

Eliminating turbulent self-interaction through the parallel boundary condition in local gyrokinetic simulations

Justin Ball^{1,†}, Stephan Brunner¹ and Ajay C.J.¹

¹Ecole Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC),
CH-1015 Lausanne, Switzerland

(Received 19 December 2019; revised 24 February 2020; accepted 25 February 2020)

In this work, we highlight an issue that may reduce the accuracy of many local nonlinear gyrokinetic simulations – turbulent self-interaction through the parallel boundary condition. Given a sufficiently long parallel correlation length, individual turbulent eddies can span the full domain and ‘bite their own tails’, thereby altering their statistical properties. Such self-interaction is only modelled accurately when the simulation domain corresponds to a full flux surface, otherwise it is artificially strong. For Cyclone Base Case parameters and typical domain sizes, we find that this mechanism modifies the heat flux by approximately 40% and it can be even more important. The effect is largest when using kinetic electrons, low magnetic shear and strong turbulence drive (i.e. steep background gradients). It is found that parallel self-interaction can be eliminated by increasing the parallel length and/or the binormal width of the simulation domain until convergence is achieved.

Key words: fusion plasma, plasma nonlinear phenomena, plasma simulation

1. Introduction

Local nonlinear gyrokinetic simulations are one of the most commonly used tools to assess turbulent transport in tokamaks. They solve the gyrokinetic model (Catto 1978; Frieman & Chen 1982; Brizard & Hahm 2007; Abel *et al.* 2013), a nonlinear system of integro-differential equations that have been derived to model plasma turbulence as accurately as possible. Indeed, the primary approximation of gyrokinetics, $\rho_* \equiv \rho_i/a \ll 1$ (i.e. the tokamak minor radius a is much larger than the ion gyroradius ρ_i), is satisfied in many existing tokamaks by factors of several hundred. Unfortunately, even with its approximations, such a high-fidelity model is very computationally expensive. For this reason, it has been crucial to streamline the calculations as much as possible.

To this end, it is helpful to minimize the volume of the simulation domain and only include the minimal number of turbulent eddies needed to obtain a statistically relevant representation of turbulence. Instead of modelling full magnetic flux surfaces across most of the plasma minor radius (referred to as ‘global’ simulations), a calculation can get by with a much smaller domain (referred to as ‘local’ simulations).

[†]Email address for correspondence: justin.ball@epfl.ch




FIGURE 1. An example flux tube (thin blue and purple lines) that is one poloidal turn long with its cross-sectional shape indicated on the outboard midplane (thick black rectangle) and at both locations on the inboard midplane (thick black parallelograms). Also shown is the central flux surface of the flux tube (transparent yellow) with a toroidal wedge removed for visual clarity.

Since turbulence in tokamaks is aligned with the magnetic field and is very anisotropic, it is important for the domain of a local simulation to have similar characteristics. Thus, the domain is long in the direction parallel to the magnetic field and quite narrow in the two perpendicular directions. Full flux surfaces are not usually modelled, nor is a large fraction of the minor radius. Typically, the domain has a rectangular cross-section on the outboard midplane and is deformed into a parallelogram at other poloidal locations due to the effect of magnetic shear. Such a domain is called a ‘flux tube’ (Beer, Cowley & Hammett 1995) (see figure 1). It can enable the computational cost of local simulations to be orders of magnitude lower than global simulations, particularly when modelling large devices.

An important difference between local and global simulations concerns boundary conditions. In global simulations, the boundaries within the flux surface are straightforward – because entire flux surfaces are modelled, the physical periodicity in the toroidal and poloidal directions can be implemented directly. However, the radial boundary condition for global simulations is less obvious. Typically Dirichlet boundary conditions are applied together with ‘buffer regions’, which are radial regions that contain sources of particles and energy (Candy & Waltz 2003; Görler *et al.* 2011)¹. Hence, it is important to monitor global simulations to ensure that the results are not sensitive to the particularities of this complex and artificial boundary condition.

In contrast, all of the boundary conditions for local simulations are elegant (Beer *et al.* 1995), although they still require care (as is emphasized by this work). Since the domain of local simulations is narrow in the radial direction, the background plasma parameters (e.g. density, temperature) are well approximated by a constant value with a linear gradient to drive turbulence. This naturally reflects the inherent scale separation between fluctuations and background gradients, which is a fundamental assumption of the gyrokinetic model (Abel *et al.* 2013). One consequence is that the

¹Additionally, an artificial Krook damping operator is often used to prevent profile relaxation in the simulation region.

turbulence experiences the same drive and background plasma conditions at opposing sides of the domain, so it should be statistically identical. This is also true in the binormal direction (i.e. the direction within the flux surface and perpendicular to the magnetic field). For simplicity, instead of generating statistically identical turbulence for these boundaries, it is substituted with exactly identical turbulence by using periodic boundary conditions. Using exactly identical turbulence instead of statistically identical turbulence does have the potential to introduce unphysical correlations into the system. However, as long as the correlation length of the turbulence is much smaller than the domain, this will not occur. In practice, the validity of this can be determined by testing for convergence in the widths (radial and binormal) of the domain.

For nonlinear simulations, the boundary condition in the parallel direction is similar, but with one additional complicating factor – the effect of magnetic shear. In this work we will study the parallel boundary condition and its effect on convergence in the parallel and binormal sizes of the simulation domain. We will find that the domain length typically used by the community, one poloidal turn, can introduce unphysical turbulent correlations, directly affecting the accuracy of the results. Individual turbulent eddies can remain correlated across the entire parallel length of the domain and interact with themselves, which is unphysical unless the domain corresponds to the full flux surface. This was first discussed in the original flux tube model paper (Beer *et al.* 1995) in the context of gyrofluid simulations with adiabatic electrons. Subsequent gyrokinetic studies (Candy 2005; Waltz *et al.* 2006; Dominski *et al.* 2015) investigated the ability of this mechanism to generate localized, steady corrugations in the background profiles. More recent work considered the self-interaction of linear eigenmodes and revealed that their parallel self-interaction can even interfere with the numerical convergence of nonlinear simulations with respect to the domain size (Weigl *et al.* 2018; C.J. *et al.* 2020)². Additionally, it has been explored in gyrokinetic simulations of low magnetic shear stellarators (Faber *et al.* 2018; Martin *et al.* 2018). This paper will show that parallel self-interaction can be eliminated by lengthening the simulation domain in the parallel direction to multiple poloidal turns or making the binormal width of the domain large. It is the first work to demonstrate that these two methods produce the same result and presents a novel, simple test for the presence of parallel self-interaction (i.e. the radial dependence of the parallel correlation function shown in figure 7).

2. The parallel boundary condition

The parallel boundary condition generally used by nonlinear flux tube simulations is called the ‘twist-and-shift’ condition (Beer *et al.* 1995). As with the perpendicular boundaries, it requires the two ends of the flux tube to have statistically identical turbulence. Finding statistically identical turbulence is non-trivial because toroidicity and plasma shaping mean that the statistical properties of turbulence change with poloidal angle. Thus, to be statistically identical the two ends of the flux tube must be at the same poloidal angle χ , although, due to axisymmetry, they are still free to be at different toroidal angles ζ . This constraint requires the length of the domain to be an integer number of poloidal circuits, which we represent by N_{pol} .

²Note that the self-interaction considered in these previous works (i.e. the nonlinear interaction of a linear eigenmode with itself) is closely related to the self-interaction considered here (i.e. the interaction of a nonlinear turbulent eddy with itself), but the two are distinct. Most notably, in the absence of magnetic shear, linear eigenmodes would not exhibit self-interaction, but nonlinear eddies still could.




