

Long global gyrokinetic simulations: source terms and particle noise control

B. F. McMillan, S. Jolliet, T. M. Tran, and L. Villard

*Centre de Recherches en Physique des Plasmas,
Association Euratom-Confédération Suisse,*

Ecole Polytechnique Fédérale de Lausanne, PPB, 1015 Lausanne, Switzerland.

A. Bottino

*Max Planck Institut für Plasmaphysik,
IPP-EURATOM Association, Garching*

P. Angelino

Association Euratom-CEA, CEA/DSM/DRFC Cadarache, France

(Dated: March 19, 2008)

In global gyrokinetic simulations it takes a long time for the turbulence to reach a quasi-steady state, and quantitative predictions about the quasi-steady state turbulence have been difficult to obtain computationally. In particular, global particle-in-cell gyrokinetic simulations have been inefficient for long simulations due to the accumulation of noise. It is demonstrated that a simple Krook operator can effectively control noise; it also introduces an unphysical dissipation, which damps the zonal flows and can significantly affect simulation results even when the relaxation time is very long. However, it is possible to project out the effects of the Krook operator on the zonal flows. This permits noise accumulation to be controlled while preserving the physics of interest: simulations are then run to determine the level of quasi-steady state transport and the variation across the ensemble of turbulent dynamics. Convergence is demonstrated both in the number of computational particles and the unphysical relaxation time.

PACS numbers: 52.30.Gz, 52.65.Tt

I. INTRODUCTION

Quite long simulations are needed to reach a quasi-equilibrium state in global gyrokinetic models of ion temperature gradient (ITG) turbulence, even when global profiles of flow and temperature gradient are strongly constrained: although the turbulent cascade in the $n \neq 0$ modes rapidly reaches an apparent quasisteady state, axisymmetric fluctuations in temperature gradient and flow evolve on long timescales. There is a significant separation in timescales between the rapid growth of ITG modes and the slow evolution of axisymmetric perturbations. We can take advantage of this separation of timescales when we introduce a dissipative noise control operator in a PIC (particle in cell) simulation: a Krook operator which weakly damps the non-axisymmetric perturbations but does not damp the zonal flows can substantially improve simulation efficiency without significantly modifying results.

We explore the use of a Krook operator in ORB5, a global gyrokinetic particle-in-cell (PIC) code[1]. The use of a relaxation operator has been proposed elsewhere, notably in ref. [2]: for a simulation with an unphysical relaxation term to correctly reproduce the system dynamics, the effect of the unphysical damping must be sufficiently small. In ref. [2] the coefficient of the Krook operator was adjusted during the simulation to keep the sum of the squared weights equal to some constant W ; they then argued that the true late time steady state could be approached in the limit where $W \rightarrow \infty$. We treat the coefficient of the Krook operator as an adjustable parameter kept fixed during a simulation, and look at the limit where the coefficient, γ_S , goes to zero: these limits should be largely equivalent from the point of view of late time fluxes. We will examine how the behaviour of the simulation is affected by non-zero values of the Krook coefficient. We restrict our attention to electrostatic simulations of ITG turbulence with adiabatic electrons, but the conclusions are probably valid more generally.

The Krook operator counteracts the growth of the marker weights, bounds the computational noise, and keeps the system close to the initial equilibrium state; but it also damps all the wavemodes in the system which we wish to simulate. The damping of wavemodes caused by the Krook operator is not necessarily catastrophic because most of the turbulent structures are rapidly driven and damped (non-axisymmetric perturbations are strongly Landau damped), so that an additional small artificial damping may be acceptable. However, it is clearly important not to artificially damp long-lived structures like zonal flows: the Krook

term acts somewhat like a collisional term and may have a large impact on the zonal flow dynamics. Applying a small zonal flow damping can significantly change the overall simulation results, especially in near-marginal systems, even when the damping rates due to sources are very much smaller than the linear growth rates of the ITG modes. The importance of the preservation of the zonal flows was also noted in Ref. [2], where the cylindrical geometry considered as an example made conservation straightforward. Because there is an explicit form for zonal flow equilibria[3], it is possible to control how a source term will impact the linear zonal flow dynamics, by ensuring the operator acts only on the non-zonal flow components of the distribution function. By relaxing towards an equilibrium that includes the zonal flow structures, we can more accurately reproduce the collisionless physics.

The Krook operator can also be used to maintain the initial temperature profile, but it is often useful to project out the component which acts on the flux-surface averaged kinetic energy so that effects due to profile relaxation can be studied separately: a separate axisymmetric heat source is then applied to model the physical heating.

When we come to use this noise-control operator to perform long simulations, we focus on a medium-size plasma configuration for which it is possible to run various convergence tests, and check the dependence on initial conditions.

