

Continuous-in-time approach to flow shear in a linearly implicit local δf gyrokinetic code

Nicolas Christen ^{*1}, Michael Barnes¹, and Felix I. Parra¹

¹Rudolf Peierls Centre for Theoretical Physics, University of Oxford,
Oxford OX1 3PU, UK

April 6, 2021

Abstract

A new algorithm for toroidal flow shear in a linearly implicit, local δf gyrokinetic code is described. Unlike the current approach followed by a number of codes, it treats flow shear continuously in time. In the linear gyrokinetic equation, time-dependences arising from the presence of flow shear are decomposed in such a way that they can be treated explicitly in time with no stringent constraint on the time step. Flow shear related time dependences in the nonlinear term are taken into account exactly, and time dependences in the quasineutrality equation are interpolated. Test cases validating the continuous-in-time implementation in the code **GS2** are presented. Lastly, nonlinear gyrokinetic simulations of a JET discharge illustrate the differences observed in turbulent transport compared to the usual, discrete-in-time approach. The continuous-in-time approach is shown, in some cases, to produce fluxes that converge to a different value than with the discrete approach. The new approach can also lead to substantial computational savings by requiring radially narrower boxes. At fixed box size, the continuous implementation is only modestly slower than the previous, discrete approach.

1 Introduction

In the core region of tokamaks, gradients of plasma parameters drive turbulence at micro-scales of the order of the ion or electron gyroradius. The resulting losses of particles, momentum and heat towards the vessel wall set a limit for the densities and temperatures that can be sustained in the device, which in turn limits the rate of fusion reactions. In current and future experiments such as JET and ITER, the core plasma is, in part, heated and fuelled by injecting energetic beams of neutral atoms of hydrogen isotopes (deuterium and tritium). As the beams penetrate into the

^{*}Contact: nicolas.christen@physics.ox.ac.uk

core, they apply a torque on the plasma and make it spin toroidally. Experimental, theoretical and computational results have shown that the presence of shear in the toroidal rotation can substantially affect turbulent transport: shear in the flow perpendicular to the mean magnetic field has been found to reduce transport, while a gradient in the parallel flow can enhance transport, see e.g. [1–13].

The first implementation of flow shear for local, δf gyrokinetic simulations was developed by Hammett *et al* [14] in the linearly implicit-in-time code GS2 [15], and was later applied to others including GENE [16], GKW [17] and CGYRO [18]. In this approach, the system of gyrokinetic and Maxwell equations is expressed in the frame that rotates and shears along with the flow, allowing the equations to be Fourier analysed across the mean magnetic field. The effective radial wave number in the lab frame – a constant without flow shear – becomes time-dependent in the presence of flow shear. To retain the benefits of using a code that is implicit in time, this approach to flow shear approximates such time-dependences by combining a nearest point approximation on a fixed grid, together with a wave number re-mapping method. This ensures that the computationally expensive part of the implicit scheme only needs to be executed at initialisation, rather than at every time step. However, many experiments only sustain modest levels of flow shear, for which the radial wave number evolves very slowly: simulations then produce long periods of time where the perpendicular flow shear has no effect, separated by discrete jumps in wave number. We wish to avoid this unphysical behaviour, and consider here an improved, continuous-in-time approach to flow shear.

This work is organised as follows. First, we give a brief overview in section 2 of the gyrokinetic ordering, the coordinates that are used, and the system of Fourier-analysed equations that has to be solved. In the next section, we summarise the main algorithm of GS2, as well as the discrete-in-time approach to flow shear developed by Hammett *et al*. Following this, we present a new algorithm that treats flow shear continuously over time. Finally, in section 4, we show numerical results obtained with GS2 that test the implementation of the continuous-in-time approach, and compare it to the discrete-in-time method.

2 Model

In this section, we give a brief overview of the standard orderings assumed in δf gyrokinetic fluxtube codes such as GS2 and present the equations solved in the code. We then explain how the presence of a sheared background toroidal flow modifies the spectral representation that is used in the plane perpendicular to the magnetic field.

2.1 Orderings and equations

The δf gyrokinetic description [19, 20] hinges on scale separations in both time and space: the fluctuations of interest are chosen to be much slower than the gyromotion of particles around magnetic field lines, and the gyroradii are assumed to be much smaller than the equilibrium length scale. The distribution function f can be de-

scribed as the sum of an equilibrium piece $\langle f \rangle$ and a fluctuating piece $\delta f = f - \langle f \rangle$, where $\langle \cdot \rangle$ denotes an ensemble average. The fluctuating part is ordered to be much smaller than the equilibrium quantity:

$$\frac{\delta f}{\langle f \rangle} \sim \rho_* \equiv \frac{\rho}{L} \ll 1, \quad (1)$$

where $\rho = |\boldsymbol{\rho}| = |\hat{\mathbf{b}} \times \mathbf{v}/\Omega|$ is the gyroradius and L is the equilibrium length scale (typically the minor radius of the tokamak). Here $\hat{\mathbf{b}} = \mathbf{B}/B$ denotes the unit vector in the direction of the magnetic field \mathbf{B} , B is the norm of \mathbf{B} , \mathbf{v} is the velocity of the particle and $\Omega = eZB/mc$ is the gyrofrequency with Z the charge number of the particle, e the elementary charge, m the particle mass and c the speed of light. In the frame moving with the background plasma flow \mathbf{u} , the fluctuations described by δf are associated with a typical frequency $\omega \sim \rho_* \Omega$ [21]. The fluctuations are assumed to be elongated along field lines ($\sim L$) but much thinner across field lines ($\sim \rho$). On the other hand, the equilibrium part evolves on the energy confinement time scale $\tau_E \sim (\rho_*^2 \omega)^{-1}$ and has a typical length scale of order L both along and across field lines. To summarize the orderings:

$$\frac{1}{\rho_* \Omega} \left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right) \ln(\delta f) \sim L \hat{\mathbf{b}} \cdot \nabla \ln(\delta f) \sim \rho \nabla_{\perp} \ln(\delta f) \sim \mathcal{O}(1) \quad (2)$$

$$\frac{1}{\rho_*^3 \Omega} \left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right) \ln \langle f \rangle \sim L \nabla \ln \langle f \rangle \sim \mathcal{O}(1) \quad (3)$$

Within this asymptotic expansion in ρ_* , we restrict ourselves to the high-flow regime, in which the background flow u is first ordered to be comparable to the ion thermal speed $v_{\text{th},i} = \sqrt{2T_i/m_i}$, with T_i the ion temperature. We then consider a subsidiary expansion in the Mach number $M \equiv u/v_{\text{th},i}$ such that $\rho_* \ll M \ll 1$, and we neglect second order effects in M . In the high-flow regime, the background plasma flow is toroidal to lowest order in ρ_* , and the associated angular frequency Ω_{ϕ} is constant on a magnetic flux surface [22]: $\mathbf{u} = \Omega_{\phi}(\psi) R^2 \nabla \phi + \mathcal{O}(\rho_* v_{\text{th},i})$ with ψ the flux surface label, R the major radius and ϕ the toroidal angle.

Making use of the assumptions above, one can express the Vlasov equation order by order in $\rho_*^n \Omega f$ for $n = 0, 1, 2$ and average over the rapid gyromotion to obtain the gyrokinetic equation governing the evolution of fluctuations of the distribution function:

$$\begin{aligned} \frac{\partial g_s}{\partial t} + \left(\mathbf{u} + w_{\parallel} \hat{\mathbf{b}} + \mathbf{V}_{d,s} + \langle \mathbf{V}_E \rangle_{\mathbf{R}_s} \right) \cdot \nabla g_s = \\ - \frac{eZ_s F_{0,s}}{T_s} \left(w_{\parallel} \hat{\mathbf{b}} + \mathbf{V}_{d,s} \right) \cdot \nabla \langle \varphi \rangle_{\mathbf{R}_s} - \langle \mathbf{V}_E \rangle_{\mathbf{R}_s} \cdot \left(\nabla F_{0,s} + \frac{Im_s w_{\parallel}}{BT_s} F_{0,s} \nabla \Omega_{\phi} \right) \end{aligned} \quad (4)$$

Here, the subscript s labels the species, $\langle \cdot \rangle_{\mathbf{R}_s} = (2\pi)^{-1} \int_0^{2\pi} (\cdot) d\gamma$ denotes an average over the gyrophase γ at fixed gyrocenter $\mathbf{R}_s = \mathbf{r} - \boldsymbol{\rho}_s$ with \mathbf{r} the particle position, and derivatives are taken at fixed $(\mathbf{R}_s, \varepsilon_s, \mu_s)$, $g_s(\mathbf{R}_s, \varepsilon_s, \mu_s) = \langle \delta f_s \rangle_{\mathbf{R}_s}$ is the gyro-averaged fluctuating distribution function, $\varepsilon_s = m_s w^2/2$ is the particle kinetic energy, $\mu_s = m_s w_{\perp}^2/2B$ is the magnetic moment, φ is the fluctuating electrostatic potential, $\mathbf{w} = \mathbf{v} - \mathbf{u}$ is the particle velocity relative to the background

flow, $w_{\parallel} = \mathbf{w} \cdot \hat{\mathbf{b}}$ is its component along the magnetic field, $I(\psi) = R^2 \mathbf{B} \cdot \nabla \phi$, and $F_{0,s}$ is a Maxwellian distribution of velocities. The species-dependent guiding centre drift velocity $\mathbf{V}_{d,s} = \mathbf{V}_{B,s} + \mathbf{V}_{C,s}$ includes the ∇B and curvature drifts $\mathbf{V}_{B,s} = \hat{\mathbf{b}}/\Omega_s \times \left[w_{\perp}^2 \nabla \ln(B)/2 + w_{\parallel}^2 \hat{\mathbf{b}} \cdot \nabla \hat{\mathbf{b}} \right]$, as well as the Coriolis drift $\mathbf{V}_{C,s} = (2w_{\parallel} \Omega_{\phi}/\Omega_s) \hat{\mathbf{b}} \times (\hat{\mathbf{Z}} \times \hat{\mathbf{b}})$, with $\mathbf{w}_{\perp} = \mathbf{w} - w_{\parallel} \hat{\mathbf{b}}$ the component of \mathbf{w} perpendicular to the magnetic field, and $\hat{\mathbf{Z}}$ the unit vector in the vertical direction. The fluctuating $\mathbf{E} \times \mathbf{B}$ drift is given by $\mathbf{V}_E = c \hat{\mathbf{b}}/B \times \nabla \varphi$. For the purpose of this work, we are neglecting fluctuations of the magnetic field and the effects of collisions.

