_s = k_\perp^2 \rho_s^2 = k_\perp^2 T_s/(m_s \Omega_s^2)$. In the final step we used the integral identity (Helander & Plunk 2022)

$$\Gamma_0(b) = 2\int_0^\infty J_0^2(x\sqrt{2b}) e^{-x^2} x dx, \tag{D.2.6}$$

where $\Gamma_0(b) = I_0(b)e^{-b}$ and I_0 is the so-called modified Bessel functions of the first kind.

We can now write the quasi-neutrality equation D.2.2 as

$$\sum_s \frac{n_s e_s^2}{T_s}[1 - \Gamma_0(b_s)]\delta\varphi_{\mathbf{k}} = \sum_s e_s \int g_{s\mathbf{k}} J_{0s} d\mathbf{v}. \tag{D.2.7}$$

D.3 Re-writing the Helmholtz free energy

Here we show how to transform the Helmholtz free energy with the non-adiabatic distribution function,

$$H_{\mathbf{k}} = \sum_s \left\{ T_s \int \frac{|h_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v} - \frac{n_s e_s^2}{T_s} |\delta\varphi_{\mathbf{k}}|^2 \right\}, \tag{D.3.1}$$

into a form where it features the gyro-averaged full perturbed distribution function,

$$H_{\mathbf{k}} = \sum_s \left\{ T_s \int \frac{|g_{s\mathbf{k}}|^2}{F_{0s}} d\mathbf{v} + \frac{n_s e_s^2}{T_s} [1 - \Gamma_{0s\mathbf{k}}(b_s)] |\delta\varphi_{\mathbf{k}}|^2 \right\}. \tag{D.3.2}$$

We start by using the definition of the gyro-averaged full perturbed distribution function

$$g_{s\mathbf{k}} = h_{s\mathbf{k}} - \frac{e_s F_s}{T_s} J_{0s\mathbf{k}} \delta\varphi_{\mathbf{k}}. \tag{D.3.3}$$

Plugging this into equation D.3.1 for h_{sk} gives for the first term

$$\sum_s T_s \int \left(\frac{|g_{s\mathbf{k}}|^2}{F_s} + g_{s\mathbf{k}}^* \frac{e_s}{T_s} J_{0s\mathbf{k}} \delta\varphi_{\mathbf{k}} + g_{s\mathbf{k}} \frac{e_s}{T_s} J_{0s\mathbf{k}} \delta\varphi_{\mathbf{k}}^* + \frac{e_s^2 F_s}{T_s^2} J_{0s\mathbf{k}}^2 |\delta\varphi_{\mathbf{k}}|^2 \right) d\mathbf{v}. \tag{D.3.4}$$

We use the quasi-neutrality equation in the following form (see equation D.2.7 in the previous quasi-neutrality section D.2)

$$\sum_s \frac{n_s e_s^2}{T_s}[1 - \Gamma_0(b_s)]\delta\varphi_{\mathbf{k}} = \sum_s e_s \int g_{s\mathbf{k}} J_{0s} d\mathbf{v}. \tag{D.3.5}$$

to replace the second and third term on the right-hand side of equation D.3.4,

$$\sum_s T_s \int \left(\frac{|g_{s\mathbf{k}}|^2}{F_s} + \frac{e_s^2 F_s}{T_s^2} J_{0s\mathbf{k}}^2 |\delta\varphi_{\mathbf{k}}|^2 \right) d\mathbf{v} \; + 2 \sum_s \frac{n_s e_s^2}{T_s}[1 - \Gamma_0(b_s)] |\delta\varphi_{\mathbf{k}}|^2. \tag{D.3.6}$$

We also note that we can re-write the fourth term on the right-hand side of equation D.3.4 as

$$\sum_s \frac{e_s^2}{T_s}|\delta\varphi_{\mathbf{k}}|^2 \int J_{0s\mathbf{k}}^2 F_s d\mathbf{v} = \sum_s \frac{n_s e_s^2}{T_s}\Gamma_0(b_s)|\delta\varphi_{\mathbf{k}}|^2.$$

The first term D.3.6 of the Helmholtz free energy thus becomes

$$\sum_s T_s \int \frac{|g_{s\mathbf{k}}|^2}{F_s} + \sum_s \frac{n_s e_s^2}{T_s}[2-\Gamma_0(b_s)]\delta\varphi_{\mathbf{k}}|^2.$$

such that including the second term of equation D.3.1, the full Helmholtz free energy becomes

$$H_{\mathbf{k}} = \sum_s T_s \int \frac{|g_{s\mathbf{k}}|^2}{F_s} d\mathbf{v} + \sum_s \frac{n_s e_s^2}{T_s}[1-\Gamma_0(b_s)]\delta\varphi_{\mathbf{k}}|^2.$$

which is the same as equation D.3.2 and thus concludes the proof.

Bibliography

Abel, I G, G G Plunk, E Wang, M Barnes, S C Cowley, W Dorland, and A A Schekochihin (Oct. 2013). Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows. In: *Reports on Progress in Physics* 76.11, p. 116201.

Albert, Christopher G., Sergei V. Kasilov, and Winfried Kernbichler (2020a). Accelerated methods for direct computation of fusion alpha particle losses within, stellarator optimization. In: *Journal of Plasma Physics* 86.2, p. 815860201.

Albert, Christopher G., Sergei V. Kasilov, and Winfried Kernbichler (2020b). Symplectic integration with non-canonical quadrature for guiding-center orbits in magnetic confinement devices. In: *Journal of Computational Physics* 403, p. 109065. ISSN: 0021-9991.

Anderson, F. Simon B., Abdulgader F. Almagri, David T. Anderson, Peter G. Matthews, Joseph N. Talmadge, and J. Leon Shohet and (1995). The Helically Symmetric Experiment, (HSX) Goals, Design and Status. In: *Fusion Technology* 27.3T, pp. 273–277.

Angelino, P., X. Garbet, L. Villard, A. Bottino, S. Jolliet, Ph. Ghendrih, V. Grandgirard, B. F. McMillan, Y. Sarazin, G. Dif-Pradalier, and T. M. Tran (May 2009). Role of Plasma Elongation on Turbulent Transport in Magnetically Confined Plasmas. In: *Phys. Rev. Lett.* 102 (19), p. 195002.

Antonsen T. M., Jr., J. F. Drake, P. N. Guzdar, A. B. Hassam, Y. T. Lau, C. S. Liu, and S. V. Novakovskii (June 1996). Physical mechanism of enhanced stability from negative shear in tokamaks: Implications for edge transport and the L-H transition. In: *Physics of Plasmas* 3.6, pp. 2221–2223. ISSN: 1070-664X.