II. EFFECTS AND SCALING OF THE KROOK OPERATOR

Consider the effects of a simple Krook operator on the linear Vlasov equation. The evolution of the Vlasov equation can be symbolically written

$$\frac{\partial(\delta f)}{\partial t} = -\dot{\mathbf{z}}_1(\delta f) \cdot \nabla_{\mathbf{z}} f_0 - \dot{\mathbf{z}}_0 \cdot \nabla_{\mathbf{z}} \delta f - \gamma_S \delta f \quad (1)$$

where $\mathbf{z} = (R, v)$ is the phase space coordinate, $\delta f = f - f_0$ is the perturbation from the initial equilibrium, and the subscripts on $\dot{\mathbf{z}}$ give the order in $\delta f/f$. The Krook operator, which is a time varying source term $S_K(\mathbf{x}, \mathbf{v}, t)$, causes a growth rate shift $-\gamma_S$ of the eigenmodes of δf . This damping of the unstable modes will clearly affect the turbulence: mixing length arguments would suggest a suppression of transport by a factor of order of the unphysical change in growth rate. However, other effects like the damping of zonal flows could lead to larger unphysical effects on transport. Later, we will demonstrate a method for preventing zonal flow damping by projecting out a component of the source term (conservation of low

order moments of the distribution function is not sufficient to ensure flow conservation, because the zonal flow equilibria have a somewhat subtle phase space structure). It is also natural to project out the heating effect of the relaxation operator, and have a separate heating operator, but in this section we will consider a Krook term which also acts on the temperature fluctuations.

We wish to use the Krook operator to keep the marker weights relatively small. The weights must at least express the low order moments of the particle distribution which represent the ITG modes and the zonal flows. Given a plasma of a certain size in gyroradius units, with $1/\rho^* = a/\rho$ (a is the minor radius and ρ the thermal gyroradius), gyrokinetic simulations typically find density fluctuations of the ions $\delta n/n \sim \rho^*$ (which is the maximum allowed in the formalism). The scaling of density fluctuations implies that typical values of $\delta f/f \gtrsim \rho^*$, and sets a minimum size for the marker weights. The marker weights are usually much larger than this (and may grow secularly) because we also resolve changes to the temperature profile and fine scale filaments in phase space.

For GyroBohm scaling, the global heat diffusion time $\tau_R \sim O(\rho^{*-2}a/c_s)$ (where c_s is the ion sound speed), and the departure from the initial temperature profile $\delta T/T \sim 1/\gamma_S \tau_R$. In order to keep the δf due to the temperature fluctuations small, it is sufficient to require $\delta T/T \lesssim \rho^*$, so $\gamma_S \gtrsim \rho^*(c_s/a)$.

For small ρ^* , a larger relaxation rate is needed to control noise accumulation due to filamentation than to control noise accumulation due to changes in the temperature profile. We would like to have a substantial proportion of the marker weight representing the current state of the ITG turbulence, the dynamics of which are largely driven by the low order moments of the distribution function, rather than having most of the weight representing the fine scale structure due to filamentation. Consider the power spectrum of fluctuations in velocity space: we can model the filamentation process as an entropy flux, transferring fluctuation power from the low order momentum-space moments driven by the instability to high order moments. In this process, the spatial structures are converted into velocity space structures by ballistic collisionless flow. The entropy flux to the high order moments can be approximated by the product of the square weight $\langle \delta f^2 \rangle_l$ in the low order moments and a typical filamentation rate γ_f ; an estimate for the weight in the high order moments $\langle \delta f^2 \rangle_h \sim \langle \delta f^2 \rangle_l \gamma_f/\gamma_S$ is given by balancing collisionless damping against dissipation by the Krook operator. Because our goal is that the total squared weight be of the same

order as that in the low order moments, we need the relaxation coefficient to scale with the filamentation rate. The filamentation rate is comparable to the inverse eddy time or the linear growth rates, all of which are $O(c_s/a)$. Note that $\gamma_S \propto c_s/a$ only describes how the Krook coefficient should scale with system parameters; the absolute value must still be much smaller than typical linear growth rates to achieve reasonable simulation accuracy.

We now develop a rough estimate for the signal to noise power ratio (which we use as a proxy for the unphysical effects of noise) as a function of the Krook relaxation coefficient. In order to time-evolve the gyrokinetic simulation, we evaluate $M = N'_\varphi N_s (2\Delta m + 1)$ integrals of the gyroaveraged density, using a Monte-Carlo algorithm, where N'_φ is the number of toroidal Fourier modes resolved, N_s is the number of radial intervals, and the range of poloidal modes resolved is $m \in [nq - \Delta m, nq + \Delta m]$ (the non-field-aligned modes are filtered[1]). The distribution function is represented by a finite number of weights $w_i = \delta f p_i$ at positions (x_i, v_i) , representing a volume of phase space p_i , and integrals are calculated from a numerical distribution function $\delta f(x, v) = \sum_i \delta(x - x_i, v - v_i) w_i$ (note that the ORB5 implementation uses some scale factors which are irrelevant to this discussion). We ignore the gyroaveraging for the moment: this mostly just reduces the noise by some numerical factor because the ITG modes are at $k\rho_i \lesssim 1$ (we are interested in scaling rather than an accurate estimate); the M density integrals are then $d_j = \int d\mathbf{x} d\mathbf{v} g_j(x) \delta f$, where we have normalised the basis functions g_j via $\int d\mathbf{x} g_j^2(x) \equiv 1$. For each density integral, a Monte-Carlo evaluation [4] gives a standard error (noise) of

$$d_{\text{noise}} = \left(\frac{\sum_i w_i^2}{N} \right)^{1/2} \quad (2)$$

where the sum is over the N markers in the system (we will suppress the particle index i for the rest of this section), assuming the markers are distributed evenly in space according to the spatial density of modes and weight. We consider the decomposition of the weights w by taking moments of the distribution function: the component w_d of each weight w represents the density fluctuation in configuration space. As long as the decomposition is orthogonal, the component w_d can be considered to contribute a proportion $\sum w_d^2 / \sum w^2$ to the total squared weight of the markers. We assume that in the turbulent steady state the total square weight in the low order moments ($\sum w_l^2$) of the distribution is proportional to $\sum w_d^2$ so that $\sum w_l^2 = C^2 \sum w_d^2$ with C some constant larger than one, but of order unity.

