ABSTRACT

Title of Dissertation: TURBULENT AND COLLISIONAL TRANSPORT IN OPTIMIZED STELLARATORS

Michael Francis Martin
Doctor of Philosophy, 2020

Dissertation Directed by: Professor William Dorland
Department of Physics

The stellarator is a fusion energy concept that relies on fully three-dimensional shaping of magnetic fields to confine particles. Stellarators have many favorable properties, including, but not limited to, the ability to operate in steady-state, many optimizable degrees of freedom, and no strict upper limit on the plasma density. Due to the three-dimensional character of stellarators, theoretical and computational studies of stellarator physics are challenging, and they also possess some disadvantages compared with tokamaks. Namely, particle confinement and impurity control are problems in generic stellarator magnetic fields that must be addressed with optimized magnetic fields. Further, simulations will require a substantial increase in grid points because of the three-dimensional structure, leading to more expensive computations. This thesis will address both topics, by first exploring the behavior of impurity particle transport in optimized stellarators, and then introducing a boundary condition to reduce the cost of stellarator turbulence simulations.

Impurity temperature screening is a favorable neoclassical phenomenon in-

volving an outward radial flux of impurity ions from the core of fusion devices. Quasisymmetric magnetic fields lead to intrinsically ambipolar neoclassical fluxes that give rise to temperature screening for low enough $\eta^{-1} \equiv d\ln n/d\ln T$. Here we examine the impurity particle flux in a number of approximately quasisymmetric stellarator configurations and parameter regimes while varying the amount of symmetry-breaking in the magnetic field. Results indicate that achieving temperature screening is possible, but unlikely, at low-collisionality reactor-relevant conditions in the core. Further, in configurations optimized for quasisymmetry, results suggest that neoclassical fluxes are small compared with a gyro-Bohm estimate of turbulent fluxes.

Calculating these turbulent fluxes is generally done by solving the gyrokinetic equation in a flux tube simulation domain, which employs coordinates aligned with the magnetic field lines. The standard "twist-and-shift" formulation of the boundary conditions [7] was derived assuming axisymmetry and is widely used because it is efficient, as long as the global magnetic shear is not too small. A generalization of this formulation is presented, appropriate for studies of non-axisymmetric, stellarator-symmetric configurations, as well as for axisymmetric configurations with small global shear. The key idea of this generalization is to rely on integrated local shear, allowing one significantly more freedom when choosing the extent of the simulation domain in each direction. Simulations of stellarator turbulence that employ the generalized parallel boundary conditions allow for lower resolution to be used compared with simulations that use the (incorrect, axisymmetric) standard parallel boundary condition.

TURBULENT AND COLLISIONAL TRANSPORT IN OPTIMIZED STELLARATORS

by

Michael Francis Martin

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor William Dorland, Chair/Advisor
Dr. Matthew Landreman, Co-Advisor
Professor Thomas Antonsen
Professor Ramani Duraiswami
Professor Adil Hassam

Dedication

To the loving memory of my grandfather, Lawrence “Larry” Deal

Acknowledgments

This thesis would not have been possible without the guidance and support of so many colleagues, friends, and family over the years.

I would first and foremost like to thank my advisor Matt Landreman, to whom I owe so much of what I've accomplished as a grad student based on his exceptional guidance over the years. It has been a privilege to be able to work with such an accomplished, knowledgeable, and kind plasma physicist who always makes time to talk science. I will always appreciate our many enlightening and instructive conversations that have been central to my progress as a scientist.

I would also like to thank my co-advisor Bill Dorland. There has rarely been a time when I don't come away energized after a conversation with Bill, as he is always overflowing with ideas, advice, and a refreshing sense of humor. It has been a pleasure to learn from and work with one of the foremost experts in turbulence simulations.

Thank you to my committee members, Adil Hassam, Thomas Antonsen, and Ramani Duraiswami. I would also like to thank Greg Howes, my advisor at the University of Iowa, as well as Jim Drake, for both being excellent instructors for the courses I took on plasma physics.

Thank you to all my fellow grad students at the University of Maryland, as well as the other members of the plasma theory group. It was a pleasure to be a part of such a large and engaged group of plasma theorists. Thank you to the grad students at the University of Iowa, especially those who endured Math Methods with

me. Thank you as well to my friends and roommates, Alex Simon, Jon Schenck, and Quinn Colville.

Finally, without family and loved ones, it can difficult to accomplish anything in life, and I want to thank all of you who have supported me over the years. Most importantly, I cannot begin to express how grateful I am to my mom, Karen Mullin, for being a rock of love and support throughout my entire life. It's been such a blessing to know that there is always someone on your side who genuinely cares and is curious about everything in your life, even something as mundane as the "Mucinex" I eat every day. Also, it's impossible to talk about love and support without mentioning my grandparents, Larry and Edna Deal, who are like second parents to me. You will both always hold a special place in my heart.

Last and certainly not least, I want to thank my best friend and girlfriend, Stella Mach, who always finds joy in discovering her name somewhere in my work. Your support through the highs and lows, as well your inspirational passion and dedication to the things that drive you have helped me strive to be the best person I can be.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

DKE	Drift-kinetic equation
GPU	Graphics Processing Unit
MCF	Magnetic confinement fusion
MHD	Magnetohydrodynamic
QA	Quasi-axisymmetric
QH	Quasi-helically-symmetric

Chapter 1: Introduction

1.1 Magnetic Confinement Fusion

Nuclear fusion as an energy source has the potential to generate clean, safe, and sustainable energy using a near-inexhaustible fuel supply. The promise of fusion energy has prompted the design of many concepts aimed at achieving fusion in a practical sense, with arguably the most successful of these being magnetic confinement fusion (MCF). MCF aims to confine a high-temperature plasma (ionized gas) using magnetic fields for a long enough time in order to allow for self-sustaining nuclear fusion reactions to occur.

The MCF devices that will be considered here achieve plasma confinement with toroidally-shaped vessels that use electromagnetic coils and/or plasma currents to generate a set of nested magnetic surfaces, or flux surfaces (see Figure 1.1). The innermost flux surface is a line that is referred to as the magnetic axis. The magnetic field lines are tangent to these flux surfaces, approximately confining charged particles to the surface, in principle. Further, the magnetic field lines have an average pitch, or twist, within the surface known as the rotational transform.

A nonzero rotational transform is required to prevent the secular motion of particles away from flux surfaces that arises due to the shaping of the magnetic field.

There are two prominent MCF concepts that use distinct approaches to generate the rotational transform, the tokamak and the stellarator, which will be discussed presently.

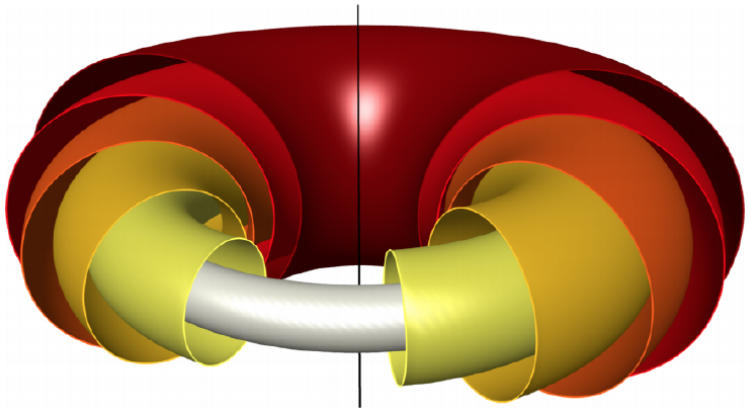

Figure 1.1: Depiction of a set of nested toroidal magnetic surfaces (flux surfaces), on which magnetic field lines will be confined. Figure courtesy of Matt Landreman.

1.1.1 Tokamaks

The most studied MCF concept to date is the tokamak (see Figure 1.2), which generates toroidal flux surfaces that have a continuous rotational symmetry in the toroidal direction (axisymmetry). By the toroidal direction, we are referring to the direction of increase of the toroidal angle, ζ, which is taken to be the long way around the torus, and the poloidal angle, θ, is the short way around the torus (see Figure 1.3). This toroidal symmetry in tokamaks leads to favorable confinement properties, and can simplify the engineering of the electromagnetic coils that are required to create the flux surfaces. Further, the ability to ignore toroidal variation has led to extensive theoretical progress in the understanding of tokamak plasmas. However, in order to produce the poloidal magnetic field component that is responsible for a

nonzero rotational transform, tokamaks require the presence of a toroidal current. The presence of a large toroidal current is a major drawback because it makes the plasma vulnerable to instabilities that lead to disruptions, which can cause a loss of confinement, and potentially damage the machine. It is also difficult to drive the required current in steady-state, forcing tokamaks toward pulsed operation.

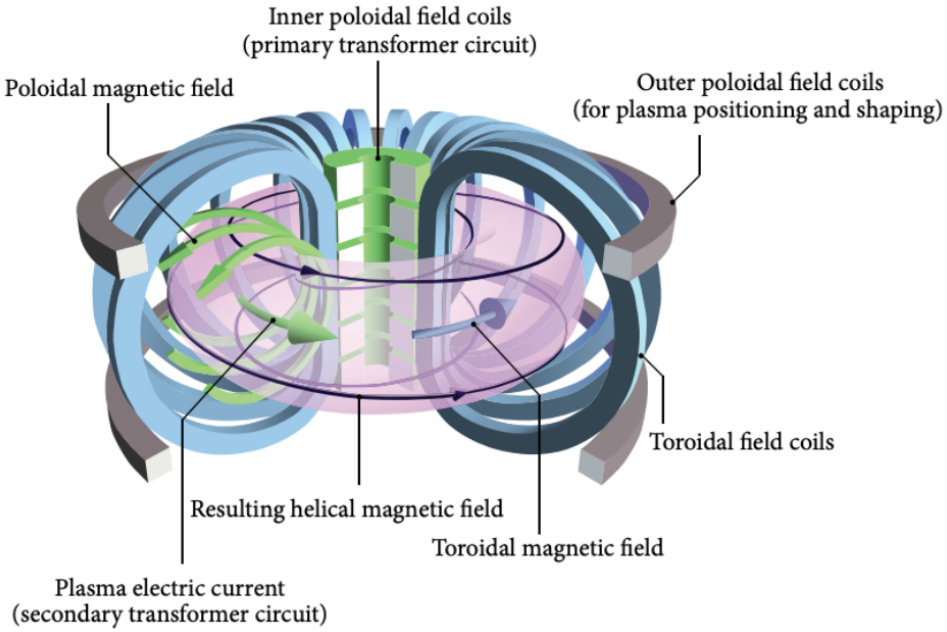

Figure 1.2: Cartoon of a tokamak experiment, where the purple torus represents a flux surface. The blue toroidal field coils generate the toroidal component of the magnetic field (blue arrow). The green transformer circuit generates the toroidal current, which is depicted by the green arrow. This current produces a poloidal magnetic field that gives the *twist* to the magnetic field that can be seen in the black line on the flux surface. Figure from [62].

1.1.2 Stellarators

Stellarators are an alternative to the tokamak, and although the concept was introduced first [91], they are not as well understood. Stellarators aim to confine particles through the use of three-dimensional flux surfaces, removing the axisymme-

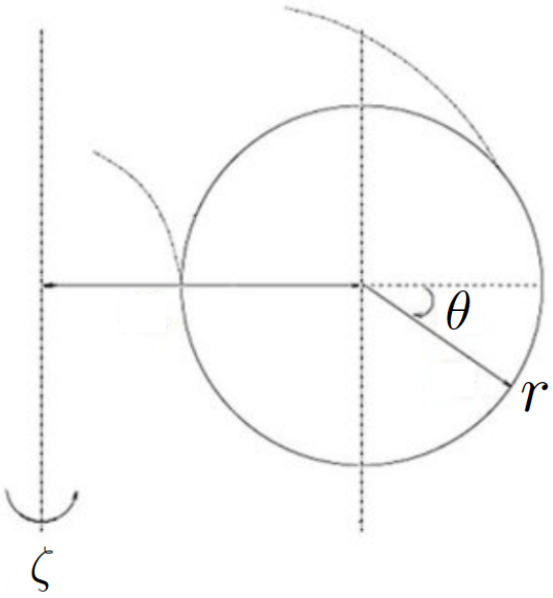

Figure 1.3: Toroidal coordinate system (r, θ, ζ).

try constraint that defines tokamaks (see Figure 1.4). A consequence of having 3-D surfaces is that stellarators possess the superior properties that they are not prone to disruptions, and are able to operate in steady-state. This is possible because of how the rotational transform is generated.

There are three different ways to produce a rotational transform [46]: an electric current (tokamaks), torsion of the magnetic axis, or rotation of the cross-section along the magnetic axis. Stellarators employ the second and third techniques by using complex electromagnetic coils to create the flux surfaces and fields, and thus do not require a plasma current. In principle, this also means that stellarators are able to operate with minimal toroidal current, making the positional equilibrium more robust as it will not be so reliant on the state of the plasma.

Further, the variation of the rotational transform across flux surfaces, which is related to the (global) magnetic shear is often made to be small in stellarators. This

avoids the low-order rational surfaces that lead to magnetic islands, and is enabled by the robustness of the positional equilibrium.

Along with the general problems in MCF, such as turbulence, there are of course unique challenges that arise when considering 3-D surfaces. Generating this shaping requires complex coils, which can be exceptionally difficult and expensive to engineer. Confinement of particles is also of concern because of secular drifts away from flux surfaces that result from asymmetries in the magnetic field.

These problems can be overcome, however. For the case of coils, the stellarator Wendelstein 7-X (W7-X) [40] currently operates with 50 of these modular coils, and experiments are able to create the desired flux surfaces to within errors of $\delta B/B \sim 10^{-5}$ [79]. The issue of particle transport can be addressed by a special type of symmetry known as quasisymmetry, which will be discussed in detail in Chapter 2. Stellarators designed to be quasisymmetric will have improved particle confinement similar to that of tokamaks.

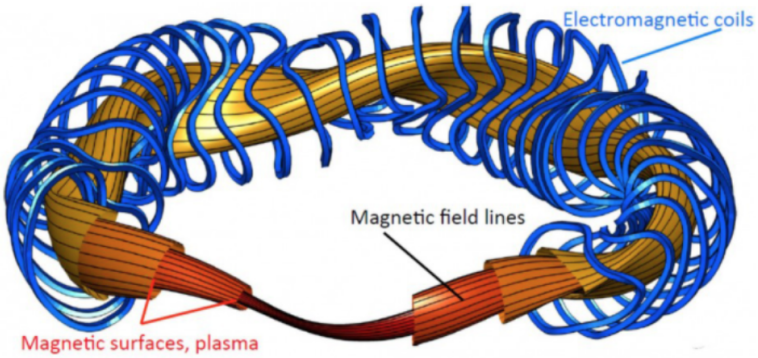

Figure 1.4: Cartoon of the W7-X stellarator showing the modular coils (blue) that produce the 3-D flux surfaces.

1.2 Transport in Plasmas

Before MCF can be realized as a practical energy source, there are a number of outstanding problems that must be overcome. Many of these outstanding issues are related to plasma transport, which covers the flow of particles and energy within the plasma. In this thesis, we will only address transport in the *core* of fusion devices. The core is defined to be the region of closed flux surfaces, where magnetic field lines ergodically cover the surface or connect with themselves, and do not intersect with the wall. In certain regions of the core, chaotic field lines or magnetic islands may be present, both of which do not lead to flux surfaces. However, in this work it is assumed that good flux surfaces exist in the regions where simulations are performed. Figure 1.5 provides a visual for the relative locations of different regions in toroidal fusion devices.

Since a long confinement time of particles and energy is central to MCF, the issues of plasma transport deserve considerable attention. For example, turbulence is a ubiquitous problem in tokamaks and stellarators that causes an increase in radial transport, which leads to decreased confinement times. There is also the issue of confining the hot alpha particle by-products of fusion reactions that are necessary for a self-sustaining reactor. However, not all particles participate in the fusion reaction. Specifically, impurity ions and the cold alpha particles (ash) that accumulate, need to be flushed from the core to avoid radiative power losses and fuel dilution. Improving the confinement time by addressing such issues in plasma transport will both enable the construction of smaller (cheaper) devices, and push

devices closer to the Lawson criterion, which is a threshold for achieving a state of self-sustaining fusion reactions.

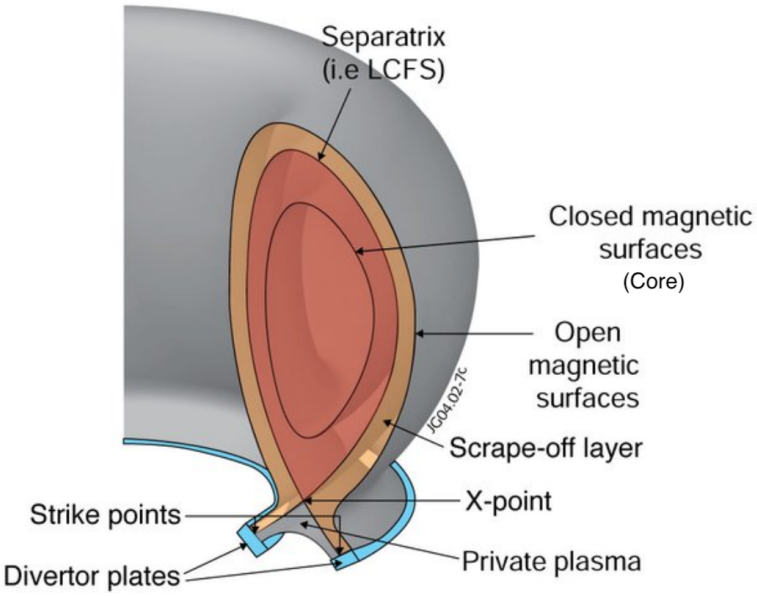

Figure 1.5: Toroidal cross section of a tokamak. The red area depicts the core region where the field lines are (ergodically) closed, leading to flux surfaces. The core is the region of interest for this thesis. Figure adapted from [96].

The study of transport in plasma revolves around solving the Fokker-Planck equation. The Fokker-Planck equation describes the evolution of the distribution function, which represents the number of particles per unit volume with position $\boldsymbol{x} = (x, y, z)$ and velocity $\boldsymbol{v} = (v_x, v_y, v_z)$. The position and velocity together $(\boldsymbol{x}, \boldsymbol{v})$ represent all possible states of a dynamical system, and is referred to as the phase space of the system. In the presence of collisions, this evolution of the distribution function in 6-D phase space is the Fokker-Planck equation and can be written as

$$\frac{df_a}{dt} = \frac{\partial f_a}{\partial t} + \boldsymbol{v} \cdot \nabla f_a + \frac{\boldsymbol{F}}{m_a} \cdot \frac{\partial f_a}{\partial \boldsymbol{v}} = \left(\frac{\partial f_a}{\partial t}\right)_{coll} = C(f_a). \tag{1.1}$$

Here, f_a is the distribution function, and $C(f_a)$ is the collision operator. Solving a

6-D equation plus time is a difficult problem irrespective of the method, making it useful to simplify the problem in some way.

In transport theory, this is done by choosing particular orderings of certain quantities in the Fokker-Planck equation to describe different physical processes. To describe what is meant by orderings, assume that any physical quantity can be expanded in some small parameter ϵ. For example, for the distribution function f_a, this would look like

$$f_a = f_{a0} + \epsilon f_{a1} + \epsilon^2 f_{a2} + \dots \tag{1.2}$$

where $\epsilon \ll 1$. Each additional term will then be $O(\epsilon)$ smaller than the previous one. Once the small parameter is selected, choices can be made about the size of each quantity. For example, if a quantity Q is known to be small, one could drop the $O(1)$ term to get $Q = \epsilon Q_1 + \epsilon^2 Q_2 + \dots$. By then substituting these expansions for each part of an equation in powers of ϵ, one can derive a set of equations at each order in ϵ. Depending on how the orderings are chosen, the resulting equations provide information about the behavior and relative importance of different physical processes.

As an example, the particle gyroradius $\rho_a = v_{ta}/\Omega_a$, where v_{ta} is the thermal velocity of species a and $\Omega_a = eB/m$ is the gyrofrequency, is one instance of a parameter that is commonly used in the ordering procedure. The gyroradius is typically small compared to larger, equilibrium-scale quantities such as the device size, and the ratio can be used as a small parameter $\rho_{*a} \equiv \rho_a/L \ll 1$, in which to take asymptotic expansions.

In this thesis, two sets of orderings are covered that lead to equations describing the primary channels of transport in fusion devices. The first leads to the drift-kinetic equation (DKE), which describes transport due to collisions and guiding-center drifts that exist in the absence of turbulence. This type of transport is termed neoclassical. The other leads to the gyrokinetic equation, which describes turbulent transport. This is concerned with the flow of particles and energy caused by instabilities that develop from fluctuating fields. The following subsections aim to cover the differences in the orderings that lead to the respective drift-kinetic and gyrokinetic equations. It should be noted here that taking an appropriate time average in the gyrokinetic ordering can actually reproduce neoclassical processes [1], indicating that turbulence and neoclassical processes can coexist. However, with the hope of both highlighting some of the properties of neoclassical and turbulent processes, as well as their differences, separate orderings will be used in the following sections.

1.2.1 Neoclassical Transport

1.2.1.1 The Drift-kinetic Ordering

To begin to describe the processes behind neoclassical transport, the following orderings are assumed: $\rho_{*a} = \rho_a/L \ll 1$, $\partial/\partial t \sim \rho_{*a}^2 v_{ta}/L$, and $\nu_a \sim v_{ta}/L$, where ν_a is the collision frequency. In the expansion of the Fokker-Planck equation, $\rho_{*a} = \rho_a/L \ll 1$ is taken to be the small parameter. It is important to note that all length scales are ordered as the equilibrium scale length L, in contrast to the gyrokinetic

ordering (Section 1.2.2.1).

The $O(\rho_*)$ terms resulting from this ordering of the Fokker-Planck equation give rise to what is known as the drift-kinetic equation (DKE) [41], which is the governing equation of neoclassical transport theory. The DKE will be discussed in more detail in Section 2.2.1.

There are a number of consequences resulting from this drift-kinetic ordering. Taking $\rho_* \ll 1$ removes information concerning the finite size of the fast particle gyration around the magnetic field, and treats the motion based on the guiding-center of its orbit (see Figure 1.6). The ordering $\partial/\partial t \sim \rho_*^2 v_{ta}/L$ assumes a slow variation in time of all physical quantities compared to all terms of $O(\rho_*)$. Thus, for the DKE, the plasma is assumed to be in a state of equilibrium (on each flux surface), where quantities such as the density and temperature, as well as their gradients are taken as constants. Finally, the choice of $\nu_a \sim v_{ta}/L$ allows for subsidiary expansions to describe different levels of collisionality. These limits are discussed more in Section 2.6.

Since neoclassical transport is strongly influenced by both collisions and trapped particles, a brief subsection has been devoted to each. This is with the hope of clarifying some of the underlying processes that lead to the calculation of neoclassical quantities, which will be covered in the final subsection.

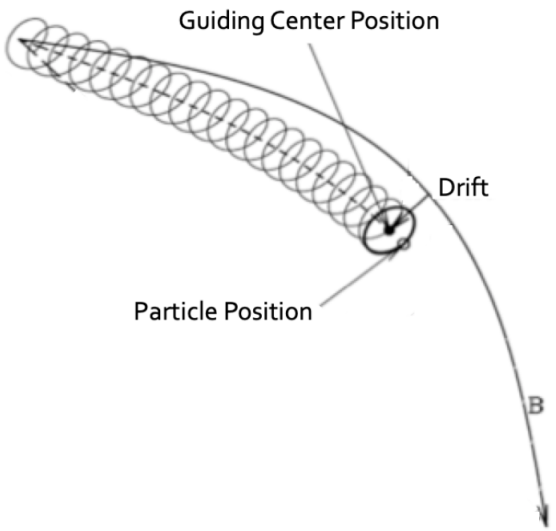

Figure 1.6: Depiction of the motion of a charged particle along a curved magnetic field. Guiding center motion concerns the motion of the *center* of the gyromotion, which will drift relative to the magnetic field due to the $\boldsymbol{E} \times \boldsymbol{B}$, ∇B, and curvature drifts. Figure adapted from [55].

1.2.1.2 The Fokker-Planck Collision Operator

The typical collision operator of the DKE (the so-called Fokker-Planck collision operator) considers the frequent small-angle collisions between charged particles. Due to the long-range forces acting between charged particles, the large-angle collisions that are seen in neutral gases are not as important in plasmas. Instead, the frequent small-angle collisions will have a much larger effect on the plasma as a whole. This behavior is standard in high-temperature plasmas with plasma parameter $\Lambda \gg 1$, where Λ is a measure of the kinetic to potential energy of the system. The small perturbations in velocity space that result from these collisions lead to simplifications and a more intuitive understanding of the collision operator. By performing an expansion in velocity space, the collision operator can be seen as a drag term plus a diffusion in velocity space [42, 60]. The drag term tends to slow

particles down, bringing the plasma as a whole toward a Maxwellian distribution.

1.2.1.3 Trapped Particles

Along with collisions, the magnetic field structure is an integral component of neoclassical transport. The variation in magnetic field strength gives rise to particles that become trapped between regions of strong magnetic fields. The cause of this can be traced back to the kinetic energy of a particle $\mathcal{E} = \frac{1}{2}m(v_\parallel^2 + v_\perp^2)$, which is conserved in the absence of an electric field. In the above equation $v_\parallel$ and $v_\perp$ represent the velocity parallel and perpendicular to the magnetic field, respectively. The kinetic energy can also be written in terms of the magnetic moment, $\mu = mv_\perp^2/2B$ as

$$\mathcal{E} = \frac{1}{2}mv_\parallel + \mu B. \tag{1.3}$$

The magnetic moment is an adiabatic invariant if the time variation of the magnetic field is small compared to the gyrofrequency ($\omega_B/\Omega \ll 1$). This is a condition that is almost always satisfied for the fluctuations relevant for transport. In this case it becomes clear that $v_\parallel$ must decrease with increasing B in order to conserve $\mathcal{E}$ and μ. Further, for a given $\mathcal{E}$, there can be a limiting value of B that will lead to $v_\parallel = 0$.

Along a field line, the parallel velocity can be expressed as

$$v_\parallel(\theta) = \sigma\sqrt{\frac{2}{m}(\mathcal{E} - \mu B(\theta))}, \tag{1.4}$$

where $\sigma = \pm 1$ represents the direction of the parallel velocity. Particles with rel-

atively large $v_\perp$ will not have the parallel velocity to overcome the increase in B along a field line and will be reflected via the mirror force when $v_\parallel = 0$ for some $\theta = \theta_0$. For the case of a tokamak, the field strength takes the form of a single well along a field line, and the reflected particle will then be reflected back at $\theta = -\theta_0$, which can be seen in Figure 1.7. The particle then becomes "trapped" in the sense that it becomes localized in the poloidal angle, and does not traverse the entire field line/flux surface. Conversely, for particles with large $v_\parallel$, the particle will not be reflected, and will traverse the entire flux surface.

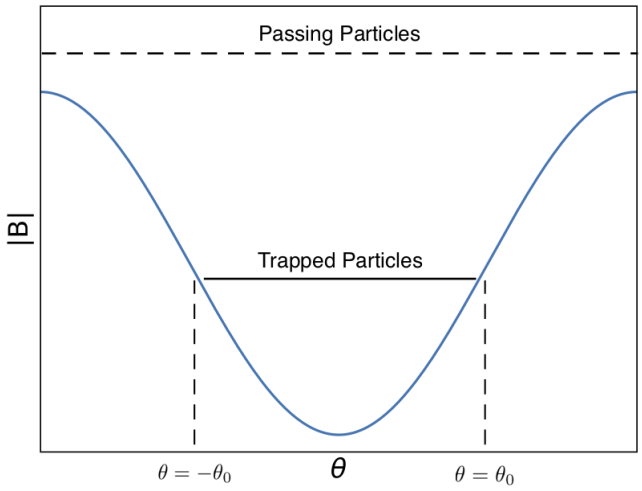

Figure 1.7: The blue curve is a plot of the magnitude of the magnetic field as a function of the poloidal angle for a tokamak. The dashed line represents a passing particle, and the solid line represents a trapped particle. Here, the locations $\theta = \pm\theta_0$ are the bounce points for the trapped particle. Note: The zero of $|B|$ is suppressed in this figure.

While the conservation principles remain the same, trapping becomes more complicated in stellarator magnetic fields (or tokamaks with magnetic field ripple), which have multiple wells of different sizes. An example of this can be seen in Figure 1.8, which shows a portion of the magnetic field as a function of θ for the W7-X

stellarator. Figure 1.8 shows that there are different *classes* of trapped particles, where the trapping occurs between different poloidal angles. More importantly, some of these regions are not symmetric about the center. Unlike stellarators, neoclassical transport in tokamaks is independent of the radial electric field, which leads to zero *net* radial motion of particles and good confinement. This net radial particle drift is zero because the drift that is experienced in moving from one bounce point to the other is equal and opposite to the drift experienced by returning to the starting bounce point in the opposite direction. In asymmetric stellarator trapping regions this cancellation does not occur. The result is secular radial motion away from the flux surface, and poor confinement of trapped particles. The exploration of techniques to mitigate this problem for impurity ions is the focus of Chapter 2.

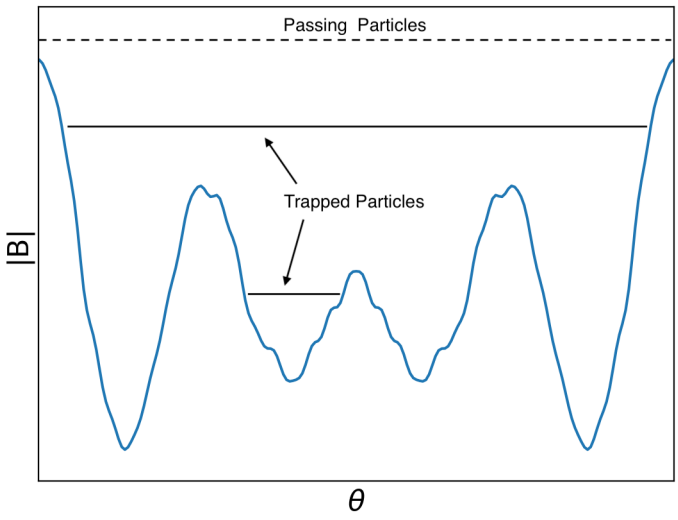

Figure 1.8: The blue curve is a plot of the magnitude of the magnetic field as a function of the poloidal angle for some stellarator. The dashed line represent a passing particle, and the solid lines represent different classes of trapped particles (i.e. localized in θ). Note: The zero of $|B|$ is suppressed in this figure.

1.2.1.4 Neoclassical Fluxes

The DKE solves for the distribution function on a flux surface. The distribution function enables the calculation of macroscopic quantities by integrating the distribution function over velocity space

$$\overline{A} \equiv \frac{1}{n_s} \int A f_s \mathrm{d}^3 v, \tag{1.5}$$

where A is some function of the particle velocity $\boldsymbol{v}$. For example, the macroscopic fluid velocity $\boldsymbol{V}_a(\boldsymbol{x}, t)$ and temperature $T_a(\boldsymbol{x}, t)$ are calculated via

$$\boldsymbol{V}_a(\boldsymbol{x}, t) = \overline{\boldsymbol{v}} \tag{1.6}$$

$$\frac{3}{2} T_a(\boldsymbol{x}, t) = \frac{\overline{m_a \boldsymbol{v}_a'^2}}{2}, \tag{1.7}$$

where $\boldsymbol{v}_a' \equiv \boldsymbol{v} - \boldsymbol{V}_a$. Using Eq 1.5, quantities such as the particle and heat flux at any given point on a flux surface can be calculated. Of principal interest is the radial component of such quantities, as it is a direct measure of how well the plasma is confined in a given magnetic field. As a function of position within a flux surface, however, these quantities can have considerable variation. It is thus more practical to obtain flux surface averages of the radial component of these quantities to determine the total transport across surfaces. The neoclassical heat flux for such an average would be

$$Q_{nc} = \left\langle \int \mathrm{d}^3 v f_s \frac{m_s v^2}{2} \boldsymbol{v}_{ds} \cdot \nabla r \right\rangle, \tag{1.8}$$

where $\boldsymbol{v}_{ds}$ is the drift velocity, and ∇r is a vector normal to the surface, whose magnitude varies within the flux surface. The notation $\langle \dots \rangle$ in Eq 1.8 represents a flux surface average, defined by

$$\langle f(r) \rangle = \frac{1}{V'(r)} \int_0^{2\pi} \int_0^{2\pi} f(r, \theta, \zeta) \sqrt{g} \, \mathrm{d}\theta \mathrm{d}\zeta, \tag{1.9}$$

where V' is the radial derivative of the volume bounded by the flux surface, and $\sqrt{g}$ is the Jacobian. The drift velocity is related to the slow motion (relative to the gyromotion) of the particle guiding center. For the work in this thesis, $\boldsymbol{v}_{ds}$ specifically refers to a combination of the motion caused by inhomogeneities in the magnetic field (Eq 2.2), and the $\boldsymbol{E} \times \boldsymbol{B}$-drift

$$\boldsymbol{v}_E = \frac{c}{B^2} \boldsymbol{E} \times \boldsymbol{B}. \tag{1.10}$$

Detailed calculations of neoclassical quantities in stellarators can be found in [93].

In summary, neoclassical transport is mainly concerned with the surface average of the radial component of quantities, which are a result of charged-particle collisions in the presence of toroidal magnetic fields.

1.2.2 Turbulent Transport

1.2.2.1 The Gyrokinetic Ordering

The assumed ordering of plasma turbulence is based on four main assumptions: strong magnetization, low frequencies, small fluctuations, and anisotropy of

these fluctuations. Strong magnetization is the equivalent of $\rho_s/L \ll 1$, and the low frequency assumption considers only frequencies well below the gyrofrequency $\omega \ll \Omega_s$. Fluctuations of fields, distributions and potentials are small compared to their equilibrium quantities $|\tilde{f}|/f \ll 1$ (fluctuations will be defined in the next section). Further, turbulent fluctuations are strongly anisotropic. They have long wavelengths parallel to the magnetic field $k_\parallel L \sim 1$, and short wavelengths perpendicular to the field $k_\perp L \gg 1$ (which can be seen in Figure 3.2). Here, $k = 2\pi/\lambda$ is the typical wavenumber for a wavelength λ. The gradients of physical quantities will then scale differently depending on the direction relative to the magnetic field. Specifically, parallel gradients vary on the equilibrium scale $\nabla_\parallel f \sim f/L$, and perpendicular gradients vary on the gyroradius scale $\nabla_\perp f \sim f/\rho$. This scale disparity is caused by rapid, sound speed communication along field lines, and diamagnetic speed communication perpendicular to field lines. The flux tube simulation domain (the topic of Chapter 3) was developed as a minimum simulation domain based on this natural scale separation.

Assuming these orderings (known as the gyrokinetic ordering) for the Fokker-Planck equation will lead to the derivation of the gyrokinetic equation (Appendix A). The derivation exploits $\omega/\Omega_s \ll 1$ to average over the fast gyromotion, which consequently leads the gyrokinetic equation to describe the evolution of a distribution of *rings* in 5-D phase space.

1.2.2.2 Fluctuating Quantities

Fluctuations in gyrokinetics and turbulence refer to a particular part of some physical quantity. For instructive purposes, one could consider the distribution function f to be comprised of an equilibrium Maxwellian part, F_0, plus a perturbation to the Maxwellian, δf, such that $f = F_0 + \delta f$. Figure 1.9 presents an example distribution function decomposed into its equilibrium and perturbed parts. The perturbed part of the distribution function is formally much smaller than the equilibrium Maxwellian part in this ordering, and this should be evident from Figure 1.9 (it is important to note that the small-scale v-dependence of Figure 1.9 is exaggerated for instructional purposes, as it is standard to assume $\partial F_0/\partial v \gg \partial(\delta f)/\partial v$).

One can further subdivide δf by defining the *fluctuations* $\tilde{f}$ of the distribution function as

$$\tilde{f} \equiv \delta f - \langle \delta f \rangle_t \,, \tag{1.11}$$

where $\langle \dots \rangle_t$ is a time average that is defined so that $\left\langle \tilde{f} \right\rangle_t = 0$. Specifically, the time average is performed on a time scale that is well-separated from the time scale of fluctuating quantities ω^{-1} and equilibrium variation of the gradients τ_E. Thus, for an intermediate time T satisfying $\omega^{-1} \ll T \ll \tau_E$, the time average as defined in [1] is

$$\langle \delta f \rangle_t \equiv \frac{1}{T} \int_{t-T/2}^{t+T/2} \delta f \, \mathrm{d}t'. \tag{1.12}$$

The averaged part of the perturbation $\langle \delta f \rangle_t$ is the quantity that is calculated in neoclassical transport. Since the above discussion is generally true for any physical

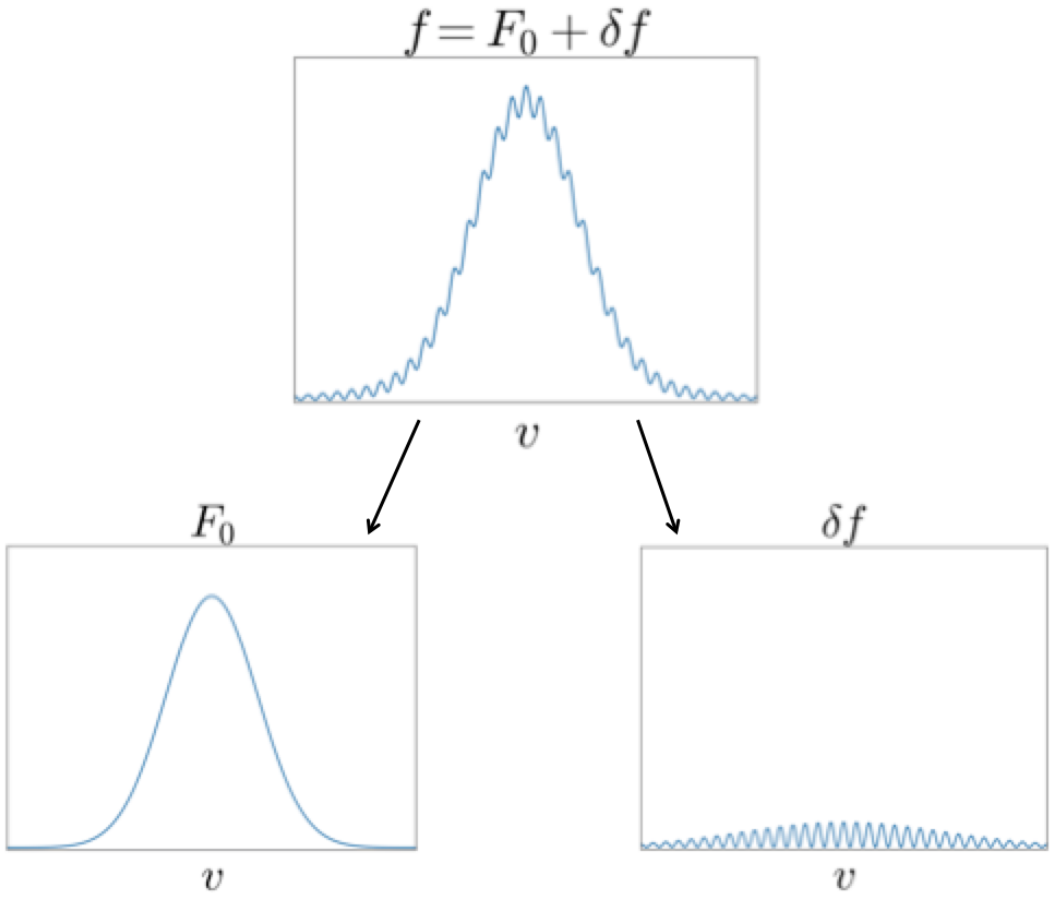

Figure 1.9: The upper figure represents an example of what a total distribution function might look like. The bottom figures represent the Maxwellian (left) and perturbed distribution (right), which together create the upper total distribution function. Each figure is plotted as a function of velocity. The small-scale v-dependence of δf is not typical (usually $\partial F_0/\partial v \gg \partial \delta f/\partial v$), and is meant for instructional purposes.

quantity (apart from the assumption of a Maxwellian for the equilibrium part), it should then be apparent that neoclassical processes ignore fluctuating quantities and involve only the steady-state behavior of small deviations to the equilibrium part of some quantity.