Bader, A., D.T. Anderson, M. Drevlak, B.J. Faber, C.C. Hegna, S. Henneberg, M. Landreman, J.C. Schmitt, Y. Suzuki, and A. Ware (Oct. 2021). Modeling of energetic particle transport in optimized stellarators. In: *Nuclear Fusion* 61.11, p. 116060.

Bähner, J.-P., J.A. Alcusón, S.K. Hansen, A. von Stechow, O. Grulke, T. Windisch, H.M. Smith, Z. Huang, E.M. Edlund, M. Porkolab, and et al. (2021). Phase contrast imaging measurements and numerical simulations of turbulent density fluctuations in gas-fuelled ECRH discharges in Wendelstein 7-X. In: *Journal of Plasma Physics* 87.3, p. 905870314.

Ball, Justin (2016). GS2 analytic geometry specification. `https://bitbucket.org/gyrokinetics/wikifiles/raw/master/JRB/GS2GeoDoc.pdf`. Accessed: 2023-10-27.

Ball, Justin, Felix I Parra, and Michael Barnes (Feb. 2016). Poloidal tilting symmetry of high order tokamak flux surface shaping in gyrokinetics. In: *Plasma Physics and Controlled Fusion* 58.4, p. 045023.

Ball, Justin, Felix I Parra, Michael Barnes, William Dorland, Gregory W Hammett, Paulo Rodrigues, and Nuno F Loureiro (July 2014). Intrinsic momentum transport in up–down asymmetric tokamaks. In: *Plasma Physics and Controlled Fusion* 56.9, p. 095014.

Bañón Navarro, A., P. Morel, M. Albrecht-Marc, D. Carati, F. Merz, T. Görler, and F. Jenko (Sept. 2011a). Free energy balance in gyrokinetic turbulence. In: *Physics of Plasmas* 18.9, p. 092303. ISSN: 1070-664X.

Bañón Navarro, A., P. Morel, M. Albrecht-Marc, D. Carati, F. Merz, T. Görler, and F. Jenko (Jan. 2011b). Free Energy Cascade in Gyrokinetic Turbulence. In: *Phys. Rev. Lett.* 106 (5), p. 055001.

Barnes, M., F.I. Parra, and M. Landreman (2019). stella: An operator-split, implicit–explicit δf-gyrokinetic code for general magnetic field configurations. In: *Journal of Computational Physics* 391, pp. 365–380. ISSN: 0021-9991.

Barnes, Michael, David Dickinson, William Dorland, Peter Alec Hill, Joseph Thomas Parker, Colin Malcolm Roach, Maurizio Giacomin, Noah Mandell, Ryusuke Numata, Tobias Schuett, Stephen Biggs-Fox, Nicolas Christen, Jason Parisi, George Wilkie, Lucian Anton, Justin Ball, Jessica Baumgaertel, Greg Colyer, Michael Hardman, Joachim Hein, Edmund Highcock, Gregory Howes, Adrian Jackson, Michael T. Kotschenreuther, Jungpyo Lee, Huw Leggate, Adwiteey Mauriya, Bhavin Patel, Tomo Tatsuno, and Ferdinand Van Wyk (July 2024). GS2 v8.2.0. Version 8.2.0.

Bauer, Frances, Octavio Betancourt, and Paul Garabedian (2012). Magnetohydrodynamic equilibrium and stability of stellarators. Springer Science & Business Media.

Beer, M. A., S. C. Cowley, and G. W. Hammett (July 1995). Field-aligned coordinates for nonlinear simulations of tokamak turbulence. In: *Physics of Plasmas* 2.7, pp. 2687–2700. ISSN: 1070-664X.

Belli, E. A., G. W. Hammett, and W. Dorland (Sept. 2008). Effects of plasma shaping on nonlinear gyrokinetic turbulence. In: *Physics of Plasmas* 15.9, p. 092303. ISSN: 1070-664X.

Biggs-Fox, Stephen Neil (Apr. 2022). “Spatial structure of micro-instabilities in tokamak plasmas: zonal flows and global effects in local gyrokinetic simulations”. PhD thesis. University of York.

Biglari, H., P. H. Diamond, and P. W. Terry (Jan. 1990). Influence of sheared poloidal rotation on edge turbulence. In: *Physics of Fluids B: Plasma Physics* 2.1, pp. 1–4. ISSN: 0899-8221.

Bindel, David, Matt Landreman, and Misha Padidar (May 2023a). Direct Optimization of Fast-Ion Confinement in Stellarators. In: *Plasma Physics and Controlled Fusion* 65.6, p. 065012.

Bindel, David, Matt Landreman, and Misha Padidar (2023b). Understanding trade-offs in stellarator design with multi-objective optimization. In: *Journal of Plasma Physics* 89.5, p. 905890503.

Boozer, Allen H. (Jan. 2005). Physics of magnetically confined plasmas. In: *Rev. Mod. Phys.* 76 (4), pp. 1071–1141.

Brower, D. L., W. A. Peebles, S. K. Kim, N. C. Luhmann, W. M. Tang, and P. E. Phillips (July 1987). Observation of a high-density ion mode in tokamak microturbulence. In: *Phys. Rev. Lett.* 59 (1), pp. 48–51.

Buller, S. (2024). Exploring the space of compact quasi-axisymmetric equilibria with multi-objective optimization. *International Stellarator and Heliotron Workshop*.

Burrell, K. H. (May 1997). Effects of E×B velocity shear and magnetic shear on turbulence and transport in magnetic confinement devices. In: *Physics of Plasmas* 4.5, pp. 1499–1518. ISSN: 1070-664X.

Camargo, Suzana J., Dieter Biskamp, and Bruce D. Scott (Jan. 1995). Resistive drift-wave turbulence. In: *Physics of Plasmas* 2.1, pp. 48–62. ISSN: 1070-664X.

Carralero, D., J. M. Estrada T. García-Regaña, E. Sánchez, T. Windisch, A. Alonso, E. Maragkoudakis, C. Brandt, K. J. Brunner, C. Gallego-Castillo, P. Kornejew, K. Parks, K. Rahbarnia, and H. Thienpondt (Apr. 2025). First experimental observation of zonal flows in the optimized stellarator Wendelstein 7-X. In: *Phys. Rev. Res.* 7 (2), p. L022009.

Choi, D-I and Wendell Horton (1980). Weakly localized two-dimensional drift modes. In: *Physics of Fluids* 23.2, pp. 356–365.

Conlin, Rory, Patrick Kim, Daniel W. Dudt, Dario Panici, and Egemen Kolemen (2024). Stellarator optimization with constraints. In: *Journal of Plasma Physics* 90.5, p. 905900501.