The signal representing the density fluctuation is of typical size

$$d_{\text{signal}} \sim \left(\frac{\sum w_d^2}{M} \right)^{1/2} \quad (3)$$

The weight in the low order moments is transferred into the higher order moments of the distribution by phase mixing/filamentation (with some characteristic rate γ_f which is of order c_s/a), and then damped by the Krook operator. By balancing the Krook damping against the increase in squared weight due to filamentation, we produce the estimate $\sum w^2 / \sum w_d^2 \sim C^2 \gamma_f / \gamma_S$. This leads to an estimate for the signal-to-noise power ratio,

$$\text{Signal/Noise} = N \gamma_S / M \gamma_f C^2. \quad (4)$$

For a simulation without a relaxation operator, the ratio of signal to noise decreases secularly with time as $N/M t_{\text{sim}} \gamma_f C^2$: because the Krook relaxation time can be much shorter than the simulation time, one can achieve a reasonable signal to noise ratio in a simulation with a Krook operator with a much smaller number of markers than in a standard PIC simulation, and reduce the computational requirements considerably.

If γ_S / γ_f is kept constant, the total number of markers required to achieve a fixed signal to noise target therefore scales with M , which in turn scales with $1/\rho^{*2}$, as $N_\varphi, N_s \propto 1/\rho^*$, and Δm is independent of ρ^* . The simulation time per timestep is approximately proportional to the number of markers (per-particle computations generally dominate), and therefore will also scale like $1/\rho^{*2}$ (this scaling is as good as can be expected for any gyrokinetic code given that the size of a field aligned grid also scales like $1/\rho^{*2}$).

III. LINEAR ZONAL FLOWS DYNAMICS WITH SOURCE TERMS.

ITG turbulence is mediated strongly by zonal flows (ZFs), and these flows can be sensitive to small changes in simulation properties, like the introduction of collisionality, or the artificial damping introduced by a relaxation operator. Even source terms which are intended only to keep the temperature profile near a target profile may affect the zonal flow dynamics. Gyrokinetic codes must accurately resolve the correct linear physics of zonal flows, and this is usually tested explicitly during the benchmarking of a code. In this section, we consider the predictions of linear theory for the behaviour of long timescale axisymmetric flows in simulations with sources, and demonstrate how to introduce a Krook operator (the source

S_K) which does not damp zonal flows. We also briefly discuss sources used in other global codes, and their impacts on linear stability of zonal flows.

We introduce a scheme for ZF conservation in the presence of sources based on an explicit representation of the ZF structure as calculated in ref. [3]. Ref. [3] gives the long-time zonal flow potential ϕ in terms of integrals of the sources $R = \int dt S$. For zonal flows of long wavenumber compared to the banana width, but small compared to system scales, this reduces to a relatively simple set of relations. To maintain ϕ constant (given an initial pure ZF state) it is sufficient to ensure (at each radial position) that

$$\left\langle \int d\mathbf{v} \left[v_{\parallel}/B - \overline{(v_{\parallel}/B)} \right] S(\vec{R}, \vec{v}) \right\rangle = 0 \quad (5)$$

and

$$\left\langle \int d\mathbf{v} S(\vec{R}, \vec{v}) \right\rangle = 0 \quad (6)$$

where the angle brackets $\langle . \rangle$ represent a flux surface average, and the bar over a quantity indicates an orbit (bounce) average. Note that the second condition is just that the source does not change the flux surface averaged density.

Given a source term S which perturbs the zonal flows, we can construct a source which does not perturb the flux-surface averaged potential by adding a second source term, which we choose based on the ZF solution in ref. [3],

$$S_{\text{corr}} = g_1(s) \left[v_{\parallel}/B - \overline{(v_{\parallel}/B)} \right] f_0 + g_2(s) f_0. \quad (7)$$

with g_1 and g_2 chosen such that the two integral relations above hold (in the implementation we only require that they are true on average across a set of radial bins). The correction scheme described here appears somewhat artificial because it does not have an obvious physical motivation (unlike, say, conserving various low order moments of f on a flux surface): it has been specifically designed to project out zonal flow effects from the Krook operator, and to do this we must capture the phase space structure of the zonal flows.

We performed simulations with an initially sinusoidal density perturbations to evaluate the effect of this Krook operator on the residual zonal flows. Figure 1 shows a typical plot of the radial electric field at a particular radial point during such a test; the rapid oscillation is the geodesic acoustic mode (GAM), and the late time flows are identified as ZFs. Without the ZF correction, the zonal flows are rapidly damped, unlike in the case with no Krook operator or a corrected operator, where the flows persist on the timescale shown. Figure 2

shows the asymptotic damping rates (for long time and small values of the relaxation rate) relative to the relaxation rate γ_S as a function of the radial ZF wavenumber (in units of the inverse banana width). As γ_S becomes comparable to c_s/a the relative ZF damping rate decreases. Uncorrected simulations have anomalous ZF damping rates approximately equal to the relaxation rate. Because the zonal flow conservation property is based on the ordering $k_r \rho_i q (R/r)^{1/2} \lesssim 1$, it is to be expected that the ZF correction is less effective at large radial wavenumber. Below $k_r \rho_i q (R/r)^{1/2} = 0.5$ the correction reduces the unphysical damping by a factor of more than 10.