The system is closed with the quasineutrality equation:

$$\sum_s Z_s \int d^3 w \langle g_s \rangle_{\mathbf{r}} = \sum_s \frac{Z_s^2 e}{T_s} \left(n_s \varphi - \int d^3 w \langle \langle \varphi \rangle_{\mathbf{R}_s} \rangle_{\mathbf{r}} F_{0,s} \right) \quad (5)$$

where n_s is the density of particles, and the brackets $\langle \cdot \rangle_{\mathbf{r}}$ denote an average over the gyrophase at fixed particle position.

2.2 Shearing frame coordinates

Solving the gyrokinetic equation over a machine-sized domain can be computationally very expensive. Instead, **GS2** and other ‘‘local’’ codes focus on a thin, filament-like simulation domain known as a flux tube, which extends along a magnetic field line as it wraps around the torus. In directions perpendicular to \mathbf{B} , the flux tube is assumed to be wide enough (several ρ) for the fluctuations on one side of the domain to be decorrelated from the ones on the opposite side. Ideally the flux tube is kept much narrower than the equilibrium length scale L so that equilibrium quantities and their gradients can be considered as constant across \mathbf{B} in the simulation. Formally, it corresponds to the limit $\rho_* \rightarrow 0$. This local approximation ensures that fluctuations at either end of the perpendicular domain are not only decorrelated but also statistically identical. It is further assumed that two such instances of turbulence can be set exactly equal to one another without affecting the statistical properties of the simulation domain. Hence, periodic boundary conditions are enforced [23] perpendicular to \mathbf{B} .

In the local approximation, the toroidal angular frequency is given by $\Omega_{\phi}(\psi) \simeq \Omega_{\phi,0} + \Omega'_{\phi,0}(\psi - \psi_0)$, where the subscript 0 indicates that a quantity is evaluated along the central field line in the simulation domain and the prime denotes differentiation with respect to the flux label ψ . In the absence of flow shear ($\Omega'_{\phi,0} = 0$), **GS2** works in the co-rotating frame (θ, x, y) , where the poloidal angle θ serves as the coordinate parallel to the field, x is the radial coordinate and y is the coordinate perpendicular to the field line within the flux surface:

$$x = \frac{q_0}{r_{\psi,0} B_r} (\psi - \psi_0), \quad (6)$$

$$y = \frac{1}{B_r} \frac{\partial \psi}{\partial r_{\psi}} \Big|_{r_{\psi,0}} (\alpha - \alpha_0), \quad (7)$$

where B_r is a reference magnetic field strength, $q(\psi)$ is the safety factor, r_ψ is the half-diameter of a flux surface at the height of the magnetic axis, and $\alpha = \phi - q\vartheta - \Omega_{\phi,0}t$ is the field line label with

$$\vartheta(\theta, \psi) = \frac{1}{q} \int_0^\theta \frac{\mathbf{B} \cdot \nabla \phi}{\mathbf{B} \cdot \nabla \theta} \Big|_{\psi, \theta'} d\theta'. \quad (8)$$

Note that the x and y coordinates are not orthogonal.

In the presence of flow shear ($\Omega'_{\phi,0} \neq 0$), the new terms appearing in the gyrokinetic equation through $\mathbf{u} \cdot \nabla$ are not periodic in x and so the periodic solution obtained by Fourier transforming in (x, y) is not consistent with the equation. The solution adopted in GS2 [14] is to restore periodicity by working in the shearing frame (θ, x, y^*) , with

$$y^* = y - x\gamma_E t, \quad (9)$$

where the shearing rate is defined as

$$\gamma_E = \frac{r_{\psi,0}}{q_0} \frac{\partial \Omega_\phi}{\partial r_\psi} \Big|_{r_{\psi,0}}. \quad (10)$$

From here on, the star superscript will be used to indicate that the associated quantity is time dependent in the presence of flow shear. It is not to be confused with the usual notation for complex conjugates.

2.3 Spectral representation

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

$$\Phi(t, x, y^*, \theta) = \sum_{k_x, k_y} \hat{\Phi}_{\mathbf{k}}(t, \theta) e^{ik_x x + ik_y y^*}, \quad (11)$$

where we define the wave vector $\mathbf{k} = k_x \nabla x + k_y \nabla y^*$, with wave numbers given by $k_x = j_x \Delta k_x$, $k_y = j_y \Delta k_y$, grid-spacings by $\Delta k_x = 2\pi/L_x$, $\Delta k_y = 2\pi/L_y$, L_x and L_y are the perpendicular sizes of the computational domain, and with the integer indices respectively spanning $-(N_x - 1)/2 \leq j_x \leq (N_x - 1)/2$ and $0 \leq j_y < N_y$. Note that the exponent can also be written in terms of (x, y) as $ik_x^*(t)x + ik_y y$ where $k_x^*(t) \equiv k_x - k_y \gamma_E t$ is the effective, time-dependent radial wave number in the lab-frame. With this spectral representation in the shearing frame, time dependences due to flow shear appear when either a gradient or a gyro-average is applied to a fluctuating quantity:

$$\nabla_\perp \Phi = \sum_{k_x, k_y} i\mathbf{k}^*(t) \hat{\Phi}_{\mathbf{k}} e^{ik_x x + ik_y y^*}, \quad (12)$$

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

$$\langle \Phi(\mathbf{r}) \rangle_{\mathbf{R}_s} = \sum_{k_x, k_y} J_0^*(t) \hat{\Phi}_{\mathbf{k}} e^{ik_x x + ik_y y^*}, \quad (13)$$

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

2.4 Fourier transformed equations

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

$$\begin{aligned} & \frac{\partial \hat{g}_{\mathbf{k},s}}{\partial t} + w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{\partial \hat{g}_{\mathbf{k},s}}{\partial \theta} + i \mathbf{V}_{d,s} \cdot \mathbf{k}^* \hat{g}_{\mathbf{k},s} + \hat{\mathcal{N}}_{\mathbf{k}} = \\ & - \frac{e Z_s F_{0,s}}{T_s} \left(w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{\partial}{\partial \theta} + i \mathbf{V}_{d,s} \cdot \mathbf{k}^* \right) (J_0^* \hat{\varphi}_{\mathbf{k}}) - i k_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_{\psi}} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) J_0^* \hat{\varphi}_{\mathbf{k}}. \end{aligned} \quad (14)$$

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

$$\hat{\mathcal{N}}_{\mathbf{k}} = \frac{c q_0}{r_{\psi,0} B_r^2} \frac{\partial \psi}{\partial r_{\psi}} \Big|_{r_{\psi,0}} \mathcal{F}_{\mathbf{k}} \left[\{g, \langle \varphi \rangle_{\mathbf{R}_s}\}_{x,y} \right], \quad (15)$$

where the Poisson bracket is defined as

$$\{g, \langle \varphi \rangle_{\mathbf{R}_s}\}_{x,y} = \frac{\partial g}{\partial x} \Big|_y \frac{\partial \langle \varphi \rangle_{\mathbf{R}_s}}{\partial y} \Big|_x - \frac{\partial g}{\partial y} \Big|_x \frac{\partial \langle \varphi \rangle_{\mathbf{R}_s}}{\partial x} \Big|_y, \quad (16)$$

and the derivatives are computed as:

$$\frac{\partial \langle \varphi \rangle_{\mathbf{R}_s}}{\partial x} \Big|_y = \mathcal{F}_{x,y}^{-1} [i k_x^* J_0^* \hat{\varphi}_{\mathbf{k}}], \quad (17)$$

$$\frac{\partial \langle \varphi \rangle_{\mathbf{R}_s}}{\partial y} \Big|_x = \mathcal{F}_{x,y}^{-1} [i k_y J_0^* \hat{\varphi}_{\mathbf{k}}], \quad (18)$$

$$\frac{\partial g}{\partial x} \Big|_y = \mathcal{F}_{x,y}^{-1} [i k_x^* \hat{g}_{\mathbf{k},s}], \quad (19)$$

$$\frac{\partial g}{\partial y} \Big|_x = \mathcal{F}_{x,y}^{-1} [i k_y \hat{g}_{\mathbf{k},s}], \quad (20)$$

with $\mathcal{F}_{x,y}^{-1}$ the inverse discrete Fourier transform to the shearing frame defined by (11), and with $\mathcal{F}_{\mathbf{k}}$ the associated forward transform. At this stage, it is worth pointing out that the Poisson bracket (16) is invariant under the transformation from the lab frame (x, y) to the shearing frame (x, y^*) . We choose to express the bracket using (x, y) , for reasons that will become clear in section 3.5. The quasineutrality equation in Fourier space becomes:

$$\sum_s Z_s \int d^3 w J_0^* \hat{g}_{\mathbf{k},s} = \sum_s \frac{Z_s^2 e n_s}{T_s} (1 - \Gamma_s^*) \hat{\varphi}_{\mathbf{k}} \quad (21)$$

with $\Gamma_s^* = (1/n_s) \int d^3 w F_{0,s} (J_0^*)^2$.

2.5 Parallel boundary condition

In [23], Beer *et al* introduced a parallel boundary condition, often referred to as the “twist-and-shift” boundary condition. It allows flux-tube simulations with smaller computational domains to be representative of an entire flux-surface.

To understand this boundary condition, consider a (ϕ, x) surface cutting through the torus, as shown in figure 1a. While the actual choice in codes for this surface is the inboard mid-plane, in this illustration we pick the outboard mid-plane for simplicity. As the flux-tube wraps around the torus, it periodically intersects with this surface after every 2π turn in poloidal angle. The periodicity in y^* implies that fluctuations at such an intersection with the flux-tube can be equivalently described by exact copies of any other intersection. Here, we consider two consecutive intersections, for example at $\theta = 0$ and $\theta = 2\pi$ (illustrated in figure 1b). For any fluctuating quantity Φ , one must then have

$$\Phi(t, \theta = 0, x, y^*(t, \theta = 0, x, \phi)) = \Phi(t, \theta = 2\pi, x, y^*(t, \theta = 2\pi, x, \phi)) \quad (22)$$

Notice that, to account for magnetic shear $\hat{s} = (r_{\psi,0}/q_0)dq/dr_\psi|_{r_{\psi,0}}$, the equality holds for fixed toroidal angle ϕ , not for fixed field-line label y^* . In Fourier space, (22) becomes

$$\hat{\Phi}_{k_x, k_y}(t, \theta = 0) = \hat{\Phi}_{k_x + 2\pi\hat{s}k_y, k_y}(t, \theta = 2\pi)e^{-i2\pi j_y q_0 \frac{2\pi}{L_\phi}} \quad (23)$$

with $L_\phi = L_y B_r / (d\psi/dr_\psi)|_{r_{\psi,0}}$. The extra phase factor can be set to unity by making an $\mathcal{O}(\rho_*)$ change in q_0 to lie on a flux surface where $q_0 2\pi / L_\phi \in \mathbb{Z}$ (remembering that j_y is an integer labelling the different k_y 's). In simulations, (23) effectively creates twist-and-shift chains of linked radial wave numbers.