Cziegler, I., P. H. Diamond, N. Fedorczak, P. Manz, G. R. Tynan, M. Xu, R. M. Churchill, A. E. Hubbard, B. Lipschultz, J. M. Sierchio, J. L. Terry, and C. Theiler (May 2013). Fluctuating zonal flows in the I-mode regime in Alcator C-Mod. In: *Physics of Plasmas* 20.5, p. 055904. ISSN: 1070-664X.

Cziegler, I., A. E. Hubbard, J. W. Hughes, J. L. Terry, and G. R. Tynan (Mar. 2017). Turbulence Nonlinearities Shed Light on Geometric Asymmetry in Tokamak Confinement Transitions. In: *Phys. Rev. Lett.* 118 (10), p. 105003.

Cziegler, I., G.R. Tynan, P.H. Diamond, A.E. Hubbard, J.W. Hughes, J. Irby, and J.L. Terry (June 2015). Nonlinear transfer in heated L-modes approaching the L–H transition threshold in Alcator C-Mod. In: *Nuclear Fusion* 55.8, p. 083007.

Cziegler, Istvan (2011). “Turbulence and transport phenomena in edge and scrape-off-layer plasmas”. PhD thesis. Massachusetts Institute of Technology.

Davies, R, D Dickinson, and H Wilson (Aug. 2022). Kinetic ballooning modes as a constraint on plasma triangularity in commercial spherical tokamaks. In: *Plasma Physics and Controlled Fusion* 64.10, p. 105001.

Diamond, P H, S-I Itoh, K Itoh, and T S Hahm (Apr. 2005). Zonal flows in plasma—a review. In: *Plasma Physics and Controlled Fusion* 47.5, R35.

Diamond, P. H. and Y.-B. Kim (July 1991). Theory of mean poloidal flow generation by turbulence. In: *Physics of Fluids B: Plasma Physics* 3.7, pp. 1626–1633. ISSN: 0899-8221.

Dickinson, David (Sept. 2012). “Effects of profiles on microinstabilities in tokamaks”. PhD thesis. University of York.

Drevlak, M., F. Brochard, P. Helander, J. Kisslinger, M. Mikhailov, C. Nührenberg, J. Nührenberg, and Y. Turkin (2013). ESTELL: A Quasi-Toroidally Symmetric Stellarator. In: *Contributions to Plasma Physics* 53.6, pp. 459–468.

Dudding, H.G., F.J. Casson, D. Dickinson, B.S. Patel, C.M. Roach, E.A. Belli, and G.M. Staebler (July 2022). A new quasilinear saturation rule for tokamak turbulence with application to the isotope scaling of transport. In: *Nuclear Fusion* 62.9, p. 096005.

Dudding, Harry G. (Oct. 2022). “A new quasilinear saturation rule for tokamak turbulence”. PhD thesis. University of York.

Enters, Y. W., S. Thomas, M. Hill, and I. Cziegler (July 2023). Testing image-velocimetry methods for turbulence diagnostics. In: *Review of Scientific Instruments* 94.7, p. 075101. ISSN: 0034-6748.

Evans, T. E., R. A. Moyer, P. R. Thomas, J. G. Watkins, T. H. Osborne, J. A. Boedo, E. J. Doyle, M. E. Fenstermacher, K. H. Finken, R. J. Groebner, M. Groth, J. H. Harris, R. J. La Haye, C. J. Lasnier, S. Masuzaki, N. Ohyabu, D. G. Pretty, T. L. Rhodes, H. Reimerdes, D. L. Rudakov, M. J. Schaffer, G. Wang, and L. Zeng (June 2004). Suppression of Large Edge-Localized Modes in High-Confinement DIII-D Plasmas with a Stochastic Magnetic Boundary. In: *Phys. Rev. Lett.* 92 (23), p. 235003.

Freidberg, Jeffrey P (2008). Plasma physics and fusion energy. Cambridge university press.

Freidberg, Jeffrey P. (2014). Ideal MHD. Cambridge University Press.

Frieman, E. A. and Liu Chen (Mar. 1982). Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. In: *The Physics of Fluids* 25.3, pp. 502–508. ISSN: 0031-9171.

Frisch, Uriel (1995). Turbulence: The Legacy of A. N. Kolmogorov. Cambridge University Press.

Fuchert, G., K.J. Brunner, K. Rahbarnia, T. Stange, D. Zhang, J. Baldzuhn, S.A. Bozhenkov, C.D. Beidler, M.N.A. Beurskens, S. Brezinsek, R. Burhenn, H. Damm, A. Dinklage, Y. Feng, P. Hacker, M. Hirsch, Y. Kazakov, J. Knauer, A. Langenberg, H.P. Laqua, S. Lazerson, N.A. Pablant, E. Pasch, F. Reimold, T. Sunn Pedersen, E.R. Scott, F. Warmer, V.R. Winters, R.C. Wolf, and W7-X Team (Feb. 2020). Increasing the density in Wendelstein 7-X: benefits and limitations. In: *Nuclear Fusion* 60.3, p. 036020.

Fujisawa, Akihide (Dec. 2008). A review of zonal flow experiments. In: *Nuclear Fusion* 49.1, p. 013001.

Garabedian, Paul R and Geoffrey B McFadden (2009). Design of the DEMO fusion reactor following ITER. In: *Journal of research of the National Institute of Standards and Technology* 114.4, p. 229.

Geiger, Joachim E., Arthur Weller, Michael C. Zarnstorff, Carolin Nührenberg, Andreas Horst Franz Werner, Yaroslav I. Kolesnichenko, and W7-AS Team (2004). Equilibrium and Stability of High-β Plasmas in Wendelstein 7-AS. In: *Fusion Science and Technology* 46.1, pp. 13–23.

Gil, Pedro F., Weiping Li, Julianne Stratton, Alan A. Kaptanoglu, and Eve V. Stenson (2025). Augmented Lagrangian methods produce cutting-edge magnetic coils for stellarator fusion reactors.

Gohil, P, G R McKee, D Schlossberg, L Schmitz, and G Wang (July 2008). Dependence of the H-mode power threshold on toroidal plasma rotation in the DIII-D tokamak. In: *Journal of Physics: Conference Series* 123.1, p. 012017.

González-Jerez, A., P. Xanthopoulos, J.M. García-Regaña, I. Calvo, J. Alcusón, A. Bañón Navarro, M. Barnes, F.I. Parra, and J. Geiger (2022). Electrostatic gyrokinetic simulations in Wendelstein 7-X geometry: benchmark between the codes stella and GENE. In: *Journal of Plasma Physics* 88.3, p. 905880310.