Turbulent processes treat these averaged quantities as constants, and instead involve the time evolution of fluctuations. The fluctuations are typically expressed as a series of Fourier harmonics $\tilde{f} = \sum_{\boldsymbol{k}} \tilde{f}_{\boldsymbol{k}} \exp\left(i\boldsymbol{k}\cdot\boldsymbol{x}\right)$. This makes it convenient to

study the evolution of plasma waves, using the $\boldsymbol{k}$ values as labels.

A common way to use the fluctuations to understand turbulence is through the use of correlation functions. The spatial correlation $C(z)$ of $\tilde{f}$ *along* the field line (z here represents the parallel coordinate) can be defined by

$$C(z) = \frac{\left\langle \tilde{f}(\boldsymbol{x}_\perp, z)\, \tilde{f}(\boldsymbol{x}_\perp, z=0) \right\rangle}{\left\langle \tilde{f}(\boldsymbol{x}_\perp, z=0)^2 \right\rangle}, \tag{1.13}$$

where the angled brackets denote an average over the perpendicular coordinates. This measures how well correlated the fluctuations are at different points along the field line, once a quasi-steady state is reached. A similar expression for the perpendicular coordinates can be constructed. The parallel (perpendicular) correlation length l_z ($l_\perp$) can be taken as the constant length scale in

$$C(z) \sim \exp\left(-\frac{|\Delta z|}{l_z}\right), \tag{1.14}$$

describing the typical length scale over which fluctuations are correlated. Knowledge of the correlation lengths allow one to efficiently set the dimensions of a simulation domain. This is in the sense that the fluctuations should be sufficiently decorrelated from one end of the domain to the other, which is discussed further in Chapter 4.

1.2.2.3 Important Quantities in Turbulent Transport

Analogous to neoclassical transport, turbulent transport concerns the transport of particles and energy across surfaces to diagnose the confinement properties

under various conditions. Simulations of the linearized gyrokinetic equation are a common way to study the evolution of each individual mode in the absence of interactions with other modes. Determining the growth rate is a standard goal of linear simulations, as it reveals which instabilities can be most problematic. Further, the structure of each mode along a magnetic field line can be found. This provides valuable information about where in the device such modes tend to concentrate.

However, the full physical picture that is seen with the nonlinear interaction is quite different. Growth rates, for example, which can be useful in determining *how fast* turbulence is reached, become less important at the state of turbulent saturation. For a saturated turbulent state, quantities such as the turbulent particle flux, and the fluctuation amplitudes for each individual mode are usually of interest. Perhaps the most valuable information, however, is the radial heat flux, summed over each of the Fourier modes, since we're really concerned about the total transport across surfaces. Like neoclassical transport, this provides a direct measure for how well the plasma is able to confine heat under a given set of parameters.

1.3 Optimized Stellarator Configurations

As discussed in Section 1.1.2 above, stellarators have a distinct advantage over tokamaks in that the positional equilibrium is well-maintained as the plasma state changes. This is because stellarators use external coils to generate the confining magnetic fields, and do not rely on the self-organization of plasma currents like in tokamaks. The consequence is that one can create optimized equilibria that will not

be strongly affected during operation. The equilibria can then be created to target the improvement of certain properties. This has led to efforts in the stellarator community to explore how to best optimize stellarators for future experiments or reactors [37].

While the possibility of optimizing the behavior of stellarators is exciting, the large number of possible parameters to optimize will require trade-offs. This typically leads to the choice of optimizing for a particular behavior, while still keeping other behaviors in check. For example, because stellarators have poor neoclassical transport, efforts have been made to optimize the equilibrium in order to reduce that transport. W7-X is an example of a stellarator optimized for neoclassical transport. It is also possible to optimize for properties like fast particle transport [5], turbulence [107], or even certain types of symmetry (explored in Chapter 2).

One specific optimization technique is to make stellarators quasisymmetric. Quasisymmetry is a property that appears only in a particular coordinate system known as Boozer coordinates. Such a symmetry in stellarators leads to guiding-center transport properties that are identical to tokamaks. This could eliminate the issue of poor particle confinement in stellarators. However, actual devices will deviate from perfect quasisymmetry, and the transport properties of these configurations must be studied. The topic of Chapter 2 addresses impurity transport in such quasisymmetric stellarators with "symmetry-breaking".

These cases are only a few examples of the potential of stellarator optimization. There is a vast landscape of parameter space that remains to be explored, making the future of stellarator experiments (or reactors) very hopeful.

1.3.1 Visualization of Optimized Stellarator Flux Surfaces

A number of optimized stellarator configurations are considered throughout this work. For visualization and reference, a flux surface is shown for each of these configurations in Figures 1.10-1.17 below. While nost configurations are visually distinct from one another, the Wistell-A [5] and ARIES-CS [75] configurations evolved from the HSX [3] and NCSX [109] designs, respectively. This results in similar-looking flux surfaces in Figures 1.11,1.12 and 1.13,1.14.

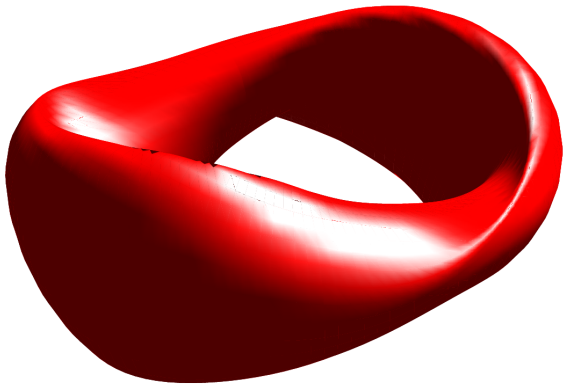

Figure 1.10: Henneberg-QA [48]

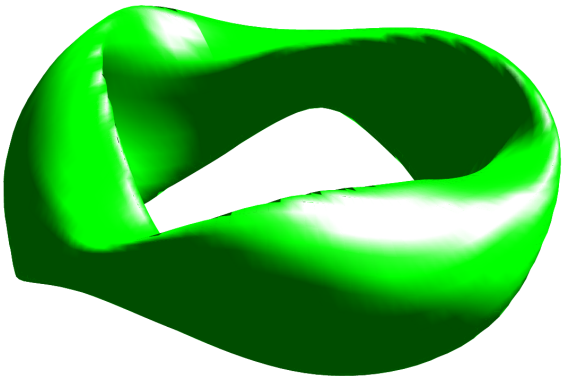

Figure 1.11: NCSX [109]

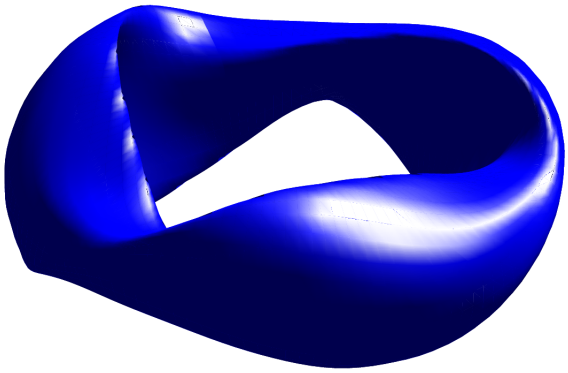

Figure 1.12: ARIES-CS [75]

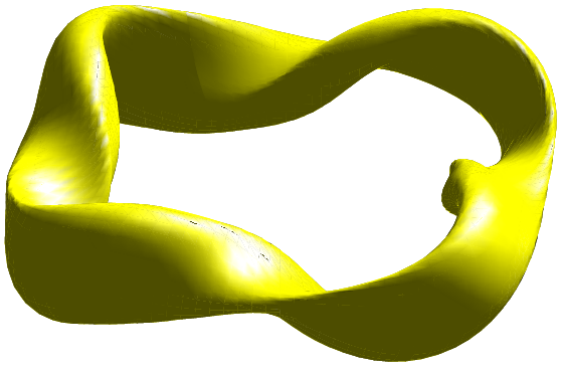

Figure 1.13: Wistell-A [5]

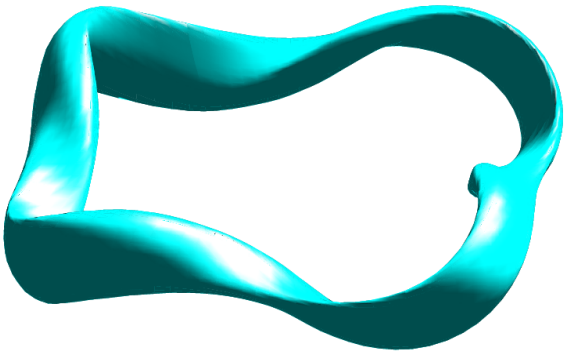

Figure 1.14: HSX [3]

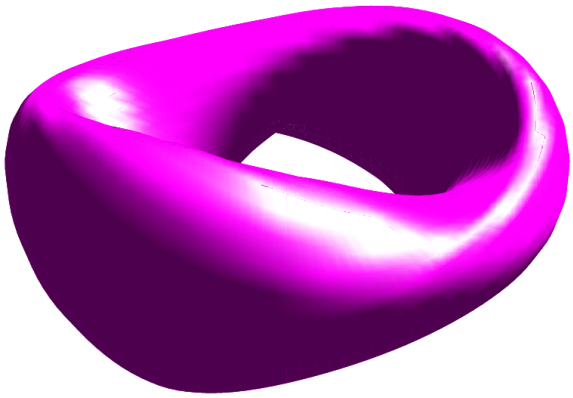

Figure 1.15: Garabedian [33]

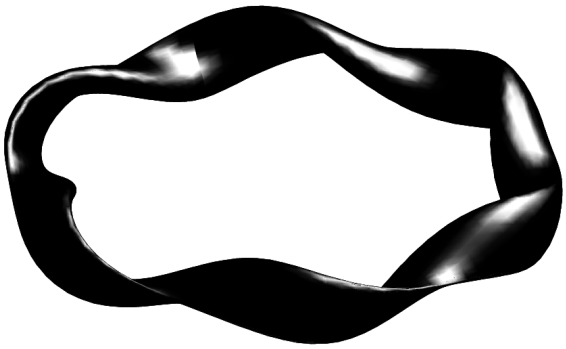

Figure 1.16: Nuhrenberg [77]

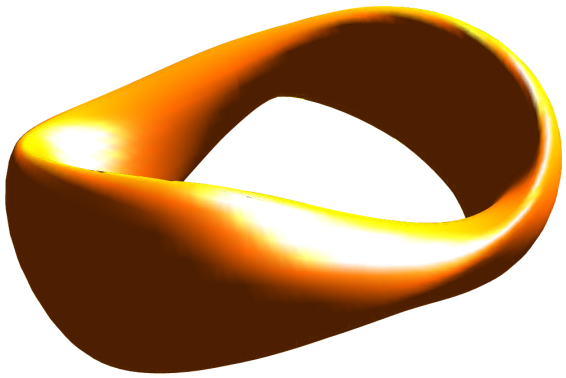

Figure 1.17: CFQS [63, 90]

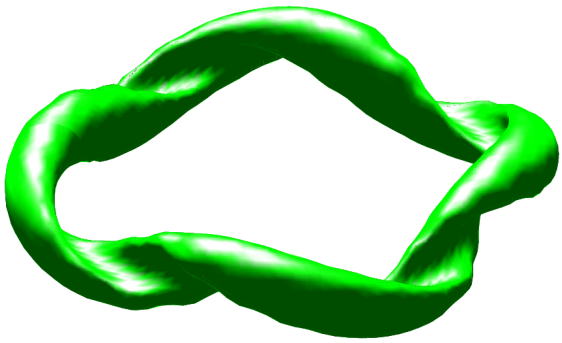

Figure 1.18: TJ-II [47]

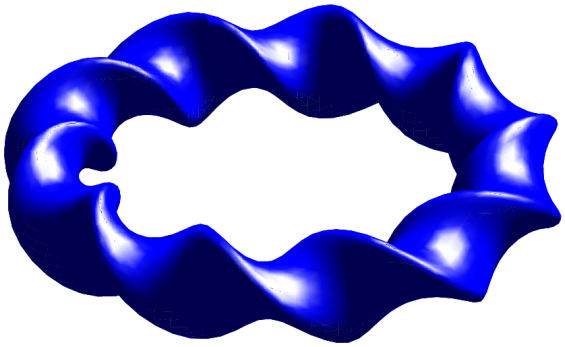

Figure 1.19: LHD [72]

Figure 1.20: W7-X [40]

1.4 Unsolved Problems in Stellarator Plasma Transport

There remain a number of open problems in stellarator transport that must be addressed [38].

Turbulence has a strong impact on plasma confinement, and it is not well-understood, especially in stellarators. Part of the problem can be traced to the high cost of turbulence simulations that have severely limited the size and scope of the runs. Therefore, to better understand and characterize turbulence, it is necessary to reduce simulation times as much as possible.

Energetic particle confinement is another transport issue that must be overcome to reach a state of self-sustaining fusion reactions. The low-collisionality of fast particles promotes the importance of achieving configurations very close to quasisymmetry to bring neoclassical transport levels closer to the respective tokamak values. Instabilities can also develop in the presence of energetic particles through wave-particle interactions that can lead to further losses.

Another foreseen problem in the stellarator community is impurity control. Impurities in the core can cause fuel dilution and radiative power losses, while impurities near the edge can reduce heat fluxes on the divertors. Specifically, the neoclassical theory of impurity particle transport in various optimized stellarators (including Φ_1 effects) needs to be better understood, along with the relative importance of turbulent vs. neoclassical processes in impurity transport.

The content in this chapter borrows heavily from the work in:
Martin M.F., Landreman M. *Journal Plasma Phys.* **86** 905860317 (2020) [67]

Chapter 2: Effects of Magnetic Field Symmetry-Breaking on Impurity Transport in Quasisymmetric Stellarators

2.1 Introduction

The ideal makeup of particles in the core of magnetic confinement fusion devices would consist exclusively of particles participating in the fusion reaction. The presence of any impurity ions can degrade fusion performance by means of fuel dilution and radiative cooling of the plasma [50, 80]. Removing impurities from the plasma core and preventing further accumulation is then of primary importance in present devices, and when designing future experiments.

Due to the symmetric nature of a tokamak, its neoclassical transport properties yield a distinct advantage over non-axisymmetric configurations because they are independent of the radial electric field, E_r, at leading order in (ρ_i/L) [42,86]. Here, ρ_i is the ion gyroradius, and L is some equilibrium scale length. In particular, if certain conditions are met (see Section 2.2.2), this absence of E_r in the transport equations leads to a property known as temperature screening [101], which guarantees an outward radial flux of impurities for large enough temperature gradients.

Conversely, unoptimized stellarator designs have been predicted to behave

poorly with regards to impurity accumulation [45, 50, 56, 57]. The lack of toroidal symmetry in the magnetic field complicates the transport quantities because of a dependence on E_r in order to maintain ambipolarity of the constituent particle fluxes. This can become an issue in reactor-relevant plasmas, which are expected to operate in the ion-root regime [65], where the negative (inward) E_r will tend to pull impurities into the core. Recent work [43], however, has found that outward impurity fluxes can be achieved in the "mixed-collisionality" regime in a stellarator (later qualified analytically with flux-surface variations in E_r by [15,16]), alleviating some of the concern.

Improving the behavior of impurities in stellarators could be addressed by contemporary stellarator design optimization, where one of the current foci is on quasisymmetric magnetic fields [77]. Quasisymmetric fields have the allure of possessing the superior transport properties of tokamaks alongside the stability of stellarators. Ideally, perfect quasisymmetry would lead to neoclassical and guiding-center transport properties identical to tokamaks [12, 81]. However, it has been shown [36] that perfect quasisymmetry can likely be achieved only on a single flux surface. Therefore, any future experiment or reactor will necessarily have some finite degree of symmetry-breaking. This makes it important to study quasisymmetric equilibria with some departure from perfect symmetry.

In this chapter, we examine the temperature screening effect using the SFINCS [61] (Stellarator Fokker-Planck Iterative Neoclassical Conservative Solver) drift-kinetic solver to calculate the impurity particle flux for a number of quasisymmetric equilibria. As we proceed, it will be necessary to distinguish between a perfectly

quasisymmetric magnetic field, and the quasisymmetry of configurations such as the National Compact Stellarator Experiment (NCSX) [109] or the Helically Symmetric Experiment (HSX) [3]. For example, the magnetic field of HSX is quasisymmetric in the sense that its quasisymmetric harmonics are dominant compared with the smaller, but non-zero, symmetry-breaking harmonics. Such a magnetic field will be referred to as the actual, or true, magnetic field of a configuration. A *perfectly* quasisymmetric magnetic field is one in which the symmetry-breaking modes are identically zero.

With this distinction, the unanswered question we would like to address is whether, in a nominally quasisymmetric stellarator with realistic deviations from perfect symmetry, the sign of the neoclassical impurity flux is outward like in tokamaks at low collisionality, or inward like in a generic stellarator.

By altering the magnitude of symmetry-breaking harmonics in the magnetic field of a given equilibrium (see Section 2.3), we are able to probe the region where temperature screening is lost. Holding the temperature constant, this effect was studied at three distinct densities, and correspondingly three distinct collisionalities (all of which are considered to be low collisionality, as defined in Section 2.6). At the lowest collisionality, no configurations are able to maintain an outward impurity flux at the true magnetic field, even for $\eta^{-1} \equiv d\ln(n_a)/d\ln(T_a) = 0$, where a refers to species. (Introducing a finite peaked density gradient for the main ions, $\eta_i^{-1} > 0$, always makes the impurity particle flux more inward, which is explained in Section 2.6.1.2). Increasing the collisionality has a favorable effect, where some configurations were even found to have an outward impurity flux. However, there

is an upper collisionality limit, beyond which temperature screening is not observed for most configurations, even in perfect quasisymmetry.

Impurity accumulation in perfect quasisymmetry with $\eta^{-1} = 0$ can either be caused by exceeding some collisionality limit, or by a dependence of the neoclassical transport on E_r, indicative of a breakdown in the intrinsic ambipolarity assumption. In the latter case, the $\boldsymbol{E} \times \boldsymbol{B}$ drift, v_E, is close to being in violation of the $v_E \sim \rho_* v_{t\alpha}$ ordering in deriving the equations solved in neoclassical codes. In Section 2.4, we examine this in further detail and calculate the resonant radial electric field, E_r^{res}, in quasisymmetric configurations. One finds that E_r^{res} is fundamentally smaller in quasi-axisymmetry (QA) than in quasi-helical symmetry (QH).

We have also compared the magnitude of the resulting neoclassical fluxes to a gyro-Bohm estimate for turbulence [22,29,103]. At reactor-relevant parameters, the neoclassical impurity particle flux did not exceed the respective turbulent flux for any impurity species or configuration. Even in the presence of a strongly peaked ($\eta^{-1} = 0.5$) density gradient, in most configurations the neoclassical impurity particle flux is less than 10% of the estimated turbulent value. This suggests that regardless of whether a configuration can achieve temperature screening, the nature of the turbulence will determine the sign of the particle flux on a surface.

The total (bulk ion + impurity) neoclassical heat flux also did not exceed the turbulent contribution. However, the ratio was larger than the analogous impurity particle flux ratio. Furthermore, the neoclassical contribution is largest near the magnetic axis, and results indicate that turbulence becomes increasingly more dominant as one moves further out radially. This is in agreement with experimen-

tal observations [19, 78] in Wendelstein 7-X (W7-X) [40] and HSX, which find that neoclassical transport is the dominant radial transport channel near the magnetic axis.

Finally, we compared the critical amount of symmetry-breaking that it takes to change the sign of the particle flux, ϵ_{sb}^{c}, to two metrics that have been used to quantify symmetry on a flux surface. These metrics are the effective helical ripple [76], ϵ_{eff}, which is a measure of neoclassical transport in the $1/\nu$ regime, and the magnitude of the symmetry-breaking terms on a flux surface, S (see Eq. 2.17). While it was found that there was some anti-correlation between ϵ_{sb}^{c} and S, there does not appear to be much of a relationship between ϵ_{eff} and ϵ_{sb}^{c}. (This should not be surprising, however, if one considers that W7-X has a very low ϵ_{eff}, yet it is far from quasisymmetry). This difference between how ϵ_{eff} and S depend on ϵ_{sb}^{c}, a quantity related to symmetry, motivates a comparison between ϵ_{eff} and S. Results indicate a configuration-specific dependence of ϵ_{eff} on S, which in many cases is non-monotonic. There is thus a disconnect between these two quantities, such that minimizing the amount of symmetry-breaking on a flux surface does not simultaneously minimize ϵ_{eff}. So while ϵ_{eff} is a useful proxy for optimizing neoclassical transport in stellarator optimization, it is a poor proxy for achieving good quasisymmetric surfaces.

This chapter is organized as follows: In Section 2.2.1, we introduce the governing equation and ordering assumptions of SFINCS in the results presented herein. In Section 2.2.2, we explain the principle of ambipolarity, the fundamentals of the temperature screening phenomenon, and the issues that arise in non-axisymmetric

geometries. In Section 2.3 we explain our approach to varying the degree of quasisymmetry on a flux surface. The quasisymmetric configurations that have been explored, and the way that these equilibria have been scaled can be found in Section 2.5. Section 2.4 explains an issue in present neoclassical stellarator codes based on the $v_E \sim \rho_* v_{ta}$ ordering, which limits the value of the radial electric field when impurities are included. Section 2.6.1 presents results on how the amount of symmetry-breaking, collisionality, and density gradient affect the behavior of the impurity particle flux for various quasisymmetric configurations. Section 2.6.2 compares the results of Section 2.6.1 to a gyro-Bohm estimate of turbulent particle and heat fluxes as a function of the impurity species, and normalized radius. Finally, Section 2.7 compares the effective helical ripple to the amplitude of symmetry-breaking terms on a flux surface.

2.2 Background

2.2.1 Drift Kinetic Equation

Neoclassical transport follows from a drift-ordering of the Fokker-Planck equation in toroidal magnetic geometry, and solving the resulting drift-kinetic equation (Eq.19 in [41]). The drift ordering assumes $\rho_{*a} = \rho_a/L \ll 1$, $v_E \sim \rho_{*a} v_{ta}$, $\partial/\partial t \sim \rho_{*a}^2 v_{ta}/L$, and $\nu_a \sim v_{ta}/L$, where ν_a is the collision frequency. The gyroradius of species a is $\rho_a = v_{ta}/\Omega_a$, the thermal velocity $v_{ta} = \sqrt{2T_a/m_a}$, with T_a and m_a the temperature and mass of species a, respectively. The gyrofrequency is $\Omega_a = Z_a e B/m_a c$, where Z_a is the species charge, B is the magnetic field magnitude,

c is the speed of light, and e is the proton charge.

Results in this chapter have been obtained by solving the drift-kinetic equation (DKE) using the SFINCS [61] code over a range of collisionality regimes, for various impurity ions. SFINCS is a radially-local DKE-solver that has been generalized to non-axisymmetry, allowing for an arbitrary number of species, fully linearized Fokker-Planck collision operator, and the capability of simulating on-surface variations in the electrostatic potential, Φ_1. The exact form of the DKE that is solved in SFINCS for this chapter (with the exception of Section 2.6.2.2) is given by Eq.(16) in [61]

$$\dot{\boldsymbol{r}} \cdot (\nabla f_{a1})_{x_a,\xi} + \dot{x}_a \left(\frac{\partial f_{a1}}{\partial x_a}\right)_{\boldsymbol{r},\xi} + \dot{\xi}_a \left(\frac{\partial f_{a1}}{\partial \xi}\right)_{\boldsymbol{r},x_a} - C_a = -\left(\boldsymbol{v}_{ma} \cdot \nabla r\right) \left(\frac{\partial F_a}{\partial r}\right)_{W_{a0}}, \tag{2.1}$$

where F_a and f_{a1} represent a Maxwellian distribution and the first-order perturbation to the distribution function, respectively. The position vector is given by $\boldsymbol{r}$, the cosine of the pitch angle is $\xi \equiv v_\parallel/v$, the velocity is $x_a \equiv v/v_{ta}$, $W_{a0} = v^2/2 + Z_a e\Phi_0/m_a$ is the lowest-order total energy, C_a is the collision operator, and $\boldsymbol{v}_{ma}$ is the magnetic drift velocity defined by

$$\boldsymbol{v}_{ma} \cdot \nabla r = \frac{m_a c}{Z_a e B^2} \left(v_\parallel^2 + \frac{v_\perp^2}{2}\right) \boldsymbol{b} \times \nabla B \cdot \nabla r. \tag{2.2}$$

The coordinate $r = \sqrt{2\psi_t/B_{av}}$ is a surface label, where $2\pi\psi_t$ is the toroidal flux, and B_{av} is some averaged magnetic field, such as the field on the magnetic axis. The electrostatic potential is split into zeroth- and first-order contributions $\Phi =$

$\Phi_0(r) + \Phi_1(r, \theta, \zeta)$, where θ and ζ are the poloidal and toroidal angles. The zeroth-order term $\Phi_0 \equiv \langle \Phi \rangle$, where $\langle \dots \rangle$ is a flux surface-average, and Φ_1 is determined from the first-order quasineutrality equation

$$\sum_a \left(-\frac{Z_a^2 e \Phi_1}{T_a} n_a + Z_a \int \mathrm{d}^3 v f_{a1} \right) = 0. \tag{2.3}$$

The time derivative terms, $\dot{\boldsymbol{r}}$, $\dot{x}_a$, and $\dot{\xi}_a$ are the phase space particle trajectories. Since SFINCS offers variations in how these trajectories are defined, we have chosen to use the "full trajectories" definition (Eq.(17) in [61]). This choice takes into account the change in potential energy as a particle drifts radially, with the corresponding change to $\dot{\xi}_a$ in order to conserve the magnetic moment, μ. Finally, note that Φ_1 effects are neglected in these phase space trajectories and in Eq 2.2.1; we will discuss how Φ_1 effect can be included in Section 2.6.2.2.

2.2.2 Ambipolarity and Temperature Screening

The property of ambipolarity can be expressed as

$$\sum_a Z_a \Gamma_a = 0, \tag{2.4}$$

where Γ_a is the radial component of the particle flux of species a

$$\Gamma_a = \left\langle \int f_a \left(\boldsymbol{v}_{ma} \cdot \nabla r \right) \mathrm{d}^3 v \right\rangle. \tag{2.5}$$

This results from the charge density being small for length scales much longer than the Debye length. Ambipolarity is then a statement that the flux surface-averaged radial current vanishes on each flux surface. While this is true in both tokamaks and stellarators, the radial electric field is set by different physical mechanisms, and E_r only affects the ambipolarity condition in stellarators. The value of the radial electric field that satisfies the ambipolarity condition in a non-axisymmetric plasma is referred to as the ambipolar radial electric field.

Neoclassical fluxes are determined from a linear combination of the equilibrium gradients in the system. The radial neoclassical impurity particle flux can be written in the form of Eq (1) in [100]:

$$\Gamma_z = -n_z \sum_a L_{11}^a \left[\frac{1}{n_a}\frac{dn_a}{dr} - \frac{Z_a e E_r}{T_a} + \delta_a \frac{1}{T_a}\frac{dT_a}{dr} \right], \tag{2.6}$$

where r is an arbitrary radial coordinate, and L_{11}^a and δ_a are transport coefficients [10,65,100] that can have a complicated dependence on E_r and the collision frequency $\nu_z = \nu_{zi} + \nu_{zz}$, where

$$\nu_{ab} = \frac{4\sqrt{2\pi} n_b Z_a^2 Z_b^2 \ln\Lambda}{3\sqrt{m_a} T_a^{3/2}}. \tag{2.7}$$

In the case of a tokamak, the toroidal symmetry causes the electric field dependence to cancel out in Eq 2.4, making transport *intrinsically ambipolar*. This, in principle, allows for rapid toroidal rotation of the plasma, as the radial electric field profile is not governed by the ambipolarity constraint, but rather by angular momentum conservation [1,42,52,89]. Moreover, for a plasma in the banana regime with $\delta_i + \delta_z >$

0, this cancellation assures a radially outward flux of the impurity species when $\eta_i^{-1} \equiv \mathrm{d}\ln n_i/\mathrm{d}\ln T_i < \eta_c^{-1}$, where η_c^{-1} is some critical ratio of the density and temperature gradients (the ratio η_i^{-1} is used here since $\nabla \ln n_i$ drives inward impurity transport, whereas $\nabla \ln n_z$ drives outward transport). This beneficial phenomenon is generally referred to as temperature screening. Specifically, we define temperature screening to be present when the flux surface-averaged impurity particle flux is positive. This definition makes clear how certain parameters affect the direction of travel of the impurities, which is ultimately the quantity of interest. However, it should be noted that temperature screening can also be defined by the sign of the temperature gradient coefficient in Eq 2.6.

The situation is less positive in stellarator geometries since the ambipolarity condition is dependent on the radial electric field, and the temperature screening effect is no longer guaranteed. For fusion-relevant, high-density plasmas, the radial electric field is directed inward [65] and will presumably act to drive higher-Z impurities into the core. It should be stated clearly here that temperature screening is a *neoclassical* effect, and its presence, or lack thereof, is independent of the turbulent fluxes.

A potential solution to this situation in stellarators lies in the design of quasisymmetric configurations. A truly quasisymmetric device, whose magnetic field varies through a fixed linear combination of Boozer angles on a flux surface, will have neoclassical and guiding-center transport properties identical to a tokamak, up to $O(\rho_{*a})$ [12, 81]. However, evidence suggests that in the absence of axisymmetry, quasisymmetry cannot be achieved exactly throughout a volume [36], meaning

that quasisymmetric devices will necessarily deviate from symmetry to some level. Therefore, it would be informative to optimization efforts to know how much breaking in the symmetry of the magnetic field can be tolerated before the temperature screening effect is lost. In the following sections, we explore the effect that magnetic field symmetry-breaking has on the impurity particle flux.

2.3 Magnetic Field Symmetry-breaking

Any magnetic field within a flux surface can be written as a sum of harmonics in the Boozer poloidal θ, and toroidal ζ angles [13]

$$B(r,\theta,\zeta) = \sum_{m,n} B_{mn}(r) \mathrm{e}^{i(m\theta - n\zeta)}. \tag{2.8}$$

Only by expressing the magnetic field in Boozer coordinates will the property of quasisymmetry become apparent. A magnetic field is considered quasisymmetric if its magnitude varies within a flux surface only through the fixed linear combination $\chi = M\theta + N\zeta$, where M, N are fixed integers, one of which may be zero. However, since perfect quasisymmetry is not practically achievable, it is possible to express the magnetic field of a quasisymmetric configuration as a sum of quasisymmetric and *non*-quasisymmetric Boozer harmonics, the latter of which will be referred to as symmetry-breaking terms. Therefore, symmetry-breaking terms with smaller magnitudes will produce better approximations to perfect symmetry.

Our approach to understanding temperature screening in stellarators exploits this fact by allowing one to adjust the amplitude of the symmetry-breaking terms

by an overall, constant scaling factor. The way we have decided to approach this is to expand in the harmonics of $1/B^2$, which can be expressed as

$$\frac{1}{B^2} = \sum_q h_q(r) \mathrm{e}^{iq\chi} + \epsilon_{sb} \sum_{m,n} h_{mn}(r) \mathrm{e}^{i(m\theta - n\zeta)}, \tag{2.9}$$

where the quantities h_q and h_{mn} are the quasisymmetric and non-quasisymmetric harmonics of the $1/B^2$ expansion, respectively. The parameter ϵ_{sb} is a scaling factor (fixed for a given simulation) that controls the amplitude of the symmetry-breaking terms. The special case of $\epsilon_{sb} = 0$ denotes a truly quasisymmetric field, while $\epsilon_{sb} = 1$ corresponds to the original magnetic field that one would get from an equilibrium code. By running simulations with ϵ_{sb} between these values, one can gain further insight into how temperature screening is affected under magnetic fields with varying degrees of symmetry-breaking.

In the context of MHD, artificially scaling the magnetic field with $\epsilon_{sb} \neq 1$ will lead to a plasma that no longer satisfies the equations of an MHD equilibrium. However, if we consider the work of Garren/Boozer [36], it is likely that the construction of a *single* quasisymmetric flux surface is possible. Then, an arbitrarily quasisymmetric magnetic field could be constructed on one of many flux surfaces that, in principle, will satisfy an MHD equilibrium. Since this applies to only one flux surface, our approach prevents the scaling of multiple flux surfaces simultaneously.

To understand our choice of expanding $1/B^2$, it is important to recognize that artificially scaling the magnetic field of an MHD equilibrium can potentially become

problematic if large currents are introduced near rational surfaces [14,46]. This can be seen in the expression for the parallel current [46]

$$J_{\parallel} \sim \frac{h_{mn}}{r - r_{mn}} \frac{dp}{dr}. \tag{2.10}$$

Since our simulations will always assume a finite pressure gradient, the h_{mn} modes must vanish on rational surfaces to avoid an infinite Pfirsch-Schlüter current. Therefore, scaling h_{mn} modes as opposed to B_{mn} modes will guarantee that such currents will not appear in this altered equilibrium.

2.4 Resonant Radial Electric Field Considerations

In tokamaks, it is well known that rapid plasma rotation is possible in the toroidal direction as a result of symmetry. If one then assumes the ordering of $v_E \sim v_{ta}$, then radial electric fields are capable of producing sonic flows. The radial electric field that corresponds to sonic rotation is known as the resonant electric field, which in axisymmetry is $E_r^{res} = r\iota B v_{ta}/(Rc)$ [11]. We take the radial electric field here to be defined by $E_r \equiv -d\Phi/dr$.

Constraints on the symmetry of the magnetic field, however, prevent ordering the flow velocity with the thermal speed in generic stellarators [44], as well as perfectly quasisymmetry ones [95]. The form of the drift kinetic equation that is solved in SFINCS uses the $v_E \sim \rho_{*a} v_{ta}$ ordering to avoid the symmetric restrictions to the magnetic field that result from sonic flows. From the SFINCS ordering, the vast majority of parameter regimes, geometries, and species, yield ambipolar ra-

dial electric fields, E_r^a, that are considerably smaller than the resonant electric field magnitude. However, the $m_a^{-1/2}$ dependence of the resonant electric field can cause the ordering to break down for heavy impurities under certain conditions. Solving this issue completely would demand a reordering to derive a new form of the drift kinetic equation. We do not attempt to tackle this problem here, but leave it to future work.

It is also very interesting and relevant to note that quasi-helically-symmetric (QH) configurations produce a considerably larger gap between E_r^a and E_r^{res} than quasi-axisymmetric (QA) configurations for otherwise identical plasma parameters. The relative size of these electric fields for a given simulation can be found by deriving an analogous expression for the resonant radial electric field in quasisymmetry.

If we start by assuming the $v_E \sim v_{ta}$ ordering, then the v_E and parallel streaming terms will be of the same order

$$v_\parallel \hat{\boldsymbol{b}} \cdot \nabla \chi \sim \frac{c}{B^2} \boldsymbol{E} \times \boldsymbol{B} \cdot \nabla \chi, \tag{2.11}$$

where $\chi = M\theta - N\zeta$. The contravariant and covariant representations of the magnetic field are, respectively,

$$\boldsymbol{B} = \nabla \psi_t \times \nabla \theta + \iota \nabla \zeta \times \nabla \psi_t \tag{2.12}$$

$$\boldsymbol{B} = L \nabla \psi_t + I \nabla \theta + G \nabla \zeta, \tag{2.13}$$

where $2\pi\psi_t$ is the toroidal flux, $L = L(\psi_t, \theta, \zeta)$ is some scalar, and as detailed in [46]

$$\int_{S_\zeta} \boldsymbol{J} \cdot \nabla\zeta \sqrt{g}\, \mathrm{d}\theta\, \mathrm{d}\psi_t = \frac{c}{2} I(\psi_t) \tag{2.14}$$

$$\int_{S_\theta} \boldsymbol{J} \cdot \nabla\theta \sqrt{g}\, \mathrm{d}\zeta\, \mathrm{d}\psi_t = \frac{c}{2} G(\psi_t), \tag{2.15}$$

where $1/\sqrt{g} = (\nabla\psi_t \times \nabla\theta)\cdot\nabla\zeta$ is the Jacobian, and S_ζ and S_θ correspond to surfaces where $\zeta = const$ and $\theta = const$, respectively. Solving for E_r when $v_E \sim v_\parallel \sim v_{th}$ yields an expression for the resonant electric field

$$E_r^{res} \sim \left| \frac{r v_{ta} B^2}{c} \frac{M\iota - N}{MG + NI} \right|. \tag{2.16}$$

Typically, $I \ll G$, so the QH devices examined in the following sections (all of which have $M = 1$, $|N| \geq 4$) will have $(E_r^{res})_{QH} \simeq |(\iota - N)/\iota|\, (E_r^{res})_{QA}$. This allows one to run neoclassical codes at larger E_r^a before the ordering breakdown is reached. For this reason, only QH results are available in some SFINCS simulations with steep gradients and/or heavier impurity ions. As a workaround for QA, we have only considered cases where $E_r^a/E_r^{res} < 1/3$, in order to be sufficiently far from the resonance to avoid unreliable results due to the breakdown of $v_E \sim \rho_{*a} v_{ta}$.

2.5 Magnetic Field Configurations

Throughout the remainder of this chapter, we aim to provide results that are general to a wide range of quasisymmetric stellarators. We have therefore chosen eight distinct quasisymmetric stellarator configurations (summarized in table 2.1)

Quasisymmetric Stellarators			
Configuration	QS Type	N_{fp}	Aspect Ratio
Henneberg [48]	QA	2	3.40
NCSX [109]	QA	3	4.37
ARIES-CS [75]	QA	3	4.56
Wistell-A [5]	QH	4	6.94
HSX [3]	QH	4	10.17
Garabedian [33]	QA	2	2.60
Nuhrenberg-Zille [77]	QH	6	11.76
CFQS [63, 90]	QA	2	4.35

Table 2.1: Quasisymmetric stellarator configurations that have been studied in this work. QA-quasi-axisymmetric, QH-quasi-helically-symmetric, N_{fp}-Number of Field Periods.

to examine, some of which were designed to be QA, and the others QH. Here, we have used the C09R00 equilibrium from NCSX, the Nuhrenberg configuration from figure 1 and table 1 in [77], and the quasi-helically-symmetric configuration of HSX. HSX is the only configuration in this list that has been constructed to date. To allow for a fair comparison between devices, each device was scaled to the minor radius, a, and on-axis magnetic field, B_0, of the Henneberg et al QA configuration [48], $a = 0.602$m and $B_0 = 2.10$T.