As we recall above, the parallel boundary condition for flux-tubes hinges on perpendicular periodicity, and hence on the assumption that the box is wider than a perpendicular correlation length. This in turn means that, for a flux-tube to represent a whole flux-surface, the parallel correlation length of fluctuations cannot be larger than 2π when projected onto θ . In the present work, we always assume that the parallel correlation length is shorter than 2π , in agreement with [24].

3 GS2 algorithm

We showed in Section 2 that the presence of flow shear introduces new time dependences in the Fourier-transformed gyrokinetic and quasineutrality equations. These time dependences would in principle be straightforward to include in codes that solve for $\hat{g}_{\mathbf{k},s}$ with an explicit time advance algorithm: the only modification required would be to update the coefficients $\mathbf{V}_{d,s} \cdot \mathbf{k}^*$ and J_0^* at every time step. Tests conducted with GS2 indicate that this increases runtime by $\sim 25\%$. However, GS2 uses an algorithm that is implicit in time. Unlike explicit algorithms, A-stable implicit algorithms have no stability criterion limiting the size of their time step (known as a Courant-Friedrichs-Lewy, or CFL condition). They can therefore lead to significant savings in computing time by taking larger time steps. Unfortunately, this also introduces additional challenges to the numerical treatment of flow shear.




Figure 1: In [1a](#), the shaded grey area corresponds to the outboard mid-plane. Looking down onto this plane from above in [1b](#), the edges of the flux-tube are shown in solid blue lines as it intersects with the plane at $\theta = 0$ and $\theta = 2\pi$. The dashed blue lines indicate exact copies of the intersection at $\theta = 0$.

In this section, we present an overview of the implicit algorithm used in [GS2](#), including a widely-used method for treating flow shear in flux tube simulations. We then propose a new, more accurate approach for treating flow shear.

3.1 Discretised equations

In [Section 2.5](#), we stress that the parallel boundary condition [\(23\)](#) creates for every k_y a set of linearly independent twist-and-shift chains of connected θ -segments. Here, we use bold-font vector notation to denote discrete arrays whose indices span every (θ, k_x) of a single twist-and-shift chain. Matrices with two such indices are denoted using calligraphic upper case letters. The nonlinear term is treated explicitly in time and is irrelevant for the present section. Therefore, references to the gyrokinetic equation in this section really mean the gyrokinetic without the nonlinear term. Discretising in time and θ , the gyrokinetic equation [\(14\)](#) can be written schematically as

$$\mathcal{A}_{n+1}^* \hat{\mathbf{g}}_{n+1} + \mathcal{B}_n^* \hat{\mathbf{g}}_n = \mathcal{C}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} + \mathcal{D}_n^* \hat{\boldsymbol{\varphi}}_n, \quad (24)$$

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

$$\mathcal{E}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} = \mathcal{W}_{n+1}^* \hat{\mathbf{g}}_{n+1} \quad (25)$$

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

3.2 Overview of the GS2 algorithm

An obvious way to solve the discretised system of equations would be to isolate $\hat{\varphi}_{n+1}$ in the quasineutrality equation (25), insert the resulting expression into the gyrokinetic equation (24) and find $\hat{\mathbf{g}}_{n+1}$. However, the velocity integration represented by \mathcal{W}_{n+1}^* leads to a dense matrix that would need to be inverted to obtain $\hat{\mathbf{g}}_{n+1}$. As this would be computationally costly, GS2 uses an alternative method [15] related to a Green's function approach. The general idea is to determine how the distribution function responds when the potential $\hat{\varphi}_n$ is set to zero and the potential $\hat{\varphi}_{n+1}$ is set to a Kronecker delta in θ and k_x . By doing this for every θ and k_x , we will show that one can compute $\hat{\varphi}_{n+1}$ without knowing $\hat{\mathbf{g}}_{n+1}$, and then solve the gyrokinetic equation (24) for the distribution function.

Since (24) is a linear equation, $\hat{\mathbf{g}}_{n+1}$ can be written as a sum of two terms (where $\hat{\mathbf{g}}_{n+1}^{\text{old}}$ only depends on "old" fluctuating quantities from the current time step n , while $\hat{\mathbf{g}}_{n+1}^{\text{new}}$ also depends on the "new" potential from time step $n+1$):

$$\hat{\mathbf{g}}_{n+1} = \hat{\mathbf{g}}_{n+1}^{\text{old}} + \hat{\mathbf{g}}_{n+1}^{\text{new}}, \quad (26)$$

which satisfy respectively

$$\mathcal{A}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{old}} + \mathcal{B}_n^* \hat{\mathbf{g}}_n = (\mathcal{C}_{n+1}^* + \mathcal{D}_n^*) \hat{\varphi}_n \quad (27)$$

and

$$\mathcal{A}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{new}} = \mathcal{C}_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}}, \quad (28)$$

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

$$\hat{\varphi}_{n+1}^{\text{dif}} = \hat{\varphi}_{n+1} - \hat{\varphi}_n. \quad (29)$$

Equation (27) can be solved directly to find $\hat{\mathbf{g}}_{n+1}^{\text{old}}$ since all other quantities are evaluated at the previous time step and are therefore known. One can then introduce a Green's function \mathcal{G}^* such that

$$\mathcal{A}_{n+1}^* \mathcal{G}_{n+1}^* = \mathbb{1}. \quad (30)$$

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

$$\hat{\mathbf{g}}_{n+1}^{\text{new}} = \mathcal{G}_{n+1}^* \mathcal{C}_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}} \equiv \left(\frac{\delta \hat{\mathbf{g}}_{n+1}^{\text{new}}}{\delta \hat{\varphi}_{n+1}^{\text{dif}}} \right)_{n+1}^* \hat{\varphi}_{n+1}^{\text{dif}} \quad (31)$$

where we defined the "response matrix" $\delta \hat{\mathbf{g}}_{n+1}^{\text{new}} / \delta \hat{\varphi}_{n+1}^{\text{dif}}$. The matrix can be computed by setting $\hat{\varphi}_{n+1}^{\text{dif}}$ to a Kronecker delta function in (θ, k_x) , solving (28) for $\hat{\mathbf{g}}_{n+1}^{\text{new}}$, filling the appropriate column of $\delta \hat{\mathbf{g}}_{n+1}^{\text{new}} / \delta \hat{\varphi}_{n+1}^{\text{dif}}$ with this result and iterating for every (θ, k_x) in the connected twist-and-shift chain. In the absence of flow shear, \mathcal{A} , \mathcal{B} , \mathcal{C} , \mathcal{D} and \mathcal{G} are independent of time, and one only needs to go through this process once per simulation (unless the time step size changes).

Inserting expressions (29), (26) and (31) into the discrete quasineutrality equation (25) yields

$$\hat{\varphi}_{n+1}^{\text{dif}} = (\mathcal{M}^*)_{n+1}^{-1} (\mathcal{W}_{n+1}^* \hat{\mathbf{g}}_{n+1}^{\text{old}} - \mathcal{E}_{n+1}^* \hat{\varphi}_n), \quad (32)$$

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

$$\mathcal{M}_{n+1}^* = \mathcal{E}_{n+1}^* - \mathcal{W}_{n+1}^* \left(\frac{\delta \hat{\mathbf{g}}^{\text{new}}}{\delta \hat{\boldsymbol{\varphi}}^{\text{dif}}} \right)_{n+1}^*. \quad (33)$$

Once $(\mathcal{M}^*)_{n+1}^{-1}$ is known, $\hat{\boldsymbol{\varphi}}_{n+1}$ can be obtained from (32) and $\hat{\mathbf{g}}_{n+1}$ can then be determined by solving the full gyrokinetic equation (24).

3.3 Challenges in the presence of flow shear

Based on the previous sections, we can identify two important issues appearing in the GS2 algorithm when flow shear is non-zero. First, consider a simulation with a fixed, finite set of radial wavenumbers k_x in the shearing frame. As the simulation advances in time, the associated wavenumbers in the lab frame $k_x^*(t, k_x) = k_x - k_y \gamma_E t$ would all tend to $\pm\infty$ depending on the sign of γ_E . In other words, the simulation would eventually contain only structures that have been highly sheared in the lab frame, and would hence be of little interest for transport studies where structures elongated in the radial direction ($k_x^* \sim 0$) play an important role. The second issue with flow shear is related to the computational cost of a simulation. As discussed in Section 3.2, the response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\boldsymbol{\varphi}}^{\text{dif}}$ required to compute $\hat{\boldsymbol{\varphi}}_{n+1}$ is independent of time if $\gamma_E = 0$, in which case it has to be computed only once at the beginning of the simulation, and every time the size of the time step changes. But if $\gamma_E \neq 0$, this matrix becomes time-dependent and would have to be re-computed at every time step, leading to prohibitively slow simulations.

3.4 Discrete-in-time approach to flow shear

We now describe how the two challenges associated with flow shear for the GS2 algorithm are addressed by a re-mapping approach developed by Hammett *et al* [14]. This approach has previously been implemented in GS2 and other gyrokinetic codes including GENE [16], GWK [17] and CGYRO [18]. We then highlight the disadvantages associated with the discrete nature of this implementation.