Goodman, A.G., K. Camacho Mata, S.A. Henneberg, R. Jorge, M. Landreman, G.G. Plunk, H.M. Smith, R.J.J. Mackenbach, C.D. Beidler, and P. Helander (2023). Constructing precisely quasi-isodynamic magnetic fields. In: *Journal of Plasma Physics* 89.5, p. 905890504.

Goodman, Alan G., Pavlos Xanthopoulos, Gabriel G. Plunk, Håkan Smith, Carolin Nührenberg, Craig D. Beidler, Sophia A. Henneberg, Gareth Roberg-Clark, Michael Drevlak, and Per Helander (June 2024). Quasi-Isodynamic Stellarators with Low Turbulence as Fusion Reactor Candidates. In: *PRX Energy* 3 (2), p. 023010.

Greene, J. M. and J. L. Johnson (Nov. 1962a). Erratum: Stability Criterion for Arbitrary Hydromagnetic Equilibria. In: *The Physics of Fluids* 5.11, pp. 1488–1488. ISSN: 0031-9171.

Greene, John M. and John L. Johnson (May 1962b). Stability Criterion for Arbitrary Hydromagnetic Equilibria. In: *The Physics of Fluids* 5.5, pp. 510–517. ISSN: 0031-9171.

Greenwald, Martin (July 2002). Density limits in toroidal plasmas. In: *Plasma Physics and Controlled Fusion* 44.8, R27.

Hartwell, G. J., S. F. Knowlton, J. D. Hanson, D. A. Ennis, and D. A. Maurer (2017). Design, Construction, and Operation of the Compact Toroidal Hybrid. In: *Fusion Science and Technology* 72.1, pp. 76–90.

Hasegawa, Akira and Masahiro Wakatani (Feb. 1983). Plasma Edge Turbulence. In: *Phys. Rev. Lett.* 50 (9), pp. 682–686.

Helander, P, C D Beidler, T M Bird, M Drevlak, Y Feng, R Hatzky, F Jenko, R Kleiber, J H E Proll, Yu Turkin, and P Xanthopoulos (Nov. 2012). Stellarator and tokamak plasmas: a comparison. In: *Plasma Physics and Controlled Fusion* 54.12, p. 124009.

Helander, P. and G.G. Plunk (2022). Energetic bounds on gyrokinetic instabilities. Part 1. Fundamentals. In: *Journal of Plasma Physics* 88.2, p. 905880207.

Helander, Per (July 2014). Theory of plasma confinement in non-axisymmetric magnetic fields. In: *Reports on Progress in Physics* 77.8, p. 087001.

Henneberg, S A, M Drevlak, and P Helander (Dec. 2019). Improving fast-particle confinement in quasi-axisymmetric stellarator optimization. In: *Plasma Physics and Controlled Fusion* 62.1, p. 014023.

Henneberg, S A and G G Plunk (May 2025). Variety of coil sets for the compact stellarator-tokamak hybrid. In: *Plasma Physics and Controlled Fusion* 67.6, p. 065013.

Henneberg, S. A. (Mar. 2024). A Compact Stellarator-Tokamak Hybrid. Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting 2024. URL: `https://www.simonsfoundation.org/event/simons-collaboration-on-hidden-symmetries-and-fusion-energy-annual-meeting-2024/`.

Henneberg, S. A. and G. G. Plunk (June 2024). Compact stellarator-tokamak hybrid. In: *Phys. Rev. Res.* 6, p. L022052.

Henneberg, S.A., M. Drevlak, C. Nührenberg, C.D. Beidler, Y. Turkin, J. Loizu, and P. Helander (Jan. 2019). Properties of a new quasi-axisymmetric configuration. In: *Nuclear Fusion* 59.2, p. 026014.

Henneberg, S.A., P. Helander, and M. Drevlak (2021). Representing the boundary of stellarator plasmas. In: *Journal of Plasma Physics* 87.5, p. 905870503.

Hirshman, S. P. and J. C. Whitson (Dec. 1983). Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. In: *The Physics of Fluids* 26.12, pp. 3553–3568. ISSN: 0031-9171.

Holland, C. (June 2016). Validation metrics for turbulent plasma transport. In: *Physics of Plasmas* 23.6, p. 060901. ISSN: 1070-664X.

Holland, C., G. R. Tynan, R. J. Fonck, G. R. McKee, J. Candy, and R. E. Waltz (May 2007). Zonal-flow-driven nonlinear energy transfer in experiment and simulation. In: *Physics of Plasmas* 14.5, p. 056112. ISSN: 1070-664X.

Holland, C., G. R. Tynan, G. R. McKee, and R. J. Fonck (Oct. 2004). Investigation of the time-delay estimation method for turbulent velocity inference. In: *Review of Scientific Instruments* 75.10, pp. 4278–4280. ISSN: 0034-6748.

Howard, N.T., C. Holland, A.E. White, M. Greenwald, and J. Candy (Dec. 2015). Multi-scale gyrokinetic simulation of tokamak plasmas: enhanced heat loss due to cross-scale coupling of plasma turbulence. In: *Nuclear Fusion* 56.1, p. 014004.

Howlett, Lena, István Cziegler, Simon Freethy, Hendrik Meyer, and the MAST team (Mar. 2023). L-H transition studies on MAST: power threshold and heat flux analysis. In: *Nuclear Fusion* 63.5, p. 052001.

Iiyoshi, A., A. Komori, A. Ejiri, M. Emoto, H. Funaba, M. Goto, K. Ida, H. Idei, S. Inagaki, S. Kado, O. Kaneko, K. Kawahata, T. Kobuchi, S. Kubo, R. Kumazawa, S. Masuzaki, T. Minami, J. Miyazawa, T. Morisaki, S. Morita, S. Murakami, S. Muto, T. Mutoh, Y. Nagayama, Y. Nakamura, H. Nakanishi, K. Narihara, K. Nishimura, N. Noda, S. Ohdachi, N. Ohyabu, Y. Oka, M. Osakabe, T. Ozaki, B.J. Peterson, A. Sagara, S. Sakakibara, R. Sakamoto, H. Sasao, M. Sasao, K. Sato, M. Sato, T. Seki, T. Shimozuma, M. Shoji, H. Suzuki, Y. Takeiri, K. Tanaka, K. Toi, T. Tokuzawa, K. Tsumori, K. Tsuzuki, K.Y. Watanabe, T. Watari, H. Yamada, I. Yamada, S. Yamaguchi, M. Yokoyama, R. Akiyama, H. Chikaraishi, K. Haba, S. Hamaguchi, M. Iima, S. Imagawa, N. Inoue, K. Iwamoto, S. Kitagawa, J. Kodaira, Y. Kubota, R. Maekawa, T. Mito, T. Nagasaka, A. Nishimura, C. Takahashi, K. Takahata, Y. Takita, H. Tamura, T. Tsuzuki, S. Yamada, K. Yamauchi, N. Yanagi, H. Yonezu, Y. Hamada, K. Matsuoka, K. Murai, K. Ohkubo, I. Ohtake, M. Okamoto, S. Satoh, T. Satow, S. Sudo, S. Tanahashi, K. Yamazaki, M. Fujiwara, and O. Motojima (Sept. 1999). Overview of the Large Helical Device project. In: *Nuclear Fusion* 39.9Y, p. 1245.