FIGURE 2. An example flux surface showing two poloidal turns of a magnetic field line (red arrows) and a global perturbation (grey and white stripes) with toroidal and poloidal mode numbers of 9 and 21 respectively. Also shown is a $N_{\text{pol}} = 2$ flux tube (thin blue and purple lines) with the parallel boundary condition applied between the two ends (thick black horizontal lines). As shown, χ is taken to be a straight-field line poloidal angle.

While N_{pol} is free to be any positive integer, it has become overwhelmingly standard in the fusion community to use $N_{\text{pol}} = 1$ (e.g. figure 1). While this is likely due to the desire to minimize the computational cost of simulations, there also have been concerns about the validity of flux tubes with $N_{\text{pol}} > 1$. For example, such simulations are often not ‘globally consistent’ (Scott 1998). This means that, when $N_{\text{pol}} > 1$, the narrow flux tube can permit Fourier modes that do not exist on the full, doubly periodic flux surface. However, in the true $\rho_* \rightarrow 0$ limit for which gyrokinetics is valid, the charged particles have no way of knowing that they are on a doubly periodic flux surface. Particles can never communicate information between the different poloidal turns by moving in the toroidal direction because the distance separating the different poloidal turns is proportional to the minor radius a , while the turbulent correlation length is proportional to the ion gyroradius ρ_i (see figure 2). Thus, the distance separating the poloidal turns is an asymptotically large number of turbulent correlation lengths, so the particles will perceive the perpendicular direction within the flux surface as infinite in the $\rho_* \equiv \rho_i/a \rightarrow 0$ limit. This means that it is not relevant which modes are allowed on a doubly periodic flux. While we should still be concerned about respecting any periodicity in the parallel direction (i.e. do not use a domain with $N_{\text{pol}} = 2$ to model a $q = 1$ surface), violating global perpendicular periodicity will only introduce an error that is small in $\rho_* \ll 1$. This is intuitive – global consistency should be unimportant in local simulations. Another concern is that flux tubes with $N_{\text{pol}} > 1$ include multiple locations at the same poloidal angle, which must all have the same value of the zonal flows (i.e. fluctuations with a binormal wavenumber $k_y = 0$). Fortunately, it appears that the physics of the gyrokinetic model ensures that this is fairly well satisfied in the simulations we have run (e.g. figure 3). However, we should point out that we have no proof that this must always be the case, so it remains an open issue.

Unlike with the perpendicular boundary conditions, substituting statistical periodicity with exact periodicity across the parallel boundary is made more complicated by the effect of magnetic shear. In local simulations, the safety factor profile is approximated




FIGURE 3. The zonal amplitude of the turbulent electrostatic potential ϕ as a function of the time t , as calculated by a gyrokinetic simulation with the parameters of tables 1 and 2 using kinetic electrons, a binormal domain width $L_y = 125\rho_i$, and $N_{\text{pol}} = 3$. The colour (i.e. black or grey) indicates two different values of the radial wavenumber (i.e. $k_x\rho_i \in \{0.05, 0.1\}$ respectively) and the symbol (i.e. line, circle or cross) indicates the three different toroidal locations at the outboard midplane (i.e. $\chi \in \{-2\pi, 0, 2\pi\}$). Here, the reference velocity $v_r \equiv \sqrt{T_r/m_r}$, temperature T_r and mass m_r , are taken to be the ion values and e is the proton charge.

Coordinate	Grid range		Number of grid points	
	Adiabatic	Kinetic	Adiabatic	Kinetic
x/ρ_r	$[-100, 100)$	$[-65, 65)$	$256N_{\text{pol}}$	$128N_{\text{pol}}$
y/ρ_r	$[-L_y/2, L_y/2)/\rho_r$	$[-L_y/2, L_y/2)/\rho_r$	$64L_y/L_{y0}$	$32L_y/L_{y0}$
χ	$[-N_{\text{pol}}\pi, N_{\text{pol}}\pi)$	$[-N_{\text{pol}}\pi, N_{\text{pol}}\pi)$	$18N_{\text{pol}}$	$18N_{\text{pol}}$
v_{\parallel}/v_s	$[-3, 3]$	$[-3, 3]$	32	32
$\sqrt{\mu}/(T_s/B_r)$	$(0, 3)$	$(0, 3)$	8	8
$t/(a/v_r)$	$[500, 1500]$	$[150, 800]$	time step < CFL limit	

TABLE 1. The nominal GENE (Jenko *et al.* 2000) coordinate grids used for the scans in L_y and N_{pol} , where $L_{y0} \equiv 125\rho_r$ is a common value used by the community. Note that all grids are equally spaced, the reference length a is the tokamak minor radius, the reference magnetic field B_r is the toroidal field in the flux tube at R_0 , R_0 is the average major radius of the flux tube, and the CFL condition is explained in Courant, Friedrichs & Lewy (1967). The other reference quantities indicated by the subscript r are taken to be the ion quantities (except for the electron-scale simulation in figure 6) and $v_s = \sqrt{2T_s/m_s}$ is the thermal velocity of species s .

to be linear across the flux tube, i.e.

$$q(x) = q_0 + \frac{dq}{dx}(x - x_0) = q_0 \left(1 + \hat{s} \frac{x - x_0}{x_0} \right), \quad (2.1)$$

where q is the safety factor, $x \equiv r$ is the minor radial coordinate within the domain, $\hat{s} \equiv (x_0/q_0) dq/dx$ is the magnetic shear and the subscript 0 indicates the quantity is

Parameter	Value	Parameter	Value
Minor radius of flux tube, x_0/a	0.54	Major radius, R_0/a	3.0
Safety factor, q_0	1.4	Magnetic shear, \hat{s}	0.8
Temperature gradient, a/L_{Ts}	2.3	Density gradient, a/L_{ns}	0.733
Ion- e^- mass ratio, m_i/m_e	3672	Ion- e^- temperature ratio, T_i/T_e	1.0
Effective ion charge, Z_{eff}	1.0	Collision frequency, $\nu_{ss'}$	0.0
Fourth-order v_{\parallel} hyperdiffusion, $\epsilon_{v_{\parallel}}$	0.2	Fourth-order z hyperdiffusion, ϵ_c	0.05

TABLE 2. The nominal Cyclone Base Case (Dimits *et al.* 2000) parameters used for the scans in L_y and N_{pol} . Simulations with kinetic electrons use $\epsilon_{v_{\parallel}} = \epsilon_c = 0.5$ (Pueschel, Dannert & Jenko 2010) and a plasma beta of $\beta = 0.001$ for numerical reasons.

evaluated at the centre of the radial domain. This means that the two ends of the flux tube are deformed into parallelograms that have opposite tilts. Hence, the shapes do not overlap nicely, which must be resolved by the parallel boundary condition (see figure 4). In essence, the parallel boundary condition determines how the magnetic field lines are connected between these two parallelograms. The twist-and-shift condition typically implemented by gyrokinetic codes starts by using axisymmetry to justify toroidally shifting the two parallelograms until they are centred on one another (Beer *et al.* 1995). Then, again using axisymmetry, we can make periodic copies of the flux tube, shifted in the toroidal direction, and completely cover the opposing parallelogram (see figure 4). Thus, any turbulent structure that extends past one end of the domain will be copied into the other end while maintaining the twisting caused by magnetic shear³.