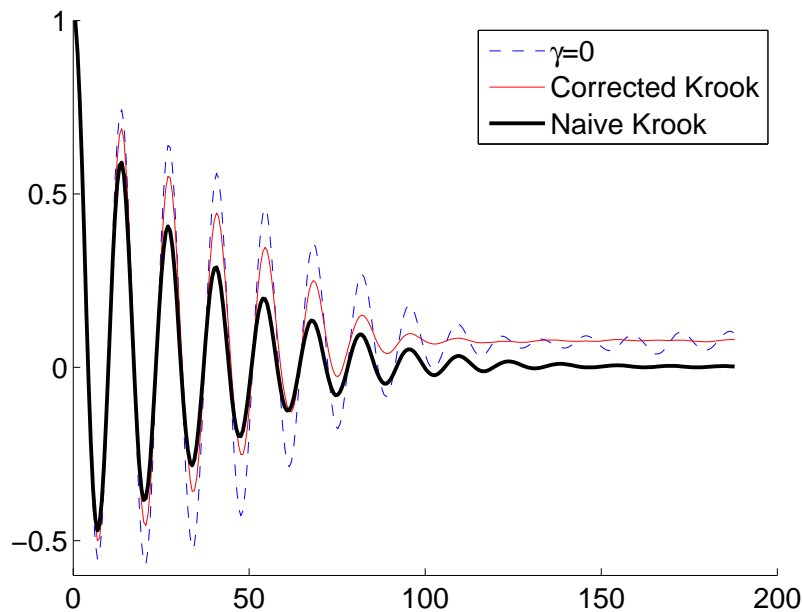


FIG. 1: Zonal flow evolution in simulation with a relaxation operator. The dashed blue curve shows the evolution of the electric field at $s = 0.5$ for a simulation with no relaxation operator. The thin solid red curve shows a simulation with the correction added to the Krook operator. The thick solid black curve shows the result with a Naive Krook operator.

The correction technique does not prevent the artificial damping of GAMs by the Krook operator. The GAM damping is about the same as in an uncorrected simulation. The GAMs in any case are linearly damped so their dynamics would not be expected to be as sensitive to an additional small damping as the zonal flow dynamics. The relaxation rate in the simulation used to produce figure 2 ($\gamma_S = 0.06c_s/a$) was set five times higher than typical for our nonlinear simulations, in order to demonstrate that the ZF correction is effective;




FIG. 2: Anomalous zonal flow damping rate (compared to the relaxation rate) as a function of zonal flow radial wavenumber in a simulation with a corrected relaxation operator.

the artificial GAM damping in nonlinear simulations is generally very small.

In order to control the temperature profile independently of the noise accumulation, we can also project out the heating effect of the Krook operator, and introduce an axisymmetric heat source S_H into ORB5. In this case the Krook operator S_K must not influence the total kinetic energy on each flux surface, so

$$\left\langle \int d\mathbf{v} v^2 S(\vec{R}, \vec{v}) \right\rangle = 0. \quad (8)$$

We can ensure this by adding a third term to S_{corr} , $g_3(s)v^2 f_0$ (flux surface density is already conserved). The heat source S_H damps the components of the perturbed distribution $\delta f(\epsilon, s)$ on a timescale γ_H . To ensure that the source does not introduce a density perturbation, a correction is applied on each flux surface, and we then have

$$S_H = -\gamma_H \left[\delta f(\epsilon, s) - f_0(\epsilon, s) \frac{\int d\mathbf{v} \delta f(\epsilon, s)}{\int d\mathbf{v} f_0(\epsilon, s)} \right]. \quad (9)$$

In the numerical implementation, $\delta f(\epsilon, s)$ is calculated by binning the marker weights on an (ϵ, s) grid. We do not apply radial smoothing to our heat source, unlike the codes GYRO and GTC. Because the source S_H is even in v_{\parallel} and conserves density on flux surfaces, it is

not expected to affect the linear dynamics of long-wavelength zonal flows. This is unlike the source of GYRO, which was expected to damp some of the system-scale flows[5]. The source used in the GTC runs of Ref. [6] on the other hand, is also even in v_{\parallel} and conserves density on each flux surface, so is expected not to linearly damp long wavelength zonal flows.