Imbert-Gérard, Lise-Marie, Elizabeth J. Paul, and Adelle M. Wright (2024). An Introduction to Stellarators: From Magnetic Fields to Symmetries and Optimization. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Jenko, F. and W. Dorland (Nov. 2002). Prediction of Significant Tokamak Turbulence at Electron Gyroradius Scales. In: *Phys. Rev. Lett.* 89 (22), p. 225001.

Jenko, F., W. Dorland, M. Kotschenreuther, and B. N. Rogers (May 2000). Electron temperature gradient driven turbulence. In: *Physics of Plasmas* 7.5, pp. 1904–1910. ISSN: 1070-664X.

Joiner, N. and W. Dorland (June 2010). Ion temperature gradient driven transport in tokamaks with square shaping. In: *Physics of Plasmas* 17.6, p. 062306. ISSN: 1070-664X.

Kappel, John, Matt Landreman, and Dhairya Malhotra (Jan. 2024). The magnetic gradient scale length explains why certain plasmas require close external magnetic coils. In: *Plasma Physics and Controlled Fusion* 66.2, p. 025018.

Kaptanoglu, Alan A., Matt Landreman, and Michael C. Zarnstorff (June 2025). Optimization of passive superconductors for shaping stellarator magnetic fields. In: *Phys. Rev. E* 111 (6), p. 065202.

Kaptanoglu, Alan A., Alexander Wiedman, Jacob Halpern, Siena Hurwitz, Elizabeth J. Paul, and Matt Landreman (Mar. 2025). Reactor-scale stellarators with force and torque minimized dipole coils. In: *Nuclear Fusion* 65.4, p. 046029.

Kirk, A, I T Chapman, T E Evans, C Ham, J R Harrison, G Huijsmans, Y Liang, Y Q Liu, A Loarte, W Suttrop, and A J Thornton (Nov. 2013). Understanding the effect resonant magnetic perturbations have on ELMs. In: *Plasma Physics and Controlled Fusion* 55.12, p. 124003.

Klinger, T, A Alonso, S Bozhenkov, R Burhenn, A Dinklage, G Fuchert, J Geiger, O Grulke, A Langenberg, M Hirsch, G Kocsis, J Knauer, A Krämer-Flecken, H Laqua, S Lazerson, M Landreman, H Maaßberg, S Marsen, M Otte, N Pablant, E Pasch, K Rahbarnia, T Stange, T Szepesi, H Thomsen, P Traverso, J L Velasco, T Wauters, G Weir, T Windisch, and The Wendelstein 7-X Team (Oct. 2016). Performance and properties of the first plasmas of Wendelstein 7-X. In: *Plasma Physics and Controlled Fusion* 59.1, p. 014018.

Kreyszig, Erwin author (1959). Differential geometry. eng. Mathematical expositions no. 11. Toronto: University of Toronto Press. ISBN: 148758945X.

Kriete, D. M., G. R. McKee, R. J. Fonck, D. R. Smith, G. G. Whelan, and Z. Yan (Sept. 2018). Extracting the turbulent flow-field from beam emission spectroscopy images using velocimetry. In: *Review of Scientific Instruments* 89.10, 10E107. ISSN: 0034-6748.

Kruskal, M. D. and R. M. Kulsrud (July 1958). Equilibrium of a Magnetically Confined Plasma in a Toroid. In: *The Physics of Fluids* 1.4, pp. 265–274. ISSN: 0031-9171.

Landreman, M., S. Buller, and M. Drevlak (Aug. 2022). Optimization of quasisymmetric stellarators with self-consistent bootstrap current and energetic particle confinement. In: *Physics of Plasmas* 29.8, p. 082501.

Landreman, Matt (Nov. 2021). Data for the paper "Magnetic fields with precise quasisymmetry". Version v2.0. Zenodo.

Landreman, Matt (July 2022). Data for the paper "Optimization of quasisymmetric stellarators with self-consistent bootstrap current and energetic particle confinement". Version v1.1. Zenodo.

Landreman, Matt and Peter J. Catto (Mar. 2012). Omnigenity as generalized quasisymmetrya). In: *Physics of Plasmas* 19.5, p. 056103. ISSN: 1070-664X.

Landreman, Matt, Jong Youl Choi, Caio Alves, Prasanna Balaprakash, Michael Churchill, Rory Conlin, and Gareth Roberg-Clark (2025). How does ion temperature gradient turbulence depend on magnetic geometry? Insights from data and machine learning. In: *Journal of Plasma Physics* 91.4, E120.

Landreman, Matt and Rogerio Jorge (2020). Magnetic well and Mercier stability of stellarators near the magnetic axis. In: *Journal of Plasma Physics* 86.5, p. 905860510.

Landreman, Matt, Bharat Medasani, Florian Wechsung, Andrew Giuliani, Rogerio Jorge, and Caoxiang Zhu (2021). SIMSOPT: A flexible framework for stellarator optimization. In: *Journal of Open Source Software* 6.65, p. 3525.

Landreman, Matt and Elizabeth Paul (2022). Magnetic fields with precise quasisymmetry for plasma confinement. In: *Physical Review Letters* 128.3, p. 035001.

Levinton, F. M., M. C. Zarnstorff, S. H. Batha, M. Bell, R. E. Bell, R. V. Budny, C. Bush, Z. Chang, E. Fredrickson, A. Janos, J. Manickam, A. Ramsey, S. A. Sabbagh, G. L. Schmidt, E. J. Synakowski, and G. Taylor (Dec. 1995). Improved Confinement with Reversed Magnetic Shear in TFTR. In: *Phys. Rev. Lett.* 75 (24), pp. 4417–4420.

Liu, Haifeng, Akihiro SHIMIZU, Mitsutaka ISOBE, Shoichi OKAMURA, Shin NISHIMURA, Chihiro SUZUKI, Yuhong XU, Xin ZHANG, Bing LIU, Jie HUANG, Xianqu WANG, Hai LIU, Changjian TANG, Dapeng YIN, Yi WAN, and CFQS team (2018). Magnetic Configuration and Modular Coil Design for the Chinese First Quasi-Axisymmetric Stellarator. In: *Plasma and Fusion Research* 13, pp. 3405067–3405067.