2.6 Results

The results generated below employ the full linearized Fokker-Planck collision operator of SFINCS, with two ion species. The main ions are taken to be hydrogen in each of the simulations, while the charge and mass of the impurity ion can vary between runs. Unless otherwise specified (2.6.2.2), the electrostatic potential is taken to be constant on a flux surface $\Phi = \Phi_0(r)$. It is assumed here that ion

temperatures are equivalent, $T_i = T_z$, due to the fast equilibration time. The choice of the temperature and density profiles in the following results is based on the modeling of an ECRH-heated, W7-X plasma in Fig 5 of [98]. The density gradient, however, is not determined from Fig 5 in [98], but rather chosen so as to give particular values of η^{-1}. Further, the density gradient is taken to be equivalent between ion species $\nabla \ln n_i = \nabla \ln n_z$ (or equivalently $\eta^{-1} = \eta_i^{-1} = \eta_z^{-1}$), meaning that the profile of Z_{eff} is flat.

The recent work of [43] has shown that temperature screening can be achieved in reactor-relevant, *mixed-collisionality* plasmas (highly-collisional impurities and low-collisionality bulk ions) at large normalized radius $r_N = 0.88$. With the increased temperature in the core, however, it is possible that the bulk ions and impurities in reactor-grade plasmas will both have low collisionalities, depending of course on the particular impurity ion. The picture for temperature screening becomes more pessimistic as collisionality decreases, which can be seen in Fig 1 and 2 in [43]. For our purposes of understanding how much symmetry-breaking can be tolerated prior to losing this effect, we choose to study collisionalities below the region of temperature screening in [43], in order to ensure that this transition will be observed in at least some cases.

Specifically, in the results that follow, both the ions and impurities will fall into what is generally considered the low-collisionality regime in stellarators $\nu_*^a \ll 1$, where we define $\nu_*^a \equiv \nu_a R/v_{ta}$. The low-collisionality regime can be further subdivided into the plateau, $1/\nu$, and $\sqrt{\nu}$ regimes, as derived for realistic aspect ratio in [16]. The plateau regime is defined by $\epsilon^{3/2} \ll \nu_*^a \ll 1$, the $1/\nu$ regime by

$\epsilon^{-1}\rho_{*a} \ll \nu_*^a \ll \epsilon^{3/2}$, and the $\sqrt{\nu}$ regime by $\nu_*^a \ll \epsilon^{-1}\rho_{*a}$, where $\epsilon \equiv a/R$. We further explicitly define the equilibrium length scale in $\rho_{*a} \equiv \rho_a/R$ to be consistent with the use of the major radius in the definitions of [16].

2.6.1 Impact of Magnetic Field Symmetry-Breaking on Impurity Particle Flux

2.6.1.1 Flat Density Profile: $\eta^{-1} = 0$

As one increases the magnitude of ϵ_{sb} from 0 up to the true magnetic field, the transport due to the helical wells will also increase. It is not clear a priori exactly how this incremental breaking of symmetry will change the impurity particle flux. However, it is clear that for many cases of interest there should be some critical value, ϵ_{sb}^c, where the radial impurity particle flux, Γ_z, changes sign, which will depend on the particular magnetic equilibrium.

In this section, we examine the ϵ_{sb} dependence of Γ_z for each of the configurations in table 2.1 in select parameter regimes. It should be understood that the magnitude of Γ_z is less important than the sign in this section.

The E_r for each simulation was chosen to be the ambipolar E_r for the $\epsilon_{sb} = 1$ case, considering that E_r becomes progressively less important in calculating radial fluxes as the magnetic field approaches symmetry (this can be seen from Figure 2.6, which is discussed at the end of this section). It also becomes difficult to accurately calculate the radial electric field for small values of ϵ_{sb}. We choose fully-ionized carbon, C^{6+}, as the impurity in Figures 2.1 and 2.3 in order to avoid proximity to

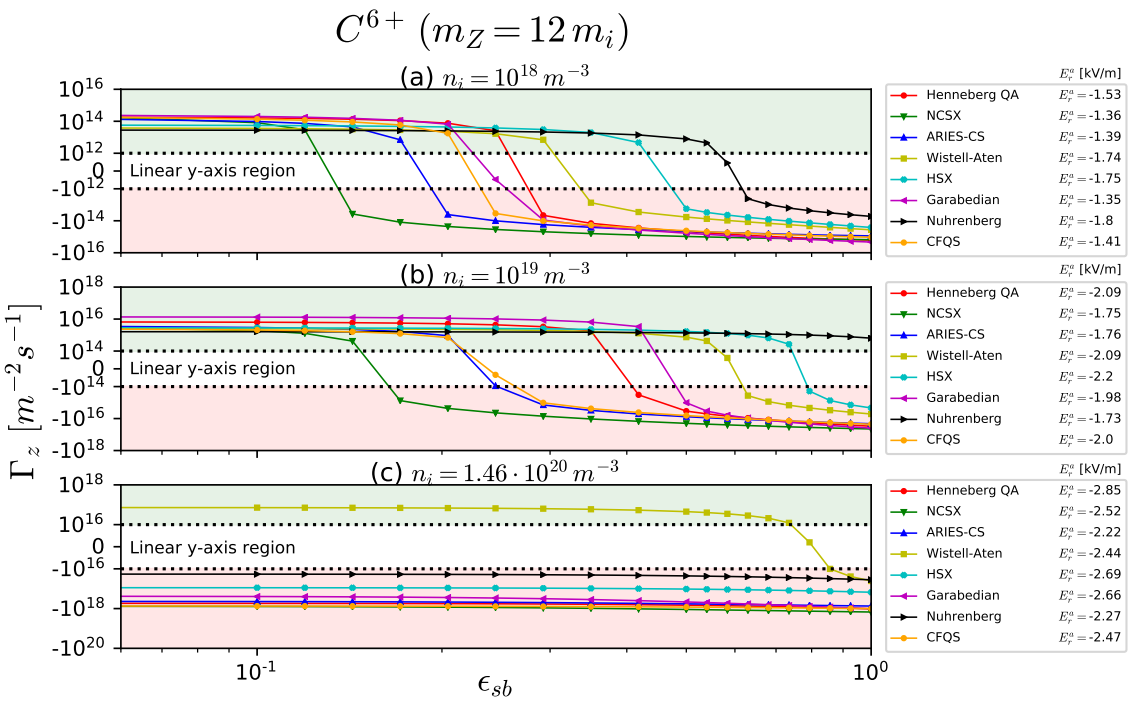

Figure 2.1: (Color online) The impurity particle flux at $\eta^{-1} = 0$ for C^{6+} is plotted as a function of the symmetry-breaking amplitude at a normalized radius of $r_N = 0.25$. $T = 4\,\mathrm{keV}$, $dT/dr = -0.97\,\mathrm{keV/m}$, and $Z_{eff} = 2$ were kept constant for all subplots. The upper, green-shaded region denotes positive Γ_z (impurity screening). The lower, red-shaded region corresponds to negative Γ_z (impurity accumulation). The normalized C^{6+} gyroradius is $\rho_{*z} = 4.17 \cdot 10^{-3}\epsilon$, and the collisionalities for each subplot are (a) $\nu_*^z = 2.26{\cdot}10^{-4}\epsilon^{-1}$, (b) $\nu_*^z = 2.26{\cdot}10^{-3}\epsilon^{-1}$, and (c) $\nu_*^z = 3.29{\cdot}10^{-2}\epsilon^{-1}$.

the resonant electric field in all configurations (see Section 2.4). Finally, we take $\alpha = \sum_{a \neq i} n_a Z_a^2/(n_i Z_i^2) = 1$, corresponding to $Z_{eff} = 2$.

In Figure 2.1(a)-(c), we have plotted results at $r_N \simeq 0.25$ for all devices at increasing values of collisionality, which is achieved by varying the ion density at constant temperature. Here, and in the results that follow, $r_N \equiv r/a = \sqrt{\psi_t/\psi_a}$, where a is the minor radius, and $2\pi\psi_a$ is the toroidal flux at the last closed flux surface (computed in VMEC [53]). At this radial location we take $T_i = T_z = 4\,\mathrm{keV}$ and $dT_i/dr = dT_z/dr = -0.97\,\mathrm{keV/m}$. In the lowest-collisionality case of Figure 2.1(a) (with $n_i = 10^{18}\,\mathrm{m}^{-3}$), there is a similar ϵ_{sb} dependence of Γ_z for each of the

devices, regardless of the type of quasisymmetry (QA or QH). For a magnetic field with near-perfect quasisymmetry ($\epsilon_{sb} = 10^{-2}$), the resulting Γ_z is positive, indicating a presence of the temperature screening effect. As ϵ_{sb} is increased, Γ_z decreases until eventually changing sign at some value of $\epsilon_{sb} < 1$.

The first thing that can be understood from this plot is that at this collisionality, none of the devices that were studied displayed temperature screening at the actual magnetic field, $\epsilon_{sb} = 1$. However, the value of ϵ_{sb}^c where temperature screening is lost will depend on the magnetic configuration. In the case of Nuhrenberg-Zille, for example, the transition occurs at $\epsilon_{sb}^c \simeq 0.6$, which is essentially saying that the symmetry-breaking terms must be $\sim 60\%$ of their actual values to ensure temperature screening under these conditions. Toward the left side of the plot, the NCSX transition occurs at $\epsilon_{sb}^c \simeq 0.1$, requiring the symmetry-breaking terms to be $\sim$10x smaller.

In Figure 2.1(b), the same plot as in Figure 2.1(a) is constructed, however, the density (and hence collisionality) has been increased by an order of magnitude. First, it should be remarked that at this collisionality, the Nuhrenberg-Zille configuration actually achieves temperature screening at $\epsilon_{sb} = 1$. While this is the only such configuration to do so, it is also true that ϵ_{sb}^c has increased for each configuration from the respective values in Figure 2.1(a). Since ϵ_{sb}^c can approximate closeness to quasisymmetry, it follows that increasing the collisionality appears to improve the "effective quasisymmetry" of a flux surface, as it relates to impurity transport.

The meaning behind our use of the term "effective quasisymmetry", which is only applicable for impurity transport and not bulk particles, can nevertheless be

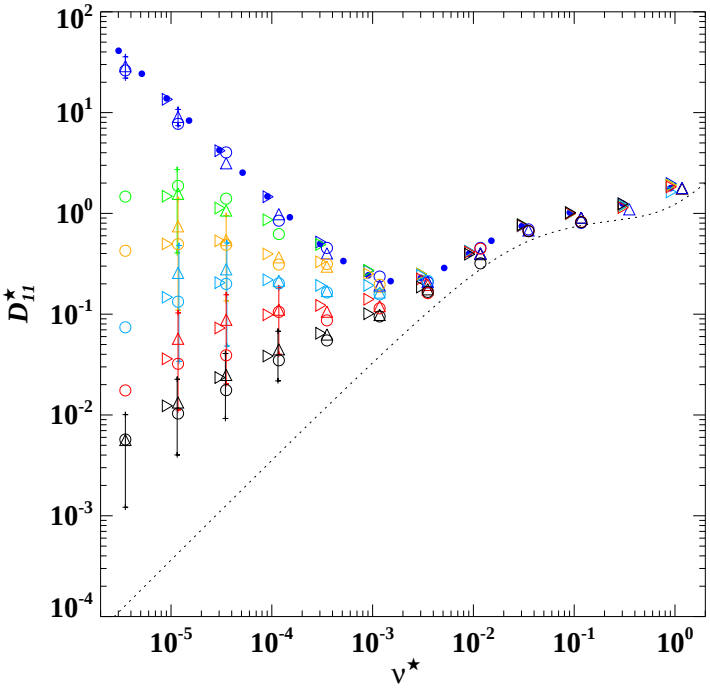

Figure 2.2: (Color online) Definitions of normalizations, markers/colors, and details can be found in Figure (14) of [10]. The D_{11} transport coefficient is plotted as a function of collisionality for HSX geometry. Here, the colors represent different v_E values. DKES [51,99] results are depicted by triangles (△), NEO-2 [59,76] by filled circles (•), and Monte-Carlo results are plotted using open circles (◯) [97], and right-point triangles (▷) [2]. The dotted line is a simulation with equivalent perfect helical symmetry and $E_r = 0$.

understood from a figure in [10] looking at bulk ion transport (which has been reused in Figure 2.2 with permissions). For comparison, the normalized ion collisionality of Figure 2.1(a) corresponds to $\nu^* \sim 10^{-5}$ in Figure 2.2. At low collisionality, there is a difference (depending on E_r) between the D_{11} coefficient (describing radial transport) for HSX and the perfectly quasisymmetric case, indicating a sensitivity of the particles to the exact structure of the magnetic field. As collisionality is increased, this difference becomes less pronounced as the contribution to transport from helically trapped particles decreases. At a high-enough collisionality, D_{11} in

Figure 2.2 is about the same for perfect quasisymmetry as it is for HSX, regardless of E_r. Noting that the magnetic trapping structures can be quite different in perfect symmetry and a nominally quasisymmetric field, the similarities in D_{11} indicate a decreased sensitivity of the particles to the exact structure of the magnetic field at higher collisionalities. Thus, in the context of quasisymmetry, increasing the collisionality brings the transport closer to symmetric levels, "effectively" increasing quasisymmetry.

However, there is a limit to the beneficial impacts of increasing the density, as can be seen in Figure 2.1(c), where $n_i = 1.46 \cdot 10^{20}\,\mathrm{m}^{-3}$. Aside from the Wistell-A configuration, all of the other configurations at near-perfect quasisymmetry do not display an outward impurity flux. There are two possible explanations for why this might take place in perfect symmetry. First, for QA configurations, it is possible that E_r^a and E_r^{res} are close enough that ambipolarity no longer holds, and the higher-order E_r terms [94] become important. To explain this effect in QH, one must recall that temperature screening in axisymmetry is not predicted at high collisionalities [87], except for cases where $\alpha \to 0$ for collisional ions *and* impurities. In Appendix C, it can be seen that beyond some critical ν_*^{ii} in axisymmetry, the impurity flux becomes negative for most η^{-1}. Figure 2.1(c) is thus indicating that we are hovering around that critical collisionality where temperature screening is not possible, even in perfect symmetry. It is interesting to note here that in Figure 2.1(a) and 2.1(b), the impurities are mostly in the $\sqrt{\nu}$ or $1/\nu$ regime. However, in Figure 2.1(c), all QH configurations are well into the plateau regime, and most of the QA configurations have transitioned to the plateau regime as well. The ions are

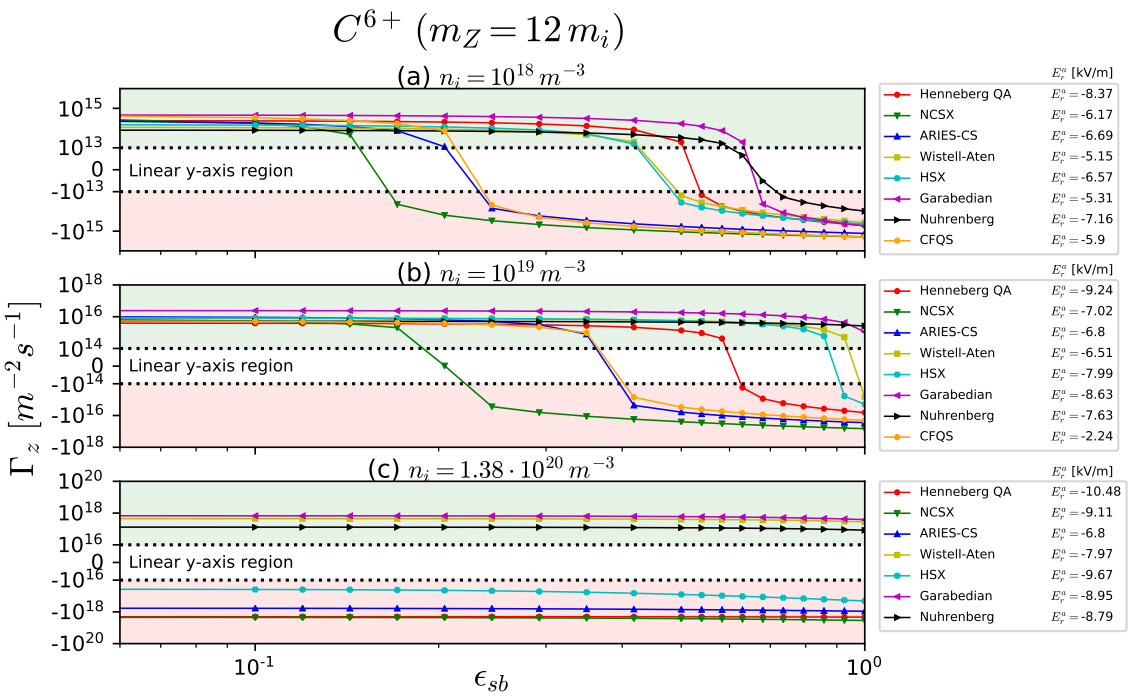

Figure 2.3: (Color online) The impurity particle flux at $\eta^{-1} = 0$ for C^{6+} is plotted as a function of the symmetry-breaking amplitude at a normalized radius of $r_N = 0.50$. $T = 3.3\,\text{keV}$, $dT/dr = -4.78\,\text{keV/m}$, and $Z_{eff} = 2$ were kept constant for all subplots. The upper, green-shaded region denotes positive Γ_z (impurity screening). The lower, red-shaded region corresponds to negative Γ_z (impurity accumulation). The normalized C^{6+} gyroradius is $\rho_{*z} = 3.79 \cdot 10^{-3}\epsilon$, and the collisionalities for each subplot are (a) $\nu_*^z = 3.32{\cdot}10^{-4}\epsilon^{-1}$, (b) $\nu_*^z = 3.32{\cdot}10^{-3}\epsilon^{-1}$, and (c) $\nu_*^z = 4.58{\cdot}10^{-2}\epsilon^{-1}$.

in the $\sqrt{\nu}$ regime in all but a couple cases at the highest collisionality.

The situation is less pessimistic in Figure 2.3, where we look at $r_N = 0.50$ with $T = 3.3\,\text{keV}$ and $dT/dr = -4.78\,\text{keV/m}$. The trends are largely similar to Figure 2.1 at each collisionality, however, there are a handful of cases where temperature screening can be seen at $\epsilon_{sb} = 1$. Furthermore, at the highest collisionality, the Wistell-A, Garabedian, and Nuhrenberg configurations have an outward impurity particle flux for all ϵ_{sb}. The collisionality at $r_N = 0.50$ is only slightly higher than at $r_N = 0.25$, and $\eta^{-1} = 0$ at both r_N, so the differences at these radii are likely caused by the distinct magnetic field modes, B_{mn}.

The differences between Figures 2.1 and 2.3 can be understood by looking at the magnitude of symmetry-breaking terms

$$S \equiv \sqrt{\sum_{m,n \neq mN} B_{mn}^2 / B_{00}^2}, \tag{2.17}$$

as a function of r_N in Figure 2.4. The summation in Eq 2.17 includes all modes that are not an integer multiple of the dominant magnetic helicity χ (i.e. $M = 1$, $N = 4$ for HSX). If we consider the curves for Henneberg QA and Garabedian, we can compare the difference in the values of ϵ_{sb}^c in Figure 2.1(a) and Figure 2.3(a). In moving from $r_N = 0.25$ to $r_N = 0.50$ in Figure 2.4, the symmetry-breaking amplitude for Henneberg QA and Garabedian decreases by ~ 4. The corresponding increase in ϵ_{sb}^c from $r_N = 0.25$ to $r_N = 0.50$ is $\sim 2 - 3$ for both configurations. If we were to then consider the respective CFQS curves (Figure 2.1(a) and Figure 2.3(a)), there is an increase in S between these radii of ~ 2, where a *decrease* in ϵ_{sb}^c is observed. This presents a connection between the closeness to quasisymmetry of a flux surface, and the realization of temperature screening. The remaining configurations have a difference in S of less than a factor of two at these radii, and S is also larger at $r_N = 0.50$. Unlike the connection between S and the change in ϵ_{sb}^c for Henneberg QA, Garabedian, and CFQS, the change in ϵ_{sb} is *positive* (although small) for the remaining configurations. This could potentially be accounted for by a complicated dependency on collisionality, E_r, and the aspect ratio.

The arguments in this section are also relevant when considering higher-Z impurities. In Figure 2.5, a similar plot to Figures 2.1 and 2.3 has been constructed

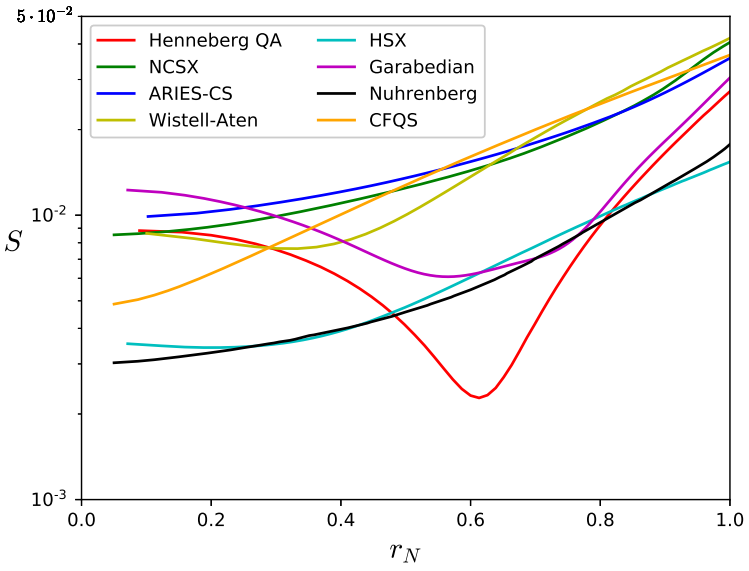

Figure 2.4: (Color online) The amplitude of the symmetry-breaking terms, S, are plotted as a function of normalized radius, r_N.

for Cr^{24+}, where results from both $r_N = 0.25$ and $r_N = 0.50$ have been consolidated into Figure 2.5. To restate from Section 2.4, only configurations where $E_r^a/E_r^{res} < 1/3$ have been considered. One difference between Cr^{24+} and C^{6+} is that the value of ϵ_{sb}^c *decreases* for the QA configurations between Figure 2.5(a) and 2.5(b). This behavior is likely due to the increase in collisionality from Figures 2.1 and 2.3, (where the impurities are now in the plateau regime in Figure 2.5(b) and collisional in 2.5(c)), meaning that the critical density where temperature screening is lost has decreased. Further, the behavior of Wistell-A in Figure 2.5 is interesting compared with other configurations. From Figure 2.4, the values of S at $r_N = 0.25$ and $r_N = 0.50$ are quite similar, however, there is clear *increase* in ϵ_{sb}^c that may not be related to the amplitude of symmetry-breaking terms.

As mentioned at the beginning of this section, E_r^a at $\epsilon_{sb} = 1$ has been used as the E_r value for all ϵ_{sb} simulations. In Figure 2.6, we use the Wistell-A curve

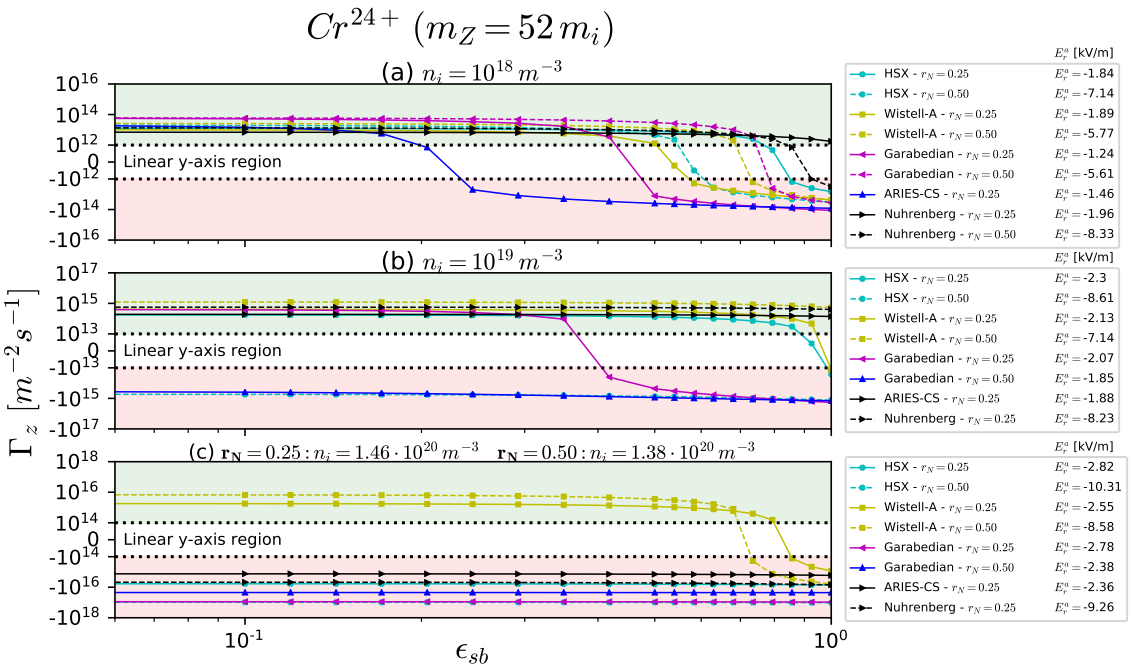

Figure 2.5: (Color online) The impurity particle flux at $\eta^{-1} = 0$ for Cr^{24+} $(m_z = 52\, m_i)$ is plotted as a function of the symmetry-breaking amplitude. For $\boldsymbol{r_N = 0.25}$: $T = 4.0\,\mathrm{keV}$, $dT/dr = -0.97\,\mathrm{keV/m}$, and $\rho_{*z} = 2.17 \cdot 10^{-3}\epsilon$ with collisionalities (a) $\nu_*^z = 3.61 \cdot 10^{-3}\epsilon^{-1}$, (b) $\nu_*^z = 3.61 \cdot 10^{-2}\epsilon^{-1}$, and (c) $\nu_*^z = 0.53\,\epsilon^{-1}$. At $\boldsymbol{r_N = 0.50}$: $T = 3.3\,\mathrm{keV}$, $dT/dr = -4.78\,\mathrm{keV/m}$, and $\rho_{*z} = 1.97 \cdot 10^{-3}\epsilon$ with collisionalities (a) $\nu_*^z = 5.31 \cdot 10^{-3}\epsilon^{-1}$, (b) $\nu_*^z = 5.31 \cdot 10^{-2}\epsilon^{-1}$, and (c) $\nu_*^z = 0.73\,\epsilon^{-1}$. $Z_{eff} = 2$ was kept constant for all subplots. The upper, green-shaded region denotes positive Γ_z (impurity screening). The lower, red-shaded region corresponds to negative Γ_z (impurity accumulation).

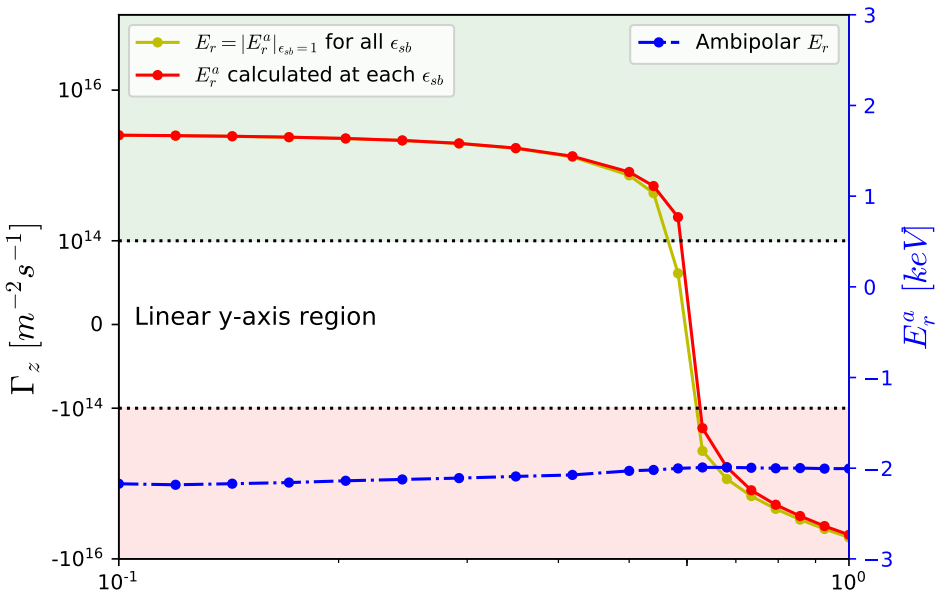

Figure 2.6: (Color online) (left y-axis) The impurity particle flux is plotted as a function of the symmetry-breaking amplitude in Wistell-A using the parameters of Figure 2.1(b). The yellow curve is identical to the yellow curve of Figure 2.1(b). The red curve uses the calculated E_r^a value for each ϵ_{sb}. (right y-axis) The blue dashed line is the E_r^a for each point in the red curve.

from Figure 2.1(b) and compare it to one generated with E_r^a calculated at each ϵ_{sb}, as evidence for our argument in making this approximation. From the dashed blue curve indicating the calculated E_r^a value, there is little variation as ϵ_{sb} decreases. More importantly, this variation also has a minor impact on Γ_z, as evinced by the red curve in Figure 2.6, which utilizes the E_r values from the dashed blue curve at each ϵ_{sb}.

The overall results in Figures 2.1, 2.3, and 2.5 indicate that for reactor-relevant plasma parameters, temperature screening could be achievable in certain configurations. However, as it is unlikely that the density profile will be completely flat, it is imperative to understand how Γ_z varies with η^{-1}.

2.6.1.2 Finite Peaked Density Gradients: $\eta^{-1} > 0$

Peaking of the main ion's density profile drives an inward neoclassical impurity flux. This result is shown for axisymmetry or quasisymmetry in Appendix C. In non-symmetric stellarators, the situation is complicated by not only the presence of the radial electric field as a driving gradient in the impurity flux, but the fact that L_{11} depends on E_r, and in different ways depending on the collisionality regime. An exact analytic solution for Γ_z is therefore intractable in most cases. However, it is possible to approximate the solution by using a similar procedure to that used in [100], but generalizing for arbitrary Z_{eff}. We start with an expression for the particle flux of an arbitrary species that is valid far from quasisymmetry at low collisionality

$$\Gamma_a = -n_a L_{11}^a \left[\frac{1}{n_a}\frac{dn_a}{dr} - \frac{Z_a e E_r}{T_a} + \delta_a \frac{1}{T_a}\frac{dT_a}{dr} \right]. \tag{2.18}$$

It is possible to then explicitly solve the ambipolarity condition $\sum_a Z_a \Gamma_a = 0$ for E_r. If we take $T_i = T_z = T_e$, it is possible to drop the contribution of Γ_e to the ambipolar condition since $L_{11}^e \ll L_{11}^i$ (though this approximation does not hold at very low collisionalities, as described in [100]). Finally, we take the temperature (T_a'/T_a) gradients to be equivalent for the bulk ions and impurities. This allows one to solve for the radial electric field

$$\frac{eE_r}{T} = \frac{L_{11}^i \frac{n_i'}{n_i} + \frac{\alpha}{Z} L_{11}^z \frac{n_z'}{n_z} + A\frac{T'}{T}}{L_{11}^i + \alpha L_{11}^z}. \tag{2.19}$$

By plugging this back into the expression for Γ_z

$$\Gamma_z = n_z L_{11}^z \left[\frac{L_{11}^i (Z \frac{n_i'}{n_i} - \frac{n_z'}{n_z})}{L_{11}^i + \alpha L_{11}^z} + C \frac{T'}{T} \right], \tag{2.20}$$

where A and C are scalars that are not relevant to the following discussion, and $L_{11}^a > 0$ in all cases so that the diffusive flux is always opposite dn_a/dr. In the approximation that L_{11}^a, A, and C (which are implicitly dependent on the pressure gradient through E_r) do not vary strongly with n'/n, and assuming Z_{eff} = constant ($n_i'/n_i = n_z'/n_z$), Eq 2.20 indicates that for $n'/n < 0$, the density gradient will have an unfavorable effect on impurity accumulation, and one that worsens as Z increases. If one were to instead assume that the ions alone determine E_r, yielding a peaked Z_{eff} profile, this would lead to $n_z' = (\alpha/Z) n_i'$ if terms proportional to T' are neglected. Therefore, for $\alpha = 1$, we have $|n_z'| \ll |n_i'|$ and a slightly stronger, yet similarly adverse effect on impurity accumulation when $n'/n < 0$.

This is evident in Figure 2.7 in the context of how Γ_z is affected by ϵ_{sb}. Each curve in the figure was calculated with the Wistell-A configuration at various η^{-1}. The red curve is the $\eta^{-1} = 0$ case (identical to the Wistell curve in Figure 2.1(c)), where the degree of quasisymmetry is nearly good enough to retain temperature screening at such parameters. If a small density gradient $\eta^{-1} = 0.03$ is introduced, ϵ_{sb}^c decreases by nearly a factor of 2. Any further increase in η^{-1} pushes the plasma to the point where even perfect quasisymmetry cannot support temperature screening. This makes the situation of temperature screening even more pessimistic, because even if $\Gamma_z > 0$ in a particular collisionality regime, simply introducing a density

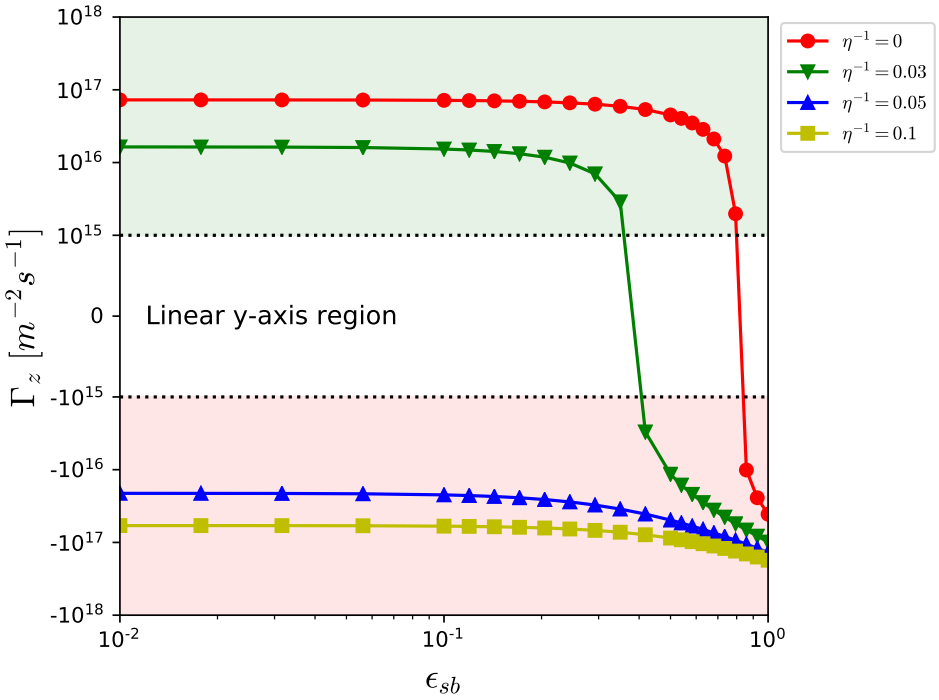

Figure 2.7: (Color online) The impurity particle flux for C^{6+} in Wistell-A is plotted as a function of ϵ_{sb}, the scaling factor for the symmetry-breaking terms, at $r_N = 0.25$. Each curve represents a different relative density gradient, but the physical parameters are otherwise identical to those of Figure 2.1(c). The upper, green-shaded region denotes positive Γ_z (impurity screening). The lower, red-shaded region corresponds to negative Γ_z (impurity accumulation).

gradient can flip the sign.

In all likelihood, there will be some finite density gradient in a reactor-relevant plasma, likely corresponding to an inward flux of impurities. It is then of interest to see how the magnitude of Γ_z changes, relative to its value at $\eta^{-1} = 0$, as the strength of the density gradient is increased. In Figure 2.8, the ratio $\Gamma_z/|\Gamma_z|_{\eta^{-1}=0}$ is plotted as a function of η^{-1} for various configurations, where each simulation was calculated at the true magnetic field, and used its own E_r^a. In every case shown in Figure 2.8(a), Γ_z is negative and a decreasing function of η^{-1}, indicating that increasing the strength of the peaked density gradient will intensify impurity accumulation.

In a scenario where the length scale of the density gradient is only twice that of the temperature gradient ($\eta^{-1} = 0.5$), the enhancement in Γ_z at this radius can be increased by a factor of ~ 20. The picture appears to at least slightly worsen at $r_N = 0.50$ in Figure 2.8(b), where the enhanced accumulation has close to doubled from the values in Figure 2.8(a) in most cases.

2.6.2 Comparison to Gyro-Bohm Turbulence Estimate

In this section, we use results from the parameter scans in the previous section to compare the neoclassical particle flux, Γ_z, and heat flux, $Q_{total} = Q_i + Q_z$, to a gyro-Bohm estimate for turbulent transport, $\Gamma_z^{gb} \sim n_z D_{gb} |\nabla T|/T$, and $Q_{total}^{gb} \sim D_{gb}|\nabla T|(n_i + n_z)$. In these expressions, $D_{gb} = \rho_*^2 v_{ti} a$ is the gyro-Bohm diffusion coefficient, where we have taken the minor radius to be the relevant length scale. The gyro-Bohm estimate is not a substitute for turbulent fluxes obtained from solving the gyrokinetic equation, but rather an order of magnitude estimate of the turbulence.