The equation for the evolution of the weights is then

$$\frac{dw_i}{dt} = -p_i \left. \frac{\partial f_0}{\partial \mathbf{z}} \right|_{\mathbf{z}_i} \cdot \dot{\mathbf{z}}_i - \gamma_S w_i + p_i S_{\text{CORR}}(\mathbf{z}, t) + p_i S_H(\mathbf{z}, t) \quad (10)$$

IV. NONLINEAR SIMULATIONS WITH A ZF PRESERVING KROOK OPERATOR.

Several simulations were run using parameters designed to match the CYCLONE base case[7] near mid radius. The physical parameters are identical to those described in Ref. [1] with $a/\rho_{s0} = 175$, $L_n = 0.45R_0$, and $a/R_0 = 0.36$, with s_0 the mid-radius point. The nominal temperature gradient scale length is $L_T = 0.145R_0$. The number of radial, poloidal and toroidal grid intervals are $N_S = 128$, $N_\chi = 512$ and $N_\varphi = 256$ respectively. The field aligned filter has a width of $\delta n \pm 5$. Only modes with $|n| \leq 120$ are kept in the filter, so the number of resolved toroidal modes $N'_\varphi = 241$. The timestep is $0.29(a/c_s)$. Unless otherwise stated, 1.5×10^8 markers are used. The logarithmic temperature gradient profile was flat over most of the simulation region (figure 3). The Krook operator was corrected using the technique of section III to prevent artificial zonal flow damping, and the relaxation rate was set to $1.26 \times 10^{-2} c_s/a$.

The initial linear growth rates are reduced because of the Krook operator, and given by $\gamma \rightarrow \gamma - \gamma_S$ as predicted. After the linear growth stage, the simulations follow the well known pattern of saturation through zonal flow generation: low n modes are generated as the linearly unstable modes approach their peak energies, and the $n = 0$ mode rapidly becomes the dominant component. The initially linearly unstable modes then decrease significantly in energy, and a nonlinear downshift of mode energy occurs, so that the dominant non-axisymmetric toroidal mode numbers at late times are around half the mode number of the most unstable linear mode.

Even though classical drift wave turbulence is deterministic, the detailed evolution of an experiment or a simulation may change wildly with small variations in initial conditions or parameters. It is not even clear that long-time averages of quantities like heat flux will




FIG. 3: Relative temperature gradient (R/L_T) versus radius at the beginning (thin line) of a simulation with a source term, and averaged over the last half of the simulation ($t \in [850, 1700](a/c_s)$, thick line).

converge to the same values given different initial conditions: for example, some turbulent systems are known to be subcritical, and have both a quiescent state and a turbulent state that can persist at long times. In fact, we are not usually interested in the behaviour of turbulence with precise parameters and very specific initial conditions, but rather the generic behaviour of turbulent systems in some parameter range; we are interested in the ensemble averaged behaviour of the turbulence. For some turbulent systems the time average asymptotes to the ensemble average (this is known as the ergodic hypothesis) but it is not known whether this is the case for ITG turbulence, and if it is the case, we do not know how rapidly this occurs. With the aim of exploring at least a certain portion of the parameter space, we ran several simulations with different quasi-random initial conditions (the marker weights are initialized using the same Hammersley sequence, but we use a parameter C to modify the sequence $a'_i = \text{mod}[a_i + C, 1.0]$ and vary the initialisation). Figure 4 shows the fluxes in the three simulations, averaged over the radial range $s = [0.52, 0.72]$ and using a moving time window. We use the time averaging window length suggested by ref. [8] of $500(a/c_s)$: at time t in the plots the averaging window is $[t - 500(a/c_s), t]$. The unit used

for the flux is the Dimits flux $\chi_D = \chi_{GB} a / L_n \sim 0.8 \chi_{GB}$, where the gyroBohm diffusivity is defined as $\chi_{GB} = \rho^2 (c_s / a)$. Ref. [8] argues that there is a transient period in their simulations after which this moving average of the flux asymptotes to the long term average, and the transient period is finished after $400a/c_s$, but our simulations the transients are longer, and there are substantial variations at late time: we suggest that the length of the transient and the window length required for good statistics are not universal. Although the use of a time averaging technique greatly reduces the variability in the flux signal, there remain large fluctuations from the ensemble trend, and these occur even at late time. To quantify the simulation-to-simulation variability, we find the standard deviation of the time-windowed flux average (ignoring the first $200c_s/a$ before nonlinear saturation) away from the ensemble trend. Here we somewhat arbitrarily take a quadratic fit to the data to represent the ensemble trend. The normalised standard deviation away from the fit is $0.13\chi_D$; the fit and lines two standard deviations above and below the fit are shown in figure 4. Despite the very small sample of simulations, this fitting and variance measurement procedure at least gives a rough quantification of variability within the ensemble: we can then assess whether changing some other parameter has had a significant impact on flux levels, by examining whether the moving-averaged values of the flux lie in the error bounds. In general the results of other gyrokinetic codes will also be sensitive to intrinsic variability, but this is often not discussed, and results are often given without quantification of repeatability. We point out that the quadratic fit generally trends downwards, and it is not clear that the ensemble has converged to some late time steady state.