Maeyama, S, M Sasaki, K Fujii, T Kobayashi, R O Dendy, Y Kawachi, H Arakawa, and S Inagaki (Apr. 2021). On the triad transfer analysis of plasma turbulence: symmetrization, coarse graining, and directional representation. In: *New Journal of Physics* 23.4, p. 043049.

Maeyama, S., Y. Idomura, T.-H. Watanabe, M. Nakata, M. Yagi, N. Miyato, A. Ishizawa, and M. Nunami (June 2015). Cross-Scale Interactions between Electron and Ion Scale Turbulence in a Tokamak Plasma. In: *Phys. Rev. Lett.* 114 (25), p. 255002.

Maeyama, S., T.-H. Watanabe, Y. Idomura, M. Nakata, A. Ishizawa, and M. Nunami (May 2017). Cross-scale interactions between turbulence driven by electron and ion temperature gradients via sub-ion-scale structures. In: *Nuclear Fusion* 57.6, p. 066036.

Mandell, N.R., W. Dorland, I. Abel, R. Gaur, P. Kim, M. Martin, and T. Qian (2024). GX: a GPU-native gyrokinetic turbulence code for tokamak and stellarator design. In: *Journal of Plasma Physics* 90.4, p. 905900402.

Manz, P., G. S. Xu, B. N. Wan, H. Q. Wang, H. Y. Guo, I. Cziegler, N. Fedorczak, C. Holland, S. H. Müller, S. C. Thakur, M. Xu, K. Miki, P. H. Diamond, and G. R. Tynan (July 2012). Zonal flow triggers the L-H transition in the Experimental Advanced Superconducting Tokamak. In: *Physics of Plasmas* 19.7, p. 072311. ISSN: 1070-664X.

Marinoni, Alessandro, O Sauter, and S Coda (2021). A brief history of negative triangularity tokamak plasmas. In: *Reviews of Modern Plasma Physics* 5.1, p. 6.

Martin, M F, M Landreman, P Xanthopoulos, N R Mandell, and W Dorland (July 2018). The parallel boundary condition for turbulence simulations in low magnetic shear devices. In: *Plasma Physics and Controlled Fusion* 60.9, p. 095008.

Martin, Y R, T Takizuka, A Author, and B Buthor (July 2008). Power requirement for accessing the H-mode in ITER. In: *Journal of Physics: Conference Series* 123.1, p. 012033.

Maurer, D. (Oct. 2015). Plasma disruption avoidance using non-axisymmetric shaping with stellarator fields. *International Stellarator and Heliotron Workshop*.

McDermott, Rachael Marie (2009). “Edge radial electric field studies via charge exchange recombination spectroscopy on the Alcator C-Mod tokamak”. PhD thesis. Massachusetts Institute of Technology.

Medasani, Bharat, Florian Wechsung, Matt Landreman, Elizabeth Paul, Rogerio Jorge, Caoxiang Zhu, Rahul Gaur, AlexWiedman, Andrew Giuliani, Dominik Stańczak-Marikin, Zhisong Qu, daringli, and tmqian (Jan. 2023). hiddenSymmetries/simsopt: v0.12.2. Version v0.12.2.

Mercier, C (1962). Critère de stabilité d’un système toroïdal hydromagnétique en pression scalaire. In: *Nucl. Fusion Suppl* 2, p. 801.

Mercier, Claude (Sept. 1964). Equilibrium and stability of a toroidal magnetohydrodynamic system in the neighbourhood of a magnetic axis. In: *Nuclear Fusion* 4.3, p. 213.

Meyer, Hendrik and on behalf of the STEP Plasma Team (2024). Plasma burn—mind the gap. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 382.2280, p. 20230406.

Miller, R. L., M. S. Chu, J. M. Greene, Y. R. Lin-Liu, and R. E. Waltz (Apr. 1998). Noncircular, finite aspect ratio, local equilibrium model. In: *Physics of Plasmas* 5.4, pp. 973–978. ISSN: 1070-664X.

Moyer, R. A., G. R. Tynan, C. Holland, and M. J. Burin (Sept. 2001). Increased Nonlinear Coupling between Turbulence and Low-Frequency Fluctuations at the $L-H$ Transition. In: *Phys. Rev. Lett.* 87 (13), p. 135001.

Mynick, H, P Xanthopoulos, B Faber, M Lucia, M Rorvig, and J N Talmadge (Aug. 2014). Turbulent optimization of toroidal configurations. In: *Plasma Physics and Controlled Fusion* 56.9, p. 094001.

Najmabadi, F., A. R. Raffray, C. Author, and D. Author (2008). The ARIES-CS Compact Stellarator Fusion Power Plant. In: *Fusion Science and Technology* 54.3, pp. 655–672.

Nakata, M., T.-H. Watanabe, and H. Sugama (Feb. 2012). Nonlinear entropy transfer via zonal flows in gyrokinetic plasma turbulence. In: *Physics of Plasmas* 19.2, p. 022303. ISSN: 1070-664X.

Nakata, Motoki and Seikichi Matsuoka (2022). Impact of Geodesic Curvature on Zonal Flow Generation in Magnetically Confined Plasmas. In: *Plasma and Fusion Research* 17, pp. 1203077–1203077.

Nelson, A. O., C. Vincent, H. Anand, J. Lovell, J. F. Parisi, H. S. Wilson, K. Imada, W. P. Wehner, M. Kochan, S. Blackmore, G. McArdle, S. Guizzo, L. Rondini, S. Freiberger, and C. Paz-Soldan (2024). First Access to ELM-free Negative Triangularity at Low Aspect Ratio.

Nemov, V. V., S. V. Kasilov, W. Kernbichler, and M. F. Heyn (Dec. 1999). Evaluation of $1/\nu$ neoclassical transport in stellarators. In: *Physics of Plasmas* 6.12, pp. 4622–4632. ISSN: 1070-664X.

Nies, Richard, Elizabeth J. Paul, Dario Panici, Stuart R. Hudson, and Amitava Bhattacharjee (2024). Exploration of the parameter space of quasisymmetric stellarator vacuum fields through adjoint optimisation. In: *Journal of Plasma Physics* 90.6, p. 905900620.

Nocedal, Jorge and Stephen J Wright (1999). Numerical optimization. Springer.

Panici, D., R. Conlin, D.W. Dudt, K. Unalmis, and E. Kolemen (2023). The DESC stellarator code suite. Part 1. Quick and accurate equilibria computations. In: *Journal of Plasma Physics* 89.3, p. 955890303.