2.6.2.1 Impurity Particle Flux with $\Phi = \Phi_0(r)$

Figure 2.9 examines how the neoclassical particle fluxes compare to Γ_z^{gb} as a function of the impurity ion charge for each device, using the equilibrium magnetic field, $\epsilon_{sb} = 1$. At $Z = 6$, the impurities are in the plateau regime, and by $Z = 24$ all impurities have become collisional $\nu_*^z \geq 1$ (depending on the aspect ratio, the $Z = 13$ impurities are also collisional). Only flat density profiles ($\eta^{-1} = 0$) are considered in Figure 2.9. In Figure 2.9(a), we look at the impurity particle flux for $r_N = 0.25$,

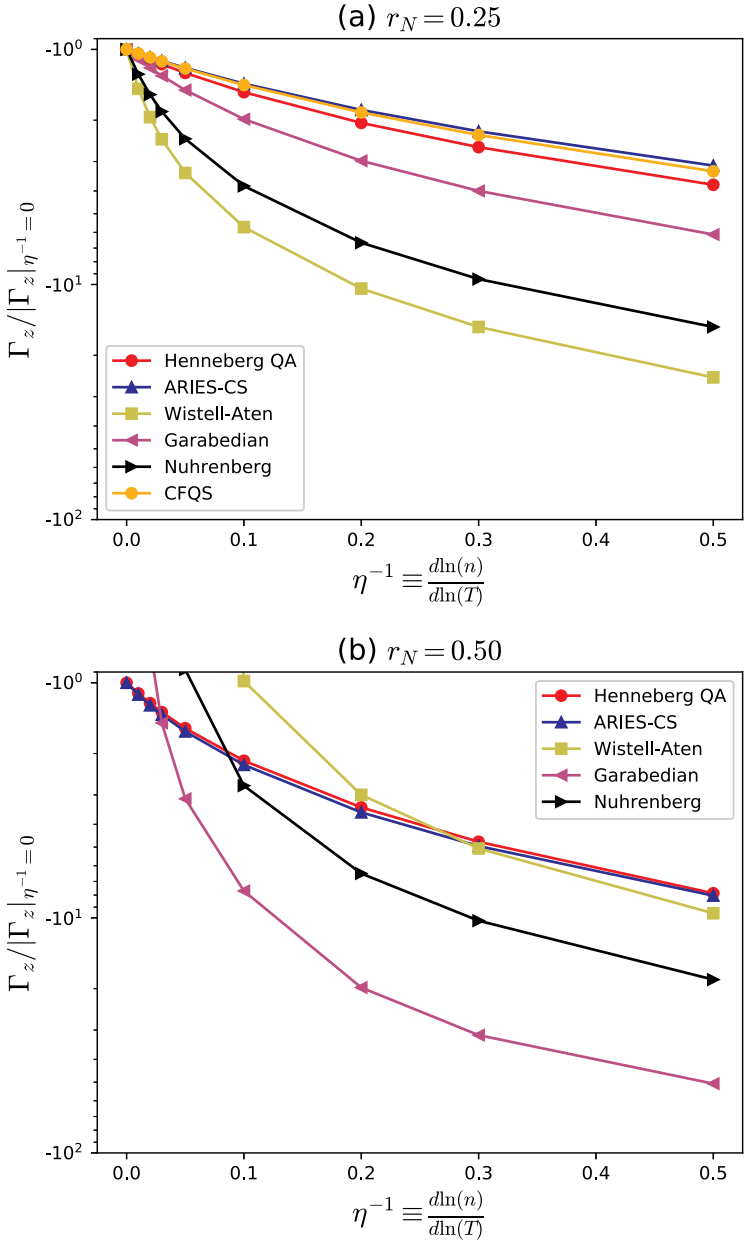

Figure 2.8: (Color online) The impurity particle flux for C^{6+} has been normalized to its magnitude at $\eta^{-1} = 0$ and plotted as a function of η^{-1} at (a) $r_N = 0.25$, and (b) $r_N = 0.50$. Every simulation was performed at the true magnetic field $\epsilon_{sb} = 1$, with the E_r^a independently calculated at every η^{-1}. The physical parameters are otherwise identical to those of Figures 2.1(c) and 2.3(c). The data points above -10^0 in (b) are those corresponding to devices with positive Γ_z at $\epsilon_{sb} = 1$, thus giving a value of $+10^0$ at $\eta^{-1} = 0$.

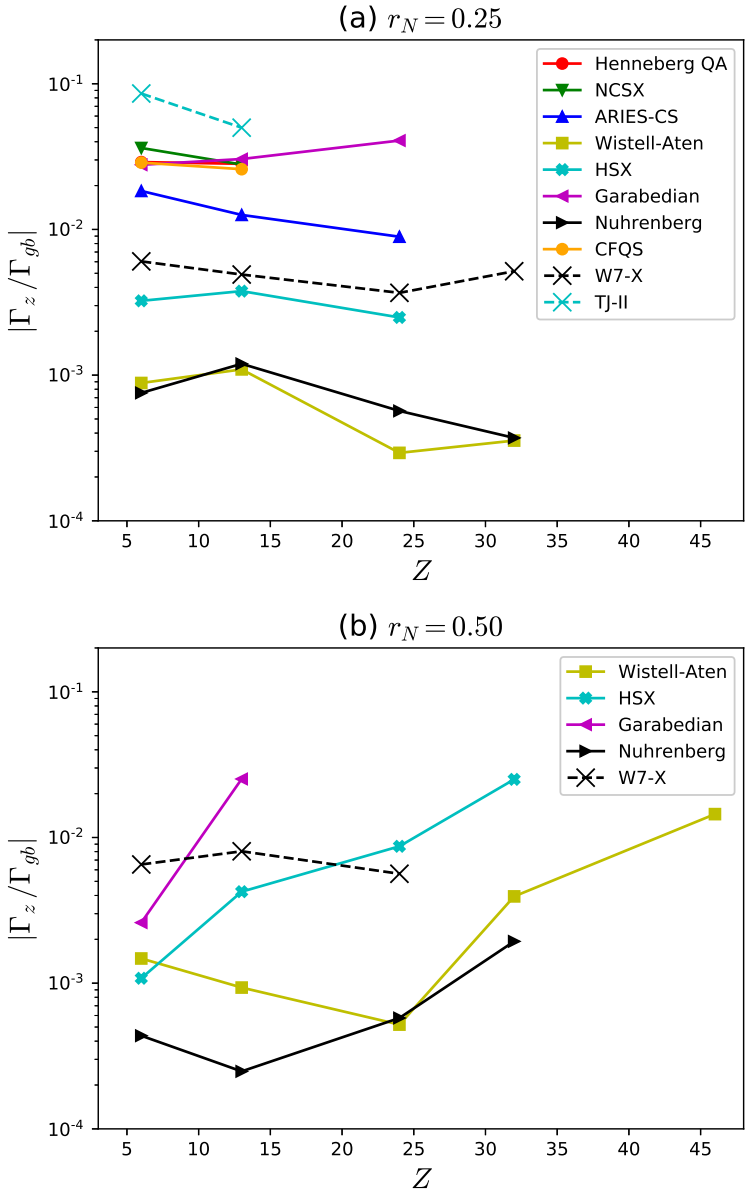

Figure 2.9: (Color online) The neoclassical impurity particle flux at $\eta^{-1} = 0$ has been normalized to a gyro-Bohm estimate of the turbulent impurity particle flux (see text). This ratio is plotted as a function of the impurity charge (and mass) for (a) $r_N = 0.25$, and (b) $r_N = 0.50$. Plasma parameters correspond to those of Figures 2.1(c) and 2.3(c). Collisionalities can be found from (a) $\nu_*^z = 9.14 \cdot 10^{-4} Z^2 \epsilon^{-1}$, and (b) $\nu_*^z = 1.27 \cdot 10^{-3} Z^2 \epsilon^{-1}$.

where the temperature gradients are weaker. At this radial location, the ratio of fluxes does not have a consistent strong trend with Z. This is observed in both QA and QH configurations, as well as the *non*-quasisymmetric TJ-II stellarator [47], and the W7-X standard configuration. Note that symmetry-breaking harmonics of B are not modified in Figure 2.9. The most striking feature of Figure 2.9(a), however, is the dominance of turbulent transport. Of the quasisymmetric configurations that were studied, the *largest* calculated neoclassical flux at $r_N = 0.25$ is only $\sim 5\%$ of the turbulent value. These small ratios indicate that regardless of whether temperature screening is present at a given collisionality, it is possible that the turbulence could control the sign of the particle flux.

At $r_N = 0.50$ in Figure 2.9(b), the overall sensitivity of this ratio to the impurity species in QA is unclear since the larger gradients push E_r^a close enough to E_r^{res} that results are unreliable (see Section 2.4). Only configurations with at least two points have been shown in Figure 2.9, eliminating all but one QA configuration. Apart from W7-X, there is an eventual point for each configuration at which further increase in Z corresponds to an increase in the relative importance of neoclassical fluxes. Even with this increase in the ratio, the neoclassical contribution to the radial particle flux is $< 3\%$ of the turbulent value.

It should be reiterated here that these results have been generated with a flat density profile. While it is still unknown exactly how the density profiles will behave in a reactor, it is likely that $|\eta^{-1}| > 0$. From Figure 2.8, it can then be inferred how this neoclassical to turbulence ratio will change if a peaked density gradient is introduced. In a non-ideal scenario, where $\eta^{-1} = 0.5$, the ratio could increase

by more than a factor of 10, depending on the configuration. At $r_N = 0.25$, this would still only lead to the neoclassical flux being $\sim 10\%$ of the turbulence for most configurations.

It is also of practical importance to understand how this ratio of neoclassical to turbulent particle flux varies with distance from the magnetic axis. This radial profile is shown in Figure 2.10, where the radial points $r_N = 0.15$ and $r_N = 0.40$ (profiles can be found in the caption of Figure 2.10) have been added to the previously calculated values at $r_N = 0.25$ and $r_N = 0.50$. For most but not all, the ratio tends to either decrease or remain constant as one moves out radially, indicating that turbulence becomes increasingly more important. This follows experimental observations [19,78] that show neoclassical fluxes at negligible levels when compared with turbulence far from the magnetic axis.

While these results point to reactor-relevant plasmas where turbulence is likely the dominant impurity particle transport channel, more work is needed to fully understand the significance of these findings. The most obvious step would be a more accurate value for the turbulent fluxes such as a quasilinear model or gyrokinetic simulations, so as to better quantify this neoclassical to turbulence particle flux ratio. Also, a recent study comparing neoclassical simulations and experimental fluxes from an laser blow-off injection of iron in W7-X [39] similarly found that $|\Gamma_z^{NC}/\Gamma_z^{anom}| \ll 1$. However, by separately considering the diffusive and convective contributions to the particle flux, it was found that neoclassical fluxes could still be responsible for determining the sign of the total particle transport, while the turbulence (anomalous transport) could control its magnitude.

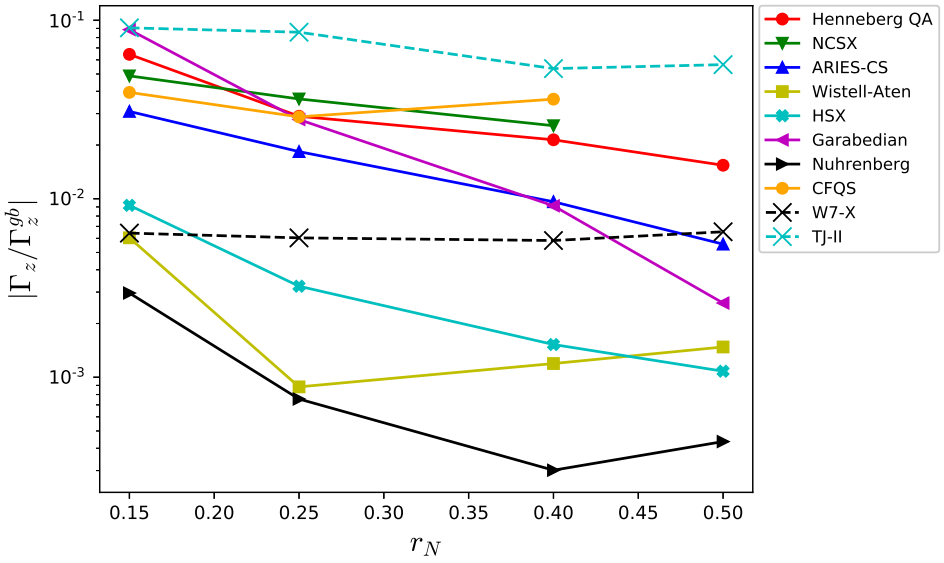

Figure 2.10: (Color online) The neoclassical impurity particle flux at $\eta^{-1} = 0$ for C^{6+} has been normalized to a gyro-Bohm estimate of the turbulent impurity particle flux (see text). This ratio is plotted as a function of the normalized radius r_N. Plasma profiles at $r_N = 0.25$ and $r_N = 0.50$ correspond to those of Figures 2.1(c) and 2.3(c), respectively. At $r_N = 0.15$: $T = 4.1\,\mathrm{keV}$, $dT/dr = -0.58\,\mathrm{keV/m}$, $n_i = 1.51 \cdot 10^{20}\,\mathrm{m}^{-3}$. At $r_N = 0.40$: $T = 3.75\,\mathrm{keV}$, $dT/dr = -3.88\,\mathrm{keV/m}$, $n_i = 1.43 \cdot 10^{20}\,\mathrm{m}^{-3}$.

2.6.2.2 Impurity Particle Flux including Φ_1 Effects

The form of the DKE in Eq. 2.1 was found by linearizing about Φ_0, assuming that it is close to a flux function, where $\Phi_1 \ll \Phi_0$. When including Φ_1 effects, one can no longer neglect the contributions of Φ_1 in the zeroth-order distribution function $f_{a0} \approx F_a[1 - Z_a e\Phi_1/T_a]$, and the energy $W_a = W_{a0} + Z_a e\Phi_1/m_a$, as was done in Eq 2.2.1. Furthermore, the radial component of the $\boldsymbol{E} \times \boldsymbol{B}$ vanishes when $\Phi = \Phi_0$, but enters the DKE for non-zero Φ_1, which would change the final term of Eq 2.2.1 to

$$- (\boldsymbol{v}_{ma} + \boldsymbol{v}_E) \cdot \nabla r \left(\frac{\partial f_{a0}}{\partial r} \right)_{W_a}, \tag{2.21}$$

where $\boldsymbol{v}_E = (c/B^2)\boldsymbol{B} \times \nabla\Phi_1$. The above replacements in the DKE will have the effect of both altering the phase space trajectories, and making the DKE nonlinear. For details on the implemented equations with Φ_1 effects see [70].

When considering Φ_1 in neoclassical transport, recent results [34] indicate that it has only a moderate impact on the particle flux for highly-charged impurities (W^{40+}), in the case of the *non*-quasisymmetric Wendelstein 7-X (W7-X) stellarator. While not quasisymmetric, W7-X is still a neoclassically optimized stellarator and will have reduced radial excursions of helically-trapped particles, limiting the size of density variations on a flux surface. Since Φ_1 is closely connected to these density fluctuations [73], it would make sense to assume (and indeed [34] has shown) that Φ_1 fluctuations are small in such configurations. In quasisymmetric experiments not deviating too far from perfect symmetry, it is reasonable to expect a similarly

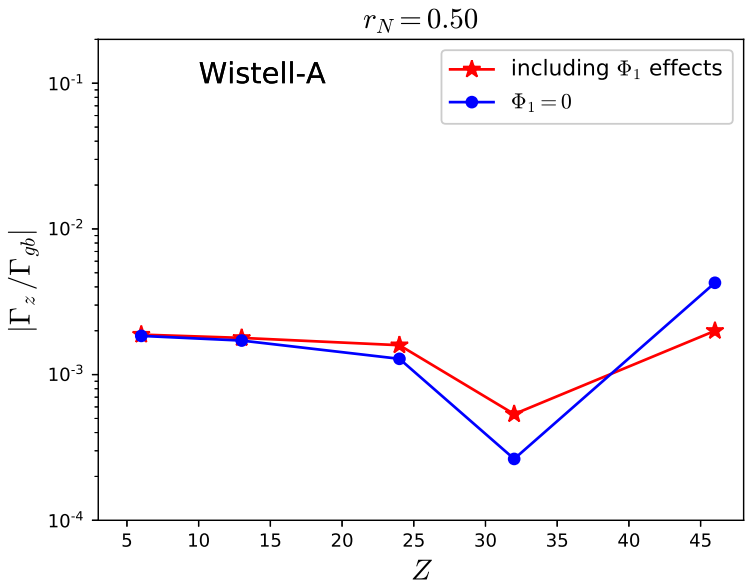

Figure 2.11: (Color online) The neoclassical impurity particle flux in Wistell-A at $\eta^{-1} = 0$ has been normalized to a gyro-Bohm estimate of the turbulent impurity particle flux (see text). A kinetic electron species has been included for all points, and plasma parameters correspond to those of Figure 2.3(c). The red curve includes Φ_1 effects in the DKE, and Φ_1 is neglected for the blue points.

small Φ_1. However, the impact of Φ_1 on the neoclassical particle flux in stellarator configurations optimized for quasisymmetry has yet to be shown. We present a first look at this behavior using the Wistell-A configuration.

In the blue curve of Figure 2.11, we recreate the Wistell-A curve from Figure 2.9(b), but now include a kinetic electron species (maintaining $\alpha = 1$ and $\Phi_1 = 0$). In plasmas where quasineutrality is satisfied, this permits one to solve the DKE with Φ_1 effects, which is shown in the red curve of Figure 2.11 (where E_r^a has been calculated at each point through the inclusion of Φ_1 effects). It is evident from Figure 2.11, that Φ_1 has a minimal effect on impurities with low charge, especially so for C^{6+} and Al^{13+}. However, with increasing Z, the difference in Γ_z with and without Φ_1 effects becomes non-negligible, differing by about a factor of 2. Also

of interest is the magnitude of Φ_1 fluctuations, which in our results range from $e|\Phi_1^{max}|/T_i = 8.9 \cdot 10^{-4}$ to $e|\Phi_1^{max}|/T_i = 1.4 \cdot 10^{-3}$, which are generally smaller than the analogous W7-X values of Figure 18 in [34].

Considering how our results differ from those for W7-X in [34], there are some differences that must be appreciated. In [34], the authors take $\alpha = 0.1$ ($Z_{eff} = 1.1$), and use this to solve a quasineutrality equation that does not consider the effect of the impurities on Φ_1. While this is reasonable for a low Z_{eff} plasma, when $\alpha = 1$ the impurity contribution will be commensurate with bulk ions in the quasineutrality equation and their effect on Φ_1 must be considered. Second, because Wistell-A is quasisymmetric, $|\Gamma_z|$ is small in the sense that it is closer to the transition between positive and negative Γ_z than the non-quasisymmetric W7-X. This could result in similarly-sized $e|\Phi_1^{max}|/T_i$ values having a comparatively stronger effect on Γ_z in Wistell-A than in W7-X. Finally, results presented in [34] employ a small but finite density gradient, where we have taken $\eta^{-1} = 0$. As we have outlined in detail in Section 6.1.2, introducing a peaked density gradient tends to have a strong influence on the impurity particle flux. Therefore, the result of introducing a density gradient alongside Φ_1 effects can be expected to modify the curves of Figure 2.11.

It is not our aim in this section to exactly quantify the differences between our results and [34]. This has been meant to both introduce new results on the effect of Φ_1 in a quasisymmetric geometry, and attempt to identify key differences between similar Φ_1 studies. Therefore, a more comprehensive study will be left to future work.

2.6.2.3 Total Heat Flux

Along with particle fluxes, it is also of great practical importance to compare the neoclassical and turbulent *heat fluxes* at different locations within the plasma. If the dominant transport channel can be identified, this can better inform future efforts in optimizing for a certain type of transport over a particular radial domain. In this section we do not distinguish between ion and impurity heat fluxes, since we primarily care about the *total* heat flux (ion+impurity) that is crossing a flux surface.

Thus, shifting our attention to the ratio of neoclassical to turbulent heat fluxes, the results in Figure 2.12 show the radial profiles of this ratio for each configuration. The overall trend is similar to Figure 2.10, except that the magnitude of this ratio is a bit higher than the respective points in Figure 2.10, especially closer to the magnetic axis.

However, it is important to mention here that unlike the impurity particle flux, we have found this ratio to be independent of η^{-1}, and the particular impurity species. So while the ratios in Figure 2.10 may appear smaller in comparison, a heavy impurity in the presence of a density gradient could change that. This is to say that these heat flux ratios are more robust over a wider range of potential reactor-relevant parameters than the impurity particle flux.

The general trend of the decreasing relative importance of neoclassical heat flux compared with turbulence with respect to radius is in agreement with experimental results [19, 78]. With that said, for $r_N \geq 0.25$, the neoclassical heat flux is, at best,

30% of the turbulent value, and in many cases this ratio is even smaller.

These magnitudes appear to be at odds with Figure 7 in [78], where neoclassical simulations (SFINCS) of an ECRH-heated W7-X experiment show that the neoclassical electron heat flux constitutes $\sim 65\%$ of the input power through the flux surface at $r_N = 0.25$. If the remaining flux is presumed to be turbulence-driven, then the neoclassical electron heat flux should be about twice the turbulent value. By comparing this neoclassical result to a gyro-Bohm estimate using $\rho_s = 3.21 \cdot 10^{-3}\,\mathrm{m}$, and $L_{T_e} \equiv (1/T_e|dT_e/dr|)^{-1} = 0.66\,\mathrm{m}$ in the expression $Q_e^{gb} \sim n_e \rho_{*s}^2 c_s a T_e L_{T_e}^{-1}$ one finds $|Q_e/Q_e^{gb}| \sim 0.05$, where Q_e is the computed neoclassical electron heat flux. The above expression uses the ion sound speed $c_s = \sqrt{T_e/m_i}$ and gyroradius $\rho_s = c_s/\Omega_i$, where $\rho_{*s} \equiv \rho_s/a$. A similar comparison can be done for HSX with Figure 13 in [19], where the neoclassical electron thermal diffusivity appears to account for $\sim 10\%$ of the experimentally measured diffusivity at $r_N = 0.25$. Using the above approximation for Q_e^{gb}, with length scales $\rho_s = 3.42 \cdot 10^{-3}\,\mathrm{m}$ and $L_{T_e} = 0.042\,\mathrm{m}$, results in $|Q_e/Q_e^{gb}| \sim 0.02$.

These inconsistencies in how well gyro-Bohm approximates the turbulence underlines the nature of gyro-Bohm as only an *estimate* of turbulence.

Setting the coefficient of D_{gb} to 1 for every configuration and set of plasma parameters is bound to yield results that can differ by an appreciable amount relative to the actual turbulent fluxes.

It should be mentioned that for the above cases, $T_e \gg T_i$, indicating that electrons will likely be important for both neoclassical and turbulent energy transport. This is in contrast to the majority of results in our work, where electrons

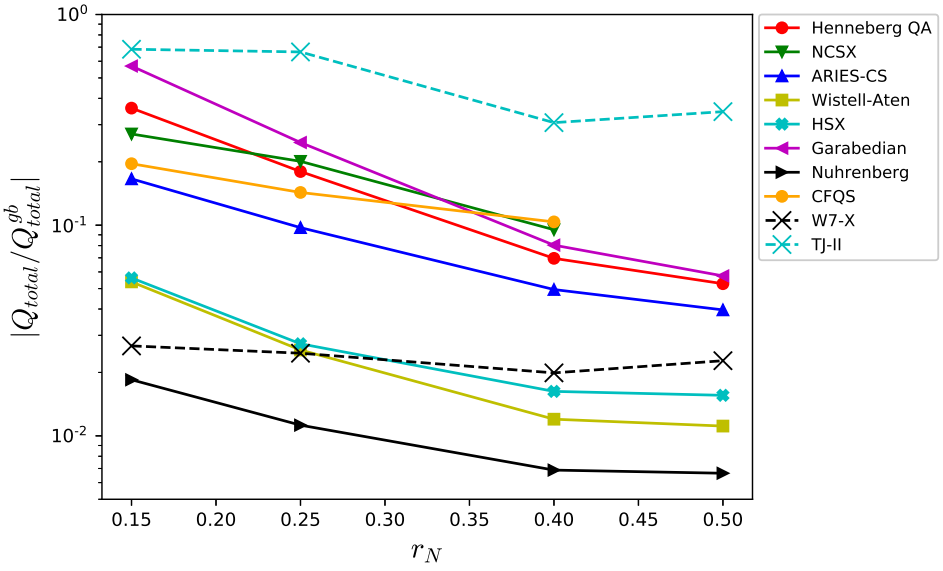

Figure 2.12: (Color online) The total (ion+impurity) heat flux at $\eta^{-1} = 0$ for C^{6+} has been normalized to a gyro-Bohm estimate of the total turbulent heat flux (see text). This ratio is plotted as a function of the normalized radius r_N. Plasma profiles are the same as for Figure 2.10.

were excluded from simulations. For plasmas with large T_e/T_i, the ion temperature gradient is no longer the relevant driving gradient, which is why c_s and ρ_s have been used in Q_e^{gb} above, in place of v_{ti} and ρ_i. While the HSX result in particular is interesting in the sense that it is the only experimental quasisymmetric stellarator data comparing transport channels, it is unclear of how relevant it is to the rest of the results in this chapter, considering that $T_e \gg T_i$.

2.7 Effective Helical Ripple as a Quasisymmetry Metric

From Section 2.6.1, we showed how there was a connection between S and ϵ_{sb}^c that helped to explain how ϵ_{sb}^c changed between the two flux surfaces that were studied. This connection can be seen more clearly in Figures 2.13(a)-(b), where the value of ϵ_{sb}^c from Figures 2.1(a) and 2.3(a) has been plotted as a function of S

on the respective surfaces for each of the configurations. For both $r_N = 0.25$ and $r_N = 0.50$, there is a visible anti-correlation between the two quantities even when considering that these configurations have very different properties. It thus seems reasonable to expect that minimizing S on a flux surface will increase ϵ_{sb}^{c}.

Along with S, the effective helical ripple, ϵ_{eff}, is sometimes taken to be a metric for quasisymmetry that could be used for stellarator optimization. ϵ_{eff}, which is a measure of neoclassical transport in the $1/\nu$ regime, was computed with the NEO code [76].

In Figure 2.14, the effective helical ripple is plotted as a function of S for each configuration. A number of these curves are multi-valued, indicating a non-monotonic change in quasisymmetry from the magnetic axis to the last closed flux surface (LCFS). To clarify the radial dependency of each curve, the open circle at the end of a curve denotes the magnetic axis, $r_N = 0$, and a closed circle the LCFS, $r_N = 1$. It can be seen [48,74] that if individual symmetry-breaking B_{mn} harmonics are plotted as a function of r_N, that the amplitude tends to increase with distance from the magnetic axis. Indeed, this trend can be seen for a handful of configurations in Figure 2.14, indicating a correlation between ϵ_{eff} and the closeness to quasisymmetry. However, this is decidedly not universal among QA configurations. Henneberg QA, for example, has a symmetry-breaking amplitude that decreases by nearly an order of magnitude from $r_N = 0 \rightarrow r_N \simeq 0.6$, and then increases again to a value at $r_N = 1$ that is larger than its value at the $r_N = 0$.

Moreover, this monotonicity in S, or lack thereof, is not necessarily tied to the value of ϵ_{eff}. Returning to Henneberg QA as an example, the initial decrease in S

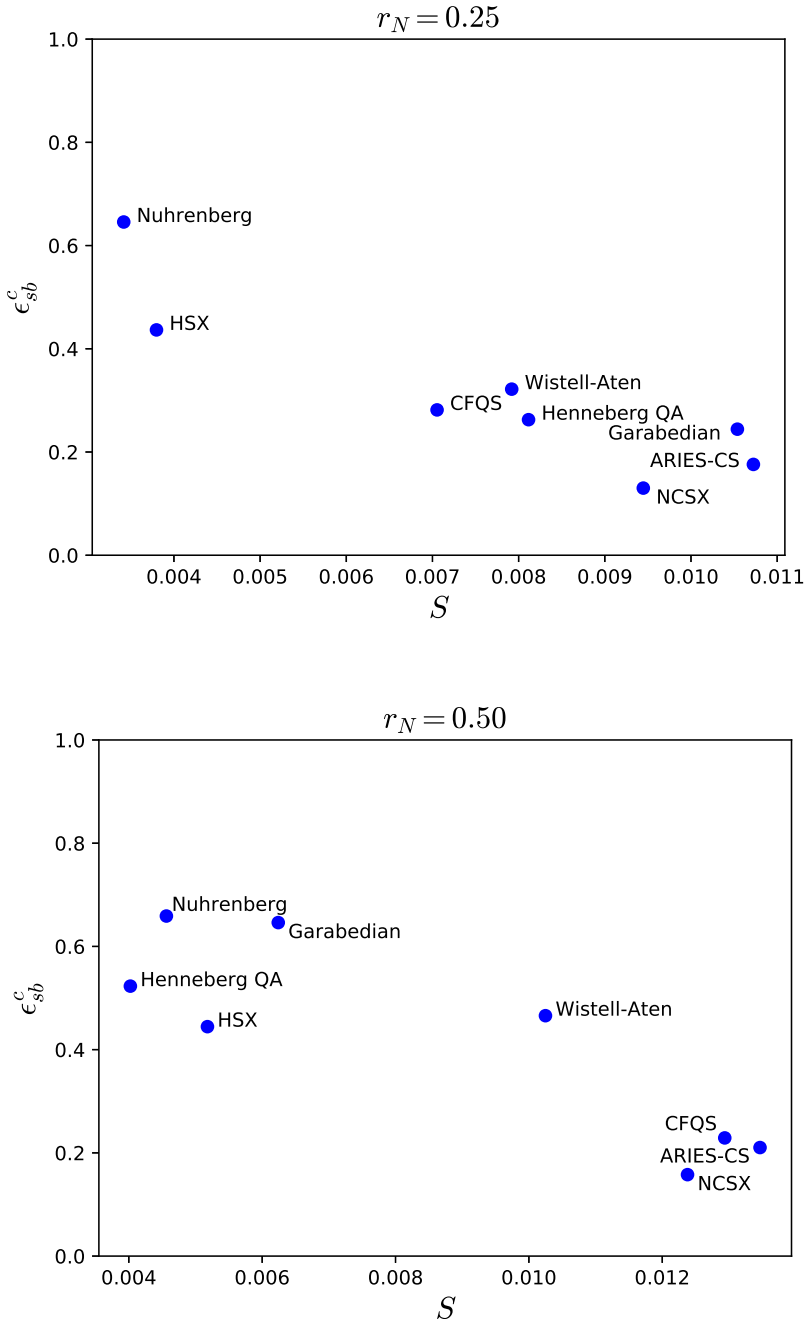

Figure 2.13: The critical symmetry-breaking parameter ϵ^c_{sb} for each configuration as a function of the corresponding S value has been plotted at (a) $r_N = 0.25$, and (b) $r_N = 0.50$, which correspond to the ϵ^c_{sb} values from Figure 2.1(a) and Figure 2.3(a), respectively.

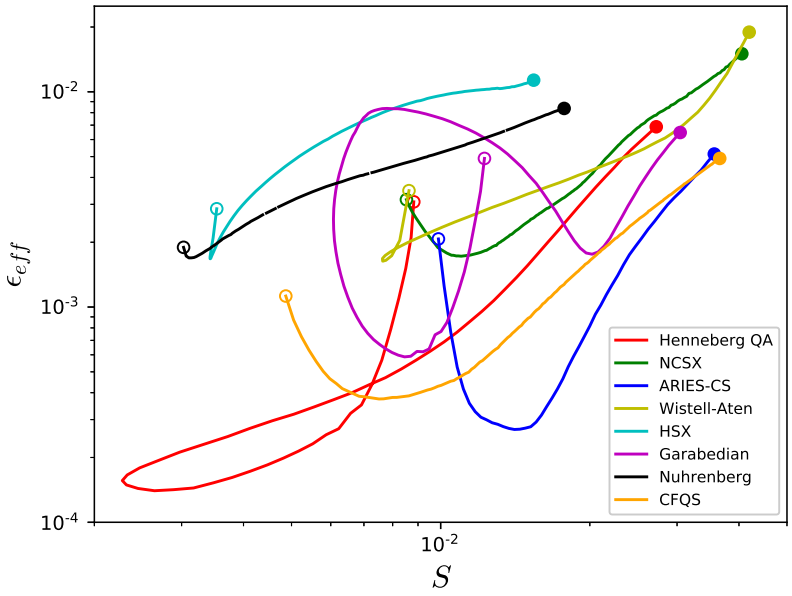

Figure 2.14: (Color online) The effective helical ripple (calculated with NEO [76]) is plotted as a function of the amplitude of the symmetry-breaking terms. The open circles here denote the value on-axis ($r_N = 0$). The closed circles correspond to the value at $r_N = 1$. These curves do not change with plasma parameters.

with r_N is accompanied with a decrease in ϵ_{eff}. Then, the subsequent increase in S corresponds to an increase in ϵ_{eff}, indicating a possible correlation between S and ϵ_{eff}. However, the behavior is different in the core of ARIES-CS, where the smallest value of S corresponds to a relatively large value of ϵ_{eff}, which decreases considerably as S is *increased*. The point here is that while there are certainly configurations where ϵ_{eff} scales with S, it is just as likely that they may not correlate well at all, and the assumption that small ϵ_{eff} indicates good quasisymmetry cannot be justified a priori. It has in fact been shown in [20] that one can achieve omnigeneity ($\epsilon_{eff} = 0$) far from quasisymmetry.

It is further interesting to note that in the cases where ϵ_{eff} does *not* scale with S, the radial location where this disagreement happens is usually within $r_N \simeq 0.5$. Above this r_N (or in some configurations, a position much closer to the magnetic axis), the scaling of ϵ_{eff} with S can be observed in every case. An interpretation of this behavior is left to future work.

2.8 Conclusions

In this work, we have examined how impurity particle flux and the temperature screening effect are influenced by varying the closeness of the magnetic field to perfect quasisymmetry. For realistic departures from symmetry, at the lowest studied collisionality (both species in the $\sqrt{\nu}$ regime) with a flat density gradient, temperature screening was not observed for any quasisymmetric configuration. However, with increasing collisionality one can see an increase in the "effective qua-

sisymmetry" of a flux surface. This can lead to temperature screening in some cases. Unfortunately, there is an upper limit to the benefits of increasing collisionality, where any further increase will lead to impurity accumulation even in perfect quasisymmetry.

When peaked density gradients are introduced, there is an overall negative effect on the impurity particle flux. Increasing the density gradient peaking ($\eta^{-1} > 0$) enhances the strength of the impurity accumulation, and also leads to a reduction in the "effective quasisymmetry". Overall, while temperature screening is technically possible at the true magnetic field in select cases, it is unlikely to be present in low-collisionality reactor-relevant regimes.

The magnitudes of these results at the true magnetic field ($\epsilon_{sb} = 1$) were then compared with a gyro-Bohm estimate for the turbulent fluxes. Even in the non-ideal scenario of $\eta^{-1} = 0.5$, the majority of configurations show neoclassical impurity particle fluxes that don't exceed 10% of the respective turbulent flux, even for highly charged impurities. However, a complete understanding of the implications of a relatively large turbulent particle flux will require further work, since determining *sign* of the particle flux may be more complicated than taking the sign of the largest transport channel [39]. In other words, while neoclassical fluxes may potentially be small, they cannot be considered irrelevant.

It was also found that when studying highly charged impurities in Wistell-A in relevant Z_{eff} plasmas, one cannot disregard the effect that including Φ_1 can have on Γ_z. Even though the effect on Γ_z is considerable, the absolute value of Φ_1 is quite small, indicating that its relationship to Γ_z is more complicated than just

considering its magnitude.

Finally, it was shown that the critical value of symmetry-breaking, ϵ_{sb}, where the impurity particle flux changes sign, appears to be anti-correlated with the amplitude of symmetry-breaking harmonics, S, on a flux surface. That this trend appears when considering configurations with widely varying properties suggests that minimizing the S on a flux surface will increase ϵ_{sb}.

The authors would like to thank S. Henneberg for providing the configuration in [48], A. Bader for providing Wistell-A, J. Schmitt for providing HSX, N. Pomphrey for providing NCSX and ARIES-CS, and S. Okamura for providing CFQS. We would also like to thank Ian Abel, Alessandro Geraldini, Rogerio Jorge, and Elizabeth Paul for useful comments and discussions. This work was supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science, under award DE-FG02-93ER54197. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

Chapter 3: Flux Tube Geometry

Calculating turbulent fluxes through simulations using the gyrokinetic equation generally require enormous computational effort. Solving the gyrokinetic equation with impurities enhances this cost since at least one additional kinetic species must be evolved. Therefore, it is important to choose the correct simulation domain and boundary conditions (Chapter 4) to minimize the overall cost.

3.1 What is a Flux Tube?

For instance, a gyrokinetic simulation of a full flux surface that uses field-line-following coordinates [7] and adiabatic electrons in stellarator geometry using the GENE code [58] currently requires on the order of 0.1 M CPU hours. The most cost-effective option is to run these codes in a flux tube (∼10-20 times faster), a simulation domain that follows a magnetic field line, is much longer than it is wide, and conserves magnetic flux throughout. The advantages of flux tube simulations and field-line-following coordinates can be seen no matter how one chooses to represent the distribution function (as $f(v)$, moments of $f(v)$, with particles, *etc.*). Figure 3.1 presents a visualization of a flux tube in the W7-X stellarator geometry.

The construction of this simulation domain starts with the expression of the

Figure 3.1: (Color online) 3D visualization of a flux tube domain in real space superimposed on a flux *surface* in the W7-X stellarator. (The perpendicular extent of the tube in was set for visualization purposes).

magnetic field in Clebsch coordinates

$$\boldsymbol{B} = \nabla\psi \times \nabla\alpha, \tag{3.1}$$

where ψ is the magnetic surface label (*e.g.* toroidal or poloidal flux), and α is a magnetic field line label. This form of $\boldsymbol{B}$ leads to $\boldsymbol{B}\cdot\nabla\psi = \boldsymbol{B}\cdot\nabla\alpha = 0$, meaning that ψ and α are constant along magnetic field lines and thus create ideal coordinates in the plane perpendicular to $\boldsymbol{B}$. The parallel coordinate z, identified with the poloidal angle θ, measures distance along the field line. These are known as *field-line-following* coordinates.

The field-line-following coordinate system is a particularly fitting choice in gyrokinetic simulations because of the anisotropic nature of turbulent fluctuations, which are elongated along a field line, and very short across it $k_\parallel/k_\perp \ll 1$ (see Figure 3.2). The perpendicular coordinate requires resolution on the gyroradius scale, while the parallel coordinate can be much more coarse-grained. For a more arbitrary chunk of plasma volume, one would need to resolve the gyroradius scale in all directions. Thus, the advantage to using field-line-following coordinates is not

solely convenient based on the structure of turbulent fluctuations, but also leads to an $O(\rho_*)$ reduction in the required resolution.

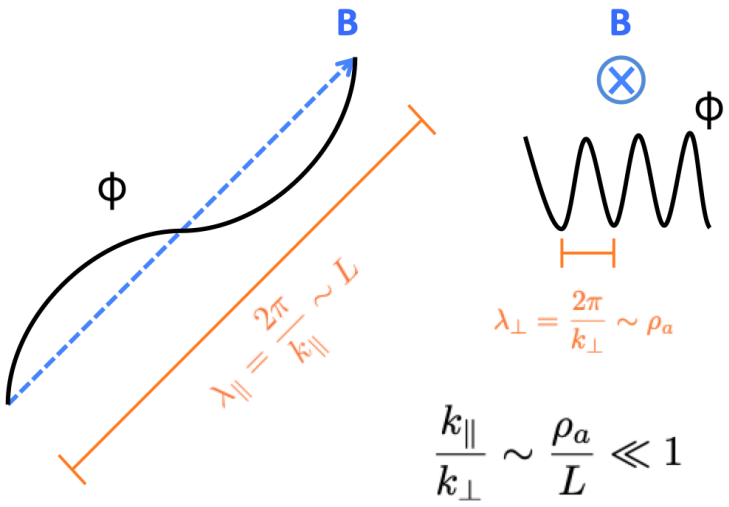

Figure 3.2: Cartoon representation of the scale of a perturbed quantity (the electrostatic potential ϕ in this case, in black) in the direction parallel and perpendicular to the magnetic field (blue).

Now since the flux tube has nonzero perpendicular extent, and ψ and α are being used as perpendicular coordinates, the logical conclusion would be that the domain simulates multiple magnetic field lines. However, flux tubes are *local* in both flux surface and magnetic field line, so it is important to understand how certain quantities are treated in the perpendicular domain.

The small-scale nature of turbulent fluctuations perpendicular to $\boldsymbol{B}$ helps to explain this. As will be discussed in 4.2, the extent of the domain is chosen to ensure that it is not shorter than the correlation length in any direction. For core plasmas, this will result in the perpendicular domain being on the order of a few ion gyroradii. Equilibrium quantities will have minimal variation across this domain and can be considered constants to lowest-order. This does not, however, restrict

variation of certain equilibrium quantities *along* the field line, such as the Jacobian

$$g^{1/2}(\psi_0, \alpha_0, z) = \left[(\nabla\psi \times \nabla\alpha) \cdot \nabla z\right]^{-1}. \tag{3.2}$$

This is because the parallel correlation length is on the order of the equilibrium scales.