One consequence of this boundary condition is that it forces at least one flux surface in the domain to have field lines that close on themselves after going through the parallel boundary just once or, equivalently, after just N_{pol} poloidal turns. Given that we have chosen to centre the parallelograms with respect to one another, the flux surface in the centre of the radial domain will always be such a surface (see figure 5). However, figures 4 and 5 show that there can be more than one – sometimes there are many. We will indicate the total number of surfaces with field lines that close on themselves after just one pass through the parallel boundary with N_{si} . The subscript *si* stands for ‘self-interaction’ because these surfaces will exhibit the strongest parallel self-interaction. Now if $N_{\text{pol}} = 1$, the parallel boundary condition will be applied after one poloidal turn, so these field lines will close after just one poloidal turn. Accordingly, we will call them ‘pseudo-integer’ surfaces – they close on themselves after one poloidal turn like a normal integer surface, but they are usually an artefact of the boundary condition and do not actually correspond to integer flux surfaces in the real device. Additionally, we see that there will be ‘pseudo-rational’ surfaces at other radial locations, where the field lines exactly close on themselves after two or more poloidal turns. For example, the ‘pseudo-half-integer’ surfaces close on themselves after two poloidal turns and will always occur midway between two neighbouring pseudo-integer surfaces. Alternatively, if N_{pol} is increased from 1 to 2, all of the pseudo-integer surfaces will be replaced by pseudo-half-integer surfaces. Thus, the lowest-order rational surfaces in an $N_{\text{pol}} = 2$ simulation will be pseudo-half-integer and there will be N_{si} of them.

³For a mathematical explanation of the parallel boundary condition in real space see appendix A or for a Fourier-space formulation see Beer *et al.* (1995).

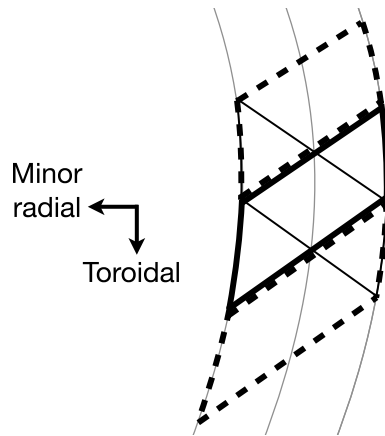


FIGURE 4. The shape of the parallel boundary (thin solid black parallelogram) on the inboard midplane of the example flux tube shown in figure 1. The parallel boundary condition is supplied by shifting the other end of the flux tube (thick solid black parallelogram) toroidally until it is centred on the first and then making periodic copies offset in the toroidal direction (two thick dashed black parallelograms). The flux surfaces are also shown (thin solid grey lines).




FIGURE 5. A Poincaré puncture plot at the outboard midplane for the flux tube from figures 1 and 4, which has $N_{\text{pol}} = 1$. Following a magnetic field line that starts at different radial positions (smallest filled circles) will take us to different toroidal locations after one poloidal circuit of the device (smaller empty circles) and two poloidal circuits (large empty circles). Two radial locations have field lines that close on themselves after passing through the parallel boundary just once, so $N_{\text{si}} = 2$.

These pseudo-rational surfaces are also involved in an important constraint that arises from the combination of the parallel and radial boundary conditions – the aspect ratio of the domain becomes discretized (i.e. the radial domain width L_x divided by the binormal domain width L_y). This can be found from a derivation in real space (see appendix A), a derivation in Fourier space (Beer *et al.* 1995) or by the following explanation. First, note that radial periodicity implies an equivalence between pairs of field lines, one on each of the two radial boundaries. In other words, any field line at $x = x_0 - L_x/2$ has a matching field line at $x = x_0 + L_x/2$ with which it is equivalent. These field lines are the same and must remain matched as we move

in the parallel direction. Otherwise, a particle moving purely along a field line would find itself transported across field lines. However, from figures 4 and 5, we see that the twist-and-shift condition, which is applied at the parallel ends of the domain, connects different field lines at different toroidal locations. This can cause problems. If we are not careful, the twist-and-shift boundary condition will connect the two field lines from a matching pair to two other field lines that are not matching. This would cause spurious cross-field transport. We must ensure that, given any matching pair of field lines, twist-and-shift always connects them to two other field lines that also form a matching pair. From figures 4 and 5, we see that this will only be achieved if

$$\Delta\zeta = N_{\text{si}}L_\zeta, \quad (2.2)$$

where $\Delta\zeta$ is the total shift in toroidal angle caused by magnetic shear along the length of the flux tube, N_{si} is some integer, and L_ζ is the width of the flux tube in toroidal angle such that $L_\zeta = 2\pi$ corresponds to the full toroidal domain. For the example shown in figure 4, we can see that $N_{\text{si}} = 2$ because the radial boundaries of the thick solid black parallelogram are offset in the toroidal direction by $1 \times L_y$ and this represents half the total shift (i.e. the shift in the range $\chi \in [-\pi N_{\text{pol}}, 0]$, but not $\chi \in [0, \pi N_{\text{pol}}]$). Additionally, figure 5 shows that, for this example, the twist-and-shift condition connects every field line on the radial boundary to itself. This trivially ensures that matching pairs remain matching as particles move across the parallel boundary. If $N_{\text{si}} = 1$, the field lines on the radial boundary would no longer close on themselves, but instead to a different matching pair offset in the y direction by $L_y/2$.

To see how equation (2.2) discretizes the domain aspect ratio, we invert the definition of the binormal coordinate

$$y(x, \zeta, \chi) \equiv \frac{x_0}{q_0}(q(x)\chi - \zeta) \quad (2.3)$$

(where χ is a straight-field line poloidal angle) and use (2.1) to calculate the toroidal shift across the sheared domain to be

$$\Delta\zeta \equiv \zeta \left(x = x_0 + \frac{L_x}{2}, y = y_0, \chi = 2\pi N_{\text{pol}} \right) - \zeta \left(x = x_0 - \frac{L_x}{2}, y = y_0, \chi = 2\pi N_{\text{pol}} \right) \quad (2.4)$$

$$= 2\pi N_{\text{pol}} \frac{dq}{dx} L_x = 2\pi N_{\text{pol}} \frac{q_0}{x_0} \hat{s} L_x, \quad (2.5)$$

where y_0 is any arbitrary binormal location. Since

$$L_y = \frac{x_0}{q_0} L_\zeta \quad (2.6)$$

follows from evaluating equation (2.3) at constant x and χ , we can use (2.2) and (2.5) to see that the possible values for the radial domain width are

$$L_x = \frac{N_{\text{si}}}{2\pi N_{\text{pol}} |\hat{s}|} L_y. \quad (2.7)$$

Thus, we find that L_x must be an integer multiple of the distance between lowest-order pseudo-rational surfaces. Importantly, we see that if the magnetic shear is very small, the radial width will be forced to be large because we cannot decrease N_{si} below 1.

This causes simulations with very low, but finite magnetic shear to become expensive. Equation (2.7) also shows that N_{si} can be used to specify the radial width of the domain. In fact, N_{si} is a common input parameter to gyrokinetic codes and goes by various names such as `nexc` in GENE, `jtwest` in GS2, `IKXSPACE` in GKW and `m_j` in GKV.

Lastly, it is important to note that the pseudo-rational surfaces created by the parallel boundary condition can be physical (Waltz *et al.* 2006; Weikl *et al.* 2018; C.J. *et al.* 2020). For example, a $q = 3$ surface really closes on itself after one poloidal turn and should be modelled using $N_{\text{pol}} = 1$, while a $q = 5/2$ surface should use $N_{\text{pol}} = 2$. Moreover, the physics occurring at these surfaces may be crucial for modelling certain behaviour such as microtearing modes (Hazeltine, Dobrott & Wang 1975; Guttenfelder *et al.* 2011). However, to accurately model the overall impact of self-interaction from rational surfaces, the spacing between pseudo-rational surfaces must also be correct. For example, the domain shown in figure 5 has $N_{\text{si}} = 2$ pseudo-integer surfaces, so, to be physical, the radial width of the simulation domain must correspond to twice the distance between the actual integer surfaces in the device⁴. To put it another way, the simulation must have the same number of ion gyroradii separating the integer surfaces as the experiment does. Unfortunately, accomplishing this turns out to require a simulation domain that corresponds to the full flux surface (as is discussed further at the end of §4). To see this, take the change in safety factor across the simulation domain $\Delta q = L_x dq/dx$ and set it equal to the number of integer surfaces in the simulation $\Delta q = N_{\text{si}}$. Then, substituting (2.6) and (2.7) with $N_{\text{pol}} = 1$ shows that $L_\zeta = 2\pi$ (i.e. the toroidal domain size is equal to the entire toroidal domain). Thus, we see that accurately modelling self-interaction on rational surfaces requires full flux surface simulations with radial domain widths on the scale of the tokamak minor radius a . Flux tubes of this size offer little computational savings compared to global simulations, but if we try to get by with a smaller domain of $L_\zeta = 2\pi/N$, it will have a radial density of pseudo-integer surfaces that is N times larger than reality.