Park, J.-K., S. M. Yang, N. C. Logan, Q. Hu, C. Zhu, M. C. Zarnstorff, R. Nazikian, C. Paz-Soldan, Y. M. Jeon, and W. H. Ko (Mar. 2021). Quasisymmetric Optimization of Nonaxisymmetry in Tokamaks. In: *Phys. Rev. Lett.* 126 (12), p. 125001.

Parra, Felix I., Michael Barnes, and Arthur G. Peeters (June 2011). Up-down symmetry of the turbulent transport of toroidal angular momentum in tokamaks. In: *Physics of Plasmas* 18.6, p. 062501. ISSN: 1070-664X.

Plunk, G. G. (2020). Perturbing an axisymmetric magnetic equilibrium to obtain a quasi-axisymmetric stellarator. In: *Journal of Plasma Physics* 86.4, p. 905860409.

Plunk, G. G., P. Helander, P. Xanthopoulos, and J. W. Connor (Mar. 2014). Collisionless microinstabilities in stellarators. III. The ion-temperature-gradient mode. In: *Physics of Plasmas* 21.3, p. 032112. ISSN: 1070-664X.

Plunk, G. G. and Per Helander (2018). Quasi-axisymmetric magnetic fields: weakly non-axisymmetric case in a vacuum. In: *Journal of Plasma Physics* 84.2, p. 905840205.

Podavini, L., A. Zocco, J.M. García-Regaña, M. Barnes, F.I. Parra, A. Mishchenko, and P. Helander (2024). Ion-temperature- and density-gradient-driven instabilities and turbulence in Wendelstein 7-X close to the stability threshold. In: *Journal of Plasma Physics* 90.4, p. 905900414.

Quénot, Georges M, Jaroslaw Pakleza, and Tomasz A Kowalewski (1998). Particle image velocimetry with optical flow. In: *Experiments in fluids* 25.3, pp. 177–189.

Ramasamy, Rohan, Haowei Zhang, Joachim Geiger, Carolin Nührenberg, Håkan M. Smith, Karl Lackner, Valentin Igochine, and the JOREK team (2025). A comparison of low-n Mercier unstable Wendelstein stellarators and quasi-interchange modes in tokamaks.

Redl, A., C. Angioni, E. Belli, O. Sauter, ASDEX Upgrade Team, and EUROfusion MST1 Team (Feb. 2021). A new set of analytical formulae for the computation of the bootstrap current and the neoclassical conductivity in tokamaks. In: *Physics of Plasmas* 28.2. 022502. ISSN: 1070-664X.

Righi, E., D.V. Bartlett, J.P. Christiansen, G.D. Conway, J.G. Cordey, L.-G. Eriksson, H.P.L. De Esch, G.M. Fishpool, C.W. Gowers, J.C.M. de Haas, P.J. Harbour, N.C. Hawkes, J. Jacquinot, T.T.C. Jones, W. Kerner, Q.A. King, C.G. Lowry, R.D. Monk, P. Nielsen, F.G. Rimini, G. Saibene, R. Sartori, B. Schunke, A.C.C. Sips, R.J. Smith, M.F. Stamp, D.F.H. Start, K. Thomsen, B.J.D. Tubbing, and N. Zornig (Mar. 1999). Isotope scaling of the H mode power threshold on JET. In: *Nuclear Fusion* 39.3, p. 309.

Ritz, Ch. P., E. J. Powers, and R. D. Bengtson (Jan. 1989). Experimental measurement of three-wave coupling and energy cascading. In: *Physics of Fluids B: Plasma Physics* 1.1, pp. 153–163. ISSN: 0899-8221.

Roberg-Clark, G. T., G. G. Plunk, P. Xanthopoulos, C. Nührenberg, S. A. Henneberg, and H. M. Smith (Sept. 2023). Critical gradient turbulence optimization toward a compact stellarator reactor concept. In: *Phys. Rev. Res.* 5 (3), p. L032030.

Roberg-Clark, G.T., P. Xanthopoulos, and G.G. Plunk (2024). Reduction of electrostatic turbulence in a quasi-helically symmetric stellarator via critical gradient optimization. In: *Journal of Plasma Physics* 90.3, p. 175900301.

Rodríguez, E., E.J. Paul, and A. Bhattacharjee (2022). Measures of quasisymmetry for stellarators. In: *Journal of Plasma Physics* 88.1, p. 905880109.

Rosenbluth, M. N. and F. L. Hinton (Jan. 1998). Poloidal Flow Driven by Ion-Temperature-Gradient Turbulence in Tokamaks. In: *Phys. Rev. Lett.* 80 (4), pp. 724–727.

Sanchez, R, S P Hirshman, A S Ware, L A Berry, and D A Spong (June 2000). Ballooning stability optimization of low-aspect-ratio stellarators*. In: *Plasma Physics and Controlled Fusion* 42.6, p. 641.

Schekochihin, A. A., S. C. Cowley, W. Dorland, G. W. Hammett, G. G. Howes, E. Quataert, and T. Tatsuno (May 2009). ASTROPHYSICAL GYROKINETICS: KINETIC AND FLUID TURBULENT CASCADES IN MAGNETIZED WEAKLY COLLISIONAL PLASMAS. In: *The Astrophysical Journal Supplement Series* 182.1, p. 310.

Schmitz, L. (Jan. 2017). The role of turbulence–flow interactions in L- to H-mode transition dynamics: recent progress. In: *Nuclear Fusion* 57.2, p. 025003.

Schuett, T M, I Cziegler, and D Dickinson (Nov. 2025a). The poloidal distribution of electrostatic zonal flow drive in strongly shaped tokamaks. In: *Plasma Physics and Controlled Fusion* 67.11, p. 115022.

Schuett, T. M., I. Cziegler, and D. Dickinson (2025b). The poloidal distribution of electrostatic zonal flow drive in strongly shaped tokamaks [Oral Contributed]. *29th EU-US Transport Task Force Workshop, Budapest, Hungary*.

Schuett, T. M. and S. A. Henneberg (2024a). Exploring novel compact qausi-axisymmetric stellarators [Poster]. *Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting, New York, USA*.

Schuett, T. M. and S. A. Henneberg (Dec. 2024b). Exploring novel compact quasi-axisymmetric stellarators. In: *Phys. Rev. Res.* 6, p. L042052.

Schuett, T. M. and S. A. Henneberg (2024c). Optimization of compact quasi-axisymetric stellarators [Oral Invited]. *International Stellarator and Heliotron Workshop, Hiroshima, Japan*.