Noise effects are diagnosed using the signal to noise diagnostic proposed in Ref [9]. This diagnostic estimates the proportion of the squared density perturbation resolved in the simulation resulting from Monte-Carlo sampling errors (due to the finite number of markers). The relative importance of many effects of physical turbulence and noise driven processes (like, for example, heat fluxes) might be expected to be dependent on this ratio. More sophisticated analyses of interactions of turbulence with noise[10] might be able to give direct estimates of how the spectral properties depend on noise levels, but we pursue the pragmatic alternative of demonstrating convergence with the number of markers in the simulation. Previous convergence tests[9] indicated that reasonable results are obtained when the signal to noise ratio (SNR) is above 10, and the simulations presented here have a SNR above 30 unless otherwise indicated. A typical trace of the signal to noise ratio is shown

in fig. 5; there is no long term trend evident in the noise level. The signal level is low in the initial part of the linear evolution because the ITG modes was allowed to grow from marker noise. We also performed convergence tests with respect to particle number (leaving all other parameters fixed). Because the system is turbulent, two simulations with slightly different parameters have exponentially divergent evolution, so that they rapidly become significantly different. The very strong sensitivity of detailed simulation evolution to parameters makes it very difficult to demonstrate absolute simulation convergence except for small and short simulations[1]. However, the simulations are not intended the evolution of a specific set of turbulent eddies, but rather to model the generic evolution of a turbulent system with certain boundary conditions. The simulation results of interest are by nature statistical quantities, like the level and variation of time-averaged late time fluxes. It is therefore sufficient to demonstrate the convergence of these averaged quantities rather than the detailed simulation history. We ran cases with 3.5, 7.5, 15 and 30×10^7 markers to look at the convergence of late-time quantities with respect to noise. The early-time instantaneous flux traces shown in figure 6 rapidly become uncorrelated, even though the initial perturbation (a smooth density perturbation we call *mode initialisation*) is the same. However, the moving time averages of the flux (Fig. 7) all fall within the variability bounds of the ensemble simulations shown earlier, indicating that changing the number of markers does not have a statistically significant effect on the flux, even for the case with the smallest number of markers. The late time SNR is roughly inversely proportional to the number of markers.

To show the effects of the Krook operator on the noise, and noise on the simulation, we performed a simulation without the Krook source term S_K but with the axisymmetric heat source S_H . In this simulation there is a late time drop in the flux (Fig. 8) as the noise accumulates (at around $t(c_s/a) = 450$), which occurs around the time where the signal to noise ratio drops below 10 (Fig. 5). Noise is expected to accumulate at a roughly constant rate (in a converged, steady state simulation), and the noise power scales with the reciprocal of the number of markers, so given that the SNR is 10 after a period $200(a/c_s)$ of saturated turbulence, we estimate that we would need 25 times as many markers to maintain the SNR above 30 until $t = 1700$: the simulation with a Krook operator is in this sense far more efficient. The average late time flux on the interval $t(c_s/a) \in [400, 900]$ is $\sim 40\%$ or $1.0\chi_D$ lower than in the simulation using the Krook operator, a much larger deviation than expected from intrinsic chaotic variation: this indicates that the simulation is clearly




FIG. 4: Temporal moving average (thick lines) over $500(a/c_s)$ of χ/χ_D on the radial range $s = [0.524, 0.724]$ for three simulations with $R/L_{T0} = 7.2$ with different initial conditions. The thick solid line is a quadratic fit to the three moving averages. The dashed lines based on the simulation-to-simulation variability bound the likely values of the flux.

affected by the removal of the Krook operator and the consequently increased noise level. The signal (in the $n \neq 0$ modes) eventually asymptotes to the noise level, indicating a noise dominated simulation, but the flux is relatively constant at late times. In fact, the much lower burstiness of the flux signal compared to noise-controlled simulations seems to be a strong indicator of simulation failure (a pure noise field is expected to have short temporal and spatial correlation lengths). The drop in flux may well be the result of unphysical perpendicular diffusion in the noise field[11] which can damp radially extended structures (such as those shown in fig. 9), even when the noise induced diffusion is much lower than the expected turbulent diffusion rate. The drop in flux occurs together with a reduction in the amplitude of the energy in the dominant toroidal mode numbers ($n \in [5, 20]$) and a large increase in the $n = 0$ energy, but the energy in the high n component of the spectrum does not change much as the flux drops. The zonal flow organises into a single radial oscillation. Streamers cannot be seen in plots of the non-zonal potential, which shows little organised structure at late time (Fig. 10), as previously noted[11].




FIG. 5: Signal to noise ratio in two simulations with $R/L_{T0} = 7.4$ and 1.5×10^8 markers. One simulation (thin line) uses a Krook operator with $\gamma_S = 1.26 \times 10^{-2}(c_s/a)$, the other (thick line) a heat source with $\gamma_H = 1.26 \times 10^{-2}(c_s/a)$.

The initial temperature gradient is relatively flat, although some departure from flatness and the nominal maximum value occur due to the effects of the canonical Maxwellian equilibrium[12]. Temperature gradients are controlled by the Krook operator during the simulation, but still evolve somewhat, and the region with a flat temperature gradient broadens (Fig. 3) at late time, as the average gradient reduces slightly. There are also significant temperature fluctuations on short radial length scales, even when temporally averaged over the last half of the simulation. There is a very clear correlation between the late time temperature gradient profile and the zonal flow structure which can be seen in figures 11 and 12. The short wavelength radial temperature fluctuations might play an important role in the nonlinear zonal flow dynamics, but our heat source is expected to damp these radial temperature fluctuations somewhat. Further study of the effect of unphysical heat sources in gyrokinetic simulation would be desirable.