3. Neglecting pseudo-rational surfaces

While the pseudo-rational surfaces created by the parallel boundary condition are unphysical unless the domain size is very large, this is not always a problem. Typically gyrokinetic simulations are used to model plasma far from low-order rational surfaces. To do this accurately we do not need to ensure that the simulation has the correct spacing of rational surfaces, we just need to ensure that their presence is negligible. Broadly speaking, this can be accomplished by verifying that the results of the simulation do not depend on the spacing of pseudo-rational surfaces. From (2.7), we see that the spacing is given by

$$\frac{L_x}{N_{\text{si}}} = \frac{L_y}{2\pi N_{\text{pol}} |\hat{\delta}|}, \quad (3.1)$$

so we can vary either L_y or N_{pol} and test for convergence. Additionally, we can specifically check the low-order pseudo-rational surfaces to see if they are exhibiting any unusual behaviour. The low-order pseudo-rational surfaces are only distinct from the other surfaces in that they close on themselves after a small number of

⁴Or more precisely, twice the distance between the integer surfaces in the device if the safety factor profile was given by (2.1).




FIGURE 6. The time-averaged zonal $E \times B$ shearing rate (i.e. $\partial^2 \phi / \partial x^2$ averaged over time and y) as a function of minor radial location within the flux tube for slab simulations (though toroidal simulations with kinetic electrons look similar (Dominski *et al.* 2015; C.J. *et al.* 2020)). Data are shown for ion-scale simulations using kinetic electrons with $m_i/m_e = 3672$ (black thick solid), kinetic electrons with $m_i/m_e = 918$ (red thick solid) and adiabatic electrons (grey thick dashed). An electron-scale simulation is shown with adiabatic ions (blue thin solid). The light grey regions indicate the portion significantly affected by pseudo-integer surfaces, while pseudo-half-integer surfaces are located half-way in between. The parameters used are given in tables 1 and 2, except $a/L_{Ts} = 10$, $a/L_{ns} = 2.0$ and several of the geometry parameters are not needed. Note the reference magnetic field B_r and reference length a become arbitrary in slab geometry.

poloidal turns. If the parallel correlation length of the turbulence is sufficiently long, an eddy can interact with itself along the field line, thereby altering its statistical properties. Hence, to ensure that their presence is not affecting things, we can verify that the flux tube has statistically homogeneous turbulence in the perpendicular plane.

For example, the observation of localized, steady $E \times B$ flow shear layers (Candy 2005; Waltz *et al.* 2006; Dominski *et al.* 2015; Weigl *et al.* 2018; C.J. *et al.* 2020) is a clear indicator that pseudo-rational surfaces may be adversely affecting the accuracy of many present-day simulations. These structures are a general consequence of self-interaction and are a fairly universal in conventional tokamak simulations with kinetic electrons. Although it is difficult to see them in tokamak simulations with adiabatic electrons, figure 6 shows that they can be found by using a slab geometry (instead of toroidal geometry). While they have been identified as an electron-scale phenomenon in linear simulations (Dominski *et al.* 2015), in nonlinear simulations they manifest as fairly narrow ion-scale structures. Figure 6 demonstrates that their width does not scale with the electron mass in slab ion-scale simulations and the same behaviour has been seen in previous toroidal ion-scale simulations (Dominski *et al.* 2017). Furthermore, figure 6 shows that similar structures can be found at electron scales in slab electron temperature gradient simulations.

These $E \times B$ flow shear layers are particularly interesting because they are constant in time – i.e. the plasma is spontaneously moving momentum around within the flux tube in order to modify the steady background flow profile. Like large-scale steady $E \times B$ flows, the tokamak will set up a parallel flow in order to convert the $E \times B$ flow layers into a purely toroidal flow (Abel *et al.* 2013). Such flow represents intrinsic rotation, which is generally prohibited by an underlying symmetry of the gyrokinetic model (Peeters *et al.* 2005; Parra, Barnes & Peeters 2011; Sugama *et al.* 2011). Thus, there must be a symmetry-breaking mechanism present. One plausible mechanism is variation in turbulence characteristics (Parra & Barnes 2015). This could be introduced

by the existence of pseudo-rational surfaces (e.g. the turbulence on integer surfaces is different than on the neighbouring irrational surfaces). Normally, this symmetry-breaking effect is small in $\rho_* \ll 1$ because the gradual change in background gradients is the usual cause for the variation in turbulence characteristics. However, the variation caused by the pseudo-integer surfaces occurs on the scale of the gyroradius, so it is not small in ρ_* . There are only two physical effects that cause variation in turbulence characteristics to break the symmetry (see Section 9 of Parra & Barnes (2015)): finite gyroradius effects and the radial magnetic drifts. Since figure 6 shows the shear layers persist in slab geometry, we believe that the finite gyroradius effect is the dominant symmetry-breaking effect. This is intuitive as the symmetry breaking occurs on a small scale and appears to be dipolar (i.e. a region of positive momentum flux is always next to a region of negative), so the radial drift orbits should average over it more effectively because they are significantly larger than the gyroradius.

Importantly, these flow shear layers are just one possible symptom of self-interaction at pseudo-rational surfaces. More generally, we can investigate spurious correlations by looking at the two-point parallel correlation function

$$C_{\parallel}(x, y, \chi_1, \chi_2) \equiv \frac{\langle \phi_{NZ}(x, y, \chi_1, t) \phi_{NZ}(x, y, \chi_2, t) \rangle_t}{\sqrt{\langle \phi_{NZ}^2(x, y, \chi_1, t) \rangle_t \langle \phi_{NZ}^2(x, y, \chi_2, t) \rangle_t}}, \quad (3.2)$$

where the subscript *NZ* signifies the non-zonal portion of the quantity and $\langle \dots \rangle_u$ indicates an average over any coordinate *u*. The quantity C_{\parallel} indicates the degree of correlation between two points χ_1 and χ_2 on the same field line (i.e. at constant *x* and *y*). Figure 7 shows the *y*-averaged correlation between adjacent outboard and inboard midplanes $\langle C_{\parallel}(x, y, \chi_1 = 0, \chi_2 = \pi) \rangle_y$ for various toroidal simulations. Note that all simulations with $N_{\text{pol}} = 1$ have sharp spikes in the parallel correlation at the locations of pseudo-integer surfaces. Thus, even though toroidal simulations with adiabatic electrons did not display flow shear layers at pseudo-integer surfaces, they do exhibit spikes in parallel correlation. We see that increasing L_y increases the spacing between pseudo-integer surfaces, in accordance with (3.1). However, we also notice that neither the width, nor the height of the spikes changes with L_y . Thus, as we increase L_y the regions affected by pseudo-integer surfaces occupy a smaller and smaller fraction of the domain. This indicates one way to eliminate the effect of pseudo-rational surfaces – increase L_y until convergence is achieved. If L_y is sufficiently large, the pseudo-rational surfaces will have a negligible impact on all volume-averaged quantities. Effectively we are ‘diluting’ away their influence.