The parallel gradients $\partial/\partial z$ of *perturbed* quantities (such as the electrostatic potential) will be assumed small compared to perpendicular gradients. This is a result of the anisotropy of plasma turbulence. It can then be shown that the spatial operators found in the gyrokinetic equation do not explicitly contain the perpendicular coordinates ψ and α. As shown in [7], the perpendicular Laplacian of some scalar function A

$$\nabla_\perp^2 A = |\nabla\psi|^2 \frac{\partial^2 A}{\partial\psi^2} + \nabla\psi \cdot \nabla\alpha \frac{\partial^2 A}{\partial\psi\partial\alpha} + |\nabla\alpha|^2 \frac{\partial^2 A}{\partial\alpha^2}, \tag{3.3}$$

is an important example of this fact. The above discussion (which is more thoroughly explored in [7]) demonstrates that physical quantities may vary along a field line, but will not depend on the location within the perpendicular plane (ψ, α).

The rotational transform

$$\iota = \iota(\psi) = \frac{\mathrm{d}\theta}{\mathrm{d}\phi} = \frac{\boldsymbol{B} \cdot \nabla\theta}{\boldsymbol{B} \cdot \nabla\phi}, \tag{3.4}$$

and the global magnetic shear

$$\hat{s} \equiv \frac{\psi_0}{q_0} \left. \frac{dq}{d\psi} \right|_{\psi=\psi_0}, \tag{3.5}$$

which are flux functions, are constant throughout the flux tube, not just the perpendicular plane.

To understand the shape of the flux tube in physical space, let us start by assuming that a field line is chosen about which to center a flux tube simulation domain in the perpendicular plane. The perpendicular coordinates of this field line are (ψ_0, α_0), which can be seen in Figure 3.3.

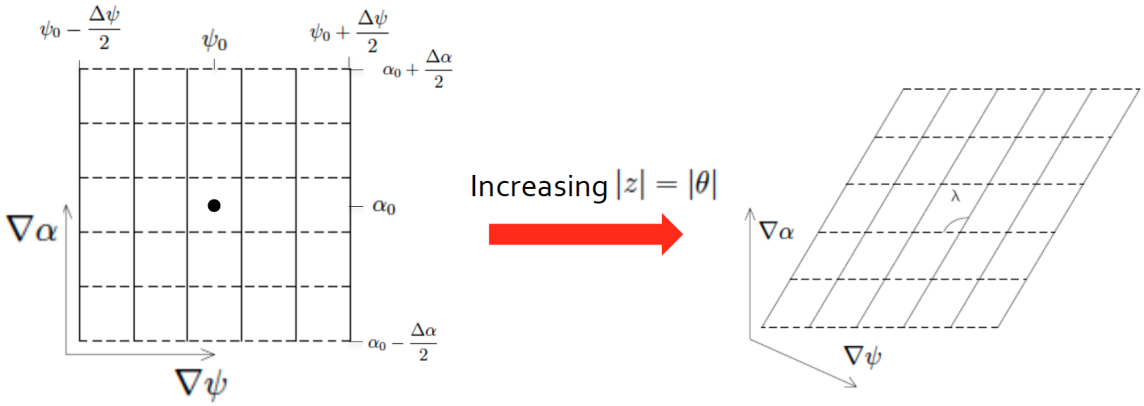

Figure 3.3: The shape of the perpendicular plane of a flux tube is defined by $\nabla\psi \cdot \nabla\alpha$. This figure is a representation of how this shape changes along a magnetic field line. This figure has been adapted from [8].

Since the magnetic field is taken to be $\boldsymbol{B} = \nabla\psi \times \nabla\alpha$, the flux tube sides are defined by the $\psi = \text{const}$ and $\alpha = \text{const}$ surfaces. This means that the scalar product of these gradients $\nabla\psi \cdot \nabla\alpha$ (which represents the non-orthogonality of the coordinates) will define the *shape* of the plane perpendicular to $\boldsymbol{B}$. From the arguments above, the quantities $\nabla\psi$ and $\nabla\alpha$ are both functions of $z = \theta$. This variation

along $\boldsymbol{B}$ is a result of magnetic shear considerations, which is the central focus of Chapter 4. The main point, however, is that the the variation in $\nabla\psi \cdot \nabla\alpha$ will cause a distortion in the shape of the perpendicular plane.

3.2 Acquiring Geometric Information

Since the main objective of the flux tube is to provide a simulation domain in which to solve the gyrokinetic equation, one needs to provide the necessary geometric information at each grid point. The specific quantities that are needed can be found in Table 3.1. For tokamaks with circular cross-sections, the calculation of these geometric quantities is relatively straightforward, as there are standard analytic formulas available. Significant complexities arise when dealing with non-standard toroidal geometries, such as shaped tokamaks or stellarators. The process involved in calculating these quantities will be detailed here.

3.2.1 Equilibrium Information

Prior to any simulations of turbulence, one must find a plasma that is in magnetohydrodynamic (MHD) equilibrium by satisfying the force balance equation

$$\nabla p = \boldsymbol{j} \times \boldsymbol{B}. \tag{3.6}$$

Perhaps the most widely used equilibrium code in the fusion community for 3D shapes is VMEC [53]. VMEC calculates an MHD equilibrium by assuming good flux surfaces, and reports the geometric and physical information needed to compute the

Geometric Information	
Quantity	Normalizing Factor
$\|B\|$	$1/B_{ref}$
$\boldsymbol{b}\cdot\nabla = \nabla_{\|}$	L_{ref}
$(\hat{s}^2/\psi_N)\|\nabla\psi\|^2$	$1/L_{ref}^2 B_{ref}^2$
$\hat{s}\nabla\psi\cdot\nabla\alpha$	$1/B_{ref}$
$\psi_N\|\nabla\alpha\|^2$	L_{ref}^2
$(2\sqrt{\psi_N}/B^3)\boldsymbol{B}\times\nabla B\cdot\nabla\alpha$	$L_{ref}^2 B_{ref}$
$2\hat{s}/(B^3\sqrt{\psi_N})\boldsymbol{B}\times\nabla B\cdot\nabla\psi$	1
$(2\sqrt{\psi_N}/B^2)\boldsymbol{B}\times\boldsymbol{\kappa}\cdot\nabla\alpha$	$L_{ref}^2 B_{ref}$
$2\hat{s}/(B^2\sqrt{\psi_N})\boldsymbol{B}\times\boldsymbol{\kappa}\cdot\nabla\psi$	1

Table 3.1: A list of geometric quantities necessary for gyrokinetic simulations in a flux tube geometry. To appropriately normalize these quantities, a scale length and reference magnetic field must also be chosen.

quantities of Table 3.1.

VMEC returns a prescribed number of Fourier harmonics for these geometric and physical quantities on each flux surface as a function of (θ_V, ζ). Here, ζ is the standard azimuthal angle in cylindrical coordinates, and θ_V is a poloidal angle, which is *not* a straight field line coordinate (this can be seen in the red curve of Figure 3.4, where θ_V is not a linear function of ζ). This poloidal angle was chosen because it reduces the number of terms required in the Fourier series. Flux surfaces are labeled by the radial coordinate $s \equiv \psi/\psi_{LCFS}$, where ψ_{LCFS} is the toroidal flux at the last closed flux surface. The angle of the sine and cosine terms in the Fourier expansions is defined as $m\theta_V - n\zeta$, where m and n denote the poloidal and toroidal mode number, respectively. For example, the entire right hand side of the expression

$$B = \sum_{m,n} B_{mn} \cos(m\theta_V - n\zeta), \tag{3.7}$$

would be provided to determine the magnetic field magnitude within a given flux surface.

VMEC also provides the function $\Lambda = \Lambda(s, \theta_V, \zeta)$ (analagous to the more general ν in Eq (8) of [7]), which is periodic in θ_V and ζ, and allows one to convert between poloidal coordinates via

$$\theta_P = \theta_V + \Lambda. \tag{3.8}$$

Here, θ_P is the poloidal angle in PEST coordinates, which is a straight-field-line

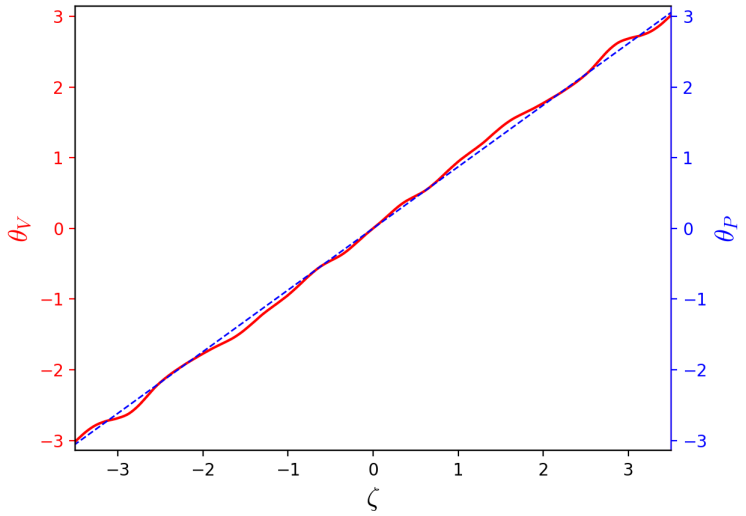

Figure 3.4: The VMEC poloidal angle θ_V is plotted in red as a function of the standard azimuthal angle ζ for the field line $\alpha = \theta_V + \Lambda - \iota\zeta = 0$. The fact that θ_V is not a linear function of ζ indicates that it is not a straight-field-line coordinate. The straight-field-line coordinate PEST coordinate θ_P is plotted in the dashed blue for the field line $\alpha = \theta_P - \iota\zeta = 0$.

angle. A plot of θ_P as a function of ζ results in a straight line, which can be seen by the dashed blue curve in Figure 3.4. PEST coordinates are realized with a straight-field-line poloidal angle in the case of the toroidal angle ζ being the standard azimuthal angle in cylindrical coordinates.

It is also possible to convert to another straight-field-line coordinate system known as Boozer coordinates, with extra steps. Using the same Λ function as well as $\iota = \iota(s)$ and a transformation function $\tilde{p} = \tilde{p}(s, \theta_V, \zeta)$ (that is periodic in $\theta_V$and ζ), it has been shown [54] that the conversion to Boozer coordinates from VMEC is

$$\theta_B = \theta_V + \Lambda + \iota\tilde{p} \tag{3.9}$$

$$\zeta_B = \zeta + \tilde{p}. \tag{3.10}$$

There is also a code BOOZ_XFORM [88] in the STELLOPT [82, 92] suite of codes, which performs this conversion. Boozer coordinates are a special set of straight-field-line coordinates, which reveal hidden symmetries in the magnetic field. This was discussed in detail in Chapter 2.

Table 3.2 lists and provides a brief description of the VMEC output quantities necessary to compute the geometric information needed for flux tube simulations. The parenthetical quantities in Table 3.2 are additional quantities for the case of equilibria that do not have up-down (or stellarator) symmetry. The covariant and contravariant components of the magnetic field, and the Jacobian, are defined in the usual sense

$$
\begin{aligned}
B^{\theta} &= \boldsymbol{B} \cdot \nabla \theta_V, \\
B^{\zeta} &= \boldsymbol{B} \cdot \nabla \zeta, \\
B_{\theta} &= \boldsymbol{B} \cdot \frac{\partial \boldsymbol{r}}{\partial \theta_V}, \\
B_{\zeta} &= \boldsymbol{B} \cdot \frac{\partial \boldsymbol{r}}{\partial \zeta}, \\
B_{s} &= \boldsymbol{B} \cdot \frac{\partial \boldsymbol{r}}{\partial s},
\end{aligned}
$$

$$
\sqrt{g} = \frac{\partial \boldsymbol{r}}{\partial s} \cdot \frac{\partial \boldsymbol{r}}{\partial \theta_V} \times \frac{\partial \boldsymbol{r}}{\partial \zeta} = \left(\nabla s \cdot \nabla \theta_V \times \nabla \zeta \right)^{-1}.
$$

VMEC Output Quantities	
Output Name	Description
`Aminor_p`	averaged minor radius at LCFS
`xm,xn`	poloidal and toroidal mode numbers
`xm_nyq,xn_nyq`	poloidal and toroidal mode numbers (Nyquist)
`iotaf`	ι on full mesh
`presf`	pressure on full mesh
`phi`	toroidal flux on full mesh
`rmnc (rmns)`	`cosmn` (`sinmn`) components of cylindrical R, full mesh
`zmns (zmnc)`	`sinmn` (`cosmn`) components of cylindrical Z, full mesh
`bmnc (bmns)`	`cosmn` (`sinmn`) components of $\|B\|$, half mesh
`gmnc (gmns)`	`cosmn` (`sinmn`) components of Jacobian, half mesh
`bsupumnc (bsupumns)`	`cosmn` (`sinmn`) components of B^θ, half mesh
`bsupvmnc (bsupvmns)`	`cosmn` (`sinmn`) components of B^ζ, half mesh
`bsubumnc (bsubumns)`	`cosmn` (`sinmn`) components of B_θ, half mesh
`bsubvmnc (bsubvmns)`	`cosmn` (`sinmn`) components of B_ζ, half mesh
`bsubsmns (bsubsmnc)`	`sinmn` (`cosmn`) components of B_s, full mesh

Table 3.2: A list and description of VMEC output quantities that are necessary to compute geometric information for flux tube simulations.

3.2.2 Calculating Geometric Quantities

The next step is to select the desired flux surface, which will allow for the calculation of flux functions such as $\iota, d\iota/d\psi$, and $dp/d\psi$, where p is the pressure. The value of ι is provided directly as a function of the flux surface on a uniform grid. The derivative quantities must then be calculated by finite difference.

At this point, quantities in the straight-field-line coordinates (ψ, θ_V, ζ) have been calculated. The next step is to calculate these quantities in the field-line-following coordinates (ψ, α, z) of [7]. Apart from the physical reason of choosing these coordinates, an advantage of transforming to (ψ, α, z) lies in the relative simplicity of the nonlinear term in the gyrokinetic equation. In the ballooning representation, for example, the nonlinear term includes an additional p sum over sections of

increasing length in the parallel coordinate ($\theta + 2\pi p$) that can become very expensive to evaluate for long flux tubes [8, 31]. The nonlinear term in field-line-following coordinates does not have such a sum, with the result being proportional to a Poisson bracket that involves only convolutions in Fourier space.

This transformation first requires the creation of poloidal and toroidal coordinates satisfying $\alpha = \theta - \iota\zeta$. Since VMEC quantities are provided on the (θ_V, ζ) grid, however, one must solve for θ_V at each θ in order to represent these quantities on the straight-field-line grid. Determining θ_V requires satisfying

$$\theta_V - \theta = \theta_V - \alpha + \iota\zeta = 0, \tag{3.11}$$

for the desired field line α at each ζ on the grid.

With the quantities now matched to the θ_V values corresponding to the straight-field-line θ values, one must then convert the VMEC Fourier representation of each quantity in Table 3.2 to their real-space representation as functions of (θ_V, ζ). This is a straightforward exercise of performing the equivalent of the sum in Eq (3.7), over the all modes that were used in VMEC.

The next step is to compute the vectors $\nabla\psi$, $\nabla\alpha$, $\nabla\theta$, and ∇B. It is helpful at this point to convert to Cartesian coordinates (X, Y, Z), by first using $X = R\cos(\zeta)$ and $Y = R\sin(\zeta)$, and taking the appropriate derivatives. One can then use the

dual relations to find gradients of the field-line-following coordinates

$$\nabla\theta_V = \frac{1}{\sqrt{g}}\left(\frac{\partial \boldsymbol{r}}{\partial\zeta}\times\frac{\partial \boldsymbol{r}}{\partial s}\right), \tag{3.12}$$

$$\nabla\zeta = \frac{1}{\sqrt{g}}\left(\frac{\partial \boldsymbol{r}}{\partial s}\times\frac{\partial \boldsymbol{r}}{\partial\theta_V}\right), \tag{3.13}$$

$$\nabla s = \frac{1}{\sqrt{g}}\left(\frac{\partial \boldsymbol{r}}{\partial\theta_V}\times\frac{\partial \boldsymbol{r}}{\partial\zeta}\right) = \left(\frac{d\psi}{ds}\right)^{-1}\nabla\psi. \tag{3.14}$$

From this information, $\nabla\alpha = \nabla(\theta_V + \Lambda - \iota\zeta)$ can be computed, with the reminder that $\iota = \iota(s)$ and $\Lambda = \Lambda(s, \theta_V, \zeta)$

$$\nabla\alpha = \left(\frac{\partial\Lambda}{\partial s} - \zeta\frac{d\iota}{ds}\right)\nabla s + \left(1 + \frac{\partial\Lambda}{\partial\theta_V}\right)\nabla\theta_V + \left(-\iota + \frac{\partial\Lambda}{\partial\zeta}\right)\nabla\zeta. \tag{3.15}$$

Finally, we can now calculate (in Cartesian coordinates) the magnetic field vector $\boldsymbol{B} = \nabla\psi\times\nabla\alpha$ and the gradients of its magnitude

$$\nabla B = \frac{dB}{ds}\nabla s + \frac{dB}{d\theta_V}\nabla\theta_V + \frac{dB}{d\zeta}\nabla\zeta. \tag{3.16}$$

The derivatives here can be calculated with either finite difference or analytic derivatives using Eq 3.7.

These gradients, appropriately normalized, provide the necessary information to calculate the required geometric quantities. Table 3.1 simply lists these quantities appropriately normalized by the unspecified reference values B_{ref} and L_{ref}. Each quantity in 3.1 is approximately constant within the perpendicular plane of a flux tube because of its narrow width. They are therefore treated as only a function of

the parallel coordinate.

3.3 Straight-field-line and Equal Arc Length Coordinates

The above quantities are calculated from VMEC to yield a straight-field-line coordinate system, which is what permits the use of field-line-following coordinates. However, while straight, the parallel (poloidal angle in this case) grid points are generally not equally-spaced along the field line. Unequally-spaced coordinates are certainly allowable, but equally-spaced grids can sometimes lead to an improvement in performance.

The subtlety of equal arc length coordinates can be understood by considering a simple equation, which considers only motion along a field line

$$\frac{\partial f}{\partial t} + v\frac{\partial \theta}{\partial l}\frac{\partial f}{\partial \theta} = 0. \tag{3.17}$$

Here, l is the equal arc coordinate, which parameterizes the distance along a field line. In most gyrokinetic codes, the poloidal angle is chosen to be the parallel coordinate, so that θ parameterizes distance along a field line. The quantity $\partial\theta/\partial l$ represents how the actual poloidal angle θ (for some definition of θ) varies with distance along the magnetic field. Due to the complicated helical nature of toroidal magnetic field lines, even an equally-spaced θ-grid will still result in $\partial\theta/\partial l = g(\theta)$, where g is some scalar function of θ. The only way to ensure that a set of grid points are equally-spaced along the field is to ensure that $\partial\theta/\partial l = \text{const}$.

As mentioned above, equal arc length coordinates are not required to run flux

tube simulations of turbulence. Many existing codes [18,27,58] use a θ-grid without equal arc lengths. The benefits to using an equally-spaced coordinate can be found when performing Fourier transforms. Having $\partial\theta/\partial l = \text{const.}$ permits the use of Fast Fourier Transfrom (FFT) libraries that require an equally-spaced grid. For codes that require FFTs [49,66], it is straightforward to interpolate quantities onto a grid that is constructed to satisfy $\partial\theta/\partial l = \text{const.}$, and still retain the straight-field-line coordinates. The majority of results in Chapter 4 are produced using the GryfX [49] code, which utilizes an equal arc length parallel coordinate.

The content in this chapter borrows heavily from the work in:
Martin M.F., et al. *Plasma Phys. Control. Fusion* **60** 095008 (2018) [68]

Chapter 4: The Parallel Boundary Condition in Turbulence Simulations

4.1 Introduction

Understanding and predicting turbulent transport in fusion devices remains one of the most pressing issues in moving fusion energy forward. In the tokamak community, microinstabilities and core turbulence have been extensively studied using an array of gyrokinetic codes [18, 27, 58, 105]. However, solving the gyrokinetic equation is a generally expensive endeavor, and the geometric complexities introduced when moving to stellarators result in commensurately more expensive computational studies. As discussed in the previous section, flux tubes are the most cost-effective option for gyrokinetic simulations, and are the focus of this chapter. Such domains use the field-line-following coordinates and boundary conditions originally developed in [7] for gyrofluid simulations. The combination of field-line-following coordinates and a flux tube domain reduces turbulence simulation runtimes by $10/\rho_*^2 \sim 10^5$ but requires an implementation of periodicity. As the coordinates are non-orthogonal and curvilinear, the flux tube domain boundaries are not manifestly periodic. For axisymmetric geometries (*e.g.,* tokamaks) the "twist-and-shift"

boundary condition [7] has been used for decades and can be expressed in logically Cartesian (x, y, z) coordinates, with (x, y) measuring appropriately normalized distances in the plane locally perpendicular to the magnetic field and z measuring distances along the magnetic field. Expressing the flux tube periodicity involves all three directions but is typically expressed as a "parallel" boundary condition, as the spatial "twist" of the magnetic field which is accumulated as one moves along a bundle of magnetic field lines is accommodated by "shifted" alignments of perpendicular-to-the-field-line Fourier modes at either end of the flux tube. The twist-and-shift boundary conditions [7,26] were designed to unwind the secular twist that arises from strong global magnetic shear, denoted here by $\hat{s}$. There are two important consequences of any physically correct twist-and-shift boundary condition. First, each (k_x, k_y) Fourier mode undergoes a shift in k_x (proportional to k_y) across the z boundary. (An equivalent condition exists for non-spectral representations.) Second, the perpendicular aspect ratio of the simulation domain, L_x/L_y, is necessarily quantized. As will be discussed in Section 4.2, all existing expressions of these constraints explicitly depend on the global magnetic shear.

Problems arise in devices such as W7-X [40], which was designed to have rotational transform with minimal radial variation, to avoid low-order rational flux surfaces. Low global shear designs are not exclusive to stellarators, however, as advanced tokamak scenarios [64] can have similarly flat q profiles, where q is the safety factor. In such geometries, the existing expression of the parallel boundary condition is inconvenient because of the intrinsically low global magnetic shear. In particular, the perpendicular aspect ratio of the simulation domain is inversely proportional to

the global shear ($L_x \propto L_y/\hat{s}$). For $\hat{s} \ll 1$, this imposes restrictive resolution requirements (a large number of x grid points) that increase computation time. Moreover, naive application of the axisymmetric formulas to non-axisymmetric geometries results in errors.

To address these shortcomings, we have generalized the parallel boundary condition for flux tube simulations. Our generalization allows for non-axisymmetric geometries and depends on local rather than global magnetic shear. In many cases of interest, the local magnetic shear can vary considerably along a flux tube, even when the global shear is weak. For 3-D equilibria, our approach requires stellarator symmetry (Section 4.4.1). In the case of axisymmetry, this generalized boundary condition reduces to conventional "twist-and-shift" when the flux tube ends are separated by an integer number of poloidal turns.

The significant variation of local magnetic shear in low *global* shear geometries presents the opportunity to optimize the flux tube length. For appropriately selected flux tube lengths, it is possible to use periodic parallel boundary conditions, or to preserve continuity in the magnetic drifts across the parallel boundary with a perpendicular aspect ratio of the simulation domain close to unity. The effects of this new boundary condition on the speed and accuracy of microinstability and turbulence simulations are explored here.

This chapter is organized as follows: In Section 4.2, we define the field-line-following coordinate system and how fluctuating quantities are represented within the simulation domain and across its perpendicular boundaries, with Section 4.3 detailing the conventional method of handling the parallel boundary. Section 4.4

presents the full derivation of the new, generalized version of the "twist-and-shift" boundary condition, including discussions of some of its useful properties. Section 4.5 discusses the characteristic behavior of shear in a stellarator flux tube geometry possessing low global shear, and how this behavior can be used to optimize the new boundary condition. Finally, Section 4.6 presents numerical studies of linear instability, secondary instability, and nonlinear turbulence, showing how relevant physical quantities depend on the parallel boundary condition choice. In the cases tested here, results indicate that even incorrect implementations of the new boundary condition do not affect the ability of a simulation to predict certain important quantities. A significant computational speedup is also observed when using the new boundary condition, compared with the conventional method.

4.2 Flux Tube Simulations

The microinstabilities that develop into the turbulence responsible for the high levels of transport observed in fusion devices are characteristically highly elongated along the magnetic field relative to their scale lengths perpendicular to the magnetic field. It is thus natural to introduce field-aligned coordinates [7] for toroidal magnetic confinement devices. Such coordinates are readily understood when the magnetic field is expressed in Clebsch form,

$$\boldsymbol{B} = \nabla\psi \times \nabla\alpha. \tag{4.1}$$

Here, ψ and α are constant on magnetic field lines, and so can be used as coordinates in the plane perpendicular to $\boldsymbol{B}$. Without loss of generality, we identify ψ as a magnetic surface label (*e.g.,* toroidal or poloidal flux) and thus think of it as the logically radial coordinate. The coordinate α is a magnetic field line label. The third coordinate z measures distances along a magnetic field line. Again without loss of generality, we identify z with the poloidal angle θ. Further details on the flux tube geometry can be found in Chapter 3.

The minimal simulation domain for a turbulence simulation should not be shorter than the correlation length in any direction. Perpendicular correlation lengths λ are observed to be on the order of a few ion gyroradii in core plasmas. It should therefore be possible to model these fluctuations in a periodic perpendicular domain of size $L_\psi \times L_\alpha$, as long as $L_\alpha/\lambda, L_\psi/\lambda$ are both large enough. We wish to find a minimal domain and so we use a periodic perpendicular domain whose lengths are measured in ion gyroradii. For any fluctuating quantity ϕ,

$$\phi(\psi, \alpha, z, t) = \phi(\psi + L_\psi, \alpha, z, t) = \phi(\psi, \alpha + L_\alpha, z, t). \tag{4.2}$$

The small perpendicular extent of the box also means that geometric quantities $(\boldsymbol{B}, \nabla\psi, \nabla\alpha)$ can be fully characterized by their local values and gradients, approximately independent of ψ, α (see Chapter 3).

The periodic perpendicular boundary conditions allow one to represent ϕ as a

Fourier series in these coordinates:

$$\phi(\psi,\alpha,z,t)=\sum_{m=-\infty}^{\infty}\sum_{n=-\infty}^{\infty}\hat{\phi}_{m,n}(z,t)\,\exp\Big(\frac{2\pi i(\psi-\psi_0)m}{L_\psi}+\frac{2\pi i(\alpha-\alpha_0)n}{L_\alpha}\Big), \quad (4.3)$$

with the constants (ψ_0,α_0) representing the center of the flux tube in the perpendicular plane. For both a simplified representation and a means to understand the steps which follow, the fluctuations can also be represented by the wavenumbers $k_\psi \equiv 2\pi m/L_\psi$ and $k_\alpha \equiv 2\pi n/L_\alpha$:

$$\phi(\psi,\alpha,z,t)=\sum_{k_\psi=-\infty}^{\infty}\sum_{k_\alpha=-\infty}^{\infty}\hat{\phi}_{k_\psi,k_\alpha}(z,t)\,\exp\Big(ik_\psi\Delta\psi+ik_\alpha\Delta\alpha\Big), \quad (4.4)$$

where $\Delta\psi=\psi-\psi_0$ and $\Delta\alpha=\alpha-\alpha_0$. The rest of this chapter concerns the conditions imposed at the ends of the domain in the parallel coordinate, z, in the context of fluctuations defined as in Eq. (4.4).

4.3 The Standard Parallel Boundary Condition

The standard parallel boundary condition [7] is based on the assumption that turbulent fluctuations should be statistically identical at two locations with the same poloidal angle (but different toroidal angle) in an *axisymmetric* geometry. It should be clear that this renders the boundary condition formally incorrect when simulating flux tubes in a stellarator, as the geometry is inherently non-axisymmetric.

Quantitatively, this assumption about turbulent fluctuations produces the fol-

lowing constraint on the fluctuating quantity ϕ:

$$\phi[\psi, \alpha(\theta + 2\pi N, \zeta), z(\theta + 2\pi N)] = \phi[\psi, \alpha(\theta, \zeta), z(\theta)]. \tag{4.5}$$

Here, we take θ and ζ to be the poloidal and toroidal angles, respectively, where magnetic field lines are straight in the (θ, ζ) plane. Further, the field line label is taken to be

$$\alpha = \zeta - q\theta, \tag{4.6}$$

where $q = q(\psi)$. By applying the above constraint to the fluctuation form (4.4), one can derive a set of conditions that must be satisfied in the simulation, namely:

$$\begin{aligned} [k_\alpha]_{z=+\pi N} &= [k_\alpha]_{z=-\pi N}\,, \\ k'_\psi \equiv [k_\psi]_{z=+\pi N} - [k_\psi]_{z=-\pi N} &= 2\pi N \frac{\mathrm{d}q}{\mathrm{d}\psi} k_\alpha, \end{aligned} \tag{4.7}$$

with N being a positive integer. Thus, by imposing the constraint in (4.5), there is a required shift in k_ψ that a (k_ψ, k_α) Fourier mode of ϕ must undergo in passing from one end of the domain to the other. This results in the standard parallel boundary condition on fluctuating quantities in flux tube simulations:

$$\phi_{k_\psi + k'_\psi, k_\alpha}\left[\theta + 2\pi N, t\right] C_{k_\alpha} = \phi_{k_\psi, k_\alpha}\left[\theta, t\right], \tag{4.8}$$

where C_{k_α} is a phase factor, $|C_{k_\alpha}| = 1$. Since we cannot retain an infinite number of modes in a simulation, the shift in k_ψ, along with the number of modes we choose

to evolve, determines the maximum $k_\perp$ value able to be resolved.

At this point, it is appropriate to introduce the coordinates (x, y), which are the standard normalization-dependent code representations of (ψ, α) that have units of length. The following normalization choices have been used in the steps that follow:

$$\frac{\mathrm{d}x}{\mathrm{d}\psi_t} = \frac{1}{aB_0\sqrt{s_0}}, \qquad \frac{\mathrm{d}y}{\mathrm{d}\alpha} = \frac{a\sqrt{s_0}}{q_0}. \tag{4.9}$$

where $x \equiv a\sqrt{\psi_t/\psi_{edge}}$, with ψ_t taken to be the toroidal flux and ψ_{edge} is the value of ψ_t at the plasma edge. In the above definitions, a is a constant representing an effective minor radius, s_0 is a flux surface label where $s \equiv \psi_t/\psi_{edge}$, and $B_0 = 2\psi_{edge}/a^2$ is the reference magnetic field. Using (4.9), we can rewrite k'_ψ in terms of x and y as

$$k'_x = 2\pi N \hat{s}\, k_y, \tag{4.10}$$

with $\hat{s} \equiv (x/q)\mathrm{d}q/\mathrm{d}x|_{x=x_0}$. It is also straightforward to show that these conditions impose a quantization on the perpendicular aspect ratio of the domain

$$\frac{L_x}{L_y} = \frac{J}{2\pi N\,|\hat{s}|}, \tag{4.11}$$

where J is a nonzero integer that can be set in the code to potentially achieve a more desirable aspect ratio. These constraints, when applied to stellarator geometries possessing low global shear, become very restrictive with respect to resolution requirements. For instance, on the $x/a = 0.357$ surface of W7-X, where $\hat{s} = -0.019$,

(4.11) corresponds to $L_x/L_y = 8.1J$. The radial extent of the simulation domain is forced to be large. Here, and in stellarator calculations below, we will take a to be the effective minor radius calculated by VMEC [53]. Good estimates of heat fluxes and other quantities of interest typically require one to resolve fluctuations with wavenumbers extending up to $k_\perp \rho_i \sim 1$. This is expensive in a radially extended domain. In a spectral decomposition, one has to use a correspondingly large number of Fourier modes. In a grid-based discretization, one has to use a large number of grid points.

4.4 The New Parallel Boundary Condition(s)

Our generalization necessarily remains consistent with stellarator symmetry, but relaxes the explicit dependence on global magnetic shear in favor of the integrated local magnetic shear.

4.4.1 Stellarator Symmetry

A flux tube demonstrating stellarator symmetry has the property that it is unchanged when rotated by 180° about an appropriate point. This symmetry implies that the magnitudes of geometric quantities are equivalent at stellarator symmetric locations. For our purposes, stellarator symmetry can be summarized by three identities,

$$|B|_{z_+} = |B|_{z_-} \qquad |\nabla\psi|_{z_+} = |\nabla\psi|_{z_-} \qquad |\nabla\alpha|_{z_+} = |\nabla\alpha|_{z_-}, \tag{4.12}$$

where $z_\pm$ indicates the two ends of the flux tube at $(\psi, \pm\theta, \pm\zeta)$.

As long as two stellarator symmetric locations are farther apart than a few correlation lengths in each direction, the fluctuations at those points will thus be indistinguishable on average. (In the absence of stellarator symmetry, the differences in magnetic geometry would not permit this assertion.) We will assert periodicity only at widely separated, stellarator symmetric points.

It is not generally known how to guarantee that a flux-tube domain is long enough, even when it is long compared with the simulation's correlation lengths. In general, for example, there is a flux of free energy along field lines in a turbulent plasma. In a simulation, this free energy flux is a form of dissipation when it is a net exhaust, but a form of noise when the net flow is into the domain. Only by simulating a full flux surface can one resolve this category of uncertainty.

4.4.2 Orthonormal Coordinates

It is convenient to construct orthonormal coordinates (u, v) to describe the plane perpendicular to the magnetic field. By doing so, we can explicitly capture the local shear information along a flux tube. In the traditional non-orthogonal coordinates $(\nabla\psi, \nabla\alpha)$, this information manifests itself in a distortion of the perpendicular plane, hiding potentially useful local magnetic shear information.

We consider a Clebsch representation of the magnetic field $\boldsymbol{B} = \nabla\psi \times \nabla\alpha$, with a field line centered on the coordinates $\psi = \psi_0$ and $\alpha = \alpha_0$. As before, we assume

ψ is a flux surface label. Notice that the following three vectors are orthonormal:

$$\hat{\boldsymbol{b}} = \frac{\boldsymbol{B}}{|\boldsymbol{B}|}, \qquad \hat{\boldsymbol{e}}_u = \frac{\nabla\psi}{|\nabla\psi|}, \qquad \hat{\boldsymbol{e}}_v = \frac{\hat{\boldsymbol{b}} \times \nabla\psi}{|\nabla\psi|}. \tag{4.13}$$

We denote the position vector of the central field line by $\boldsymbol{r}_0(z)$, where z parameterizes the position along the field line. Any point in the flux tube can be labeled with coordinates $(\psi = \psi_0 + \Delta\psi,\ \alpha = \alpha_0 + \Delta\alpha,\ z)$. The position vector $\boldsymbol{r}$ for this point in the flux tube can be written

$$\boldsymbol{r}\left(\psi, \alpha, z\right) \approx \boldsymbol{r}_0(z) + \left(\frac{\partial \boldsymbol{r}}{\partial \psi}\right)_{\alpha,z} \Delta\psi + \left(\frac{\partial \boldsymbol{r}}{\partial \alpha}\right)_{\psi,z} \Delta\alpha. \tag{4.14}$$

At the same time, we can parameterize the perpendicular plane using alternative coordinates (u, v) defined in terms of the orthonormal basis (4.13):

$$\begin{aligned} u &= (\boldsymbol{r} - \boldsymbol{r}_0) \cdot \hat{\boldsymbol{e}}_u, \\ v &= (\boldsymbol{r} - \boldsymbol{r}_0) \cdot \hat{\boldsymbol{e}}_v. \end{aligned} \tag{4.15}$$

Substituting (4.13) and (4.14) into (4.15), noting

$$\left(\frac{\partial \boldsymbol{r}}{\partial \psi}\right)_{\alpha,z} \cdot \nabla\psi = 1, \qquad \left(\frac{\partial \boldsymbol{r}}{\partial \psi}\right)_{\alpha,z} \cdot \nabla\alpha = 0, \tag{4.16}$$

and using $\boldsymbol{B} = \nabla\psi \times \nabla\alpha$ to find $(\partial\boldsymbol{r}/\partial\psi) \cdot \boldsymbol{B} \times \nabla\psi = -\nabla\psi \cdot \nabla\alpha$ and $(\partial\boldsymbol{r}/\partial\alpha) \cdot \boldsymbol{B} \times \nabla\psi = |\nabla\psi|^2$, we obtain a relation between the orthonormal coordinates (u, v)

and the standard Clebsch coordinates:

$$u = \frac{\Delta\psi}{|\nabla\psi|},$$
$$v = \frac{-(\nabla\psi\cdot\nabla\alpha)\Delta\psi + |\nabla\psi|^2\Delta\alpha}{B|\nabla\psi|}. \tag{4.17}$$

Figure 4.1 presents an example of how the two sets of coordinates parameterizing the perpendicular plane compare at an arbitrary location along the parallel coordinate, z, of the flux tube.

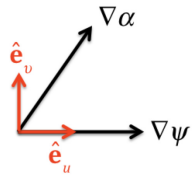

Figure 4.1: (Color Online) Vector directions in the perpendicular plane for the orthonormal $(\hat{\boldsymbol{e}}_u, \hat{\boldsymbol{e}}_v)$ and Clebsch $(\nabla\psi, \nabla\alpha)$ coordinates at an arbitrary z location where $\nabla\psi\cdot\nabla\alpha \neq 0$.

4.4.3 Boundary Condition Derivation

Using the orthonormal coordinates of (4.17), the new parallel boundary condition for flux tube simulations can be derived, assuming certain requirements are met. The fluctuations expressed as functions of (u, v) at the two ends of the flux tube should either be (i) at stellarator-symmetric locations or (ii) separated by an integer number of poloidal turns in axisymmetry. This allows us to take the fluctuating quantity ϕ of (4.4) to be equal at the two ends of a stellarator-symmetric flux tube, and determine which values of k_ψ and k_α are connected as in (4.7).