To address the issue of modes getting more and more sheared over time in the lab frame, Hammett's approach uses a wavenumber re-mapping method [14]. In practice, the Fourier series (11) of a fluctuating quantity Φ is truncated to a finite sum. At $t = 0$, we have $-K_x \leq k_x \leq K_x$ and $-K_y \leq k_y \leq K_y$. The set of Fourier coefficients that are being simulated is updated over time, such that $-K_x \lesssim k_x^* \lesssim K_x$ at all times (i.e. such that modes radially elongated in the lab frame are always included in the simulation). For each k_y there is a corresponding time $T_{\text{map}}(k_y) = \Delta k_x / |\gamma_E k_y|$ at which

$$k_x^*(t + T_{\text{map}}, k_x, k_y) = k_x^*(t, k_x, k_y) \mp \Delta k_x. \quad (34)$$

From here on, the upper sign will always correspond to $\gamma_E > 0$ and the lower to $\gamma_E < 0$. In the lab frame, $T_{\text{map}}(k_y)$ corresponds to the time required for modes with this particular k_y to regain radial periodicity on the edges of the box (for a visual illustration, see figure 4b, which will be presented in section 4.1). Over the interval

of time T_{map} , the most sheared mode in the simulation, $\hat{\Phi}_{\mp K_x, k_y}$ is considered to shear into structures that are radially fine enough to get averaged out to a negligible amplitude by the gyro-motion of particles; this mode is dropped from the simulation. Simultaneously, the mode $\hat{\Phi}_{\pm(K_x + \Delta k_x), k_y}$ has now been unsheared enough to become of interest and gets added into the simulation with zero amplitude. The Fourier sum at time t is then given by

$$\Phi(t, x, y^*, \theta) = \sum_{k_y = -K_y}^{K_y} \sum_{k_x = -K_x \pm S(t, k_y) \Delta k_x}^{K_x \pm S(t, k_y) \Delta k_x} \hat{\Phi}_{\mathbf{k}}(t, \theta) e^{ik_x x + ik_y y^*}, \quad (35)$$

where $S(t, k_y)$ is the number of times the simulation has dropped a mode with k_y and added a new one. In the code, the shift in radial wavenumber is triggered every time k_x^* crosses the mid-point between the two nearest lying $n\Delta k_x$ values, with $n \in \mathbb{Z}$. This is implemented by incrementing S at every shift, and by using a ‘‘re-mapped’’ version of the Fourier transformed quantities, denoted by a tilde:

$$\tilde{\Phi}_{k_x, k_y} = \hat{\Phi}_{k_x \pm S \Delta k_x, k_y}, \quad (36)$$

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

To address the computing time issue associated with the time dependence of the response matrix, Hammett’s approach makes a nearest grid point (NGP) approximation [14]: $k_x^*(t, k_x, k_y) \simeq \bar{k}_x(t, k_x, k_y)$, where

$$\bar{k}_x(t, k_x, k_y) = k_x \mp S(t, k_y) \Delta k_x \quad (37)$$

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

Combined with the wavenumber re-mapping, the NGP approximation gets rid of all time dependences due to flow shear in the code, and all star superscripts in sections 3.1 and 3.2 can be ignored. Therefore, the response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\boldsymbol{\varphi}}^{\text{dif}}$ in this implementation only has to be computed at initialisation and when the time step size changes, similarly to cases with $\gamma_E = 0$. As an illustration of the time dependences vanishing, the inverse Fourier transform of the radial derivative at fixed y of a quantity Φ is given by:

$$\left. \frac{\partial \Phi}{\partial x} \right|_y = \sum_{k_y = -K_y}^{K_y} \sum_{k_x = -K_x \pm S(t) \Delta k_x}^{K_x \pm S(t) \Delta k_x} k_x^*(t) \hat{\Phi}_{k_x, k_y}(t) e^{ik_x^*(t)x + ik_y y} \quad (38)$$

$$\stackrel{\text{remap}}{=} \sum_{k_y = -K_y}^{K_y} \sum_{\bar{k}_x = -K_x}^{K_x} k_x^*(t) \tilde{\Phi}_{\bar{k}_x, k_y}(t) e^{i\bar{k}_x x + ik_y y} e^{i(k_x^*(t) - \bar{k}_x)x} \quad (39)$$

$$\stackrel{\text{NGP}}{\simeq} \sum_{k_y = -K_y}^{K_y} \sum_{\bar{k}_x = -K_x}^{K_x} \bar{k}_x \tilde{\Phi}_{\bar{k}_x, k_y}(t) e^{i\bar{k}_x x + ik_y y}. \quad (40)$$

where we explicitly show all time dependences. In the first expression, k_x^* is time-dependent, and the set of k_x is changing over time. In the last expression, the set of \bar{k}_x values entering the sum is fixed, and all time dependences from flow shear vanish. However, the NGP will produce two different types of errors in simulations.

First, unphysical, discrete jumps will occur in the linear time evolution of fluctuations, with a period of $T_{\text{map}}(k_y)$. This effect vanishes in the limit of small spacing in the radial wavenumber grid, $\Delta k_x \rightarrow 0$. But it is made worse for cases with weaker flow shear, where radial wavenumbers in the lab frame k_x^* evolve very slowly in time and get approximated to the same nearest neighbour for a long period of time.

The second type of error associated with the NGP arises when Fourier transforms are performed, e.g., in the computation of the nonlinear term. A phase factor $e^{i(k_x^* - \bar{k}_x)x}$ is missed by the discrete-in-time approach to flow shear. This can be seen in the example above, where the phase factor has been approximated to unity to go from (39) to (40). This type of error has been stated in [25] not to vanish in the limit $\Delta k_x \rightarrow 0$. Following this observation, the authors added the missing phase factor in the gyrokinetic code GENE. One can however argue that given that $|k_x^* - \bar{k}_x| \sim \Delta k_x$, the error at any fixed location in (x, y) should get smaller when Δk_x is decreased. Errors at the radial edges of the box would be independent of Δk_x , since $L_x \sim 1/\Delta k_x$. But this should not affect the centre of the domain: two instances of turbulence separated by more than one correlation length should not influence one another, regardless of whether one instance suffers from the missing phase factor. This would suggest that the error associated with the missing $e^{i(k_x^* - \bar{k}_x)x}$ factor could be reduced by decreasing Δk_x , while only considering the same portion of the simulation domain and ignoring the erroneous radial edges. A simple test case presented in section 4.2 supports this argument, but nonlinear gyrokinetic simulations in section 4.3 show a more complicated picture.

The next section will present a new approach to flow shear that removes both types of errors in a linearly implicit δf gyrokinetic code.

3.5 Continuous-in-time approach to flow shear

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

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

The first novel idea is to remove time dependences from the implicit part of the algorithm by evaluating them explicitly in time. In the gyrokinetic equation (14), every term $\mathcal{L}[k_x^*] \hat{\Phi}_{\mathbf{k}}$ that depends on the wavenumber in the lab frame can be re-written as:

$$\mathcal{L}[k_x^*] \hat{\Phi}_{\mathbf{k}} = \mathcal{L}[\bar{k}_x] \hat{\Phi}_{\mathbf{k}} + (\mathcal{L}[k_x^*] - \mathcal{L}[\bar{k}_x]) \hat{\Phi}_{\mathbf{k}} \quad (41)$$

where we added and subtracted the same term, evaluated at the nearest $n\Delta k_x$ neighbour. For every term in the discretised gyrokinetic equation (24) that is evaluated at the next time step $n + 1$, we perform (41). The first term $\mathcal{L}[\bar{k}_x] \hat{\Phi}_{\mathbf{k}}$ is discretised in the usual implicit way, and hence appears in the implicit part of the

algorithm. This term is nothing but the NGP approximation that was used in the discrete-in-time implementation. In our new approach, we choose to treat the $(\mathcal{L}[k_x^*] - \mathcal{L}[\bar{k}_x]) \hat{\Phi}_{\mathbf{k}}$ term explicitly in time; i.e., to evaluate it at the current time step n . As for the explicit nonlinear term (15), there is an associated Courant-Friedrichs-Lewy condition that has to be fulfilled for the scheme to be stable. But the new explicit terms are always small compared to the corresponding implicit terms since $\rho |k_x^* - \bar{k}_x| \leq \rho(\Delta k_x)/2 \ll 1$. Hence, the new condition for numerical stability should be less stringent than the one already set by the nonlinear term. One should note, however, that if Δk_x is made large enough (much larger than any sensible choice for a nonlinear simulation), this scheme would eventually become unstable. Following the decomposition (41), the time dependences in \mathcal{A} and \mathcal{C} vanish from (27) and (28). Equations (30) and (31) then yield a time-independent response matrix $\delta \hat{\mathbf{g}}^{\text{new}} / \delta \hat{\varphi}^{\text{dif}}$.

This decomposition alone does not remove all time dependences from the implicit part of the code: to compute $\hat{\varphi}_{n+1}$, one must evaluate the right-hand side of equation (32), where \mathcal{E}^* , \mathcal{W}^* and $(\mathcal{M}^*)^{-1}$ depend on time. We update \mathcal{E}^* and \mathcal{W}^* at every time step with little extra computational cost. The time dependence in the matrix \mathcal{M}^* could be computed exactly, however inverting it at every time step would be very costly. Instead, we use linear interpolation to approximate the time dependence of its inverse $(\mathcal{M}^*)^{-1}$. Note that while there is a simple dependence of \mathcal{M}^* on time, there is no such simple dependence for the inverse $(\mathcal{M}^*)^{-1}$. At the beginning of a simulation, we compute three matrices \mathcal{M}_ℓ^{-1} with $\ell = -1, 0, +1$, defined respectively by replacing all instances of k_x^* in (33) with $\bar{k}_x + \ell \Delta k_x$. Interpolating then gives

$$(\mathcal{M}^*)^{-1} \simeq \left(1 - \frac{|k_x^* - \bar{k}_x|}{\Delta k_x}\right) \mathcal{M}_0^{-1} + \frac{|k_x^* - \bar{k}_x|}{\Delta k_x} \mathcal{M}_{\text{sgn}(k_x^* - \bar{k}_x)}^{-1}. \quad (42)$$

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

As was argued in [25], the implementation of Fourier transforms also needs to be modified to take into account the continuous time evolution of k_x^* . This has direct implications for the computation of the nonlinear term (15). With the wave number re-mapping, the discrete inverse Fourier transform (11) can be written as:

$$\mathcal{F}_{x,y}^{-1} [\hat{\Phi}_{\mathbf{k}}] = \sum_{k_y=-K_y}^{K_y} e^{i(k_x^*(t,k_y) - \bar{k}_x)x_n} e^{ik_y y_m} \sum_{\bar{k}_x=-K_x}^{K_x} \tilde{\Phi}_{\bar{k}_x, k_y} e^{i\bar{k}_x x_n}, \quad (43)$$

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

$$\mathcal{F}_{\mathbf{k}} [\Phi] = \sum_{n=-N_x}^{N_x} e^{-i(k_x^*(t,k_y) - \bar{k}_x)x_n} e^{-i\bar{k}_x x_n} \sum_{m=-N_y}^{N_y} \Phi(x_n, y_m) e^{-ik_y y_m}. \quad (44)$$

With the NGP approximation made in the discrete-in-time approach, k_x^* is approximated to \bar{k}_x and the first exponential terms in (43) and (44) are set to unity. In the continuous-in-time approach, we update the phase factor at every time step and

modify the Fourier transforms in the following way: for $\mathcal{F}_{x,y}^{-1}$, inverse transform in k_x , multiply by $e^{i(k_x^* - \bar{k}_x)x}$ and inverse transform in k_y ; for $\mathcal{F}_{\mathbf{k}}$, transform in y , multiply by $e^{-i(k_x^* - \bar{k}_x)x}$ and transform in x .

4 Comparison of the approaches to flow shear

In this section, we compare numerical results obtained with the discrete-in-time and continuous-in-time approaches. We first focus on tests of the modifications made to the linear GS2 algorithm. We then present a toy case illustrating the changes in the Fourier transforms and how they can affect the nonlinear term. Finally, a full nonlinear gyrokinetic simulation of an experiment at the Joint European Torus is shown, highlighting the impact of this work on simulating turbulent transport in the presence of flow shear.