Figure 7 also shows a second strategy to eliminate pseudo-integer surfaces – increase N_{pol} until convergence is achieved. If the flux tube is longer, a turbulent eddy has to remain correlated over a longer distance in order to ‘bite its own tail’. In other words, the $N_{\text{pol}} = 2$ has no pseudo-integer surfaces, only pseudo-half-integer surfaces. Accordingly, we see that, compared to the $N_{\text{pol}} = 1$ simulation, the $N_{\text{pol}} = 2$ case has much smaller peaks at its lowest order pseudo-rational surfaces, which are more closely spaced. This suggests that, given the parameters used for these simulations, a small number of poloidal turns should be sufficient to eliminate the effect of pseudo-rational surfaces and achieve convergence. This is consistent with the original convergence study of Beer *et al.* (1995), which found that two or three poloidal turns was sufficient to achieve convergence in gyrofluid simulations with adiabatic electrons. However, it was not obvious that kinetic and adiabatic electron simulations would behave similarly because, in linear simulations, kinetic electrons enable modes that are very extended along the magnetic field line (i.e. ‘giant tails’ in the ballooning envelope) (Hallatschek & Dorland 2005).




FIGURE 7. The parallel correlation between inboard and outboard locations following a field line at a given minor radial location within the flux tube. The base simulation has kinetic electrons with $N_{\text{pol}} = 1$ and $L_y = 250\rho_i$ (black thick solid). The other three curves are produced by changing the base simulation to either have adiabatic electrons (grey thick solid), $L_y = 500\rho_i$ (red thin solid) or $N_{\text{pol}} = 2$ (blue thick dashed). Note the $L_y = 500\rho_i$ simulation has only a single pseudo-integer surface, which occurs at $x = 0$. The $N_{\text{pol}} = 2$ simulation has no pseudo-integer surfaces, but has pseudo-half-integer surfaces at $x = \{0, \pm 25, \pm 50, \pm 75\}\rho_i$.

4. Resolution study

We will now use the gyrokinetic code GENE (Jenko *et al.* 2000; Görler *et al.* 2011) to perform a resolution study in L_y and N_{pol} , the two strategies to eliminate pseudo-integer surfaces. Figures 8(a) and 8(b) show the results using standard Cyclone Base Case parameters (Dimitis *et al.* 2000) with adiabatic electrons and kinetic electrons respectively (see table 2). The simulations with adiabatic electrons are considerably less expensive, so we are able to perform a more complete parameter scan and more rigorously ensure adequate resolution (see table 1). Regardless, both display qualitatively similar behaviour.

For adiabatic electrons, figure 8(a) shows that the fully converged value of the ion heat flux appears to be a bit less than 7 in gyroBohm units ($Q_{gB} = \rho_*^2 v_r n_i T_r$, where n_i is the ion number density). Such convergence can be achieved by simply increasing L_y to very large values while maintaining $N_{\text{pol}} = 1$. The same result can also be achieved by increasing N_{pol} , which somewhat alleviates our concerns about the validity of using $N_{\text{pol}} > 1$ (see § 2). However, this is only true when we are already using a sufficiently large value of L_y . If L_y is not large enough, the simulation will not be fully converged even if $N_{\text{pol}} \rightarrow \infty$. This is for the same reason that we need a sufficiently large domain in the x direction – self-interaction in the perpendicular direction. Completely independent of parallel self-interaction, it is still necessary to ensure that the turbulent correlation length in the y direction is much smaller than the domain width in the y direction (as explained in § 1). If this is not satisfied, convergence will be achieved as $N_{\text{pol}} \rightarrow \infty$ (because the parallel self-interaction will vanish), but it will not converge to the same value as when the binormal self-interaction is also eliminated. Note in figure 8 that the simulations that are not fully resolved in either L_y or N_{pol} have heat fluxes lower than the converged value. This indicates that both binormal and parallel self-interaction tend to reduce the ability of turbulence to transport energy. This may

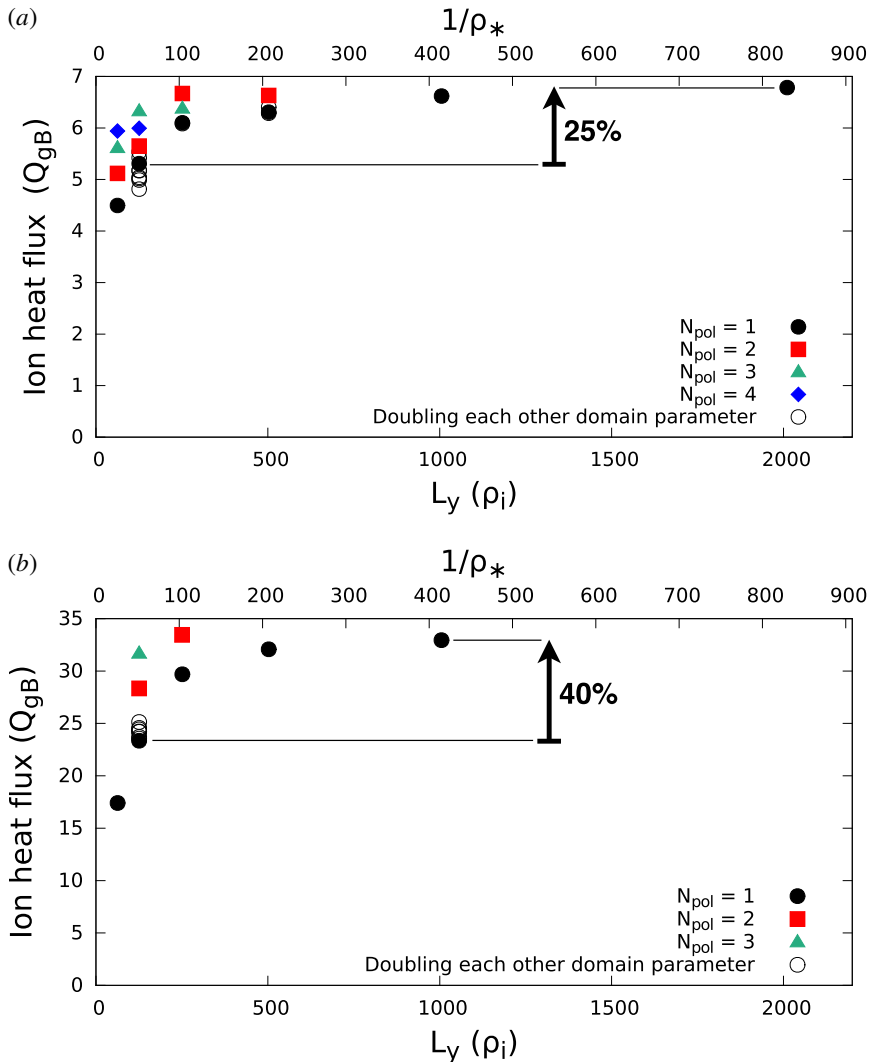


FIGURE 8. Resolution studies in L_y and N_{pol} with the parameters of tables 1 and 2 using adiabatic electrons (a) and kinetic electrons (b). The open black circles indicate simulations with $N_{pol} = 1$ that individually double the grid resolution or the domain width of the three other coordinates (i.e. x , v_{\parallel} and μ) as well as the grid resolutions in y and χ . The simulation at $L_y = 2000\rho_i$ was forced to have double the nominal x domain width (and thus double the number of x grid points) due to the aspect ratio quantization of (2.7). The top axis shows the value of ρ_* for which a given L_y corresponds to the full flux surface.

be because the correlations imposed by self-interaction reduce the degrees of freedom accessible to the turbulence, preventing it from behaving in ways that would more effectively drive transport.