We start by rearranging (4.17) to get expressions for $(\Delta\psi, \Delta\alpha)$ as functions of

(u, v):

$$\begin{aligned} \Delta\psi &= u|\nabla\psi|, \\ \Delta\alpha &= \frac{B}{|\nabla\psi|}v + \frac{\nabla\psi\cdot\nabla\alpha}{|\nabla\psi|}u. \end{aligned} \tag{4.18}$$

Substituting (4.18) into (4.4), we see that the fluctuations for each wavenumber pair (k_ψ, k_α) have the form

$$\hat{\phi}_{k_\psi,k_\alpha}(z,t)\exp\left(i|\nabla\psi|\left[k_\psi + k_\alpha\frac{\nabla\psi\cdot\nabla\alpha}{|\nabla\psi|^2}\right]u + i\left[k_\alpha\frac{B}{|\nabla\psi|}\right]v\right), \tag{4.19}$$

where terms depending on u have been collected. Identifying the (u, v) planes at the two ends of the flux tube, then the coefficients multiplying u and v in (4.19) must each match, yielding:

$$\left[k_\alpha\frac{B}{|\nabla\psi|^2}\right]_{z_+} = \left[k_\alpha\frac{B}{|\nabla\psi|^2}\right]_{z_-}, \tag{4.20}$$

$$\left[k_\psi + k_\alpha\frac{\nabla\psi\cdot\nabla\alpha}{|\nabla\psi|^2}\right]_{z_+} = \left[k_\psi + k_\alpha\frac{\nabla\psi\cdot\nabla\alpha}{|\nabla\psi|^2}\right]_{z_-}. \tag{4.21}$$

These relations hold for all stellarator-symmetric flux tubes, as well as in axisymmetric geometry where the flux tube goes around an integer number of times poloidally, such that the ends coincide in a poloidal projection. If neither (i) or (ii) are satisfied, then the magnetic geometry at the two ends of the flux tube is dissimilar and we do not expect the turbulence to be statistically similar at the two ends, so the derivation breaks down. On the other hand, if either of these two

conditions above are satisfied, we can reduce (4.20) and (4.21) using (4.12). Then (4.20) indicates that we should link a given k_α to the same k_α at the other end of the flux tube, which is consistent with Beer's result. We can then write (4.21) as

$$k_\psi^{shift} \equiv [k_\psi]_{z_+} - [k_\psi]_{z_-} = -\left(\left[\frac{\nabla\psi \cdot \nabla\alpha}{|\nabla\psi|^2} \right]_{z_+} - \left[\frac{\nabla\psi \cdot \nabla\alpha}{|\nabla\psi|^2} \right]_{z_-} \right) k_\alpha, \tag{4.22}$$

Stellarator-symmetry allows for further reduction by noting that $\nabla\psi \cdot \nabla\alpha$ is an odd function along the field line, i.e. $[\nabla\psi \cdot \nabla\alpha]_{z_+} = -[\nabla\psi \cdot \nabla\alpha]_{z_-}$:

$$k_\psi^{shift} = 2 \left(\frac{[\nabla\psi \cdot \nabla\alpha]_{z_-}}{|\nabla\psi|^2} \right) k_\alpha = -2 \left(\frac{[\nabla\psi \cdot \nabla\alpha]_{z_+}}{|\nabla\psi|^2} \right) k_\alpha \tag{4.23}$$

Equation (4.23) is our new boundary condition. We note here that the quantities $\nabla\psi \cdot \nabla\alpha$ and $|\nabla\psi|^2$ determining k_ψ^{shift} are already computed in every stellarator gyrokinetic code workflow, as they are needed to compute $k_\perp^2$. Thus, there are no new geometric quantities that need to be computed in order to use the new boundary condition. It is also possible to derive the same result if the orthonormal condition is relaxed, and $(\hat{\boldsymbol{b}}, \hat{\boldsymbol{e}}_u, \hat{\boldsymbol{e}}_v)$ are taken to be *orthogonal* vectors.

For completeness, using the change of variables employed in (4.9) for the conventional boundary condition, (4.23) can be written in terms of x and y to yield

$$k_x^{shift} = 2 \left(\frac{[\nabla x \cdot \nabla y]_{z_-}}{|\nabla x|^2} \right) k_y = -2 \left(\frac{[\nabla x \cdot \nabla y]_{z_+}}{|\nabla x|^2} \right) k_y. \tag{4.24}$$

Finally, one can directly derive a quantization condition on the aspect ratio of the

simulation domain from (4.24) to be

$$\frac{L_x}{L_y} = \frac{J}{2}\frac{|\nabla x|^2}{|\nabla x \cdot \nabla y|}, \tag{4.25}$$

where J is a nonzero integer.

4.4.4 Perpendicular Wavenumber Continuity

Our formulation manifestly produces perpendicular wavenumbers

$$k_\perp = \left(k_\alpha^2 |\nabla\alpha|^2 + 2k_\alpha k_\psi \nabla\alpha \cdot \nabla\psi + k_\psi^2 |\nabla\psi|^2\right)^{1/2}, \tag{4.26}$$

that are continuous when passing through the boundary. In contrast, when the conventional "twist-and-shift" condition is used, $k_\perp$ is continuous only in the case of axisymmetry with an integer number of poloidal turns. For the $\alpha = 0$ flux tube in W7-X running from $[-\pi, \pi]$, Figure 4.2 shows a plot of $k_\perp$ over a connected domain by linking the flux tube to itself at the boundaries $\pm\pi, \pm 3\pi$ using the boundary conditions in question. If the conventional boundary condition [7] is applied instead of (4.23), then $k_\perp$ becomes discontinuous at the boundary (as one can see in the blue curve) which may cause undesirable numerical behavior. For Figure 4.2 and results that follow, we have chosen to normalize wavenumbers to the ion gyroradius, defined to be $\rho_i \equiv v_{ti}/\Omega$, where $v_{ti} \equiv \sqrt{T_i/m_i}$ is the ion thermal velocity, and Ω is the ion cyclotron frequency.

This continuity might be important because $k_\perp$ appears in the argument of

the Bessel functions in the gyrokinetic equation, making it noteworthy that $k_\perp$ increases faster with $|\theta|$ for the new boundary condition than for the old condition. This behavior is expected, since in the old approach, k_ψ is increased by an amount proportional to the small global shear, while in the new approach, k_ψ increases by an amount related to the local shear, which is generally larger. A large rise in $k_\perp$ with $|\theta|$ is desirable because it leads to localization of the eigenfunctions and turbulence within a small number of linked domains (since the Bessel functions cause the plasma response to decrease with $k_\perp$), leading to less expensive simulations.

Alongside the plots of $k_\perp$ using the two boundary conditions in Figure 4.2, we have also plotted $k_\perp$ in what we call the 'extended domain', meaning a very long flux tube with no linkages across the tube ends. (For this figure, the extended domain represents a tube of length $\geq 10\pi$.) While the extended domain represents the true magnetic geometry, its length makes nonlinear simulations impractical, so the workaround is to use shorter domains that can be connected. The behavior of $k_\perp$ past the first connection in a linked domain will generally be different than in the extended domain, regardless of the boundary condition choice.

4.4.5 Axisymmetric Limit

Let us now show that (4.23) reduces to Beer's condition in axisymmetric geometry if the flux tube extends an integer number of times poloidally around the

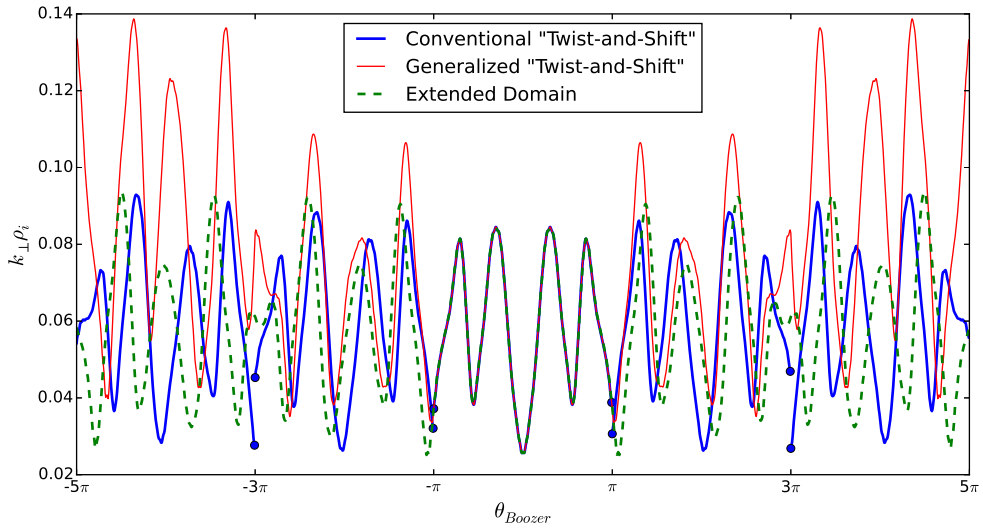

Figure 4.2: (Color online) The new boundary condition ensures that $k_\perp$ remains continuous across linked domains, whereas the condition of [7] generally leads to discontinuities at the boundaries in a stellarator. Shown here are five linked domains of parallel length $\Delta\theta = 2\pi$, so the boundaries are at $\theta = \pm\pi$ and $\pm 3\pi$. The calculation here is for the $\alpha = 0$ field line on the x/a=0.357 surface in W7-X, considering $k_x\rho_i = 0$ and $k_y\rho_i = 0.05$ in the central domain. (For other choices of $k_y\rho_i$, the curves in the figure would merely be scaled by a constant.)

torus. By using the definition of α in (4.6), we can write

$$\nabla\psi \cdot \nabla\alpha = \nabla\psi \cdot \nabla\zeta - q\nabla\psi \cdot \nabla\theta - \theta\frac{\mathrm{d}q}{\mathrm{d}\psi}|\nabla\psi|^2. \tag{4.27}$$

Due to axisymmetry, $\nabla\psi \cdot \nabla\zeta$ is the same at the forward and backward end of the flux tube. The same is true of $\nabla\psi \cdot \nabla\theta$. Therefore, these terms cancel when (4.27) is substituted into (4.22). The remaining term gives

$$\left(k_\psi^{shift}\right)_{AS} \equiv [k_\psi]_{+\pi N} - [k_\psi]_{-\pi N} = \frac{\mathrm{d}q}{\mathrm{d}\psi}[\theta(z_+) - \theta(z_-)]k_\alpha, \tag{4.28}$$

$$\left(k_\psi^{shift}\right)_{AS} = k'_\psi = 2\pi N\frac{\mathrm{d}q}{\mathrm{d}\psi}k_\alpha, \tag{4.29}$$

where N is the number of times the flux tube extends poloidally around the torus. This is equivalent to the conventional "twist-and-shift" boundary condition in (4.7).

Continuity of $k_\perp$ can be shown in axisymmetry by noting $\nabla\psi$ is the same at z_+ and z_-, and from (4.6),

$$[\nabla\alpha]_{z_+} = [\nabla\alpha]_{z_-} - 2\pi N \frac{dq}{d\psi}\nabla\psi. \tag{4.30}$$

4.5 Selecting the Flux Tube Length Using Local Magnetic Shear

In a stellarator-symmetric flux tube, $\nabla\psi\cdot\nabla\alpha$ is an odd function of z. This can be seen from the fact that $\nabla\psi\cdot\nabla\alpha$ flips sign under the replacements $(\theta \to -\theta,\ \zeta \to -\zeta)$, where now θ and ζ are any straight-field-line coordinates satisfying $\alpha = \zeta - q\theta$. This is why all terms in (4.27) generally add and allowed for the last step in producing (4.23). In particular, in a stellarator it is generally not valid to drop the $\nabla\psi\cdot\nabla\zeta$ and $\nabla\psi\cdot\nabla\theta$ terms, even if the flux tube goes an integer number of times around the torus poloidally.

As discussed in [46], the local magnetic shear is

$$S = \boldsymbol{B}\cdot\nabla\left(\frac{\nabla\psi\cdot\nabla\alpha}{|\nabla\psi|^2}\right). \tag{4.31}$$

Therefore, the shift to k_ψ in (4.23) represents the integral of the local shear along the flux tube, which makes this new boundary condition advantageous for a couple of reasons. First, k_ψ^{shift} is no longer solely dependent on a potentially restrictive constant global shear but rather a locally varying function. It is important to

note, however, that while $\hat{s}$ is not explicit in (4.23), the global shear information is contained in the geometric quantities $\nabla\psi \cdot \nabla\alpha$ and $|\nabla\psi|^2$. Second, the fact that k_ψ^{shift} depends on a function of z evaluated at flux tube ends means that the length of the tube can be chosen such that an optimal k_ψ^{shift} is obtained.

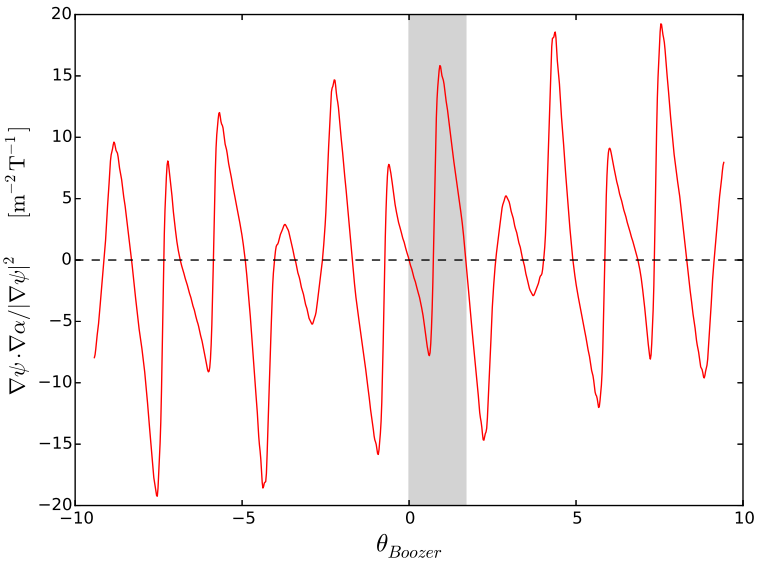

Figure 4.3: Color online) The angle between the two field-aligned coordinate directions $\nabla\psi$ and $\nabla\alpha$, along the same field line. The large departures from orthogonality $(\pi/2)$ indicate that the local shear is significant even though the global shear is small, $\hat{s} = -0.019$. Here we take $2\pi\psi$ to be the poloidal flux. The shaded region indicates the length of the tube in Figure 4.6.

Figure 4.4 gives some insight into how the local shear and the simulation domain length are related. In Figure 4.3 we have plotted the integrated local shear, which defines k_ψ^{shift} up to a constant. This curve shows that the local shear has an oscillatory form in this geometry and in fact changes sign a number of times over this domain. These frequent sign changes in the local shear, and by definition k_ψ^{shift}, provide the opportunity to make $k_\psi^{shift} = 0$ if the flux tube length is chosen such that the ends lie where the local shear vanishes. If k_ψ^{shift} vanishes this implies that $[k_\psi]_{z_+} = [k_\psi]_{z_-}$, which in combination with (4.20) assures that the parallel bound-

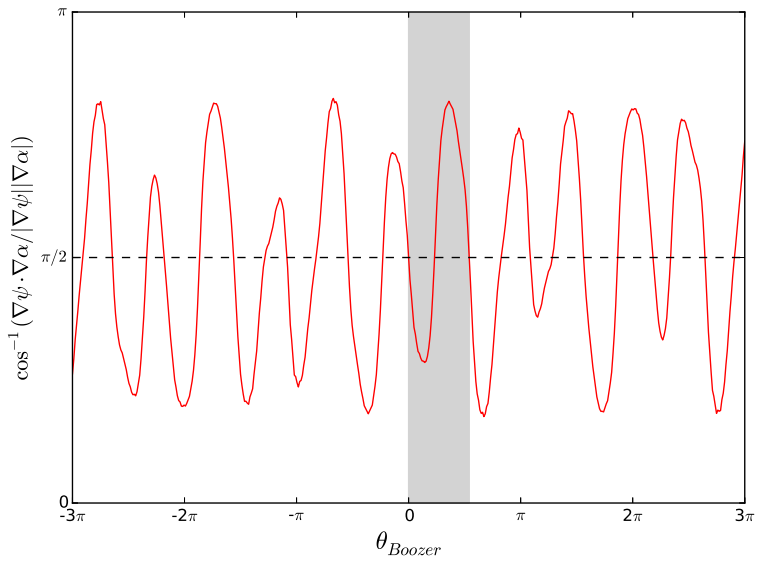

Figure 4.4: (Color online) Variation of $\nabla\psi\cdot\nabla\alpha/|\nabla\psi|^2$, the quantity arising in the new boundary condition, in the W7-X standard configuration for the surface with normalized radius x/a=0.357 and the field line $\alpha = 0$. The shaded region indicates the length of the tube in Figure 4.6.

ary condition becomes periodic. Along with improving computational efficiency, periodic boundary conditions remove the quantization on the aspect ratio of the simulation domain (4.25). Now, while in principle one could decide to minimize the length of the tube with this condition in mind by choosing flux tube ends to lie at the first zero of the local shear ($\sim \pi/3$ in this case), this boundary condition only allows for periodicity and *does not imply correct results for an arbitrarily small flux tube*, as will be discussed in the following sections.

This type of behavior that allows for periodic boundary conditions is a result of the low global shear, which permits the oscillations about zero to dominate the functional form of the integrated local shear. While small, the effect of the global shear is visible in the slight linear trend of the function. Conversely, the linear trend for a geometry with significant global shear would dominate, and reduce (or perhaps

eliminate) any zero crossings in the integrated local shear. The zero crossings that remain, if any, would then be concentrated near the center and periodic boundary conditions would be limited to shorter flux tubes. Figure 4.5 examines this effect by comparing the curve from Figure 4.3 to the same quantity for larger global shear devices, namely LHD [72] and NCSX [109]. These complications don't preclude one from using the boundary condition in high global shear geometries, but the value of optimizing the tube length is more limited.

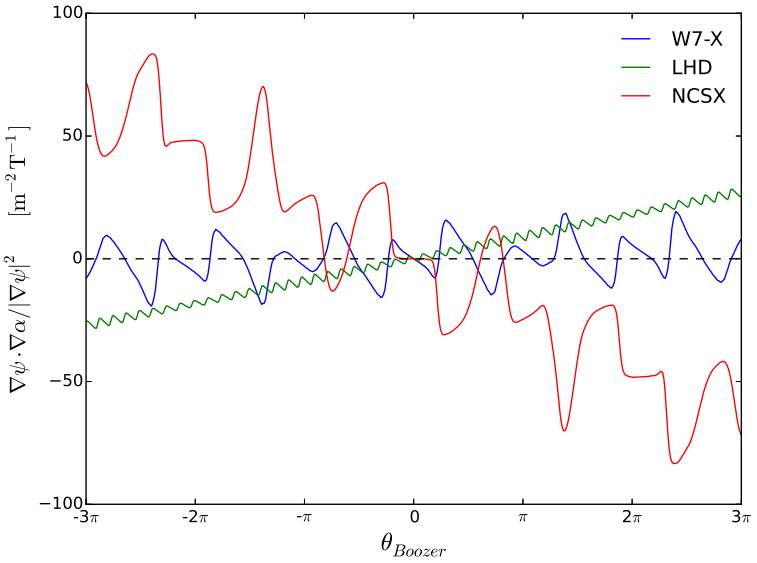

Figure 4.5: (Color online) The quantity $\nabla\psi \cdot \nabla\alpha/|\nabla\psi|^2$ over the domain $[-3\pi, 3\pi]$ for the standard equilibrium configurations of W7-X, NCSX, and LHD. Each curve denotes the $\alpha = 0$ field line at a radial position of $x/a \approx 0.36$.

To get more of a sense for what is happening physically, Figures 4.6 and 4.7 illustrate how the local shear influences the overall shape of a flux tube. Figure 4.6 shows the $\alpha = 0$ W7-X flux tube at $x/a = 0.357$ extending from the outboard midplane at $\theta = 0$ (bean cross section) to $\theta = 1.70$ and is meant to coincide with the shaded regions of Figure 4.3. The difference between high and low global shear cases is clear as the 3-dimensional shape of the flux tube constitutes a twisting-

and-untwisting of the domain, in contrast to the near monotonic twisting of high global shear flux tube. Since the condition for periodic parallel boundary conditions require $\nabla\psi \cdot \nabla\alpha = 0$ (implying a rectangular perpendicular cross section), this twisting-untwisting characteristic of the flux tube is what affords many potential lengths for which periodic boundary conditions are possible.

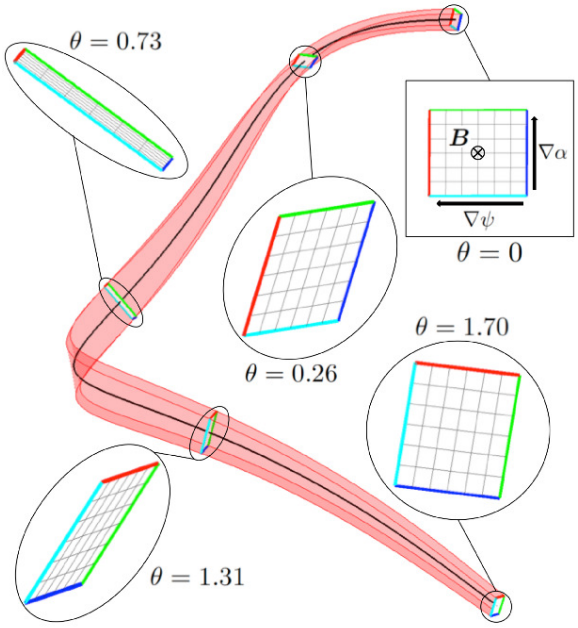

Figure 4.6: (Color online) 3D visualization of the $\alpha = 0$ W7-X flux tube in the field-line-following coordinates (ψ, α, z) at $x/a = 0.357$. The relevant field line is centered at (ψ_0, α_0) in the domain (black line), with the perpendicular boundaries located at $\psi = \psi_0 \pm \Delta\psi$ and $\alpha = \alpha_0 \pm \Delta\alpha$. The color scheme is as follows: $\psi = \psi_0 + \Delta\psi$ (red line); $\psi = \psi_0 - \Delta\psi$ (blue line); $\alpha = \alpha_0 + \Delta\alpha$ (green line); $\alpha = \alpha_0 - \Delta\alpha$ (cyan line). All cross sections are projected along the magnetic field at the given θ location. The $\theta = 0, 0.73, 1.70$ positions correspond to a vanishing of the integrated local shear (see Figure 4.3), resulting in a rectangular cross section.

Figure 4.7: (Color online) 3D visualization of the $\alpha = 0$ flux tube domain in real space superimposed on the flux *surface* at $x/a = 0.357$. (The extent of the tube in ψ,α was set for visualization purposes). This is the same flux tube from Figure 4.6, but shown from $\theta = [-0.73, 0.73]$, where $\nabla\psi \cdot \nabla\alpha = 0$ and the ends of the tube have a rectangular perpendicular cross section.

4.5.1 Magnetic Drift Continuity

The magnetic drift term in the gyrokinetic equation (Appendix A), $\boldsymbol{v}_m \cdot \nabla_\perp h$, is continuous in axisymmetry with the standard twist-and-shift condition, but the term is generally discontinuous across the parallel boundary of a flux tube in a stellarator, for both the old and new boundary conditions. It is not obvious that continuity of this term matters, since discontinuity of coefficients in a PDE does not necessarily cause the solution to be discontinuous or otherwise pathological. To investigate whether it makes a difference, it is possible to make the magnetic drift term continuous in the steps that follow. We begin by taking the ∇B-drift part of the magnetic drift term in the gyrokinetic equation (A.1), noting that $\nabla_\perp = i(k_\psi \nabla\psi + k_\alpha \nabla\alpha)$:

$$\boldsymbol{v}_m \cdot \nabla_\perp h \propto \boldsymbol{B} \times \nabla B \cdot (k_\psi \nabla\psi + k_\alpha \nabla\alpha) h. \tag{4.32}$$

Setting $\boldsymbol{v}_m \cdot \nabla_\perp h$ equal at both ends of the tube, we note that the $\nabla\alpha$- and $\nabla\psi$-components of the ∇B-drift are even and odd functions, respectively, in stellarator-symmetric flux tubes. Applying this fact along with (4.23), the resulting condition can be derived:

$$\left([k_\psi]_{z_+} + \left[\frac{\nabla\psi \cdot \nabla\alpha}{|\nabla\psi|^2} \right]_{z_+} k_\alpha \right) [\boldsymbol{B} \times \nabla B \cdot \nabla\psi]_{z_+} = 0. \tag{4.33}$$

This condition can only be satisfied for all k_ψ and k_α in the simulation when $[\boldsymbol{B} \times \nabla B \cdot \nabla\psi]_{z_+} = 0$. In other words, the only way to make the magnetic drifts continuous across the parallel boundary is to choose the flux tube length such that the radial component of the ∇B-drift vanishes at the ends. The argument here is identical for the curvature drift as well, since $\boldsymbol{B} \times \nabla B \cdot \nabla\psi = 0$ is equivalent to $\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla\psi = 0$, at any β (normalized pressure). Therefore, both the ∇B and curvature drifts become continuous at the same tube length.

Similar to the quantity $\nabla\psi \cdot \nabla\alpha$ discussed above, $\boldsymbol{B} \times \nabla B \cdot \nabla\psi$ varies significantly along a field line, and in fact has many zero-crossings regardless of the global magnetic shear, which is clear from its form in the s-alpha model, $(\boldsymbol{B} \times \nabla B \cdot \nabla\psi)_{s-\alpha} \approx (2a/R)\,\hat{s}\sin(z)$. Unfortunately, locations where $\nabla\psi \cdot \nabla\alpha$ and $\boldsymbol{B} \times \nabla B \cdot \nabla\psi$ vanish do not coincide, meaning magnetic drift continuity cannot be accompanied by appropriately enforced periodic boundary conditions. However, as one can see in Figure 4.8, there are numerous locations along a field line where the aspect ratio quantization condition (solid blue curve) approaches unity at the same time as $\boldsymbol{B} \times \nabla B \cdot \nabla\psi = 0$ occurs (vertical dashed lines). The solid blue curve demonstrates how the length of

the flux tube affects the required aspect ratio, with the vertical dashed lines representing locations where $\boldsymbol{B} \times \nabla B \cdot \nabla \psi = 0$. The effects of magnetic drift continuity in simulations are explored in the following sections.

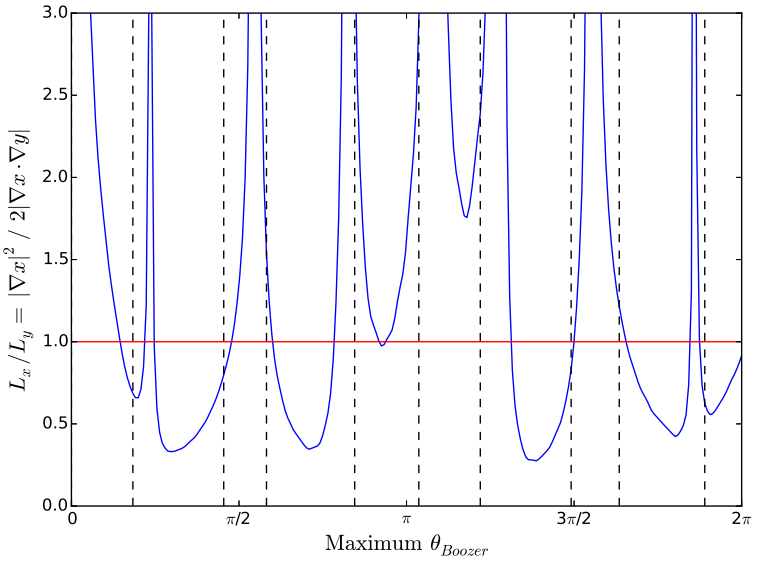

Figure 4.8: (Color online) The solid blue curve shows the quantized domain aspect ratio as a function of the flux tube's maximum θ_{Boozer}, for $J = 1$. (The flux tube length is twice the maximum θ_{Boozer}.) The horizontal red line represents the ideal $L_x/L_y = 1$ case. The dashed vertical lines correspond to flux tube lengths for which $\boldsymbol{B} \times \nabla B \cdot \nabla \psi = 0$, so the magnetic drift term is continuous.

4.6 Numerical Results

Many questions related to the boundary condition are generic with respect the representation of the distribution function. The majority of simulations we present below used the GPU-based gyrofluid code GryfX [49] (Appendix D presents some information on techniques for testing and portability of GPU-based codes). While GryfX has the option to employ a hybrid approach to simulate zonal flow dynamics with a gyrokinetic model, we have chosen to use GryfX in a pure gyrofluid config-

uration, in which all modes are evolved using the 4+2 set of gyrofluid equations in [9]. Compared with any comprehensive gyrokinetic model, a gyrofluid model is very inexpensive. The speedup that is achieved by using GryfX allows for simulations with extremely large N_x (the number of grid points in the x direction) values which would otherwise be computationally impractical in gyrokinetics. This in turn facilitates a more complete survey of the boundary condition issues.

4.6.1 Linear Convergence Results

We consider linear problems first. Linear flux tube stability analyses have been performed in W7-X geometry in the collisionless, electrostatic, adiabatic electron limit. All simulations use $\alpha = 0$ (bean cross section) flux tubes, with geometric information calculated by applying the GIST code [106] to a VMEC equilibrium. Each flux tube is located at the radial position $x/a = 0.357$, and unless otherwise specified, references to simulated perpendicular wavenumbers are normalized to ρ_i.

4.6.1.1 Growth Rate Convergence

We used the gyrofluid code GryfX [49] to investigate growth rate convergence with respect to both the number of simulated radial modes and length of the flux tube for various boundary condition choices. We assume $T_i/T_e = 1$ and equilibrium scale lengths of $a/L_T = 3.0$ and $a/L_n = 0.0$. One boundary condition considered is the conventional "Twist-and-Shift" condition; in this case the flux tube length is taken to be exactly an integer or half-integer number of poloidal turns. A second

option is the new boundary condition, applied to a flux tube with length chosen so $[\boldsymbol{B} \times \nabla B \cdot \nabla \psi]_{z_\pm} = 0$ ('Continuous Magnetic Drifts'). A third option, which we will call 'Exact periodicity', is periodicity in z imposed for a flux tube with length chosen so $[\nabla \psi \cdot \nabla \alpha]_{z_\pm} = 0$, consistent with the new boundary condition. The fourth option, which we call 'Forced Periodicity', is to impose periodicity in z for a tube length at which there is no rigorous analytic justification for doing so, since $[\nabla \psi \cdot \nabla \alpha]_{z_\pm} \neq 0$. In this case we choose the flux tube length to be exactly an integer or half-integer number of poloidal turns.

The first convergence study examines how the growth rates for each boundary condition choice change as a function of N_x for two binormal wavenumbers $k_y = 0.2, 0.5$. Each flux tube has been chosen to be $\sim$1 poloidal turn in length. The N_x we refer to in this chapter is defined to be the number of *aliased* radial modes, where the actual number of simulated (dealiased) radial modes is $\sim 2/3\, N_x$. These simulated radial wavenumbers are integer multiples of the minimum radial wavenumber, defined by $k_\psi^{\min} \equiv 1/L_\psi$, where the modes are connected via k_ψ^{shift}.

In Figure 4.9, results make clear that regardless of the chosen boundary condition, N_x has a very minimal effect on the calculated linear growth rate, and moreover, for $N_x \geq 4$ the growth rate has reasonably converged for both k_y values. This leads one to think that only a few connected domains (or in some cases only a single k_x value) are necessary to reproduce the extended domain result (i.e. the result in an extremely long flux tube), the eigenfunction of which is displayed in Figure 4.10 for the $(k_x, k_y) = (0.0, 0.2)$ mode.

In order to better understand how changing N_x and the flux tube length affect

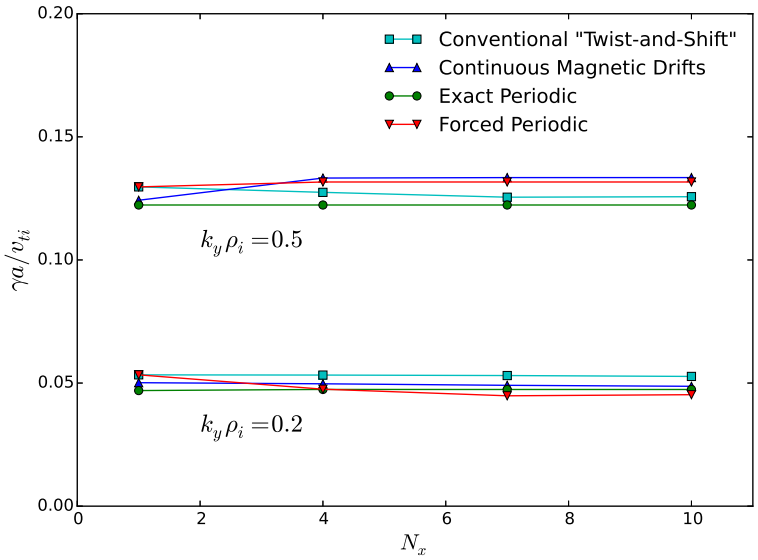

Figure 4.9: (Color online) - Growth rate for ∼1 poloidal turn flux tubes as a function of the number of simulated radial modes. The lengths of the flux tubes for each boundary condition choice: Conventional "Twist-and-Shift"/Forced Periodic $[-\pi, \pi]$, Exact Periodic $[-1.086\pi, 1.086\pi]$, Continuous Magnetic Drifts $[-1.045\pi, 1.045\pi]$.

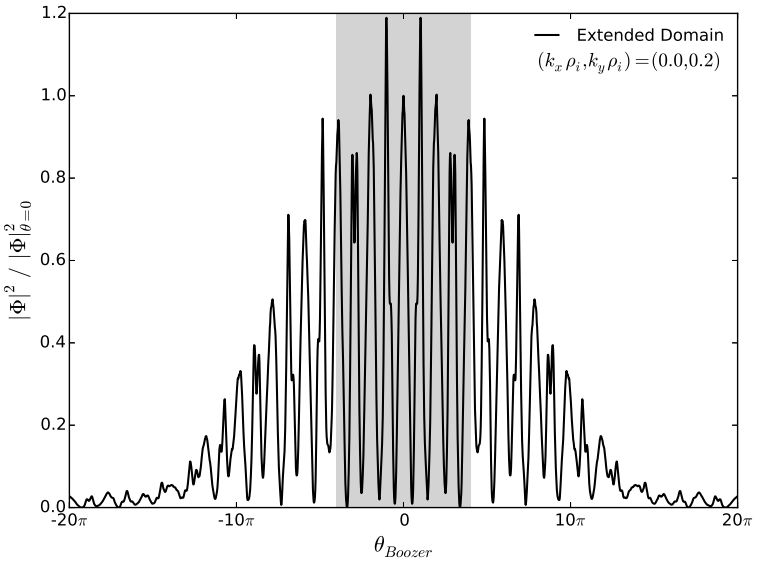

Figure 4.10: Eigenfunction for the $(k_x, k_y) = (0.0, 0.2)$ mode over the extended domain $[-20\pi, 20\pi]$. The shaded region indicates the extent of the plots in Figure 4.12.

our ability to reproduce the extended domain solution, we compare its eigenfunction to the eigenfunctions generated in shorter flux tubes employing the various boundary

condition options. To visualize this, we zoom in on the shaded region of Figure 4.10 and superimpose plots of eigenfunctions of the connected modes for flux tubes of length ∼0.5 poloidal turns ($N_x = 7$) and ∼1 poloidal turn ($N_x = 4$) in Figures 4.11 and 4.12, respectively, for each boundary condition choice.

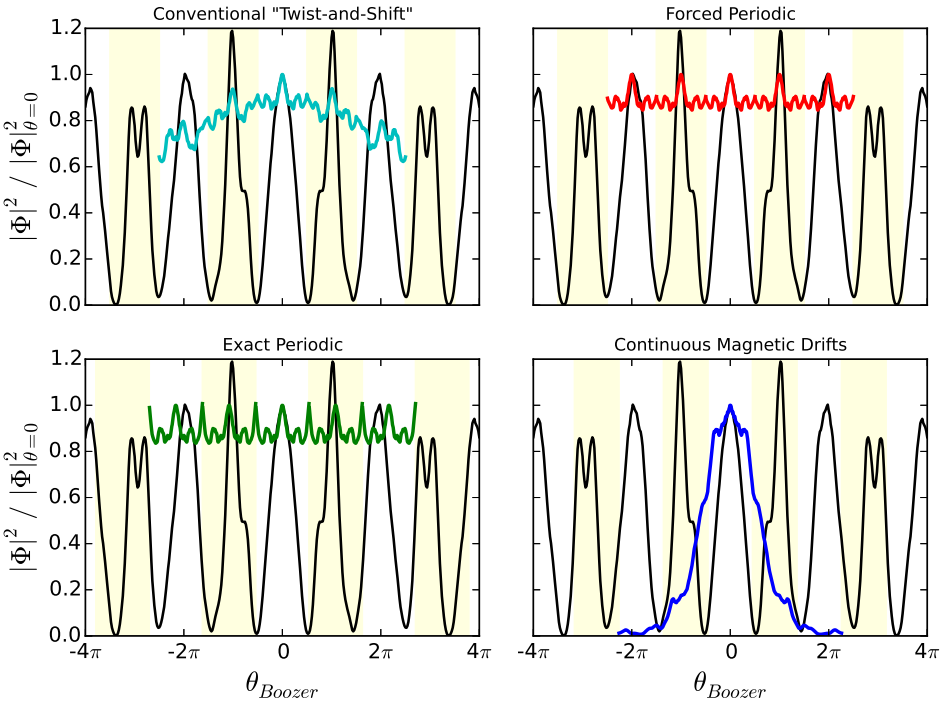

Figure 4.11: (Color Online) Using the four boundary conditions from Figure 4.9 for flux tubes of ∼0.5 poloidal turn, eigenfunctions for the connected regions ($N_x = 7$) are plotted over a portion of the extended domain eigenfunction (black line). The center region in each plot corresponds to the $(k_x, k_y) = (0.0, 0.2)$ mode, where the adjacent shaded regions have $(k_x, k_y) = (k_x^{shift}, 0.2)$, where k_x^{shift} depends on the boundary condition choice. The lengths of the flux tubes for each boundary condition choice: Conventional "Twist-and-Shift"/Forced Periodic $[-\pi/2, \pi/2]$, Exact Periodic $[-0.54\pi, 0.54\pi]$, Continuous Magnetic Drifts $[-0.45\pi, 0.45\pi]$.

The central region of each plot in Figures 4.11 and 4.12 correspond to the mode $(k_x, k_y) = (0.0, 0.2)$, which is connected at each end (in the yellow-shaded region) to the eigenfunctions of modes with the same k_y and a different k_x determined by the k_x^{shift} calculated from each boundary condition. As N_x increases, more modes with

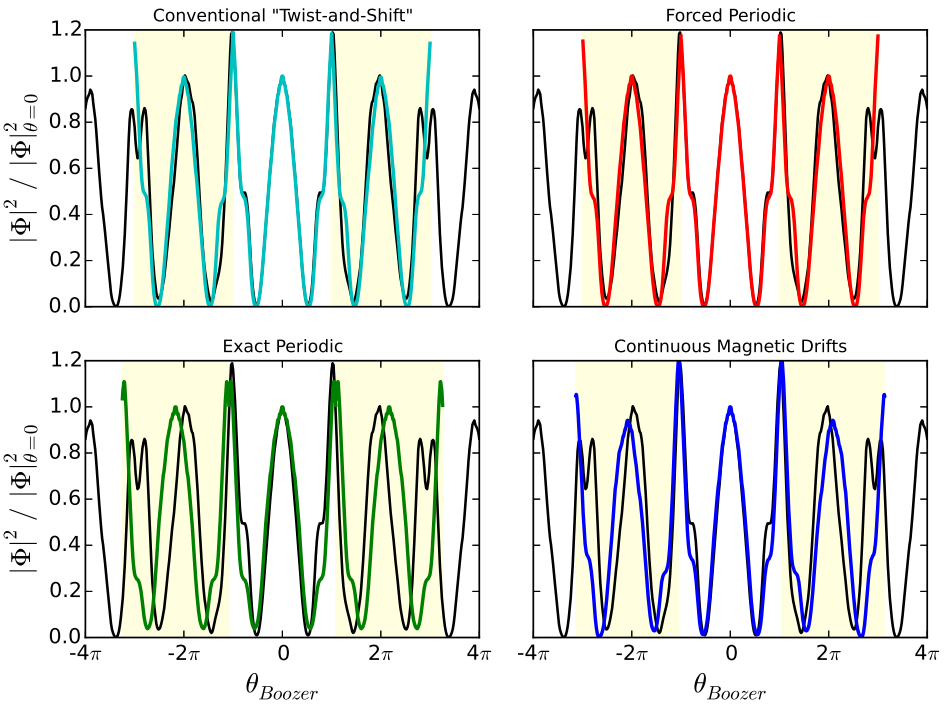

Figure 4.12: (Color online) Using the four boundary conditions from Figure 4.9 for flux tubes of ∼1 poloidal turn, eigenfunctions for the connected regions ($N_x = 4$) are plotted over a portion of the extended domain eigenfunction (black line). Exact flux tube lengths are given in the caption of Figure 4.9.

the same k_y are linked, with k_x changing by integer multiples of k_x^{shift}.