4.1 Linear simulations

We first present linear simulations of the widely used ‘‘Cyclone Base Case’’ [26], with added flow shear $\gamma_E = 0.2v_{\text{th},i}/a$. Here, a is the half-diameter of the last closed flux-surface at the height of the magnetic axis. As we explain in section 2.5, twist-and-shift chains of connected radial wave numbers can be seen as a single ballooning structure, that can extend in θ by more than 2π . At any given time, this chain corresponds to a direct-space structure that is twisted along θ by the presence of magnetic shear \hat{s} (illustrated by the sketch in figure 2a). Now, if $\gamma_E \neq 0$, the background flow shear will contribute to the twisting, enhancing it at one end in θ and decreasing it at the other end (figure 2b). After a period of time T_{map} , the wave number re-mapping is applied in order to retain the radially extended part of the structure in the simulation. As shown in figure 2c, this amounts to moving the computational domain in direct space along θ . In the absence of flow shear, the ballooning mode typically peaks at $k_x^* = 0$ and $\theta = 0$, where the mode is aligned with the ∇B and curvature drives for the ion temperature gradient instability, and where finite gyroradius stabilisation is weakest. When $\gamma_E \neq 0$, it modifies the twist of the structure, and the peak moves along the field line in θ towards the smallest k_x^* in the chain. As the peak passes from $\theta = 0$ (bad curvature region) to $\theta = \pi$ (good curvature region) it experiences a varying growthrate, resulting in a so-called Floquet mode [10] (analogous to Bloch states in condensed matter physics). The Floquet oscillation period of the growthrate is $T_F = 2\pi\hat{s}/\gamma_E$.

To highlight the effects on linear physics of the NGP approximation for radial wave numbers, we show the time evolution of the fluctuating electrostatic potential φ , averaged over θ and summed along one particular twist-and-shift chain. In figure 3, $\ln(\sum_{k_x} |\langle \hat{\varphi}_{\mathbf{k}} \rangle_{\theta}|^2)$ is plotted for several values of the grid spacing in radial wave number, Δk_x , with $\langle \cdot \rangle_{\theta}$ denoting an average over the poloidal angle. figure 3a shows the results obtained with the discrete-in-time approach to flow shear, which uses the NGP approximation. Discrete jumps are generated every time a wave number re-mapping occurs for this twist-and-shift chain. For fine enough grid-spacings, the simulation recovers from the jumps, and the time-averaged growthrate remains

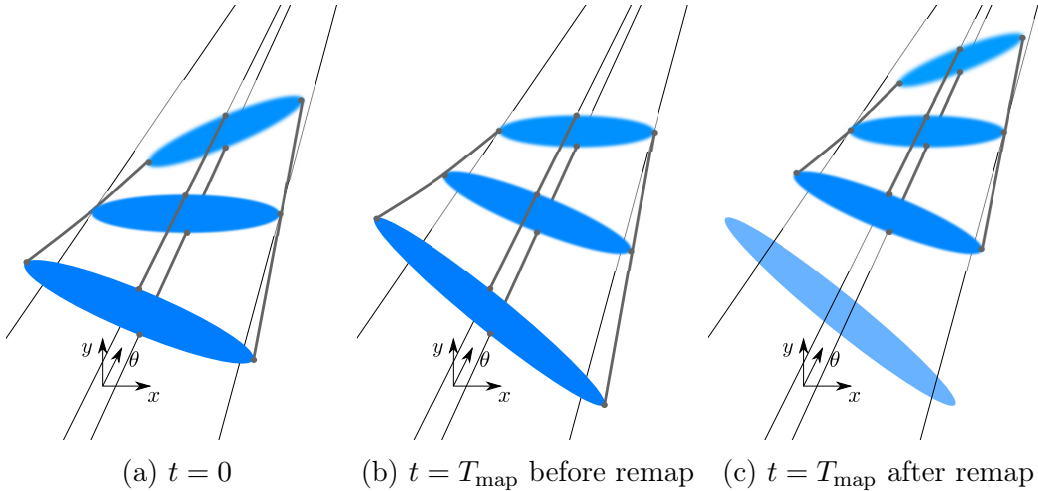


Figure 2: Cartoon illustrating how the wavenumber re-mapping amounts to moving the computational domain in θ . Initially (2a) the structure is twisted by magnetic shear. Over a period of time T_{map} , flow shear contributes to the twisting, enhancing it at one end and decreasing it at the other (2b). In order to always include radially extended structures in the simulation, the computational domain is shifted in θ via a wavenumber re-map (2c).

approximately correct. For larger Δk_x however, the simulation fails to capture the correct evolution of the Floquet mode. figure 3b shows the results with the new continuous-in-time approach: it does not generate any discrete jumps at re-maps, and reproduces the correct Floquet behaviour over a much larger range of Δk_x than the discrete-in-time algorithm. Note that typical nonlinear gyrokinetic simulations currently use $\rho(\Delta k_x) \sim 0.1$. Section 4.3 will show that our approach to capture the linear physics correctly by avoiding the NGP approximation can lead to computational savings.


4.2 Test cases for Fourier transforms

We now show two cases that test the changes made to the Fourier transforms in the continuous-in-time algorithm. The first case highlights the effect on direct-space quantities when the NGP approximation is being made in Fourier transforms. Given a field $\Psi(t, x, y)$ with the initial condition $\Psi(0, x, y) = \cos [2(\Delta k)y]$, we solve a simple shearing motion in the lab frame


$$\left. \frac{\partial \Psi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Psi}{\partial y} = 0, \quad (45)$$

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

$$\left. \frac{\partial \hat{\Psi}_{\mathbf{k}}}{\partial t} \right|_{x,y^*} = 0, \quad (46)$$



(a) Discrete-in-time



(b) Continuous-in-time

Figure 3: Time evolution of the normalised electrostatic potential for a single twist-and-shift chain, summed over the parallel coordinate.

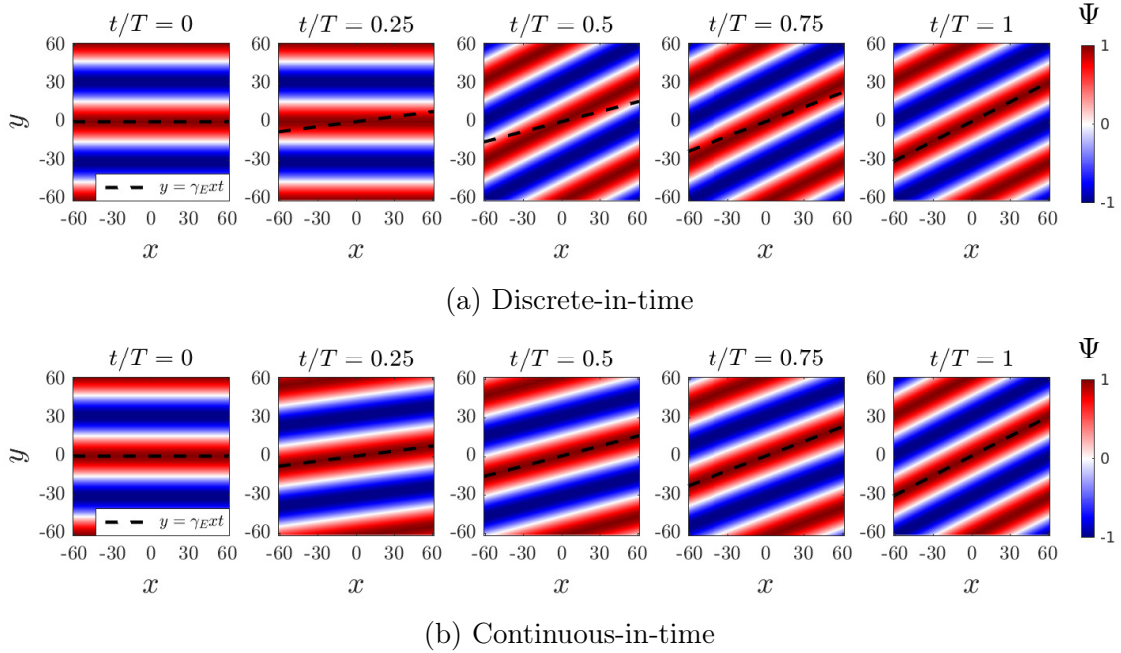


Figure 4: Time evolution of Ψ in the lab frame, according to equation (45). The discrete-in-time approach (4a) does not include the phase factor appearing in Fourier transforms, whereas the continuous-in-time approach (4b) does. Both methods become equivalent at $t = T$, when the NGP approximation is exact. The dashed black line corresponds to $y^* = 0$ and shows the tilt that flow shear should be producing.

so the only operation we have to apply at every time step is the wave number re-mapping described in section 3.4, followed by an inverse Fourier transform to visualise Ψ in direct space. In the discrete-in-time approach, the transform is given by (43) without the first phase factor. In the continuous-in-time approach, the phase factor is included. The time evolution obtained with each approach is shown in figure 4 for $t \in [0, T]$, where T is the time required for $k_x^*(t)$ to change from 0 to $-\Delta k_x$. The correct shearing motion is indicated by the line $y^* = 0$. Until the first re-map, i.e. $t < T/2$, the discrete-in-time approach can be seen to lag behind the correct shearing motion (\bar{k}_x remains 0). After the re-map (\bar{k}_x has now jumped to $-\Delta k_x$), it is ahead of the correct motion, until both coincide again at $t = T$ (i.e. when $\bar{k}_x = k_x^*$). In the continuous-in-time approach, the exact time evolution of k_x^* is included in the Fourier transform, and the numerical result coincides with the correct shearing motion at all times.