Figure 8(a) shows that, using $L_y \approx 125\rho_i$ and $N_{pol} = 1$, as is common in the community (Watanabe *et al.* 2015), under-predicts the ion heat flux by approximately 25% in adiabatic electron simulations. While this is already concerning, figure 8(b) shows that the effect of self-interaction increases when using kinetic electrons. Instead of a 25% error, we see roughly 40%. This is consistent with figure 7,

which showed that the kinetic electron simulation had more pronounced spikes in the parallel correlation function. Moreover, the effect of self-interaction can be even bigger for other physical parameters. For example, Watanabe *et al.* (2015) shows that taking Cyclone Base Case parameters and lowering the magnetic shear from $\hat{s} = 0.8$ to 0.2 increases the error in adiabatic electron simulations from 25% to at least 50%. Going further, we simulated Cyclone Base Case parameters with adiabatic electrons and $\hat{s} = 0$ and found that changing $N_{\text{pol}} = 1$ to $N_{\text{pol}} = 2$ increased the heat flux by 70%. Simulations of stellarators with $\hat{s} = 0$ have even found that turbulence can be stable when using $N_{\text{pol}} = 1$, but unstable when $N_{\text{pol}} = 2$ (see figure 11 of Faber *et al.* (2018)). Clearly, self-interaction through the parallel boundary condition can be a significant effect in local gyrokinetic simulations.

An important practical result is the relative computational cost of the two routes – increasing N_{pol} or increasing L_y . This is briefly considered by Beer *et al.* (1995) in the context of gyrofluid simulations of a TFTR experimental shot, concluding that ‘it appears that faster convergence is obtained by allowing the domain to be longer than a parallel correlation length’. However, our conclusion from the above Cyclone Base Case simulations is the opposite – increasing L_y while maintaining $N_{\text{pol}} = 1$ was the cheapest way to reach convergence. For example, the $N_{\text{pol}} = 2$, $L_y = 250\rho_i$ simulation required roughly double the computational resources of the $N_{\text{pol}} = 1$, $L_y = 1000\rho_i$ simulation. This can be understood by studying how the coordinate system grids change in response to increasing L_y versus N_{pol} (see table 1).

We see that as L_y is increased the number of y grid points must be increased proportionally. Additionally, after L_y reaches a certain value, equation (2.7) forces us to increase L_x (and thus the number of x grid points) proportionally as well. However, for the parameters of tables 1 and 2, this was not necessary until $L_y > 1000\rho_i$. For simulations with lower values of \hat{s} , for which parallel self-interaction is expected to be more of a concern, this constraint would kick in at lower values of L_y . On the other hand, when increasing N_{pol} we clearly must increase the number of parallel grid points. Moreover, as explained in detail in Watanabe *et al.* (2015), properly resolving a longer parallel domain turns out to require more radial grid points, due to the fact that longer domains are twisted more by magnetic shear. Adding more radial modes decreases the time step of explicit codes like GENE because we are allowing smaller radial scales, which makes the CFL limit more constraining. In light of this, it is not surprising that simulations with $N_{\text{pol}} > 1$ are so costly – it is necessary to increase the parallel, radial and temporal resolutions proportionally with N_{pol} . Additionally, pushing to extreme values of L_y has the added advantage of ensuring that self-interaction in the y direction is completely eliminated.

While increasing L_y was the better way to reach convergence in our simulations, there are a few important caveats that make generalization difficult. First, at lower values of \hat{s} the domain aspect ratio quantization condition (i.e. (2.7)) kicks in at lower values of L_y , making the L_y route more costly. Second, codes that partially use implicit methods to advance in time may observe different computational costs than what we have seen in GENE, which uses only explicit methods, as they are not limited by the CFL condition. Third, adding perpendicular grid points (specifically when increasing L_x and L_y at fixed resolution) is fundamentally different than adding parallel grid points because the perpendicular modes are all coupled through a Fourier transform in the nonlinear term, while only some of the parallel grid points are connected via the parallel streaming term. Thus, the relative computational and communication costs may depend strongly on the computer architecture, memory and parallelization scheme. Finally, Scott (2001) and Watanabe *et al.* (2015) each present a clever way

of potentially getting around the need to increase the radial and temporal resolutions with N_{pol} . While neither have been implemented in GENE, they both could make $N_{\text{pol}} > 1$ simulations much more affordable, potentially more affordable than large L_y simulations. The first approach, the ‘flux tube train’ (Watanabe *et al.* 2015), has been implemented in the local gyrokinetic codes GKV (Watanabe & Sugama 2006) and stella (Barnes, Parra & Landreman 2019). The other approach, the ‘shifted metric’ coordinate system (Scott 2001), has not been implemented in any local gyrokinetic code and it is currently unclear if it is possible to implement in a radially periodic domain (Told & Jenko 2010).

As discussed in §2, the parallel self-interaction occurring in local simulations is physical when L_y corresponds to a full flux surface. The same is true of self-interaction in the y direction – if we are modelling a full flux surface, the toroidal periodicity of a real device will lead to self-interaction if the toroidal correlation length is comparable to the toroidal circumference. Thus, we can also view the $N_{\text{pol}} = 1$ points in figure 8 as a scan showing how the physical effect of self-interaction weakens with decreasing ρ_* ⁵. This is indicated on the top horizontal axes in figure 8. For context, the DIII-D tokamak (Luxon 2002) has $\rho_* \approx 1/300$ and a full flux surface simulation corresponds to

$$L_y = \frac{2\pi}{q_0} \frac{1}{\rho_*} \frac{x_0}{a} \rho_i \approx 700\rho_i. \quad (4.1)$$

Therefore, figure 8 shows that self-interaction is a finite ρ_* effect that we are usually justified in eliminating when modelling present-day large tokamaks. However, it might still play an important role in smaller machines or specific parameter regimes⁶.

Lastly, we performed a temperature gradient scan to determine how self-interaction affects the critical gradient and profile stiffness in Cyclone Base Case simulations. Figure 9 shows two data sets – one with a large value of L_y , where self-interaction has been mostly diluted away, and one with a small value of L_y , where self-interaction is still significant. We see that the difference in the ion heat flux between the two values of L_y is generally reduced as the temperature gradient decreases. This leads to similar estimates of the critical gradient for both sets of simulations, giving evidence that large L_y domain widths may not be necessary for simulations near the critical gradient. However, the self-interaction present in the simulations with smaller L_y reduces the profile stiffness, more significantly when using kinetic electrons.

5. Conclusions

In this work we have shown that turbulent self-interaction through the parallel boundary condition can significantly affect the results of local gyrokinetic simulations. Such self-interaction can be physical, but only when the simulation domain corresponds to a full flux surface for the device being modelled. Using a narrow flux tube to model a large device will artificially strengthen self-interaction, which can

⁵This convergence looks quantitatively similar to the system size investigations in McMillan *et al.* (2010) and Lin *et al.* (2012).

⁶Note that there are many other finite ρ_* effects apart from the self-interaction discussed here. Some of these are contained in current global gyrokinetic codes (e.g. profile shearing), while many others have only been derived recently (Parra & Calvo 2011) and are not in any code. Thus, while going to a full flux surface simulation will allow us to properly treat self-interaction, it is no guarantee that our overall simulation will be accurate for a small machine. On the contrary, the fact that self-interaction is important suggests that many of the other finite ρ_* effects may be important too.