Schuett, T. M., S. A. Henneberg, D. Dickinson, and I. Cziegler (2025). Ion temperature gradient turbulence in compact quasi-axisymmetric stellarators [Poster]. *Simons Collaboration on Hidden Symmetries and Fusion Energy Annual Meeting, New York, USA*.

Schuett, Tobias M and Sophia A Henneberg (June 2025). Optimization of compact quasi-axisymmetric stellarators. In: *Plasma Physics and Controlled Fusion* 67.6, p. 065024.

Schuett, Tobias M. (2025). Initial version of GS2 with gyrokinetic transfer. `https://bitbucket.org/tobiasschuett/gs2-implement/src/adding-functions-in/`. commit 6c04ce3.

Singh, Rameswar, P.H. Diamond, and A.O. Nelson (Nov. 2023). Geometric dependencies of the mean E × B shearing rate in negative triangularity tokamaks. In: *Nuclear Fusion* 63.12, p. 126053.

Stroteich, Sven, Pavlos Xanthopoulos, Gabriel Plunk, and Ralf Schneider (2022). Seeking turbulence optimized configurations for the Wendelstein 7-X stellarator: ion temperature gradient and electron temperature gradient turbulence. In: *Journal of Plasma Physics* 88.5, p. 175880501.

Sudo, S., Y. Takeiri, H. Zushi, F. Sano, K. Itoh, K. Kondo, and A. Iiyoshi (Jan. 1990). Scalings of energy confinement and density limit in stellarator/heliotron devices. In: *Nuclear Fusion* 30.1, p. 11.

Sugama, H. and W. Horton (Aug. 1995). Neoclassical and anomalous transport in axisymmetric toroidal plasmas with electrostatic turbulence. In: *Physics of Plasmas* 2.8, pp. 2989–3006. ISSN: 1070-664X.

Sugama, H., T.-H. Watanabe, and M. Nunami (Nov. 2009). Linearized model collision operators for multiple ion species plasmas and gyrokinetic entropy balance equations. In: *Physics of Plasmas* 16.11, p. 112503. ISSN: 1070-664X.

Takizuka, T (Apr. 2004). Roles of aspect ratio, absolute B and effective Z of the H-mode power threshold in tokamaks of the ITPA database. In: *Plasma Physics and Controlled Fusion* 46.5A, A227.

Terry, J. L., A. von Stechow, S. G. Baek, S. B. Ballinger, O. Grulke, C. von Sehren, R. Laube, C. Killer, F. Scharmer, K. J. Brunner, J. Knauer, S. Bois, and W7-X Team (Sept. 2024). Realization of a gas puff imaging system on the Wendelstein 7-X stellarator. In: *Review of Scientific Instruments* 95.9, p. 093517. ISSN: 0034-6748.

Terry, P. W. (Jan. 2000). Suppression of turbulence and transport by sheared flow. In: *Rev. Mod. Phys.* 72 (1), pp. 109–165.

Wagner, F., G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke, G. Haas, M. Huang, F. Karger, M. Keilhacker, O. Klüber, M. Kornherr, K. Lackner, G. Lisitano, G. G.

Lister, H. M. Mayer, D. Meisel, E. R. Müller, H. Murmann, H. Niedermeyer, W. Poschenrieder, H. Rapp, H. Röhr, F. Schneider, G. Siller, E. Speth, A. Stäbler, K. H. Steuer, G. Venus, O. Vollmer, and Z. Yü (Nov. 1982). Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak. In: *Phys. Rev. Lett.* 49 (19), pp. 1408–1412.

Wakatani, Masahiro and Akira Hasegawa (Mar. 1984). A collisional drift wave description of plasma edge turbulence. In: *The Physics of Fluids* 27.3, pp. 611–618. ISSN: 0031-9171.

Weller, A., S. Sakakibara, K. Y. Watanabe, K. Toi, J. Geiger, M. C. Zarnstorff, S. R. Hudson, A. Reiman, A. Werner, C. Nührenberg, S. Ohdachi, Y. Suzuki, H. Yamada, The W7-AS Team, and The LHD Team (2006). Significance of MHD Effects in Stellarator Confinement. In: *Fusion Science and Technology* 50.2, pp. 158–170.

Winsor, Niels, John L. Johnson, and John M. Dawson (Nov. 1968). Geodesic Acoustic Waves in Hydromagnetic Systems. In: *The Physics of Fluids* 11.11, pp. 2448–2450. ISSN: 0031-9171.

Xanthopoulos, P., F. Merz, T. Görler, and F. Jenko (July 2007). Nonlinear Gyrokinetic Simulations of Ion-Temperature-Gradient Turbulence for the Optimized Wendelstein 7-X Stellarator. In: *Phys. Rev. Lett.* 99 (3), p. 035002.

Xanthopoulos, P., A. Mischchenko, P. Helander, H. Sugama, and T.-H. Watanabe (Dec. 2011). Zonal Flow Dynamics and Control of Turbulent Transport in Stellarators. In: *Phys. Rev. Lett.* 107 (24), p. 245002.

Xanthopoulos, P., H. E. Mynick, P. Helander, Y. Turkin, G. G. Plunk, F. Jenko, T. Görler, D. Told, T. Bird, and J. H. E. Proll (Oct. 2014). Controlling Turbulence in Present and Future Stellarators. In: *Phys. Rev. Lett.* 113 (15), p. 155001.

Xu, M., G. R. Tynan, C. Holland, Z. Yan, S. H. Muller, and J. H. Yu (2009). Study of nonlinear spectral energy transfer in frequency domain. In: *Physics of Plasmas* 16.4, p. 042312.

Yan, Z., G. R. McKee, R. Fonck, P. Gohil, R. J. Groebner, and T. H. Osborne (Mar. 2014). Observation of the $L-H$ Confinement Bifurcation Triggered by a Turbulence-Driven Shear Flow in a Tokamak Plasma. In: *Phys. Rev. Lett.* 112 (12), p. 125002.

Zarnstorff, M C, L A Berry, A Brooks, E Fredrickson, G-Y Fu, S Hirshman, S Hudson, L-P Ku, E Lazarus, D Mikkelsen, D Monticello, G H Neilson, N Pomphrey, A Reiman, D Spong, D Strickler, A Boozer, W A Cooper, R Goldston, R Hatcher, M Isaev, C Kessel, J Lewandowski, J F Lyon, P Merkel, H Mynick, B E Nelson, C Nuehrenberg, M Redi, W Reiersen, P Rutherford, R Sanchez, J Schmidt, and R B White (Nov. 2001). Physics of the compact advanced stellarator NCSX. In: *Plasma Physics and Controlled Fusion* 43.12A, A237.

Zohm, Hartmut (2015). Magnetohydrodynamic stability of tokamaks. John Wiley & Sons.