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

two fields, $\Psi(t, x, y)$ and $\Phi(t, x, y)$, with initial conditions

$$\Psi_0(x, y) = \cos [2(\Delta k)y], \quad (47)$$

$$\Phi_0(x, y) = \frac{C}{\sigma^2} e^{-(x^2+y^2)/(2\sigma^2)}, \quad (48)$$

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

$$\left. \frac{\partial \Psi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Psi}{\partial y} + \{\Psi, \Phi\}_{x,y} = 0, \quad (49)$$

$$\left. \frac{\partial \Phi}{\partial t} \right|_{x,y} + \gamma_E x \frac{\partial \Phi}{\partial y} = 0. \quad (50)$$

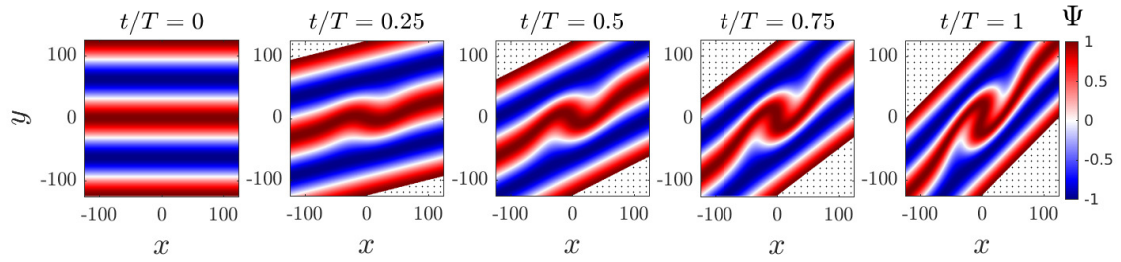
The analytical solution for Ψ can be computed as being

$$\Psi(t, x, y) = \cos [2(\Delta k)Y(t, x, y)], \quad (51)$$

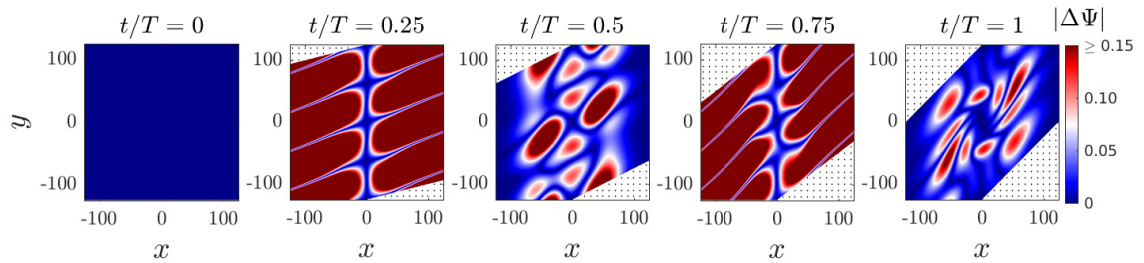
where we defined $Y(t, x, y) = y^* \cos(\Phi_0 t / \sigma^2) - x \sin(\Phi_0 t / \sigma^2)$. Numerically, equations (49) and (50) are solved in Fourier space, and the Poisson bracket is computed pseudo-spectrally. Note that the maximum amplitude of the analytical solution is $|\Psi| = 1$.

In figure 5a, we plot the analytical solution (51) over a time period $T = 1/\gamma_E$. We indicate with dots areas of the box where analytical and numerical results should not be compared, because the numerical solution is periodic in x and y , while the analytical one is not. In figure 5b, we show the difference $|\Delta \Psi| = |\Psi_{\text{num}} - \Psi|$ between the analytical solution and the numerical result, obtained with the discrete-in-time algorithm. Although the numerical solution looks more correct at $t = T/2$ (when the NGP approximation is exact for half of the values of k_y) and $t = T$ (when it is exact for all values of k_y), errors in the pseudo-spectral computation of the Poisson bracket at $0 < t < 1/\gamma_E$ have altered the time-evolution of the simulation. In figure 5c, we see that the continuous-in-time approach deviates from the analytical solution by less than 0.01 at any time, compared to differences of up to 1.4 observed with the discrete-in-time scheme.


Finally, figure 6 illustrates how the error $|\Delta \Psi|$ is affected when Δk is reduced. A smaller Δk (i.e. a larger simulation domain) does not alter the pattern or the maximum amplitude of the error produced by the missing phase factor. It does however reduce the error within the original simulation domain. This simple test case therefore suggests that the error from the missing phase factor in local gyrokinetic codes would not vanish when simply reducing Δk_x , which agrees with the observations made in [25]. This would also suggest that to make this error vanish, only the original portion of the simulation domain should be considered, while simultaneously reducing Δk_x . We find in section 4.3 that nonlinear simulations do not seem to agree with this last point.



(a) Analytical time evolution of Ψ .




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




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


Figure 5: Analytical solution to equations (49) and (50) for Ψ , and comparison to numerical solutions obtained with the discrete- and continuous-in-time approaches. At $t = T$, the NGP approximation is exact, making both approaches equivalent for that time step.




(a) Discrete-in-time, $\Delta k_x = 0.025$.
Identical to $t/T = 0.25$ in figure 5b.



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



(c) Continuous-in-time, $\Delta k_x = 0.025$.
Identical to $t/T = 0.25$ in figure 5c.



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

Figure 6: $|\Delta\Psi| = |\Psi_{\text{num}} - \Psi|$ for the nominal Δk_x and for Δk_x halved, with both discrete- and continuous-in-time approaches.

| | | |
|-------------------|---------------------------------------|---|
| I_p | 2.6MA | Plasma current |
| B_T | 2.9T | Vacuum toroidal field on axis |
| P_{NBI} | 17MW | Neutral beam heating power |
| R_ψ | $3.06a$ | $[\max(R) + \min(R)]/2$ for this flux surface |
| r_ψ | $0.508a$ | $[\max(R) - \min(R)]/2$ for this flux surface |
| $ q_0 $ | 1.43 | flux-surface averaged safety factor |
| \hat{s} | 0.574 | flux-surface averaged magnetic shear |
| κ | 1.36 | Miller elongation [29] |
| $d\kappa/dr_\psi$ | $0.146/a$ | elongation gradient |
| δ | 0.0571 | arcsin of Miller triangularity [29] |
| $d\delta/dr_\psi$ | $0.129/a$ | gradient of GS2 triangularity |
| γ_E | $-0.063v_{\text{th},i}/a$ | background flow shear rate |
| Ω_ϕ | $-0.08v_{\text{th},i}/a$ | background flow angular frequency |
| n_i/n_e | 1.0 | ion to electron density ratio |
| $1/L_{n_i}$ | $0.602/a$ | inverse ion density gradient length |
| $1/L_{n_e}$ | $0.602/a$ | inverse electron density gradient length |
| T_e/T_i | 0.855 | electron to ion temperature ratio |
| $1/L_{T_i}$ | $1.759/a$ | inverse ion temperature gradient length |
| $1/L_{T_e}$ | $1.551/a$ | inverse electron temperature gradient length |
| ν_{ii} | $2.6 \times 10^{-4}v_{\text{th},i}/a$ | ion collisionality |
| ν_{ee} | $0.02v_{\text{th},i}/a$ | electron collisionality |
| L_x | $77.5\rho_i$ | box size in x |
| L_y | $69.8\rho_i$ | box size in y |
| K_x | $3.49/\rho_i$ | largest radial wave number |
| K_y | $0.99/\rho_i$ | largest wave number in y |
| N_θ | 33 | number of θ grid points |
| N_ε | 16 | number of ε_s grid points |
| N_λ | 27 | number of pitch angles |

Table 1: Simulation parameters for the JET discharge #68448. The gradient length of any quantity Q is given by $L_Q = 1/[d \log(Q)/dr_\psi]$.

4.3 Nonlinear simulations

Having validated the implementation of the continuous-in-time approach in sections 4.1 and 4.2, we now present nonlinear, electrostatic gyrokinetic simulations performed with the local code GS2. They show that our new approach to flow shear can lead to computational savings compared to the discrete-in-time approach, and that differences in fluxes can persist if Fourier transforms are not evaluated correctly when computing the nonlinear term of the gyrokinetic equation. The plasma parameters (presented in table 1) are taken from the JET discharge #68448 for which a substantial part of the heating was provided by neutral beam injection. The data for this discharge is documented in the JETPEAK database [27]. The simulations include two kinetic species (deuterium ions and electrons), numerical hyperviscosity [28], and collisions.

For a species s , the turbulent heat flux Q_s is given by

$$Q_s = \left\langle \int d^3\mathbf{v} \frac{m_s v^2}{2} \delta f_s \mathbf{V}_E \cdot \nabla \psi \right\rangle_\psi \quad (52)$$

where $\langle \cdot \rangle_\psi$ denotes volume averages over the flux-tube. Figure 7 shows how Q_i differs in simulations using the new continuous-in-time approach to flow shear, compared to simulations with the discrete-in-time approach. In figure 7a, Δk_x is the only parameter that is being changed (along with the number of k'_x s so as to keep the maximum wavenumber fixed). It can be seen that the two approaches to flow shear do not converge to the same heat flux. In this particular case, the value from the continuous-in-time method is lower by about 20%. A simulation was performed by only adding the phase factor to the discrete-in-time approach (blue cross at $L_x \sim 150\rho_i$), and another simulation was performed by removing the phase factor from the continuous approach (red dot at $L_x \sim 150\rho_i$): both clearly show that it is this phase factor that explains the difference in convergence, and that the changes to the linear parts of the code have little impact at large L_x (as one would expect from section 4.1). At lower values of L_x in figure 7a, the main result is that the continuous-in-time approach stays closer to its converged value for smaller box sizes than the discrete approach. This agrees with the results of figure 3 which show that the continuous algorithm is able to reproduce the linear physics correctly with a larger Δk_x (smaller L_x) than the discrete-in-time approach. A simulation was performed by adding the phase factor to the discrete-in-time approach (blue cross at $L_x \sim 20\rho_i$), and another simulation was performed by removing the phase factor from the continuous approach (red dot at $L_x \sim 20\rho_i$): they indicate that, at small L_x , the changes to the linear parts of the code can play a bigger role than the error from the missing phase factor. Surprisingly, we note that the phase factor appears to have almost no impact in the simulations with $L_x \sim 20\rho_i$, unlike at $L_x \sim 150\rho_i$. At $L_x \sim 20\rho_i$, we observe that the size of the simulation domain is becoming comparable to the radial extent of the eddies. And from figure 6, we know that the error from the missing phase factor is strongest at the radial edges of the simulation. When the box is small enough, we therefore expect the missing phase factor to be unable to substantially affect a whole eddy: only the parts of the eddy that are close to the radial edges of the box are strongly impacted. In contrast, when the box becomes larger, entire eddies are affected by the phase factor. However, additional work would be needed to determine if this could explain why simulations with smaller boxes are less affected by the phase factor. Figure 7b is meant to highlight the cases with smaller L_x . The fluxes are this time plotted against $T_{\text{map}}(k_y)/T_F$, that is the number of wavenumber re-mappings that occur per Floquet period (meaning that the smallest L_x is now on the right of the figure). T_{map} is evaluated at the wavenumber k_y with the largest amplitudes in the simulation. All the points are identical to those on the previous figure, except for those in square boxes for which both Δk_x and Δk_y were changed (again, keeping the maximum wavenumbers fixed). The conclusions are the same as for figure 7a.

It is interesting to note that the exact mechanism through which the added phase factor decreases the converged heat flux remains unclear. The only notable difference when adding the phase is that the amplitude of the box-scale zonal mode

($k_x = \Delta k_x$, $k_y = 0$) decreases, as is illustrated by figure 8. We previously showed in figures 6a and 6b that omitting the phase factor produces errors that are larger at the radial edges of the simulation domain. However it is not yet clear how this error relates to the box-scale zonal mode, and how this mode might affect the level of transport observed in simulations.