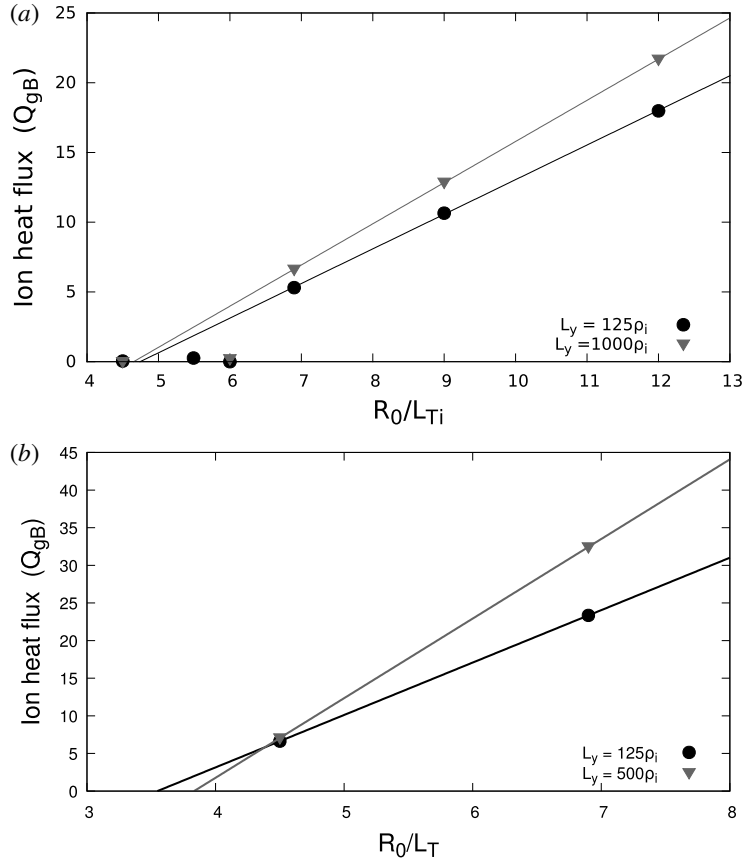


FIGURE 9. A critical gradient study using adiabatic electrons (a) or kinetic electrons (b) for a domain width that mostly eliminates self-interaction (grey triangles) and another for a domain width that allows significant self-interaction (black circles). All simulations use the parameters of tables 1 and 2 and best-fit linear trendlines are also shown.

reduce the heat flux. This implies that, if the parallel self-interaction is significant, the simulation will only properly model it for a single finite value of ρ_* . However, if self-interaction is negligible, the simulation is appropriate for any sufficiently small value of ρ_* . In other words, self-interaction is a finite ρ_* effect.

Thus, to achieve the true $\rho_* \ll 1$ limit assumed in deriving gyrokinetics, one would like to completely eliminate parallel self-interaction. We have shown in figure 8 that this can be achieved by increasing the width of the simulation domain in the binormal direction L_y and/or lengthening the simulation domain using N_{pol} . Currently available results suggest that self-interaction is stronger for kinetic electrons, low values of magnetic shear and strong turbulence drive (i.e. steep background gradients). This may be because these parameters tend to increase the parallel correlation length of the turbulence. To verify that self-interaction has been eliminated, one should pay special attention to convergence in the binormal domain width and check for spikes in the parallel correlations function (e.g. figure 7). Additionally, one should take care when using a resolution study done with adiabatic electrons to justify domain resolutions for simulations using kinetic electrons.

Acknowledgements

The authors would like to thank B. McMillan, M. Hardman, B. Faber, A. Peeters, F. Parra and M. Barnes for useful discussions pertaining to this work. This work was supported by the EUROfusion – Theory and Advanced Simulation Coordination (E-TASC). This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement no. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. We acknowledge the CINECA award under the ISCRA initiative, for the availability of high performance computing resources and support. Lastly, this work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID s863 and s956.

Appendix A. Derivation of aspect ratio quantization

To see why the aspect ratio L_x/L_y of a flux tube is quantized, we must start with the boundary conditions. The binormal boundary condition is

$$\phi(x, y + L_y, \chi) = \phi(x, y, \chi), \quad (\text{A } 1)$$

where y is defined by (2.3). Here we give the turbulent electrostatic potential ϕ as an example, but the same boundary conditions also hold for the electromagnetic perturbed fields as well as the gyrokinetic distribution function. The radial boundary condition is

$$\phi(x + L_x, y, \chi) = \phi(x, y, \chi), \quad (\text{A } 2)$$

where it is important that we are holding y constant in applying radial periodicity. Lastly, the parallel condition is more complicated and is called the ‘twist-and-shift’ (Beer *et al.* 1995). It is defined by taking parallel periodicity at constant toroidal angle ζ , according to

$$\phi(x, y(x, \zeta, \chi + 2\pi N_{\text{pol}}), \chi + 2\pi N_{\text{pol}}) = \phi(x, y(x, \zeta, \chi), \chi). \quad (\text{A } 3)$$

This relation is also depicted in figure 4. Using the definition of y given by (2.3) and linearizing the safety factor profile, we can see that

$$y(x, \zeta, \chi + 2\pi N_{\text{pol}}) = y(x, \zeta, \chi) + 2\pi N_{\text{pol}} x_0 + 2\pi N_{\text{pol}} \hat{s}(x - x_0). \quad (\text{A } 4)$$

Thus, equation (A 3) can be written as

$$\phi(x, y + 2\pi N_{\text{pol}} \hat{s}(x - x_0), \chi + 2\pi N_{\text{pol}}) = \phi(x, y, \chi), \quad (\text{A } 5)$$

where we have assumed that $2\pi N_{\text{pol}} x_0/L_y$ is very close to an integer $N_q \in \mathbb{Z}$ and then applied the binormal periodicity (i.e. (A 1)) N_q times. This is valid because the minor radial location of the centre of the flux tube x_0 is much larger than the binormal width L_y , so a slight radial shift ensures that the assumption is satisfied and does not affect the simulation.

The flux tube aspect ratio quantization condition fundamentally arises from the combination of the parallel and radial boundary conditions. As explained in the main text, if field lines are not connected properly across these two boundaries, parallel motion can lead to an artificial jump in the binormal direction. We begin by writing

the radial boundary condition given by (A 2). Then, replace both sides of the equation using the parallel boundary condition given by (A 5) to find

$$\begin{aligned} & \phi(x + L_x, y + 2\pi N_{\text{pol}} \hat{s}(x + L_x - x_0), \chi + 2\pi N_{\text{pol}}) \\ &= \phi(x, y + 2\pi N_{\text{pol}} \hat{s}(x - x_0), \chi + 2\pi N_{\text{pol}}). \end{aligned} \quad (\text{A } 6)$$

Next, we again use the radial boundary condition given by (A 2) to replace the right side of (A 6) and show

$$\begin{aligned} & \phi(x + L_x, y + 2\pi N_{\text{pol}} \hat{s}(x + L_x - x_0), \chi + 2\pi N_{\text{pol}}) \\ &= \phi(x + L_x, y + 2\pi N_{\text{pol}} \hat{s}(x - x_0), \chi + 2\pi N_{\text{pol}}). \end{aligned} \quad (\text{A } 7)$$

This expression can only be true if the $2\pi N_{\text{pol}} \hat{s} L_x$ term in the second argument on the left side can be removed using repeated application of binormal periodicity (given by (A 1)). Thus, we require that

$$2\pi N_{\text{pol}} |\hat{s}| L_x = N_{\text{si}} L_y, \quad (\text{A } 8)$$

for some integer $N_{\text{si}} \in \mathbb{Z}$. This is the flux tube aspect ratio quantization condition given by (2.7). If it is not satisfied, field line identity will not be consistent across the parallel boundary of the flux tube and spurious cross-field transport will be introduced into the simulation. A particle travelling purely along a field line can find itself on a different field line after passing through the parallel boundary, thereby enabling fictitious correlations between distant regions within the flux tube. Note that this is a separate issue from fictitious correlations that can be introduced by pseudo-integer surfaces (i.e. the primary focus of this paper). By applying (2.5) and (2.6) to (A 8), we can derive (2.2) and show that the radial domain width must correspond to an integer number of lowest-order pseudo-rational surfaces.