By comparing the various boundary conditions (colored lines) to the solid black curve of the extended domain in Figure 4.11 for the ∼0.5 poloidal turn flux tube, it is apparent that none of the boundary condition options reliably model the form of the extended domain eigenfunction. Furthermore, the connected eigenfunction of the Continuous Magnetic Drifts case is distinctly more narrow than the other options. This seemingly peculiar structure is based on how the k_x dependence of $k_\perp$ changes at connection points based on k_x^{shift}. For comparatively larger k_x^{shift}, $k_\perp$ will increase accordingly at each connection (see Figure 4.2), leading to more localized eigenfunctions. Moreover, shorter flux tubes will have more connections per unit

length, causing this increase in $k_\perp$ to occur more frequently, introducing further localization. This larger k_x^{shift} (and slightly shorter flux tube) in the Continuous Magnetic Drifts case is the cause of its comparatively narrow eigenfunction, relative to the smaller shift resulting from the conventional method, and $k_x^{shift} = 0$ for the two periodic boundary conditions. Such disagreement among the boundary condition choices in addition to the poor reconstruction of the extended domain eigenfunction might lead one to expect that growth rate results would be inaccurate, with the continuous magnetic drift result being the biggest outlier.

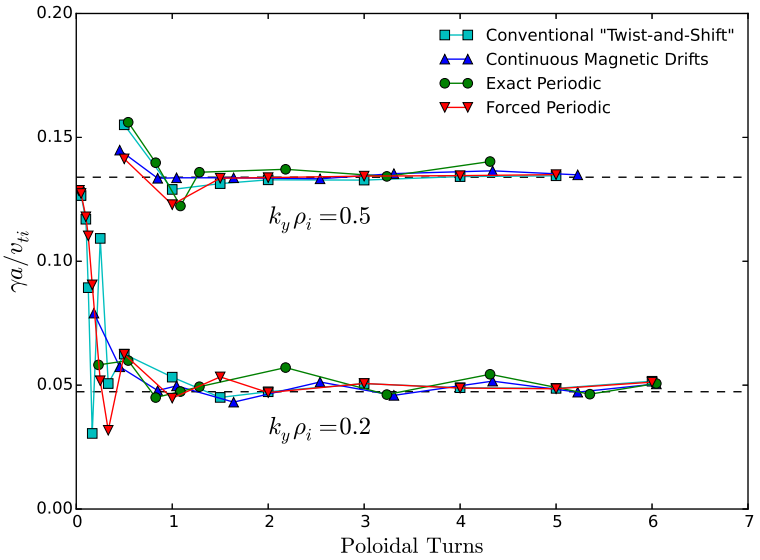

Figure 4.13: (Color online) Linear growth rates as a function of flux tube length for the binormal wavenumbers $k_y = 0.2, 0.5$, using the boundary conditions described in Section 4.6.1.1. The dashed lines represent the *true* results, obtained from the extended domain simulations. In each simulation, the number of θ grid points, N_z, is scaled proportionally with flux tube length, maintaining a fixed θ resolution for each run.

However, in Figure 4.13, which shows the growth rate as a function of flux tube length, the relevant data points at ∼0.5 poloidal turns do not support this line of thinking. The growth rates from each boundary condition are closely clustered

and are within $\sim 30\%$ of the true result.

This same comparison of connected eigenfunctions has been done for flux tubes of $\sim$1 poloidal turn in Figure 4.12. Unlike $\sim$0.5 poloidal turn flux tubes, agreement with the functional form of the extended domain eigenfunction is quite good, and the boundary condition options have only minor differences between one another. As one might expect, the growth rate is closer to the true result from Figure 4.13, but the extreme contrast between Figures 4.11 and 4.12 make it surprising that the growth rates with $\sim$0.5 poloidal turn flux tubes even come close to the true result. An interpretation of this result is given in Appendix B.

The conclusion from the results in this section is that the parallel boundary condition has a seemingly insignificant effect on linear growth rates, as long as the flux tube length exceeds some minimum value. However, it remains to be shown in Section 4.6.3 how these findings translate to nonlinear simulations when the modes become coupled.

4.6.1.2 Linear Zonal Flow Response

Due to the importance of zonal flows in the saturation of turbulence, representing the response as accurately as possible is advantageous in simulations. For this reason, understanding the behavior of the dynamic zonal flow response and Rosenbluth-Hinton (RH) [85] residual values as the flux tube length is varied is desirable. It is important to note here that $k_y = 0$ modes are self-periodic for both the conventional and generalized twist-and-shift boundary conditions. Hence, the

choice of boundary condition affects $k_y = 0$ modes only through the tube length, with no effect on the linkages at the parallel boundary. For calculations in this section, the gyrokinetic code GS2 was used in lieu of GryfX for the purpose of avoiding the closure approximations of the gyrofluid set of equations, which have historically had difficulties in matching zonal flow responses well [25]. The GS2 normalizations are slightly different from GryfX, so for this section we normalize results to the GS2 ion gyroradius, $\rho_{i,\mathrm{GS2}} = \sqrt{2}\rho_i$, and thermal velocity, $v_{ti,GS2} = \sqrt{2}v_{ti}$.

For the figures in this section, the results are produced from flux tube lengths chosen such that the "Continuous Magnetic Drifts" (blue) and "Exact Periodic" (red) boundary condition options are applicable, which correspond to $[\mathbf{B} \times \nabla B \cdot \nabla\psi]_{z_\pm} = 0$ and $[\nabla\psi \cdot \nabla\alpha]_{z_\pm} = 0$, respectively. The flux tubes where $[\mathbf{B} \times \nabla B \cdot \nabla\psi]_{z_\pm} = 0$ are of particular interest, as linear studies [69, 71] reveal a dependence on the radial bounce-averaged magnetic drift of the zonal flow residual in stellarators, a quantity that vanishes in axisymmetry. The bounce-average of $\mathbf{B} \times \nabla B \cdot \nabla\psi$ will thus be performed between two points where this term vanishes, making it possible that such flux tube lengths could result in unique zonal flow behavior compared with other tube lengths.

Numerical studies have shown that in stellarator geometries, the dynamic response of zonal flows has a central role in the regulation of turbulent transport [108]. In Figure 4.14, this linear response is plotted for $k_x\rho_{i,GS2} = 0.15, 0.4$ at flux tube lengths of $\sim$1 poloidal turn. For both wavenumbers in Figure 4.14, the response is nearly identical for both flux tube types. This leads to the expectation that the nonlinear effect of the zonal flows will likely be quite similar for both boundary

condition options, which is confirmed by the results of Section 4.6.3.

The long-time zonal flow behavior was also studied, where the RH residual value has been calculated for $k_x\rho_{i,\mathrm{GS2}} = 0.15, 0.4$ in a variety of flux tube lengths, in Figure 4.15. For calculations done with flux tubes less than a full poloidal turn, there is an apparent downward trend in the residual for both flux tube types as the length is increased. For longer flux tubes the residual values have an oscillatory behavior with an amplitude approaching some constant value as the length is increased. The results of Figure 4.15 demonstrate that although $[\mathbf{B} \times \nabla B \cdot \nabla\psi]_{z_\pm} = 0$ in the blue curve, it appears to have a minor effect relative to flux tubes where this quantity does not vanish at the ends of the domain.

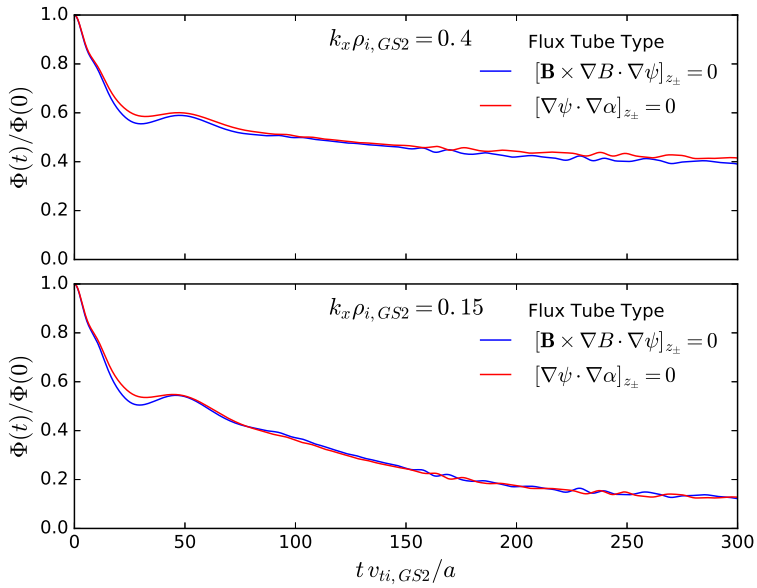

Figure 4.14: (Color online) Gyrokinetic (GS2) simulations of the linear zonal flow response for flux tubes lengths of ∼1 poloidal turn for radial wavenumbers $k_x\rho_{i,GS2} = 0.15, 0.4$, where $\rho_{i,GS2} = \sqrt{2}\rho_i$ and $v_{ti,GS2} = \sqrt{2}v_{ti}$. The blue and red curves for each radial wavenumber correspond to tube lengths satisfying the conditions $[\mathbf{B} \times \nabla B \cdot \nabla\psi]_{z_\pm} = 0$ or $[\nabla\psi \cdot \nabla\alpha]_{z_\pm} = 0$, respectively.

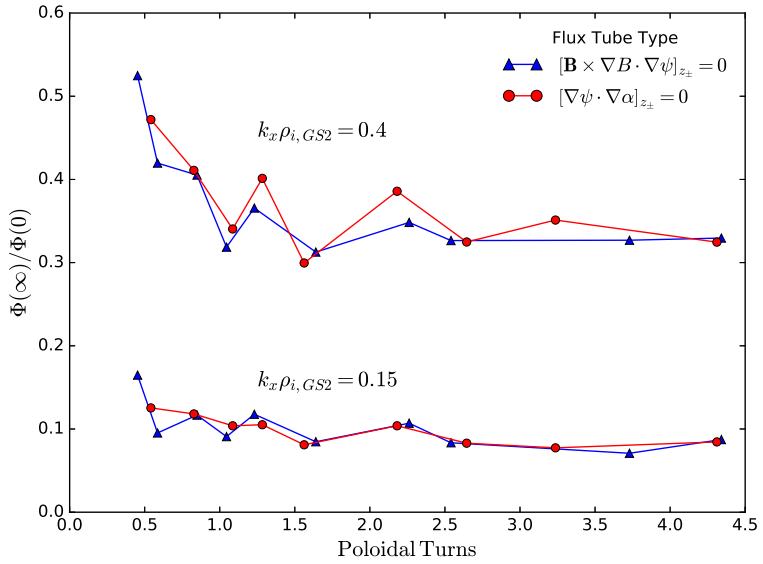

Figure 4.15: (Color online) Gyrokinetic (GS2) calculations of the Rosenbluth-Hinton residual (fit to an exponential decay model) as a function of flux tube length for radial wavenumbers $k_x\rho_{i,\mathrm{GS2}} = 0.15, 0.4$, where $\rho_{i,\mathrm{GS2}} = \sqrt{2}\rho_i$. The blue and red curves for each radial wavenumber correspond to tube lengths satisfying the conditions $[\mathbf{B}\times\nabla B\cdot\nabla\psi]_{z_\pm} = 0$ or $[\nabla\psi\cdot\nabla\alpha]_{z_\pm} = 0$, respectively.

4.6.2 Secondary Instability

We demonstrate in this section that there is a case, the evolution of a secondary instability, in which the discontinuity associated with an incorrectly applied boundary condition could have an effect on results.

The nonlinear generation of zonal flows in plasmas are due in part to a Kelvin-Helmholtz-like secondary instability [83] that develops from the primary Ion-Temperature-Gradient (ITG)-driven radial streamers. The primary ITG instability (characterized by $k_x = 0, k_y \neq 0$) is nonlinearly coupled to a $k_x \neq 0, k_y = 0$ mode through a three-wave interaction with another unstable ($k_x \neq 0, k_y \neq 0$) mode, sometimes referred to as the pump wave.

Returning to GryfX simulations, we address the behavior of the aforemen-

tioned modes involved in the generation of the secondary instability when subjected to flux tubes of ∼1 poloidal turn, employing the "Exact" and "Forced" periodic boundary conditions. For the simulation performed here, we begin with a short linear setup run to initialize the primary mode. The simulation is then nonlinearly restarted with a primary mode amplitude so large that the nonlinear term dominates the equation, allowing one to study the interactions among only the three aforementioned modes. The resulting eigenfunctions are plotted in Figure 4.16, using the finite wavenumbers $k_x = k_y = 0.2$.

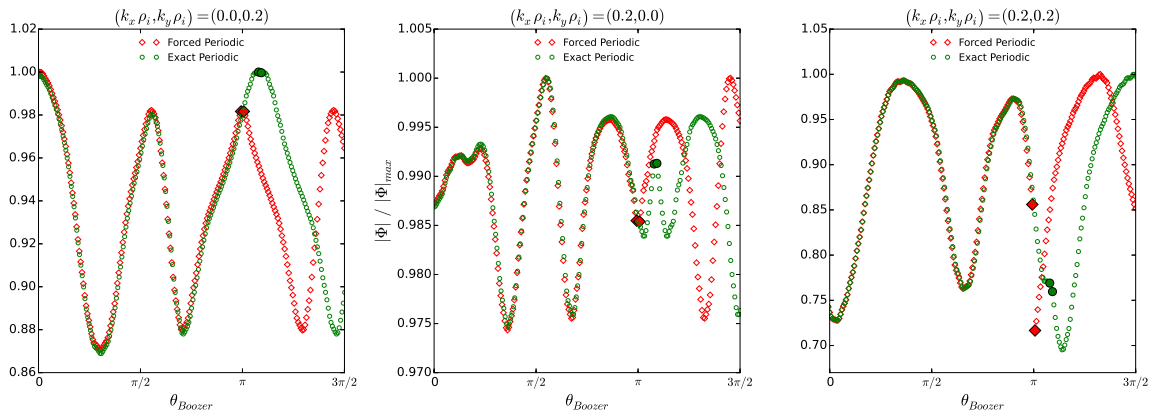

Figure 4.16: (Color online) Eigenfunctions for the primary mode (left), $k_y = 0$ mode (center), and pump wave (right) in ∼1 poloidal turn flux tubes. The larger, filled markers denote grid points at the ends of the domain, for the "Exact" (circles) and "Forced" (diamonds) periodic cases. Note the suppressed zeros and different vertical axes in each figure.

One can see from the $k_y = 0$ (center) and primary mode (left) eigenfunctions that the particular boundary condition does not affect the continuity of these modes across the connected domain. However, the boundary condition appears to have a notable effect on the pump wave (right) in the form of a discontinuity in the eigenfunction across the connections at $\theta_{Boozer} \approx \pm\pi$. Such behavior arises through the discontinuity in $k_\perp$ between connected domains when using a boundary condition

option that is not strictly valid, such as incorrectly enforcing periodicity as we have done here.

This discontinuity in Figure 4.16.c can be understood by considering the dominant terms in the gyrokinetic equation (Appendix A) with k corresponding to the pump wave:

$$\frac{\partial h_k}{\partial t} + \sum_{k',k''} \frac{c}{B} \{\langle \phi_{k'} \rangle, h_{k''}\} + \dots \quad , \tag{4.34}$$

where the angled brackets $\langle ... \rangle$ denote a gyroaveraging operation performed at constant guiding center, and $\{\cdot, \cdot\}$ is the Poisson bracket. For simulations using forced periodicity, even if $\phi_{k'}$ and $h_{k''}$ are continuous, the discontinuity in $k_\perp$ will cause the Bessel function $J_0(k_\perp \rho_i)$ involved in $\langle ... \rangle$ to be discontinuous, causing h_k to develop a discontinuity. Thus, incorrectly enforcing periodicity, or otherwise improperly using conventional "twist-and-shift" in the parallel direction introduces errors. However, the consequences resulting from introducing these discontinuities are not well understood, and further study is warranted to quantify the full effect it may have on a given simulation.

4.6.3 Nonlinear Results

We now turn to discussion of the nonlinear behavior associated with the various boundary conditions, where unless otherwise stated, results pertain to W7-X geometry under the conditions stated in Section 4.6.1, with all simulations performed in the gyrofluid approximation.

As mentioned in Section 4.2, fully resolving the dominant k-space fluctuation region is required to correctly calculate heat flux values. The ability of a simulation to satisfy these resolution requirements is directly tied to the radial wavenumber shift of the parallel boundary condition. If a minimally resolved run requires simulating up to some particular k_ψ value, having a small k_ψ^{shift} (based on the boundary condition) clearly forces one to include a large number of simulated modes. The reverse situation holds true for a large k_ψ^{shift}.

Figure 4.17 presents the fluctuation spectrum for a one poloidal turn flux tube with $N_x = 96$ using the conventional "twist-and-shift" boundary condition alongside the spectrum for an unoptimized case of the new boundary condition. The contrast between the two figures shows unambiguously that a large portion of the fluctuation region is not captured with the conventional boundary condition for this number of radial modes. The spectrum found with the new boundary condition indicates a localized region of larger relative amplitude in the center, suggesting the simulation contains the most important fluctuations. On the other hand, the reduced k_ψ range in the conventional case does not allow for enough modes to sufficiently capture this region, and the fluctuation amplitudes become artificially large. Such a case requires one to include more radial modes, and a calculation of transport coefficients at this resolution leads to inaccurate results.

The difference in resolution capabilities between the boundary condition is presented in Figure 4.18 by comparing the saturated heat flux as a function of the number of simulated radial modes, for the boundary condition variations as described in Section 4.6.1.1. Each simulation in Figure 4.18 uses a flux tube that is

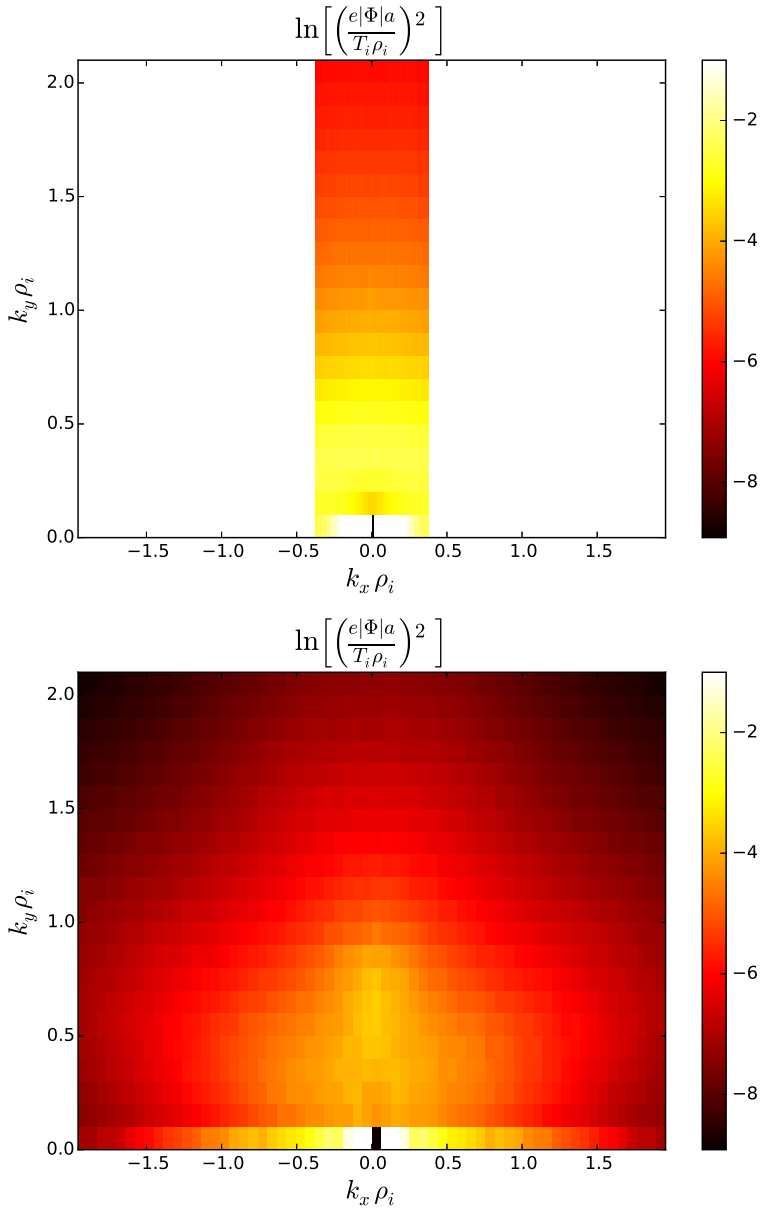

Figure 4.17: (Color online) - 2D fluctuation spectra with $N_x = 96$ using (top) conventional "twist-and-shift" covering $k_x = [-0.38, 0.38]$ and (bottom) generalized "twist-and-shift" covering $k_x = [-1.96, 1.96]$. The increased k_x range in (bottom) permits fluctuation localization in the domain, while artificially high fluctuations result (top) due to the lack of resolution.

$\sim$1 poloidal turn in length, with exact lengths given in the caption of Figure 4.12.

This figure shows a stark difference in how quickly the results converge with N_x to the correct heat flux (somewhere between 3.5-4.5), based on the chosen boundary

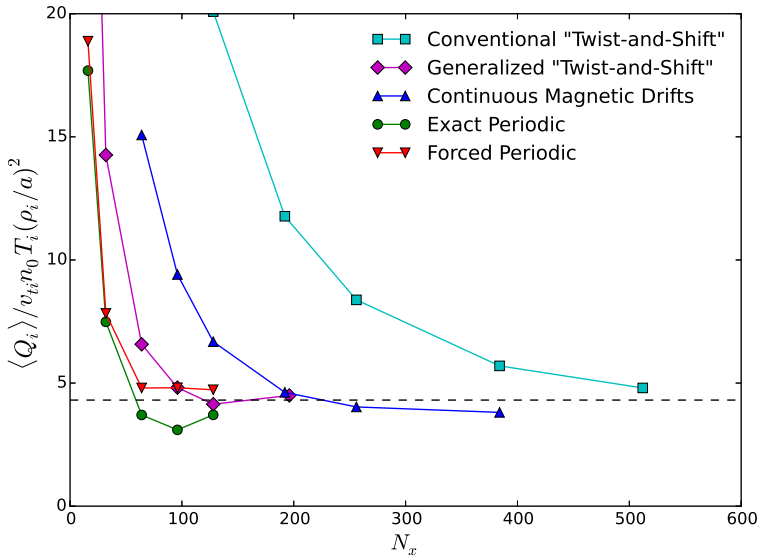

Figure 4.18: (Color online) Saturated heat flux in W7-X as a function of the radial resolution for various boundary condition choices using flux tubes of ∼1 poloidal turn. The dashed line is calculated as the average of the heat flux for the rightmost data point of each boundary condition.

condition. For example, heat flux convergence requires $N_x \simeq 96$ for an unoptimized case of the new boundary condition, compared with $N_x \simeq 512$ for the conventional boundary condition. Such a drastic decrease in required resolution leads to a reduction in computational time of ∼7x in GryfX.

It should be emphasized here that k_x^{shift} and the domain aspect ratio are directly related, in the sense that larger aspect ratios will produce smaller k_x^{shift} values. So while results converged with $N_x \simeq 96$ using the unoptimized boundary condition, for a flux tube of one poloidal turn ($L_x/L_y = 1.59$), there is no guarantee that every flux tube length will give an aspect ratio in reasonable proximity to one (which can be seen in Figure 4.8). For a poorly chosen flux tube length with respect to the aspect ratio, it may be that convergence with the new boundary condition is slower with respect to N_x, than with conventional "twist-and-shift" ($L_x/L_y = 8.14$).

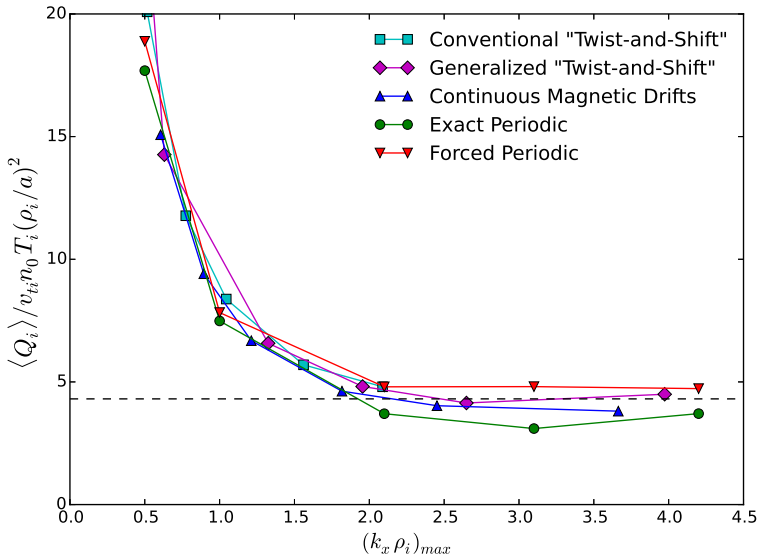

Figure 4.19: (Color online) The same data as Figure 4.18, plotted as a function of the maximum simulated k_x value. Exact flux tube lengths are given in the caption of Figure 4.9. The dashed line is calculated as the average of the heat flux for the rightmost data point of each boundary condition.

However, this also suggests that the required $N_x \simeq 192$ for convergence of the continuous magnetic drifts flux tube in Figure 4.18 is not necessarily directly related to continuity of the magnetic drifts, but may just be a consequence of the domain aspect ratio.

Another interesting result is the behavior of the two periodic cases ('forced' and 'exact'). Surely the most noteworthy outcome pertaining to these runs is the fact that simulations where periodicity is incorrectly enforced converge to the same saturated heat flux as the exact periodicity runs. This behavior is consistent with the linear calculations in Figure 4.13 in the sense that the particular choice of boundary condition is irrelevant if the flux tube has sampled "enough" of the geometry. Further, for both the exact and forced periodic simulations, we observe a convergence to the correct heat flux using even fewer radial modes than boundary conditions not

employing periodicity. This provides some evidence that simply applying periodicity may be the optimal choice, even when the theory and continuity properties dictate that it is not strictly valid.

Apart from the differences we find in the required N_x between the boundary condition choices, the previous point made regarding the importance of simulating a large enough region of k-space can be further appreciated by plotting the same heat flux data of Figure 4.18, but instead as a function of the maximum simulated k_x value in Figure 4.6.3. In doing this, we see that the all heat flux curves nearly overlap, demonstrating that irrespective of which boundary condition is used and how large N_x may need to be, heat flux convergence is ultimately determined by the range of k-space that is being simulated.

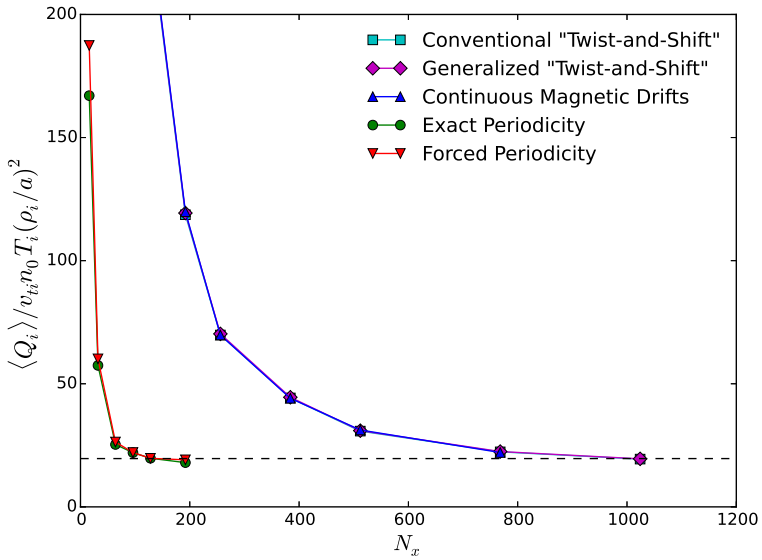

Figure 4.20: (Color online) - Axisymmetric saturated heat flux calculations, where the boundary conditions used in each curve follow Figure 4.18, with the following exception: conventional and generalized "twist-and-shift", and the flux tube producing continuous magnetic drifts become equivalent in axisymmetry as discussed in Section 4.4.5. The exact periodic flux tube extends from $[-1.13\pi, 1.13\pi]$ to satisfy $[\nabla\psi \cdot \nabla\alpha]_{z_\pm} = 0$. The dashed line is calculated as the average of the heat flux for the rightmost data point of each boundary condition.

While results in Figure 4.18 are based on a W7-X geometry, one should expect to see consistent behavior in axisymmetry for similar global shear values based on the conditions discussed above that set the resolution requirements. In Figure 4.20 we present the same study as in Figure 4.18 for a flux tube in a VMEC-generated axisymmetric geometry, designed to have a global shear value, $\hat{s} = -0.018$, close to that of W7-X. First of all, the conventional and generalized cases overlap exactly, as one would expect based on how the boundary condition simplifies in axisymmetry. Beyond this, the simulations model the W7-X case of Figure 4.18 quite well in reference to the required resolution for convergence to the correct saturated heat flux value.

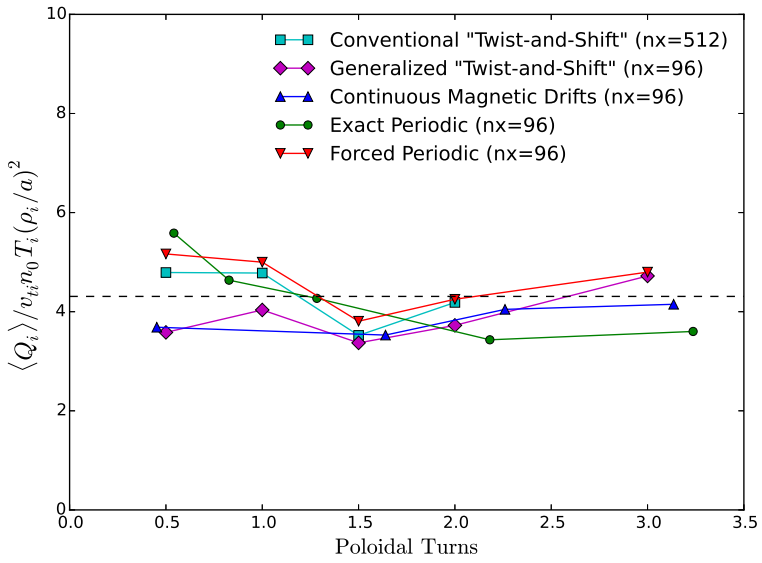

Figure 4.21: (Color online) - Saturated heat flux in W7-X as a function of flux tube length for the various boundary condition choices detailed in Section 4.6.1.1. The dashed line is calculated as the average of the heat flux of the rightmost data point for each boundary condition choice in Figure 4.18.

Finally, to tie in the linear studies of how flux tube length affects results, the heat flux is calculated for the boundary condition options over a range of flux tube

lengths in Figure 4.21. Within statistical fluctuations, the results appear converged to $\langle \bar{Q}_i \rangle \sim 4$ for arguably all flux tube lengths > 0.5 poloidal turns, which is consistent with the linear growth rate results of Section 4.6.1. There is then no reason to expect any significant change in behavior in the limit of the flux tube becoming arbitrarily long.

4.7 Conclusions

The "twist-and-shift" parallel boundary condition used for gyrokinetic simulations of turbulence in axisymmetric equilibria has been generalized for non-axisymmetric geometries and for configurations with low global magnetic shear. Twist and shift boundary conditions are associated with field-line-following coordinates in flux tube simulations. When the variation of local magnetic shear is strong compared with the global magnetic shear, the flux tube twists and untwists as one moves along the field lines. Our generalization takes advantage of this phenomenon by using the integrated local magnetic shear to determine the boundary conditions instead of relying only on the global magnetic shear. As a result, a considerably smaller periodic computational domain can be identified and additional opportunities for optimization of the simulation domain are exposed. The conventional boundary condition of [7] is a perfect subset of our generalized formalism.

Linear stability analyses of W7-X stellarator equilibria have been undertaken using a variety of boundary condition options. The growth rates and frequencies are found to be insensitive to the details of the boundary conditions as long as the

simulation domain is sufficiently long in the direction of the magnetic field. We observe a weak dependence of the calculated eigenvalues on the parallel extent L of the simulation domain as long as $L \gtrsim 2\pi v_i/\omega$. This rule of thumb is consistent with Fourier decompositions of the eigenmodes along the field line. A flux tube that extends at least one poloidal turn was found to be sufficiently long for the W7-X case we examined [see Fig. (4.13)].

In general, nonlinear simulations that are used to estimate turbulent fluxes of heat, *etc.,* are very expensive and are the primary targets of our development of improved boundary conditions. We have surveyed the behavior of secondary instabilities (which can be highly elongated along the magnetic field) and zonal flows in this context. Although we identify cases for which an incorrect (ungeneralized) boundary condition introduces potentially significant parallel discontinuities in the secondary pump waves, we do not observe further serious consequences (such as numerical instability). The importance of this finding will presumably depend on the details of any given numerical discretization. Zonal flows are strictly periodic along the field line, and are therefore not directly affected by the generalization of the boundary condition. Because our approach allows the use of shorter sections of field line, however, we examined the sensitivity of key zonal flow properties to the extent of the flux tube. We found that one poloidal turn is evidently long enough to produce consistent short- and long-time zonal flow responses in the W7-X configuration we examined.

The ideal computational domain for a nonlinear problem can be as small as a few correlation lengths in each direction. When the global magnetic shear is small,

standard "twist-and-shift" boundary conditions force one to use a flux tube that can be considerably longer in the radial direction. Our generalization of the boundary condition makes it possible to hew more closely to the ideal in non-axisymmetric configurations, and we have observed approximately order-of-magnitude speed-ups as a result. Once converged, nonlinear heat flux simulations seem to be essentially unaffected by further details of the boundary conditions, even as one uses the flexibility enabled by our formalism to satisfy additional continuity properties at the boundaries.

The authors acknowledge illuminating conversations about this topic with Per Helander, Gabe Plunk, and Ben Faber. This work was supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Science, under Award Number DE-FG02-93ER54197.

Chapter 5: Conclusions and Future Work

5.1 Conclusions

Stellarators are a promising concept for achieving nuclear fusion as an energy source. They possess many advantages over the conventional tokamak that make them a viable alternative for MCF. As with any concept, however, many challenges must still be overcome, and the goal of this thesis has been to address two of them. Both topics revolve around the neoclassical and turbulent transport of heat and particles in the core of fusion devices.

5.1.1 Neoclassical Impurity Transport in Quasisymmetric Stellarators

The symmetry of tokamaks leads to the beneficial property that heavy impurity ions are automatically expelled from the core of fusion plasmas for large temperature gradients because the average radial motion is unaffected by the radial electric field. This phenomena is referred to as temperature screening. In stellarators, the direction of the impurity transport direction depends on the radial electric field, which is expected to be directed inward for reactor-relevant plasmas, and will

tend to pull impurities into the core. Therefore, stellarators are not predicted to display temperature screening.

In Chapter 2, the behavior of impurity ions in quasisymmetric magnetic fields has been explored. Although perfectly quasisymmetric magnetic fields have neoclassical transport properties identical to tokamaks, actual quasisymmetric concepts or experiments will have a finite amount of symmetry-breaking. In the eight quasisymmetric devices that were studied in this thesis, results indicate that while temperature screening of impurities is possible in select parameters regimes, it is unlikely to be present over a wide range of reactor-relevant plasma parameters. However, this is not to say that it is impossible to achieve temperature screening in future optimized stellarator concepts.

By comparing this neoclassical result to an estimate of turbulence, the magnitude of the neoclassical impurity particle flux is considerably smaller than the turbulent counterpart. The consequences of this disparity in magnitudes are not clear, but the relative size of the turbulence does indicate that perhaps there are other factors beyond the neoclassical temperature screening effect that will control the direction of impurity particle transport in stellarators.

5.1.2 Reducing the Cost of Turbulence Simulations

Turbulence simulations are computationally intensive even for axisymmetric tokamak geometries. The added complexity of stellarators will magnify this cost because of the larger number of grid points required. The boundary condition of

Chapter 4 serves to reduce the cost of turbulence simulations in stellarators.

The standard parallel boundary condition in flux tube simulations of turbulence was originally developed for geometries that are axisymmetric and have order unity levels of global magnetic shear. The derivation and consequences of a modified version of this parallel boundary condition, which permits non-axisymmetric geometries and low global shear, has been covered in Chapter 4.

The new boundary condition allows for the selection of a flux tube length based on local magnetic shear information, which will determine the aspect ratio of the perpendicular plane. One can then optimize the length to either allow for a periodic parallel boundary condition or to find an aspect ratio as close to unity as possible. Domains with aspect ratio close to one will minimize the required resolution in those coordinates, whereas periodic boundary conditions will serve to further decrease the computation time.

Simulations of W7-X using the GryfX code with the new parallel boundary condition show a reduction of a factor of seven in compute time, compared with otherwise identical simulations with the conventional boundary condition. By reducing the cost, one can increase both the quantity and quality of simulations, with the hope of furthering the understanding of turbulence in stellarators.

5.2 Future Work

There are number of natural extensions to the work in this thesis that would be worthwhile to explore.

Firstly, there should be further testing of the parallel boundary condition to understand its performance impact fully. The work in Chapter 4 tests the boundary condition using a hybrid gyrokinetic-gyrofluid code, and the logical next test would be to use a fully gyrokinetic code. Further, the above results explored the ITG instability in the electrostatic limit, and it would be enlightening to understand the behavior when considering electromagnetic effects, or for trapped electron mode (TEM) or kinetic ballooning mode (KBM) instabilities. Other (low-shear) stellarators should be considered, in particular HSX, which has global shear values that are similar to W7-X. Recent simulations of TEM-turbulence in HSX [30] have shown that very long flux tubes are required to produce saturated electron heat fluxes. It could be worthwhile to compare performance of these simulations when using the generalized version of the parallel boundary condition. Finally, a comparison should be done between flux tube and flux surface simulations to better understand when full-surface simulations are necessary.

On the topic of neoclassical impurity transport, the effect of Φ_1 requires further work. As a reminder, the majority of results in Chapter 2 assume the electrostatic potential is a flux function $\Phi = \Phi_0(r)$, and ignore the first-order contribution $\Phi_1(r, \theta, \zeta)$, which includes variation within a flux surface. As mentioned in Section 2.6.2.2, Φ_1 is expected to be small in quasisymmetric configurations not deviating too far from symmetry. In Figure 2.11, however, it was shown that including Φ_1 effects leads to a factor of two change in the impurity particle flux for large values of Z in the QHS Wistell-A configuration. It is thus important to conduct a comprehensive study to fully understand how the inclusion of Φ_1 effects with highly-

charged impurities will alter the neoclassical transport in different quasisymmetric stellarators.