Figure 12 shows the formation of zonal flows on both large and small scales. The standard interpretation of zonal flow drive is through the radial derivative of the Reynolds stress: the Reynolds stress has a global-scale quasi-steady profile due to the quasi-equilibrium profile, as




FIG. 6: Instantaneous value of χ/χ_D on the radial range $s = [0.524, 0.724]$ for four simulations with $R/L_{T0} = 7.2$ but different numbers of computational markers.

well as small scale time-fluctuating component due to random turbulent fluctuations. Large scale flows due to the turbulence gradient are absent in radially homogeneous turbulence simulations, and might need to be explicitly damped for the purposes of benchmarking against flux tube codes. Even though the growth timescale of system-scale flows is expected to become much longer than typical gyrokinetic timescales in large systems, this is evidently not the case here with $a/\rho = 140$: system-scale flows are seen almost immediately after saturation.

Plots of the flux versus time and radius (figure 13) can be roughly described as a background averaged flux, plus diagonally aligned bursts. This appears to be the result of an avalanche process which will be considered in detail elsewhere.

We also performed simulations without projecting out the effects of the Krook operator on zonal flows. Four simulations were performed with 300, 150, 75 and 37.5 million markers with a naive Krook operator: all of these simulations appear to be well converged in particle number, and we treat them here as an ensemble of equivalent simulations. The moving-average flux traces are shown in figure 14, together with the 2σ bounds from the ensemble of simulations using a ZF conserving Krook operator. Near the beginning of the simulation




FIG. 7: Moving time average (over $500(a/c_s)$) of χ/χ_D on the radial range $s = [0.524, 0.724]$ (thin solid curves) for three simulations with $R/L_{T0} = 7.2$ but different numbers of computational markers. The thick solid line gives the fit from figure 4, and the dashed curves are the variability bounds from figure 4.

($t = 800a/c_s$) the averaged flux levels are close to those in the zonal flow conserving ensemble but towards the end ($t = 1450a/c_s$) all the average flux levels are above the 2σ bounds. It is therefore clear that the ZF correction makes a significant difference to late time flux levels. Qualitatively, there is a clear decrease in the long radial wavelength components of the zonal electric field when the zonal flow conservation is turned off: long wavelength flows no longer build up. It appears that the fluxes reach a quasi-steady state much more quickly when we introduce this artificial zonal flow damping which suggests that the long timescales required to reach steady state in the earlier simulations are a result of the buildup of large scale zonal flows. Note that we are considering the collisionless limit, but physical tokamaks are at least somewhat collisional, and even small values of collisionality will damp the large scale flows. In a simulation with collisions the ZF correction may be less important.

Convergence of the simulations with respect to the relaxation time was tested by allowing the relaxation time to become long while keeping the temperature gradients and noise levels fixed. We used the heat source S_H to fix the temperature profile fixed, while projecting out




FIG. 8: Average χ/χ_D on the radial range $s = [0.524, 0.724]$ for two simulations with $R/L_{T0} = 7.2$ (solid lines) and their temporal moving average (dashed lines) over $500(a/c_s)$. One simulation (thin line) uses a Krook operator with $\gamma_S = 1.26 \times 10^{-2}(c_s/a)$, the other (thick line) a heat source with $\gamma_H = 1.26 \times 10^{-2}(c_s/a)$.

the heating effects of the Krook term S_K , as described in section III. Average noise levels in the simulation are expected to be proportional to $1/N\gamma_S$, so we scaled the number of simulation markers inversely with the relaxation rate γ_S to keep noise levels fixed. Simulations were run with $\gamma_S = 0.63, 1.26, 2.52$ and $5.04 \times 10^{-2}(c_s/a)$, and γ_H fixed at $1.26 \times 10^{-2}(c_s/a)$ (for $\gamma_S = 0$ the initial linear ITG growth rate is $0.14(c_s/a)$ at midradius). Moving time averages of the diffusivity in each of these simulations are plotted in figure 15, along with the variability range from figure 4. In the limit where the relaxation parameter is small, we have several simulations with different values of γ_S but similar dynamics, which are quantitatively difficult to distinguish. Even though the three flux traces for $\gamma_S \leq 2.52 \times 10^{-2}(c_s/a)$ do not lie entirely within the variability bounds, there is no clear trend with γ_S , and the departures are on the boundary of statistical significance. We therefore conclude that the Krook operator has at most a small effect for $\gamma_S \leq 2.52 \times 10^{-2}(c_s/a)$. For the simulation with $\gamma_S = 5.04 \times 10^{-2}(c_s/a)$, the heat flux is very small, and the simulation appears to be marginally stable. As a rule of thumb, we suggest that $\gamma/\gamma_S \gtrsim 10$ is probably sufficient to

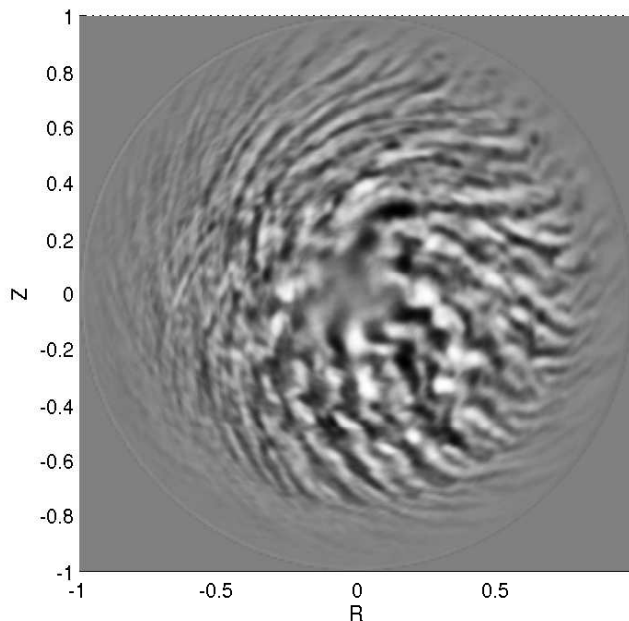


FIG. 9: Non-zonal electric potential (as greylevels) on the plasma cross section in the late-time saturated phase of the simulation with $R/L_{T0} = 7.2$ (solid lines) and a Krook operator with $\gamma_S = 1.26 \times 10^{-2}(c_s/a)$.