REFERENCES

- ABEL, I., PLUNK, G., WANG, E., BARNES, M., COWLEY, S., DORLAND, W. & SCHEKOCHIHIN, A. 2013 Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport, and energy flows. *Rep. Prog. Phys.* **76**, 116201.
- BARNES, M., PARRA, F. & LANDREMAN, M. 2019 *stella*: An operator-split, implicit–explicit δf -gyrokinetic code for general magnetic field configurations. *J. Comput. Phys.* **391**, 365.
- BEER, M., COWLEY, S. & HAMMETT, G. 1995 Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Phys. Plasmas* **2** (7), 2687.
- BRIZARD, A. & HAHM, T. 2007 Foundations of nonlinear gyrokinetic theory. *Rev. Mod. Phys.* **79** (2), 421.
- C.J., AJAY, BRUNNER, S., MCMILLAN, B., BALL, J., DOMINSKI, J. & MERLO, G. 2020 How the self-interaction mechanism affects zonal flow drive and the convergence of turbulent transport in flux-tube simulations. *J. Plasma Phys.* (in preparation).
- CANDY, J. 2005 Beta scaling of transport in microturbulence simulations. *Phys. Plasmas* **12** (7), 072307.
- CANDY, J. & WALTZ, R. 2003 An Eulerian gyrokinetic-Maxwell solver. *J. Comput. Phys.* **186** (2), 545.
- CATTO, P. 1978 Linearized gyro-kinetics. *Plasma Phys.* **20** (7), 719.
- COURANT, R., FRIEDRICHS, K. & LEWY, H. 1967 On the partial difference equations of mathematical physics. *IBM J. Res. Dev.* **11** (2), 215–234.

- DIMITS, A., BATEMAN, G., BEER, M., COHEN, B., DORLAND, W., HAMMETT, G., KIM, C., KINSEY, J., KOTSCHENREUTHER, M., KRITZ, A. *et al.* 2000 Comparisons and physics basis of tokamak transport models and turbulence simulations. *Phys. Plasmas* **7**, 969.
- DOMINSKI, J., BRUNNER, S., GOERLER, T., JENKO, F., TOLD, D. & VILLARD, L. 2015 How non-adiabatic passing electron layers of linear microinstabilities affect turbulent transport. *Phys. Plasmas* **22** (6), 062303.
- DOMINSKI, J., MCMILLAN, B., BRUNNER, S., MERLO, G., TRAN, T.-M. & VILLARD, L. 2017 An arbitrary wavelength solver for global gyrokinetic simulations. Application to the study of fine radial structures on microturbulence due to non-adiabatic passing electron dynamics. *Phys. Plasmas* **24** (2), 022308.
- FABER, B., PUESCHEL, M., TERRY, P., HEGNA, C. & ROMAN, J. 2018 Stellarator microinstabilities and turbulence at low magnetic shear. *J. Plasma Phys.* **84** (5), 905840503.
- FRIEMAN, E. & CHEN, L. 1982 Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *Phys. Fluids* **25**, 502.
- GÖRLER, T., LAPILLONNE, X., BRUNNER, S., DANNERT, T., JENKO, F., MERZ, F. & TOLD, D. 2011 The global version of the gyrokinetic turbulence code GENE. *J. Comput. Phys.* **230** (18), 7053.
- GUTTENFELDER, W., CANDY, J., KAYE, S., NEVINS, W., WANG, E., BELL, R., HAMMETT, G., LEBLANC, B., MIKKELSEN, D. & YUH, H. 2011 Electromagnetic transport from microtearing mode turbulence. *Phys. Rev. Lett.* **106** (15), 155004.
- HALLATSCHKE, K. & DORLAND, W. 2005 Giant electron tails and passing electron pinch effects in tokamak-core turbulence. *Phys. Rev. Lett.* **95**, 055002.
- HAZELTINE, R., DOBROTT, D. & WANG, T. 1975 Kinetic theory of tearing instability. *Phys. Fluids* **18** (12), 1778–1786.
- JENKO, F., DORLAND, W., KOTSCHENREUTHER, M. & ROGERS, B. 2000 Electron temperature gradient driven turbulence. *Phys. Plasmas* **7** (5), 1904.
- LIN, Z., ETHIER, S., HAHM, T. & TANG, W. 2012 Verification of gyrokinetic particle simulation of device size scaling of turbulent transport. *Plasma Sci. Technol.* **14** (12), 1125.
- LUXON, J. 2002 A design retrospective of the DIII-D tokamak. *Nucl. Fusion* **42** (5), 614.
- MARTIN, M., LANDREMAN, M., XANTHOPOULOS, P., MANDELL, N. & DORLAND, W. 2018 The parallel boundary condition for turbulence simulations in low magnetic shear devices. *Plasma Phys. Control. Fusion* **60** (9), 095008.
- MCMILLAN, B., LAPILLONNE, X., BRUNNER, S., VILLARD, L., JOLLIET, S., BOTTINO, A., GÖRLER, T. & JENKO, F. 2010 System size effects on gyrokinetic turbulence. *Phys. Rev. Lett.* **105** (15), 155001.
- PARRA, F. & BARNES, M. 2015 Intrinsic rotation in tokamaks: theory. *Plasma Phys. Control. Fusion* **57** (4), 045002.
- PARRA, F., BARNES, M. & PEETERS, A. 2011 Up-down symmetry of the turbulent transport of toroidal angular momentum in tokamaks. *Phys. Plasmas* **18** (6), 062501.
- PARRA, F. & CALVO, I. 2011 Phase-space Lagrangian derivation of electrostatic gyrokinetics in general geometry. *Plasma Phys. Control. Fusion* **53** (4), 045001.
- PEETERS, A., ANGIONI, C. *et al.* 2005 Linear gyrokinetic calculations of toroidal momentum transport in a tokamak due to the ion temperature gradient mode. *Phys. Plasmas* **12** (7), 072515.
- PUESCHEL, M., DANNERT, T. & JENKO, F. 2010 On the role of numerical dissipation in gyrokinetic Vlasov simulations of plasma microturbulence. *Comput. Phys. Commun.* **181** (8), 1428–1437.
- SCOTT, B. 1998 Global consistency for thin flux tube treatments of toroidal geometry. *Phys. Plasmas* **5** (6), 2334–2339.
- SCOTT, B. 2001 Shifted metric procedure for flux tube treatments of toroidal geometry: avoiding grid deformation. *Phys. Plasmas* **8** (2), 447–458.
- SUGAMA, H., WATANABE, T., NUNAMI, M. & NISHIMURA, S. 2011 Momentum balance and radial electric fields in axisymmetric and nonaxisymmetric toroidal plasmas. *Plasma Phys. Control. Fusion* **53** (2), 024004.
- TOLD, D. & JENKO, F. 2010 Applicability of different geometry approaches to simulations of turbulence in highly sheared magnetic fields. *Phys. Plasmas* **17** (4), 042302.

- WALTZ, R., AUSTIN, M., BURRELL, K. & CANDY, J. 2006 Gyrokinetic simulations of off-axis minimum- q profile corrugations. *Phys. Plasmas* **13** (5), 052301.
- WATANABE, T.-H. & SUGAMA, H. 2006 Velocity-space structures of distribution function in toroidal ion temperature gradient turbulence. *Nucl. Fusion* **46** (1), 24.
- WATANABE, T.-H., SUGAMA, H., ISHIZAWA, A. & NUNAMI, M. 2015 Flux tube train model for local turbulence simulation of toroidal plasmas. *Phys. Plasmas* **22** (2), 022507.
- WEIKL, A., PEETERS, A., RATH, F., SEIFERLING, F., BUCHHOLZ, R., GROSSHAUSER, S. & STRINTZI, D. 2018 The occurrence of staircases in ITG turbulence with kinetic electrons and the zonal flow drive through self-interaction. *Phys. Plasmas* **25** (7), 072305.