Considering the large parameter space of reactor-relevant plasmas, it could be useful to perform further numerical work with more parameter scans. This could include scans over quantities such as the collisionality, temperature gradient (with fixed density gradient), Z_{eff} value, or temperature ratio. Furthermore, having an analytic theory around magnetic field symmetry-breaking in such a large parameter space would prove useful as a means of providing further insight into the results of Chapter 2.

Finally, the natural connection to the work in Chapters 2 and 4 is to run turbulence simulations with impurities in quasisymmetric stellarators (there have not been turbulent impurity transport simulations in quasisymmetric stellarators, but there has been recent numerical [35] and experimental [39] work studying this in W7-X). This will allow for a considerably more accurate ratio of neoclassical to turbulent particle fluxes than was presented in Section 2.6.2. While turbulence simulations that include impurities are commensurately more computationally intensive, the generalized parallel boundary condition could yield a performance enhancement.

Appendix A: Gyrokinetic Equation

In the electrostatic, collisionless limit, the gyrokinetic equation that describes an arbitrary species is given by:

$$\frac{\partial h}{\partial t} + \boldsymbol{v}_{\parallel} \cdot \nabla h + \Big(\langle \boldsymbol{v}_{\boldsymbol{E}} \rangle + \boldsymbol{v}_m \Big) \cdot \nabla_{\perp} h + \langle \boldsymbol{v}_{\boldsymbol{E}} \rangle \cdot \nabla_{\perp} F_M = \frac{q F_M}{T} \frac{\partial \langle \phi \rangle}{\partial t}. \tag{A.1}$$

The evolved quantity, h, represents the non-adiabatic part of the fluctuating distribution function. The quantity F_M represents the background Maxwellian equilibrium distribution. The angled brackets $\langle ... \rangle$ denote a gyroaveraging operation performed at constant particle guiding center, where the quantity $\langle \boldsymbol{v}_{\boldsymbol{E}} \rangle$ is the gyroaverage of the $\boldsymbol{E} \times \boldsymbol{B}$ velocity. The fluctuating electric field $\boldsymbol{E}$ is defined by the gradient of the fluctuating electrostatic potential, ϕ, which is self-consistently calculated under the assumption of quasineutrality. The magnetic drifts are contained in $\boldsymbol{v}_m = \hat{\boldsymbol{b}} \times \Big[\frac{v_{\perp}^2}{2\Omega} \nabla \log B + \frac{v_{\parallel}^2}{\Omega} (\hat{\boldsymbol{b}} \cdot \nabla \hat{\boldsymbol{b}}) \Big]$, where $\hat{\boldsymbol{b}}$ is the unit vector along the magnetic field.

Appendix B: Minimum Flux Tube Length

Here, we present an interpretation of the Section 4.6.1 results of linear growth rate convergence with respect to flux tube length. In particular, we attempt to understand the minimum flux tube length for linear convergence (which from Figures 4.11, 4.12, and 4.13 appears to be ∼1 poloidal turn) using two approaches.

First, one can consider the wave period $2\pi/\omega$ of the $(k_x, k_y) = (0.0, 0.2)$ mode where the real frequency is $|\omega|_{k_y=0.2} \approx 0.1\ v_{ti}/a$, and estimate the distance $d_i \approx 35$m in which a thermal ion travels in that time. This distance can be compared with the length of a one poloidal turn flux tube $L_{[-\pi,\pi]} = \int_{-\pi}^{\pi} \mathrm{d}l/(\hat{\boldsymbol{b}} \cdot \nabla\theta)$, whereupon one finds that $d_i \sim L_{[-\pi,\pi]}$. Our numerical results are therefore consistent with the hypothesis that the linear growth rates are converged when the flux tube length is at least $\sim 2\pi v_{ti}/|\omega|$. Physically, this hypothesis is plausible since thermal particles should sample the correct geometry for the relevant timescale of the mode.

A second approach for understanding the minimum flux tube length for convergence is to examine the mode's parallel structure in Fourier space. The Fourier transform over the parallel coordinate of $\Phi(z)$ for the $(k_x, k_y) = (0.0, 0.2)$ mode has been plotted in Figure B.1 for a range of flux tubes lengths. It is apparent that the power in $|\Phi_k|$ is strongly peaked around $k_z a \approx 0.08$, which corresponds to a length

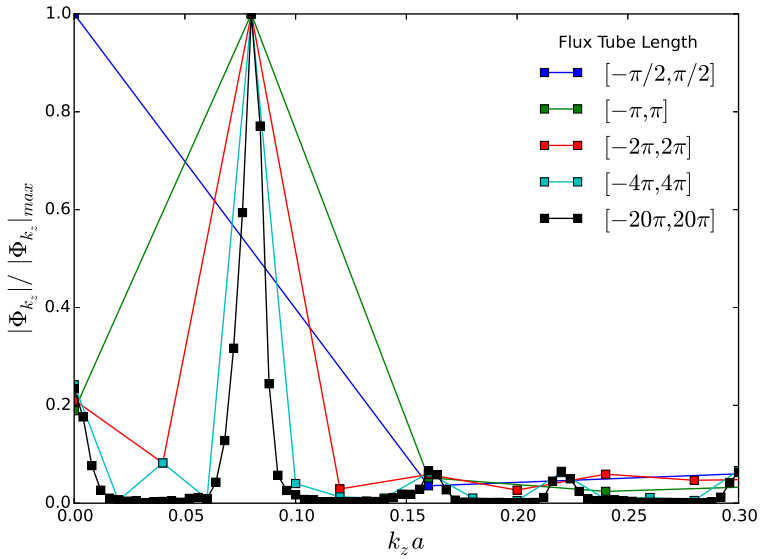

Figure B.1: (Color online) - Fourier transform of the electrostatic potential along the parallel coordinate for various flux tube lengths, using the Forced Periodic boundary condition.

of ∼1 poloidal turn. Since results in Figure 4.13 show that growth rates are mostly converged at one poloidal turn for $k_y = 0.2$, it appears that linearly, a single data point near $k_z a \approx 0.08$ is enough to resolve this spike and get the correct growth rate. Figure B.1 also provides some explanation for the $\sim 30\%$ error in growth rates for ∼0.5 poloidal turn flux tubes, since power at $k_z a \approx 0.08$ is binned with $k_z a = 0$ due to lack of resolution. The growth rate trend becomes erratic and unreliable as the flux tube length is decreased further, and at small enough lengths the growth rate is quite inaccurate, which can be seen in Figure 4.13. Also in line with resolving the spike at $k_z a \approx 0.08$ is that as the flux tube is increased past one poloidal turn, the growth rate converges to the true result. We note here that data points in Figure 4.13 using the conventional "twist-and-shift" boundary condition for flux tubes with lengths < 0.5 poloidal turns yield nearly identical results to using the new boundary condition at an unoptimized length.

Appendix C: Impurity Temperature Screening in Axisymmetry and Quasisymmetry

A number of papers [23,32,84,87] have examined neoclassical transport quantities in the axisymmetric limit and can provide a more intuitive look at the properties of a plasma that can affect temperature screening. These results are also applicable in perfect quasisymmetry. In [23], an expression for the bulk ion (taken to be hydrogen here) particle flux in the presence of electrons and a single impurity was derived, using a momentum-conserving, pitch-angle scattering collision operator, where all species are taken to be in the banana regime

$$\Gamma_i = -C_i \left\{ \frac{T_i}{T_e} \left[\frac{n_i'}{n_i} - \left(\frac{0.09 + 0.5\alpha}{0.53 + \alpha} \right) \frac{T_i'}{T_i} \right] - \frac{T_z}{ZT_e} \left[\frac{n_z'}{n_z} - 0.17 \frac{T_z'}{T_z} \right] \right\}. \tag{C.1}$$

Here, C_i is a positive coefficient that is independent of thermodynamic gradients, Z is the impurity charge, and $\alpha = n_z Z^2/n_i = Z_{eff} - 1$ represents the effect of the impurities on the transport coefficients. Typically, one can take $T_i = T_z = T$, and we further assume the profiles of each species to be equal: $T_i'/T_i = T_z'/T_z = T'/T$, and $n_i'/n_i = n_z'/n_z = n'/n$. An expression for the *impurity* particle flux can then

easily be obtained by satisfying the ambipolarity condition, $\sum_a Z_a \Gamma_a = 0$. Based on $\Gamma_e \sim \sqrt{m_e/m_i}\,\Gamma_i$, if we take $\Gamma_e \simeq 0$, then $\Gamma_z \simeq -\Gamma_i/Z$. The impurity particle flux can now be shown to be

$$\Gamma_z = \frac{C_i}{Z}\left\{\frac{T}{ZT_e}\left[L_a\frac{n'}{n} - L_b\frac{T'}{T}\right]\right\}, \tag{C.2}$$

where $L_a \equiv Z - 1$ and $L_b \equiv \{[Z(0.09 + 0.5\alpha)/(0.53 + \alpha)] - 0.17\}$. L_b is always positive for $Z > 1$.

If it is assumed that both density and temperature profiles are peaked, then achieving a positive (outward) Γ_z is based on two properties: Z_{eff}, and more importantly, the ratio $\eta^{-1} \equiv d\ln(n)/d\ln(T)$. In the absence of a density gradient, the term in square brackets in C.2 will always be positive and lead to temperature screening. However, based on the α, there is some critical $\eta_c^{-1}(\alpha) \sim 0.4$ where $\eta^{-1} > \eta_c^{-1}$ will always lead to impurity accumulation. This effect can be seen in Figure C.1, where SFINCS has been used to calculate the impurity particle flux for C^{6+} over a range of collisionalities for an axisymmetric geometry model $B(\theta) = B_0(1 + \epsilon\cos\theta)$, with $\iota = 0.689$ and $\alpha = 1$. It is also clear that temperature screening is only accessible up to some maximum collisionality, regardless of η^{-1}. However, such high collisionalities put not only impurities, but also bulk ions, into the Pfirsch-Schlüter regime, which is generally too collisional to be relevant in the core of reactor scenarios.

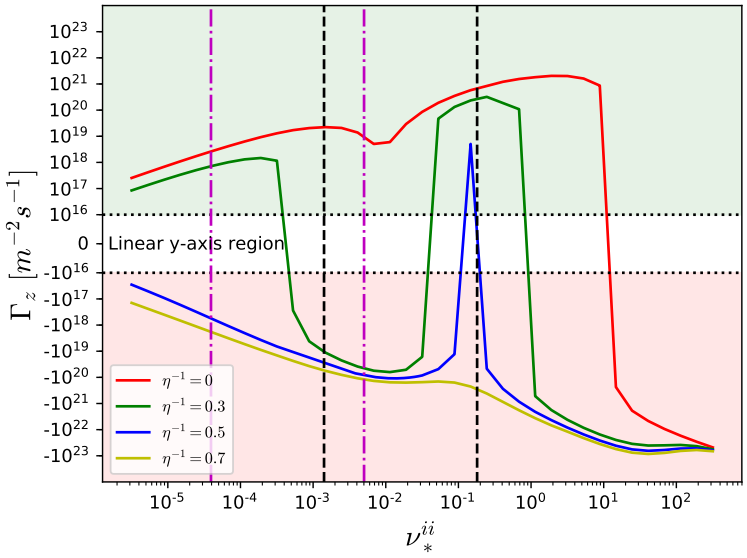

Figure C.1: (Color online) The impurity particle flux for C^{6+} is plotted in an axisymmetric geometry as a function of the normalized ion-ion collisionality $\nu_*^{ii} \equiv \nu R/v_{ti}$, where $\epsilon = r_N(a/R)$. In this plot, $r_N = 0.25$, and a/R=0.16. The vertical lines signify the transitions between collisionality regimes for ν_*^{ii} (black dashed) and ν_*^{zz} (magenta dotted). The transitions are as follows: $\nu_*^{ii} = 1.41 \cdot 10^{-3}$ (main ion banana-plateau), $\nu_*^{ii} = 0.18$ (main ion plateau-Pfirsch-Schlüter), $\nu_*^{ii} = 3.93 \cdot 10^{-5}$ (impurity banana-plateau), and $\nu_*^{ii} = 5.03 \cdot 10^{-3}$ (impurity plateau-Pfirsch-Schlüter). The upper, green-shaded region denotes positive Γ_z (impurity screening). The lower, red-shaded region corresponds to negative Γ_z (impurity accumulation).

Appendix D: Testing and Portability of GPU-based Codes

In recent years, it has become increasingly more common to use graphics processing units (GPUs) for the purpose of scientific computing. GPUs enable the acceleration of scientific codes by using single instruction, multiple data (SIMD) commands to parallelize parts of the code. Utilizing GPUs for code sections that perform arithmetic operations on large amounts of data can result in a significant reduction in computing time. The presence of more than one kind of processor (i.e. CPU + GPU/FPGA/etc.) does, however, require additional steps to create production-level code that can be tested, and used across multiple different architectures.

Central to the writing of GPU-based codes is handling the communication between the host (CPU) and device (GPU). On pure CPU architectures, it is straightforward to allocate and deallocate memory, create functions, and perform operations without having to specify anything about *where* these processes are executed. Since the principal advantage to GPUs is executing a set of instructions in parallel, the device must have knowledge of these functions, and these functions will need to access arrays, which are allocated somewhere. An important point in host-device communication is that the transfer from host-to-device or device-to-host is an ex-

pensive action, and should be minimized. The solution to this is two-fold. First, memory should be allocated on the device to provide quick access to array elements. Second, the instructions to be performed must be able to be interpreted by the device, requiring what are known as kernels and device functions, which are compiled specifically for the GPU.

There is a subtle difference between kernels and device functions that is central to this chapter. Device functions can simply be thought of as a regular function that is compiled to be executed within kernels, on the GPU. Kernels are *initializers*. When called, they provide instructions for *how* to parallelize a set of instructions through the use of a coordinate-based grid of threads, where a thread is the parallel execution unit. The following bit of code gives an example for how a kernel is called in CUDA C syntax

```
dim3 block = dim3(4, 4, 4);

dim3 grid = dim3(1, 1, 512);

kernel<<<grid, block>>>(input_array);
```

The variables `block` and `grid` are arrays that can be 1-D, 2-D, or 3-D, which describe the layout of threads within a kernel, where a grid is a collection of blocks (along the `z` coordinate in this example). Threads within a block are limited (1024 threads for NVIDIA Volta architecture) since they are expected to reside within a single processor core, while grid dimensions can be considerably larger. In the above code, a block is a cube with sides of length four threads, and the grid has 512 of these blocks. The body of a kernel executes for each thread, which has a unique identifier

that depends on the values of `block` and `grid`. This identifier can be found, for example, via

```
// ---------------------------------------
int index = threadIdx.x + blockIdx.x*blockDim.x;
// ---------------------------------------
```

The variables `threadIdx`, `blockIdx`, and `blockDim` are automatically-generated CUDA variables for use within kernels. In principle, a kernel could execute identically to a device function, but there is reason to avoid this for the benefit of testing, and if necessary, code portability.

The ability to test functions, conditional statements, and other sections of the code becomes increasingly important as the size and complexity of the code increases. There are numerous testing packages that one might want to import in order to ensure these parts of the code are functioning properly. However, these packages cannot interpret device-based code, which includes any type of kernel. This is where the importance of device functions comes in. Device functions can be called from within the body of the kernel just as a regular host-based function would be called within a CPU code. This similarity leads to the option of creating "host-device" functions, which are compiled separately for the host and device so that the function can be interpreted on both. The ability to create these host-device functions presents an advantageous design paradigm for kernels.

Without the use of host-device functions, it would be difficult to test the behavior of a kernel. If one were to instead fill the body of a kernel with small host-

device function calls (which could be inlined), then each part of the kernel could be tested without a decrease in performance. For example, consider a kernel call that initializes a grid of threads with the objective of manipulating an array with three separate sets of instructions (syntax for each of the examples is specific to CUDA C)

```
// ----------------------------------------
__global__ void kernel(double *array) {
  int index = threadIdx.x + blockIdx.x*blockDim.x;
  array[index] = array[index] + 5;
  array[index] = array[index] * 10;
  array[index] = array[index] - 3;
}
// ----------------------------------------
```

Within the body of the kernel above, we have the three sets of instructions performed serially on each thread (which is done in parallel with the other threads). There is nothing wrong with how this kernel executes, and assuming there is nothing incorrect in the instructions, it will produce the desired result. However, if something is wrong in the instructions, there is no easy way of finding out the cause based on the output, since that bit of code cannot be tested. If we were to instead make these three sets of instructions into host-device functions, this would give one the ability to test each function separately without altering the functionality of the kernel. The altered version of this kernel where three host-device functions are defined, can be

seen below

```
// ----------------------------------------
__host__ __device__ void add_5(double arr_element)
{
  arr_element += 5;
}
__host__ __device__ void multiply_10(double arr_element)
{
  arr_element *= 10;
}
__host__ __device__ void subtract_3(double arr_element)
{
  arr_element -= 3;
}
__global__ void kernel(double *array) {
  int index = threadIdx.x + blockIdx.x*blockDim.x;
  add_5(array[index]);
  multiply_10(array[index]);
  subtract_3(array[index]);
}
// ----------------------------------------
```

For a more practical example, the calculation of the quantity $\Gamma_0(b) \equiv I_0(b)\mathrm{e}^{-b}$,

which appears in gyrokinetic/gyrofluid turbulence simulations, will be adapted slightly to allow for unit testing. Here, $I_0(b) = J_0(ib)$, where J_0 is the zeroth-order Bessel function, and i represents the imaginary unit. The kernel, host-device function, and Catch2 unit tests (https://github.com/catchorg/Catch2) can be seen below. For simplicity, only the line that calls the Γ_0 function in the kernel is shown.

```
// ----------------------------------------
__host__ __device__ float g0(float b) {


  float tol = 1.e-7;
  float tk, b2, b2sq;
  float g, x, xi, err;


  if (b < tol) {return 1.0;}
  b2 = 0.5 * b;
  b2sq = b2 * b2;
#ifdef __CUDA_ARCH__
  tk = __expf(-b);
#else
  tk = exp(-b);
#endif}
  g = tk;
  x = 1.;
```

```
  err = 1.;


  while (err > tol) {
    xi = 1./x;
    tk  = tk * b2sq * xi * xi;
    g += tk;
    x  += 1.;
    err = abs(tk/g);
  }


  if(g<tol) {g=tol;}
  return g;
}


// ----------------------------------------
__global__ void calc_phiavgdenom(float* PhiAvgDenom,
  float* kperp2, float* jacobian, specie* species,
  float ti_ov_te) {
  ...
  pfilter2 += s.dens*s.z*s.zt
  *(1. - g0(kperp2[(idx + i*nx)*nyc]*s.rho2));
  ...
}
```

```
// ---------------------------------------
// Unit Testing for the function g0(float b)
// g0( b < tol ) should return 1
CHECK( g0(1.e-8) == 1 );
// g0( b < 0 ) should return 1
CHECK( g0(-0.1) == 1 );

// Verify that result are accurate to within 1%
CHECK( g0(0.001) == Approx(0.999001).epsilon(0.01) );
CHECK( g0(0.01) == Approx(0.990075).epsilon(0.01) );
CHECK( g0(0.1) == Approx(0.907101).epsilon(0.01) );
CHECK( g0(0.5) == Approx(0.645035).epsilon(0.01) );
CHECK( g0(1.0) == Approx(0.465760).epsilon(0.01) );
CHECK( g0(10.0) == Approx(0.127833).epsilon(0.01) );
// ---------------------------------------
```

In the above bits of code, the kernel `calc_phiavgdenom` uses the host-device function `g0` (equivalent to Γ_0) to calculate a quantity `pfilter2`. The function `g0` has the qualifier `__host__ __device__` along with an `#ifdef` macro, which enables the function to be compiled separately for the host and device with processor-specific lines of code. In this case, the exponential function is evaluated with the built-in CUDA function `__expf` on the device, and with the `exp` function from the `<cmath>`

header library on the host. Including the `#ifdef` macro does not change how `g0` is evaluated when called from the kernel. It does, however, allow `g0` to be tested to ensure that it works properly when any routine calls it. The final section of code uses `CHECK` statements from Catch2 to ensure that the function returns the correct value based on the argument that is passed. To provide thorough testing, a similar process should be repeated for any function that is called within a kernel.

Another benefit to using host-device functions within kernels is for portability of codes across different platforms. For example, if one wanted to run a GPU-based code on a CPU-based architecture, the kernels would need to be completely rewritten for the instructions to be readable to a CPU. By filling the kernels with consecutive calls to device-host functions, all of the instructions are already in a format that can be understood. The only change that would need to be made is how the body of the kernel is called.

Today, access to supercomputers is becoming increasingly more common, leading many codes to be developed to run on these platforms. However, each of these systems comes with their own hardware, packages, and sometimes languages, requiring very specific builds to be able to run on them. This forces developers to rewrite large portions of code in order to utilize the parallelism of a specific system. To address this problem, there has been an effort to develop portable programming ecosystems like Kokkos [28], RAJA [6], and oneAPI. The purpose is to provide general commands that can access the parallelism of any system, so that rewriting is minimized. Such portability will allow any code to benefit from different types of parallelism with relatively minor modifications.

Bibliography

[1] Abel, I.G., Plunk, G.G., Wang, E., Barnes, M., Cowley, S.C., Dorland, W., Schekochihin, A.A. Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows. *Report Prog. Phys.* **76** 116201 (2013).

[2] Allmaier, K., Kasilov, S.V., Kernbichler, W., Leitold, G.O. Variance reduction in computations of neoclassical transport in stellarators using a δf method. *Phys. Plasmas* **15** 072512 (2008).

[3] Anderson, F.S.B., et al. The Helically Symmetric Experiment, (HSX) Goals, Design and Status. *Fusion Tech.* **27** 273 (1995).

[4] Antonsen Jr. T.M., Lane B. Kinetic equations for low frequency instabilities in inhomogeneous plasmas. *Phys. Fluids* **23** 1205 (1980).

[5] Bader, A., Drevlak, M., Anderson, D.T., Faber, B.J., Hegna, C.C., Likin, K.M., Schmitt, J.C., Talmadge, J.N. Stellarator equilibria with reactor relevant energetic particle losses. *J. Plasma Phys.* **85** 905850508 (2019).

[6] Beckingsale D.A., et al. RAJA: Portable performance for large-scale scientific applications. *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC).* https://github.com/LLNL/RAJA.

[7] Beer M.A. Cowley S.C. and Hammett G.W. Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Phys. Plasmas* **2** 2687 (1995).

[8] Beer M.A., PhD thesis, Princeton University (1995).

[9] Beer M.A., Hammett G.W. Toroidal gyrofluid equations for simulations of tokamak turbulence. *Phys. Plasmas* **3** 4046 (1996).

[10] Beidler, C.D., et al. Benchmarking of the mono-energetic transport coefficients-results from the International Collaboration on Neoclassical Transport in Stellarators (ICNTS). *Nucl. Fusion* **51** 076001 (2011).

[11] Beidler, C.D., et al. Transport modeling for W7-X on the basis of W7-AS experimental results. *Proceedings of the 17th International Toki Conference and 16th International Stellarator/Heliotron Workshop, Toki* (2008).

[12] Boozer, A.H. Transport and isomorphic equilibria. *Phys. Fluids* **26** 496 (1983).

[13] Boozer, A.H. Establishment of magnetic coordinates for a given magnetic field. *Phys. Fluids* **25** 520 (1982).

[14] Boozer, A.H. Plasma equilibrium with rational magnetic surfaces. *Phys. Fluids* **24** 1999 (1981).

[15] Buller S., Smith H.M., Helander P., Mollén A., Newton S.L., Pusztai I. Collisional transport of impurities with flux-surface varying density in stellarators. *J. Plasma Phys.* **84** 905840409 (2018).

[16] Calvo I., Parra F.I., Velasco J.L., Alonso J.A., García-Regaña, J.M. Stellarator impurity flux driven by electric fields tangent to magnetic surfaces. *Nucl. Fusion* **58** 124005 (2018).

[17] Calvo I., Parra F., Velasco J.L., Alonso J.A. The effect of tangential drifts on neoclassical transport in stellarators close to omnigeneity. *Plasma Phys. Control. Fusion* **59** 055014 (2017).

[18] Candy J., Waltz R.E. An Eulerian gyrokinetic-Maxwell solver. *J. Comput. Phys.* **186** 545 (2003).

[19] Canik, J.M., et al. Reduced particle and heat transport with quasisymmetry in the Helically Symmetric Experiment. *Phys. Plasmas* **14** 056107 (2007).

[20] Cary, J.R., Shasharina, S.G. Helical plasma confinement devices with good confinement properties. *Phys. Rev. Lett.* **78** 674 (1997).

[21] Catto P. et al. Linearized gyro-kinetics. *Plasma Phys.* **23** 639 (1981).

[22] Connor, J.W. Invariance principles and plasma confinement. *Plasma Phys. Control. Fusion* **30** 619 (1988).

[23] Connor, J.W. The neo-classical transport theory of a plasma with multiple ion species. *Plasma Phys.* **15** 765 (1973).

[24] Cowley S.C., Kulsrud R.M., Sudan R. Considerations of ion-temperature-gradient-driven turbulence. *Phys. Fluids B* **3** 2767 (1991).

[25] Dimits A.M. et al. Comparisons and physics basis of tokamak transport models and turbulence simulations. *Phys. Plasmas* **7** 969 (2000).

[26] Dimits A.M. Fluid simulations of tokamak turbulence in quasiballooning coordinates. *Phys. Rev. E* **48** 4070 (1993).

[27] Dorland W., Jenko F., Kotschenreuther M., Rogers B.N. Electron temperature gradient turbulence. *Phys. Rev. Lett.* **85** 5579 (2000).

[28] Edwards H.C., Trott C.R, Sunderland D. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing* **74** 3202 (2014).

[29] Engelmann, F. Extrapolation of tokamak energy confinement to Next Step Devices. *Plasma Phys. Control. Fusion* **32** 1101 (1990).

[30] Faber, B. et al. Stellarator microinstabilities and turbulence at low magnetic shear. *Journal Plasma Phys.* **84** 905840503 (2018).

[31] Frieman E.A., Chen L. Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *Phys. Fluids* **25** 502 (1982).

[32] Galeev, A.A., Sagdeev, R.Z. Transport Phenomena in a Collisionless Plasma in a Toroidal Magnetic System. *Soviet Phys. JETP* **26** 233 (1968).

[33] Garabedian, P.R. Stellarators with the magnetic symmetry of a tokamak. *Phys. Plasmas* **3** 2483 (1996).

[34] García-Regaña, J.M., et al. Electrostatic potential variation on the flux surface and its impact on impurity transport. *Nucl. Fusion* **57** 056004 (2017).

[35] García-Regaña, J.M., et al. Turbulent impurity transport simulations in Wendelstein 7-X plasmas. arXiv:2008.07662 (2020).

[36] Garren, D.A., Boozer, A.H. Existence of quasihelically symmetric stellarators. *Phys. Fluids B* **3** 2822 (1991).

[37] Gates D.A., et al. Recent advances in stellarator optimization. *Nucl. Fusion* **57** 126064 (2017).

[38] Gates D.A., et al. Stellarator Research Opportunities: A Report of the National Stellarator Coordinating Committee. *Journal Fusion Energy* **37** 51 (2018).

[39] Geiger, B., et al. Observation of anomalous impurity transport during low-density experiment in W7-X with laser blow-off injections of iron. *Nucl. Fusion* **59** 046009 (2019).

[40] Grieger G. et al. Physics optimization of stellarators. *Phys. Fluids B: Plasma Physics* **4** 2081 (1992).

[41] Hazeltine, R.D. Recursive derivation of drift-kinetic equation. *Plasma Phys.* **15** 77 (1973).

[42] Helander P., Sigmar D.J. *Collisional Transport in Magnetized Plasmas.* Cambridge University Press (2002).

[43] Helander P., Newton S.L., Mollén A., Smith H.M. Impurity transport in a mixed-collisionality stellarator plasma. *Phys. Rev. Lett.* **118** 115002 (2017).

[44] Helander, P. On rapid plasma rotation. *Phys. Plasmas* **14** 104501 (2007).

[45] Helander P., et al. Stellarator and tokamak plasmas: a comparison. *Plasma Phys. Control. Fusion* **54** 124009 (2012).

[46] Helander P. Theory of plasma confinement in non-axisymmetric magnetic fields. *Rep. Prog. Phys.* **77** 087001 (2014).

[47] Hender, T., et al. Studies of a Flexible Heliac Configuration. *Fusion Tech.* **13** 521 (1988).

[48] Henneberg, S.A., Drevlak, M., Nührenberg, C., Beidler, C.D., Turkin, Y., Loizu, J., Helander, P. Properties of a new quasi-axisymmetric configuration. *Nucl. Fusion* **59** 26014 (2019).

[49] Highcock E.G., Mandell N.R., Barnes M., Dorland W. Optimisation of confinement in a fusion reactor using a nonlinear turbulence model. *J. Plasma Phys.* **84** 905840208 (2018).

[50] Hirsch, M. Major results from the stellarator Wendelstein 7-AS *Plasma Phys. Controlled Fusion* **50** 053001 (2008).

[51] Hirshman, S.P., Shaing, K.C., van Rij, W.I., Beasley, C.O., Crume, E.C. Plasma transport coefficients for nonsymmetric toroidal confinement systems. *Phys. Fluids* **29** 2951 (1986).

[52] Hirshman, S.P. The ambipolarity paradox in toroidal diffusion, revisited. *Nucl. Fusion* **18** 917 (1978).

[53] Hirshman S.P., van Rij W.I., Merkel P. Three-dimensional free boundary calculations using a spectral Green's function method. *Comput. Phys. Comm.* **43** 143 (1986).

[54] Hirshman S.P. Transformation from VMEC to Boozer Coordinates. "https://princetonuniversity.github.io/STELLOPT/docs/Transformation%20from%20VME (1995).

[55] Howes G.G., et al. Astrophysical Gyrokinetics: Basic Equations and Linear Theory. *The Astrophysical Journal* **651** 590 (2006).

[56] Ida, K., et al. Observation of an impurity hole in a plasma with an ion internal transport barrier in the Large Helical Device. *Phys. Plasmas* **16** 056111 (2009).

[57] Igitkhanov, Y., Polunovsky, E., Beidler, C.D. Impurity dynamics in nonaxisymmetric plasmas. *Fusion Sci. Tech.* **50** 268 (2006).

[58] Jenko F., Dorland W., Kotschenreuther M., Rogers B.N. Electron temperature gradient driven turbulence. *Phys. Plasmas* **7** 1904 (2000).

[59] Kernbichler, W., Kasilov, S.V., Leitold, G.O., Nemov, V.V., Allmaier, L. Recent Progress in NEO-2 — A Code for Neoclassical Transport Computations Based on Field Line Tracing. *Plasma Fusion Res.* **3** S1061-1 (2008).

[60] Landau, L.D. The transport equation in the case of coulomb interactions. *Physikalische Zeitschrift der Sowjetunion* **10** 154 (1936).

[61] Landreman, M., Smith, H.M., Mollén, A., Helander, P. Comparison of particle trajectories and collision operators for collisional transport in nonaxisymmetric plasmas. *Phys. Plasmas* **21** 042503 (2014).

[62] Li, S., Jiang, H., Ren, Z., Xu, C. Optimal Tracking for a Divergent-Type Parabolic PDE System in Current Profile Control. *Abstract and Applied Analysis.* **2014** 940965 (2014).

[63] Liu, H., et al. Magnetic configuration and modular coil design for the Chinese First Quasi-Axisymmetric Stellarator. *Plasma Fusion Res.* **13** 3405067 (2018).

[64] Luce T.C. et al. Stationary high-performance discharges in the DIII-D tokamak. *Nucl. Fusion* **43** 321 (2003).

[65] Maaßberg, H., Beidler, C.D., Simmet, E.E. Density control problems in large stellarators with neoclassical transport. *Plasma Phys. Control. Fusion* **41** 1135 (1999).

[66] Mandell N.R., Dorland W., Landreman M. Laguerre-Hermite pseudo-spectral velocity formulation of gyrokinetics. *J. Plasma Phys.* **84** 905840108 (2018).

[67] Martin M.F., Landreman M. Impurity temperature screening in stellarators close to quasisymmetry. *Journal Plasma Phys.* **86** 905860317 (2020).

[68] Martin M.F., et al. The parallel boundary condition for turbulence simulations in low magnetic shear devices. *Plasma Phys. Control. Fusion* **60** 095008 (2018).

[69] Mishchenko A., Helander P., Könies A. Collisionless dynamics of zonal flows in stellarator geometry. *Phys. Plasmas* **15** 072309 (2008).

[70] Mollén A., Landreman M., Smith H.M., García-Regaña J.M., Nunami M. Flux-surface variations of the electrostatic potential in stellarators: impact on the radial electric field and neoclassical impurity transport. *Plasma Phys. Control. Fusion* **60** 084001 (2018).

[71] Monreal P. et al. Residual zonal flows in tokamaks and stellarators at arbitrary wavelengths. *Plasma Phys. Control. Fusion* **58** 045018 (2016).

[72] Motojima O. et al. Recent advances in the LHD experiment. *Nucl. Fusion* **43** 1674 (2003).

[73] Mynick, H. Calculation of the poloidal ambipolar field in a stellarator and its effect on transport. *Phys. Fluids* **27** 2086 (1984).

[74] Mynick, H.E., Boozer, A.H., Ku, L.P. Improving confinement in quasi-axisymmetric stellarators. *Phys. Plasmas* **13** 1 (2006).

[75] Najmabadi, F., et al. the ARIES-CS Compact Stellarator. *Fusion Sci. Tech.* **54** 655 (2008).

[76] Nemov, V.V., Kasilov, S.V., Kernbichler, W., Heyn, M.F. Evaluation of $1/\nu$ neoclassical transport in stellarators. *Phys. Plasmas* **6** 4622 (1999).

[77] Nührenberg, J., Zille, R. Quasi-helically symmetric toroidal stellarators. *Phys. Lett. A* **129** 113 (1988).

[78] Pablant, N.A., et al. Core radial electric field and transport in Wendelstein 7-X plasmas. *Phys. Plasmas* **25** 022508 (2018).

[79] Pederson T.S., et al. Confirmation of the topology of the Wendelstein 7-X magnetic field to better than 1:100,000. *Nature Comm.* **7** 13493 (2016).

[80] Post R.F. Impurity radiation losses from a high temperature plasma. *J. Nucl. Energy* **3** 273 (1961).

[81] Pytte, A., Boozer, A.H. Neoclassical transport in helically symmetric plasmas *Phys. Fluids* **24** 88 (1981).

[82] Reiman A., et al. Physics design of a high-beta quasi-axisymmetric stellarator. *Plasma Phys. Control. Fusion* **41** B273 (1999).

[83] Rogers B.N., Dorland W., Kotschenreuther M. Generation and stability of zonal flows in ion-temperature-gradient mode turbulence. *Phys. Rev. Lett.* **85** 25 (2000).

[84] Rosenbluth, M.N., Hazeltine R.D., Hinton F.L. Plasma transport in toroidal confinement systems. *Phys. Fluids* **15** 116 (1972).

[85] Rosenbluth M.N., Hinton F.L. Poloidal flow driven by ion-temperature-gradient turbulence in tokamaks. *Phys. Rev. Lett.* **80** 724 (1998).

[86] Rutherford, P.H. Collisional diffusion in an axisymmetric torus. *Phys. Fluids* **13** 482 (1970).

[87] Rutherford, P.H. Impurity transport in the Pfirsch-Schlüter regime. *Phys. Fluids* **17** 1782 (1974).

[88] Sanchez R., et al. Ballooning stability optimization of low-aspect-ratio stellarators. *Plasma Phys. Control. Fusion* **42** 641 (2000).

[89] Shaing, K.C. Plasma flows and radial electric field in nonaxisymmetric toroidal plasmas. *Phys. Fluids* **29** 2231 (1986).

[90] Shimizu, A., et al. Configuration property of the Chinese First Quasi-Axisymmetric Stellarator. *Plasma Fusion Res.* **13** 3403123 (2018).

[91] Spitzer L. The stellarator concept. *Phys. Fluids* **1** 253 (1958).

[92] Spong D.A., et al. J* optimization of small aspect ratio stellarator/tokamak hybrid devices. *Phys. Plasmas* **5** 1752 (1998).

[93] Sugama H., Nishimura S. How to calculate the neoclassical viscosity, diffusion, and current coefficients in general toroidal plasmas. *Phys. Plasmas* **9** 4637 (2002).

[94] Sugama H., Watanabe T.H., Nunami M., Nishimura S. Momentum balance and radial electric fields in axisymmetric and nonaxisymmetric toroidal plasmas. *Plasma Phys. Control. Fusion* **53** 024004 (2011).

[95] Sugama, H., Watanabe, T.-H., Nunami, M., Nishimura, S. Quasisymmetric toroidal plasmas with large mean flows. *Phys. Plasmas* **18** 82505 (2011).

[96] Theiler C. PhD Thesis. École Polytechnique Fédérale de Lausanne (2011).

[97] Tribaldoes, V. Monte Carlo estimation of neoclassical transport for the TJ-II stellarator. *Phys. Plasmas* **8** 1229 (2001).

[98] Turkin, Y., Beidler, C.D., Maaßberg, H., Murakami, S., Tribaldos, V., Wakasa, A. Neoclassical transport simulations for stellarators. *Phys. Plasmas* **18** 022505 (2011).

[99] van Rij, W.I., Hirshman, S.P. Variational bounds for transport coefficients in three-dimensional toroidal plasmas. *Phys. Fluids B* **1** 563 (1989).

[100] Velasco, J.L., et al. Moderation of neoclassical impurity accumulation in high temperature plasmas of helical devices. *Nucl. Fusion* **57** 016016 (2017).

[101] Wade, M.R., Houlberg, W.A., Baylor, L.R. Experimental confirmation of impurity convection driven by the ion-temperature gradient in toroidal plasmas *Phys. Rev. Lett.* **84** 282 (2000).

[102] Waltz R.E., Boozer A.H. Local shear in general magnetic stellarator geometry. *Phys. Fluids B* **5** 2201 (1993).

[103] Waltz, R.E., DeBoo, J.C., Rosenbluth, M.N. Magnetic-Field Scaling of Dimensionally Similar Tokamak Discharges. *Phys. Rev. Lett.* **65** 19 (1990).

[104] Waltz R.E., Kerbel G.D., Milovich J. Toroidal gyro-Landau fluid model turbulence simulations in a nonlinear ballooning mode representation with radial modes. *Phys. Plasmas* **1** 2229 (1994).

[105] Watanabe T.-H., Sugama H. Velocity–space structures of distribution function in toroidal ion temperature gradient turbulence. *Nucl. Fusion* **46** 24 (2005).

[106] Xanthopoulos P. et al. A geometry interface for gyrokinetic microturbulence investigations in toroidal configurations. *Phys. Plasmas* **16** 082303 (2009).

[107] Xanthopoulos P., et al. Controlling turbulence in present and future stellarators. *Phys. Rev. Lett.* **113** 155001 (2014).

[108] Xanthopoulos et al. Zonal flow dynamics and control of turbulent transport in stellarators. *Phys. Rev. Lett.* **107** 245002 (2011).

[109] Zarnstorff M.C. et al. Physics of the compact advanced stellarator NCSX. *Plasma Phys. Controlled Fusion* **43** A237 (2001).