Unlike it was suggested at the end of section 4.2, our simulations seem to indicate that considering only a smaller subset of a large simulation domain does not reduce the error from the missing phase factor. One could have imagined turbulent fluctuations in the absence of the phase factor setting up a region in the middle of the box where, on average over time, flow shear is reduced or the temperature gradient is substantially enhanced. However we observe neither of these two cases. Further work would be needed to confirm whether a different turbulent heat diffusivity $\chi_i = -Q_i/(dT_i/dr_\psi)$ is indeed obtained with and without phase factor. The diffusivity would have to be averaged over time and space for small subsets of the computational domain.

The implementation of the continuous-in-time approach can have an impact on computational costs in two ways. Firstly, the additional flexibility in the choice of Δk_x with the continuous approach can lead to substantial computational savings. In the nonlinear simulations presented here, the discrete-in-time approach needs $T_{\text{map}}/T_F \sim 0.12$ to remain converged, whereas the continuous approach is still roughly fine at $T_{\text{map}}/T_F \sim 0.18$. This corresponded to a 46% decrease in computing time. Secondly, the continuous approach requires three large matrices to be inverted at initialisation or when changing the time-step. The discrete approach only requires a single matrix inversion. In theory, this could increase the computing time by up to a factor of three if a simulation were to require its time-step to change very often. In practice however, the time-step typically changes only a few times throughout a simulation. In the many simulations we have carried out with the continuous approach, the need for three matrix inversions has never increased the computing time by more than 25%, and the typical increase is around 15%.

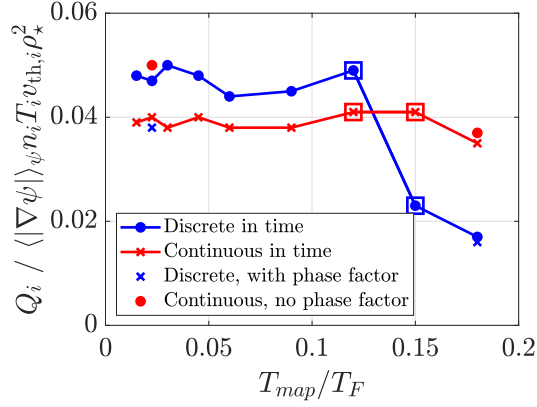
Lastly, while the converged fluxes for this JET discharge differ significantly between the discrete and continuous approaches, it is important to stress that this would by no means be the case for every nonlinear simulation that includes flow shear. In fact, the majority of simulations that we have performed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

5 Conclusions

In this work, we presented a new approach to flow shear for linearly implicit, local δf gyrokinetic codes. This method treats flow shear continuously over time, unlike the widely-used, discrete-in-time approach developed by Hammett *et al* [14]. The main complication resides in the response-matrix approach adopted by codes with an implicit-in-time linear algorithm: when $\gamma_E \neq 0$, this matrix becomes time dependent and it is prohibitively expensive to re-compute it at every time step. In the discrete-in-time approach to flow shear, this time dependence was approximated by combining a nearest grid point approximation (NGP) for radial wave numbers,




(a)



(b)

Figure 7: Turbulent heat flux, plotted for different Δk_x in 7a, and for different $T_{map}(k_y)/T_F$ in 7b (where k_y is the wavenumber with the largest amplitudes). All points in 7a and 7b are identical, except for the ones highlighted by square boxes, for which both Δk_x and Δk_y were modified.



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

Figure 8: Amplitude of the electrostatic potential for the box-scale zonal mode in solid black and for the dominant non-zonal mode in dashed red.

together with a wave number re-mapping procedure. In the new approach to flow shear, a decomposition of terms allows the time dependences in the linear gyrokinetic equation to be treated explicitly in time, without further restricting the time step size. Additionally, time dependences arising in the quasineutrality equation are approximated by linear interpolation, and Hammett’s re-mapping in wave number is applied without nearest grid-point approximation. Following these steps leads to an implicit-in-time linear algorithm, where flow shear is included continuously in time, and for which response matrices only need to be computed during initialisation. Finally, the pseudo-spectral evaluation of the nonlinear term also requires modifications to take into account the time dependence of the radial wave number, as has been argued in [25].

We implemented this new approach in the GS2 code, and showed numerical simulations that validate the changes made to both linear and nonlinear parts of the code. Linearly, the discrete-in-time approach produces discontinuous jumps at every re-map in wave number, and it fails to reproduce Floquet oscillations when the grid-spacing in radial wave number becomes too large. The continuous-in-time approach produces no such jumps, and is able to capture the correct Floquet behaviour for a wider range of resolutions in wave number. For the nonlinear term, a first test case illustrates the errors in real space quantities when a NGP approximation is made in Fourier transforms, and how the continuous-in-time approach gets rid of those errors by including a phase factor. A second test case shows the impact of these errors when the equations for time evolution include a pseudo-spectrally-computed Poisson bracket: the discrete-in-time approach accumulates errors from the NGP approximation, while the continuous-in-time approach correctly reproduces the expected analytical result. Finally, nonlinear gyrokinetic simulations of a JET discharge demonstrate that the continuous and discrete approaches can, in some cases, lead to a different converged value for the turbulent heat flux (20% difference in this example). However for the majority of nonlinear simulations that we have performed with flow shear, the two methods yield essentially identical fluxes. It was also shown that the continuous approach can lead to substantial computational savings by allowing for radially narrower boxes (46% decrease in computing time for this example). When using the same box size, the continuous implementation is somewhat slower than the discrete algorithm (about 15% slower for cases in the example presented here, and at most 25% in all simulations performed so far).

6 Acknowledgments

The authors would like to thank H Weisen and P Sirén for providing the data of the JET shot analysed in section 4.3. They would also like to thank M R Hardman, O Beeke, J Ball and P J Dellar for very useful discussions.

The GS2 branch that includes the continuous-in-time algorithm in the presence of flow shear is available at https://bitbucket.org/gyrokinetics/gs2/branch/ndc_branch with the latest commit at the time of writing being 0abdca. The corresponding “Makefiles” branch can be found at https://bitbucket.org/gyrokinetics/makefiles/branch/ndc_branch with the commit ba24979, and the

corresponding “utils” branch is available at https://bitbucket.org/gyrokinetics/utils/branch/ndc_branch with the commit 8e41f9a.

7 Funding

This work was supported by the Berrow Foundation (scholarship of NC); the Steppes Fund for Change (scholarship of NC); the Fondation H el ene et Victor Barbour (scholarship of NC); Lincoln College, Oxford; the Department of Physics of the University of Oxford; and EPSRC (grant EP/R034737/1).

The authors also acknowledge the use of the ARCHER High Performance Computer through the Plasma HEC Consortium EPSRC (grant EP/L000237/1 under project e281-gs2); the EUROfusion High Performance Computer (Marconi-Fusion, projects FUA32_MULTEI, FUA33_MULTEI and FUA34_MULTEI); and the high performance computing cluster Hydra at the Theoretical Physics subdepartment of the University of Oxford.

8 Declaration of interests

The authors report no conflict of interest.

A

In this section, we present how the system of gyrokinetic and quasineutrality equations is discretised in **GS2**. The discretisation in time and poloidal angle θ can be modified via the parameters r_t and r_θ respectively. The user can set r_t between zero (fully-implicit linear scheme) and one (fully-explicit scheme). Similarly r_θ can be set between zero (centred in θ) and one (upwinding in θ). The discrete indices used in this work are defined in Table 2. With these definitions, the time derivative of any fluctuating quantity Φ is approximated by

$$\frac{\partial\Phi}{\partial t} \approx \frac{1 + \sigma r_\theta}{2} \frac{\Phi_{n+1,j_\theta+1} - \Phi_{n,j_\theta+1}}{\Delta t} + \frac{1 - \sigma r_\theta}{2} \frac{\Phi_{n+1,j_\theta} - \Phi_{n,j_\theta}}{\Delta t}, \quad (53)$$

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

$$\frac{\partial\Phi}{\partial\theta} \approx (1 - r_t) \frac{\Phi_{n+1,j_\theta+1} - \Phi_{n+1,j_\theta}}{\Delta\theta} + r_t \frac{\Phi_{n,j_\theta+1} - \Phi_{n,j_\theta}}{\Delta\theta} \quad (54)$$

and all other terms are approximated by

$$\Phi \approx \frac{1 - \sigma r_\theta}{2} \left(r_t \Phi_{n,j_\theta} + (1 - r_t) \Phi_{n+1,j_\theta} \right) + \frac{1 + \sigma r_\theta}{2} \left(r_t \Phi_{n,j_\theta+1} + (1 - r_t) \Phi_{n+1,j_\theta+1} \right). \quad (55)$$

We now use bold-font vector notation for the potential and distribution function, to denote discrete arrays whose indices span every (θ, k_x) of a single twist-and-shift

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

Table 2: Indices used in discrete notations.

chain. Matrices with two such indices are denoted using bold upper case letters. With this notation, discretising the gyrokinetic equation (14) according to (53), (54) and (55) yields:

$$\mathcal{A}_{n+1}^* \hat{\mathbf{g}}_{n+1} + \mathcal{B}_n^* \hat{\mathbf{g}}_n = \mathcal{C}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} + \mathcal{D}_n^* \hat{\boldsymbol{\varphi}}_n, \quad (56)$$

where we have defined the matrices

$$\mathcal{A}_{n+1,j,j'}^* = \frac{1}{\Delta t} \left(\frac{1 + \sigma r_\theta}{2} \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \delta_{j,j'} \right) \quad (57)$$

$$+ (1 - r_t) w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (\delta_{j+1,j'} - \delta_{j,j'}) \quad (58)$$

$$+ i(1 - r_t) \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_{\theta+1}} \cdot \mathbf{k}_{n+1,j+1}^* \delta_{j+1,j'} \right) \quad (59)$$

$$+ \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n+1,j}^* \delta_{j,j'} \Big), \quad (60)$$

$$\mathcal{B}_{n,j,j'}^* = - \frac{1}{\Delta t} \left(\frac{1 + \sigma r_\theta}{2} \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \delta_{j,j'} \right) \quad (61)$$

$$+ r_t w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (\delta_{j+1,j'} - \delta_{j,j'}) \quad (62)$$

$$+ i r_t \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_{\theta+1}} \cdot \mathbf{k}_{n,j+1}^* \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n,j}^* \delta_{j,j'} \right), \quad (63)$$

$$\mathcal{C}_{n+1,j,j'}^* = - \frac{e Z_s F_{0,s}}{T_s} w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} (1 - r_t) (J_{0,n+1,j+1}^* \delta_{j+1,j'} - J_{0,n+1,j}^* \delta_{j,j'}) \quad (64)$$

$$- i \frac{e Z_s F_{0,s}}{T_s} (1 - r_t) \left(\frac{1 + \sigma r_\theta}{2} \mathbf{V}_{d,s,j_{\theta+1}} \cdot \mathbf{k}_{n+1,j+1}^* J_{0,n+1,j+1}^* \delta_{j+1,j'} \right) \quad (65)$$

$$+ \frac{1 - \sigma r_\theta}{2} \mathbf{V}_{d,s,j_\theta} \cdot \mathbf{k}_{n+1,j}^* J_{0,n+1,j}^* \delta_{j,j'} \Big) \quad (66)$$

$$- i k_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_\psi} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) (1 - r_t) \quad (67)$$

$$\cdot \left(\frac{1 + \sigma r_\theta}{2} J_{0,n+1,j+1}^* \delta_{j+1,j'} + \frac{1 - \sigma r_\theta}{2} J_{0,n+1,j}^* \delta_{j,j'} \right), \quad (68)$$

$$\mathcal{D}_{n,j,j'}^* = -\frac{eZ_s F_{0,s}}{T_s} w_{\parallel} \hat{\mathbf{b}} \cdot \nabla \theta \frac{1}{\Delta \theta} r_t (J_{0,n,j+1}^* \delta_{j+1,j'} - J_{0,n,j}^* \delta_{j,j'}) \quad (69)$$

$$-i \frac{eZ_s F_{0,s}}{T_s} r_t \left(\frac{1 + \sigma r_{\theta}}{2} \mathbf{V}_{d,s,j_{\theta}+1} \cdot \mathbf{k}_{n,j+1}^* J_{0,n,j+1}^* \delta_{j+1,j'} \right. \quad (70)$$

$$\left. + \frac{1 - \sigma r_{\theta}}{2} \mathbf{V}_{d,s,j_{\theta}} \cdot \mathbf{k}_{n,j}^* J_{0,n,j}^* \delta_{j,j'} \right) \quad (71)$$

$$-ik_y \frac{c}{B_r} \left(\frac{\partial F_{0,s}}{\partial r_{\psi}} + \gamma_E \frac{q_0 I m_s w_{\parallel}}{r_{\psi,0} B T_s} F_{0,s} \right) r_t \left(\frac{1 + \sigma r_{\theta}}{2} J_{0,n,j+1}^* \delta_{j+1,j'} \right. \quad (72)$$

$$\left. + \frac{1 - \sigma r_{\theta}}{2} J_{0,n,j}^* \delta_{j,j'} \right). \quad (73)$$

Finally, the quasineutrality equation (21) is evaluated at time step $n+1$ and written in the discrete form:

$$\mathcal{E}_{n+1}^* \hat{\boldsymbol{\varphi}}_{n+1} = \mathcal{W}_{n+1}^* \hat{\boldsymbol{\mathbf{g}}}_{n+1}, \quad (74)$$

where we have defined the matrix

$$\mathcal{E}_{n+1,j,j'}^* = \sum_s \frac{e^2 Z_s n_s}{T_s} (1 - \Gamma_{s,n+1,j}^*) \delta_{j,j'}. \quad (75)$$

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

$$\mathcal{W}_{n+1,j,j'}^* \Phi = \sum_s Z_s \int d^3 w J_{0,n+1,j}^* \delta_{j,j'} \Phi. \quad (76)$$

References

- [1] P. J. Catto, M. N. Rosenbluth, and C. S. Liu. Parallel velocity shear instabilities in an inhomogeneous plasma with a sheared magnetic field. *The Physics of Fluids*, 16(10):1719–1729, 1973.
- [2] M. Artun and W. M. Tang. Gyrokinetic analysis of ion temperature gradient modes in the presence of sheared flows. *Physics of Fluids B: Plasma Physics*, 4(5):1102–1114, 1992.
- [3] M. Barnes, F. I. Parra, E. G. Highcock, A. A. Schekochihin, S. C. Cowley, and C. M. Roach. Turbulent transport in tokamak plasmas with rotational shear. *Phys. Rev. Lett.*, 106:175004, Apr 2011.
- [4] F. J. Casson, A. G. Peeters, Y. Camenen, W. A. Hornsby, A. P. Snodin, D. Strintzi, and G. Szepesi. Anomalous parallel momentum transport due to exb flow shear in a tokamak plasma. *Physics of Plasmas*, 16(9):092303, 2009.
- [5] A. M. Dimits, T. J. Williams, J. A. Byers, and B. I. Cohen. Scalings of ion-temperature-gradient-driven anomalous transport in tokamaks. *Phys. Rev. Lett.*, 77:71–74, Jul 1996.

- [6] P. Mantica, D. Strintzi, T. Tala, C. Giroud, T. Johnson, H. Leggate, E. Lerche, T. Loarer, A. G. Peeters, A. Salmi, S. Sharapov, D. Van Eester, P. C. de Vries, L. Zabeo, and K. D. Zastrow. Experimental study of the ion critical-gradient length and stiffness level and the impact of rotation in the jet tokamak. *Phys. Rev. Lett.*, 102:175002, Apr 2009.
- [7] G. R. McKee, P. Gohil, D. J. Schlossberg, J. A. Boedo, K. H. Burrell, J. S. deGrassie, R. J. Groebner, R. A. Moyer, C. C. Petty, T. L. Rhodes, L. Schmitz, M. W. Shafer, W. M. Solomon, M. Umansky, G. Wang, A. E. White, and X. Xu. Dependence of the l- to h-mode power threshold on toroidal rotation and the link to edge turbulence dynamics. *Nuclear Fusion*, 49(11):115016, sep 2009.
- [8] A. G. Peeters and C. Angioni. Linear gyrokinetic calculations of toroidal momentum transport in a tokamak due to the ion temperature gradient mode. *Physics of Plasmas*, 12(7):072515, 2005.
- [9] E. J. Synakowski, S. H. Batha, M. A. Beer, M. G. Bell, R. E. Bell, R. V. Budny, C. E. Bush, P. C. Efthimion, G. W. Hammett, T. S. Hahm, B. LeBlanc, F. Levinton, E. Mazzucato, H. Park, A. T. Ramsey, G. Rewoldt, S. D. Scott, G. Schmidt, W. M. Tang, G. Taylor, and M. C. Zarnstorff. Roles of electric field shear and shafranov shift in sustaining high confinement in enhanced reversed shear plasmas on the tftr tokamak. *Phys. Rev. Lett.*, 78:2972–2975, Apr 1997.
- [10] F. L. Waelbroeck and L. Chen. Ballooning instabilities in tokamaks with sheared toroidal flows. *Physics of Fluids B: Plasma Physics*, 3(3):601–610, 1991.
- [11] R. E. Waltz, R. L. Dewar, and X. Garbet. Theory and simulation of rotational shear stabilization of turbulence. *Physics of Plasmas*, 5(5):1784–1792, 1998.
- [12] R. E. Waltz, G. D. Kerbel, J. Milovich, and G. W. Hammett. Advances in the simulation of toroidal gyro-landau fluid model turbulence. *Physics of Plasmas*, 2(6):2408–2416, 1995.
- [13] R. E. Waltz, G. M. Staebler, J. Candy, and F. L. Hinton. Gyrokinetic theory and simulation of angular momentum transport. *Physics of Plasmas*, 14(12):122507, 2007.
- [14] G. W. Hammett, W. Dorland, N. F. Loureiro, and T. Tatsuno. Implementation of large scale $\mathbf{E} \times \mathbf{B}$ shear flow in the gs2 gyrokinetic turbulence code. Poster presented at the DPP meeting of the American Physical Society, 2006.
- [15] M. Kotschenreuther, G. Rewoldt, and W. M. Tang. Comparison of initial value and eigenvalue codes for kinetic toroidal plasma instabilities. *Computer Physics Communications*, 88(2):128 – 140, 1995.
- [16] F. Jenko, W. Dorland, M. Kotschenreuther, and B. N. Rogers. Electron temperature gradient driven turbulence. *Physics of Plasmas*, 7(5):1904–1910, 2000.

- [17] A. G. Peeters, Y. Camenen, F. J. Casson, W. A. Hornsby, A. P. Snodin, D. Strintzi, and G. Szepesi. The nonlinear gyro-kinetic flux tube code gkw. *Computer Physics Communications*, 180(12):2650–2672, 2009. 40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures.
- [18] J. Candy and E. A. Belli. Spectral treatment of gyrokinetic shear flow. *Journal of Computational Physics*, 356:448 – 457, 2018.
- [19] P. J. Catto. Linearized gyro-kinetics. *Plasma Physics*, 20(7):719–722, jul 1978.
- [20] E. A. Frieman and L. Chen. Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *The Physics of Fluids*, 25(3):502–508, 1982.
- [21] H. Sugama and W. Horton. Nonlinear electromagnetic gyrokinetic equation for plasmas with large mean flows. *Physics of Plasmas*, 5(7):2560–2573, 1998.
- [22] P. J. Catto, I. B. Bernstein, and M. Tessarotto. Ion transport in toroidally rotating tokamak plasmas. *The Physics of Fluids*, 30(9):2784–2795, 1987.
- [23] M. A. Beer, S.C. Cowley, and G.W. Hammett. Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Physics of Plasmas*, 2(7):2687–2700, 1995.
- [24] M. Barnes, F. I. Parra, and A. A. Schekochihin. Critically balanced ion temperature gradient turbulence in fusion plasmas. *Phys. Rev. Lett.*, 107:115003, Sep 2011.
- [25] B. F. McMillan, J. Ball, and S. Brunner. Simulating background shear flow in local gyrokinetic simulations. *Plasma Physics and Controlled Fusion*, 61(5):055006, mar 2019.
- [26] A. M. Dimits, G. Bateman, M. A. Beer, B. I. Cohen, W. Dorland, G. W. Hammett, C. Kim, J. E. Kinsey, M. Kotschenreuther, A. H. Kritz, L. L. Lao, J. Mandrekas, W. M. Nevins, S. E. Parker, A. J. Redd, D. E. Shumaker, R. Sydora, and J. Weiland. Comparisons and physics basis of tokamak transport models and turbulence simulations. *Physics of Plasmas*, 7(3):969–983, 2000.
- [27] P. Siren, J. Varje, H. Weisen, and L. Giacomelli. Role of JETPEAK database in validation of synthetic neutron camera diagnostics and ASCOT- AFSI fast particle and fusion product calculation chain in JET. *Journal of Instrumentation*, 14(11):C11013–C11013, nov 2019.
- [28] E. A. Belli. *Studies of numerical algorithms for gyrokinetics and the effects of shaping on plasma turbulence*. PhD thesis, Princeton University, 2006.
- [29] R. L. Miller, M. S. Chu, J. M. Greene, Y. R. Lin-Liu, and R. E. Waltz R E. Noncircular, finite aspect ratio, local equilibrium model. *Physics of Plasmas*, 5(4):973–978, 1998.