obtain reasonable simulation accuracy, given that the difficulty of time averaging (among other things) precludes a very accurate comparison with other codes or experiment. The expected scaling of noise levels with $1/N\gamma_S$ was well observed, so the late-time SNR of all these simulations was around 30.

V. CONCLUSIONS

A relaxation operator is an attractively simple way to ensure that marker weights stay bounded in gyrokinetic PIC simulations and that the turbulence approaches a steady state. In order to preserve the physics of interest, the relaxation operator must act only on states which have large typical damping or forcing rates so that the additional effect of the operator is negligible in the balance. All the nonaxisymmetric perturbations are rapidly Landau damped (although they may have a positive net growth rate), but some axisymmetric perturbations are not collisionlessly damped. In particular, the zonal flows in the plasma sometimes have zero or near-zero damping (they are not linearly damped), and control the




FIG. 10: Non-zonal electric potential (as greylevels) on the plasma cross section in the late-time noise-dominated phase of the simulation with $R/L_{T0} = 7.2$ (solid lines) and a heating operator with $\gamma_H = 1.26 \times 10^{-2}(c_s/a)$.

level of non-axisymmetric turbulence. We demonstrated numerically that a small artificial damping of the zonal flows has a large effect on simulation results. We therefore require that the relaxation operator does not damp the zonal flows in the plasma. Using a Krook operator which does not damp zonal flows, we have demonstrated that simulation results are not significantly modified even for relatively large values of the relaxation rate, and that simulation times are not limited by noise accumulation. To perform the simulations shown in this paper without using a Krook operator would require at least an order of magnitude more markers; this is likely to be typical, so that the use of a zonal-flow conserving Krook operator will result in a reduction in computational requirements over a standard PIC simulation by more than an order of magnitude, without significantly modifying the results.

We note that other techniques have also been examined for noise reduction in δf PIC simulations using coarse-graining onto phase space grids[13, 14]. These schemes can be less dissipative than a Krook operator because they act preferentially on fine-scale structure in the distribution function. There are some drawbacks to the coarse-graining approach,




FIG. 11: Change in inverse temperature scale length $R/L_T - (R/L_T)|_{t=0}$ versus time and radius for a simulation with $(R/L_{T0})|_{t=0} = 7.4$.

however, including the need to handle a high resolution 5-D grid. A side-by-side comparison would be interesting to examine whether the reduction in dissipation is sufficient to justify the additional computational expense. It would also be useful to determine the level of unphysical zonal flow dissipation in a coarse-grained simulation as a function of grid resolution.

-
- [1] S. Jolliet, A. Bottino, P. Angelino, R. Hatzky, T. Tran, B. McMillan, O. Sauter, K. Appert, Y. Idomura, and L. Villard, *Computer Physics Communications* **177**, 409 (2007).
 - [2] J. A. Krommes, *Physics of Plasmas* **6**, 1477 (1999).
 - [3] M. N. Rosenbluth and F. L. Hinton, *Phys. Rev. Lett.* **80**, 724 (1998).
 - [4] A. Aydemir, *Physics of Plasmas* **1**, 822 (1994).
 - [5] R. Waltz, J. Candy, and M. Rosenbluth, *Physics of Plasmas* **9**, 1938 (2002).
 - [6] Z. Lin, T. Hahm, S. Ethier, W. Lee, J. Lewandowski, G. Rewoldt, W. Tang, W. Wang, L. Chen, and P. Diamond (2002).




FIG. 12: Radial electric field (shown as greylevels) as a function of simulation time and position in a simulation with $R/L_{T0} = 7.4$.

- [7] 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, *Physics of Plasmas* **7**, 969 (2000).
- [8] J. Candy, R. Waltz, and W. Dorland, *Physics of Plasmas* **11**, L25 (2004), ISSN 1070-664X.
- [9] A. Bottino, A. G. Peeters, R. Hatzky, S. Jolliet, B. McMillan, T. Tran, and L. Villard, *Physics of Plasmas* **14**, 010701 (2007).
- [10] J. Krommes, *Physics of Plasmas* **14**, 090501 (2007).
- [11] W. Nevins, G. Hammett, A. Dimits, W. Dorland, and D. Shumaker, *Physics of Plasmas* **12**, 122305 (2005).
- [12] P. Angelino, A. Bottino, R. Hatzky, S. Jolliet, O. Sauter, T. Tran, and L. Villard, *Physics of Plasmas* **13**, 052304 (2006).
- [13] S. Brunner, E. Valeo, and J. Krommes, *Physics of Plasmas* **6**, 4504 (1999).
- [14] Y. Chen and S. E. Parker, *Physics of Plasmas* **14**, 082301 (2007).

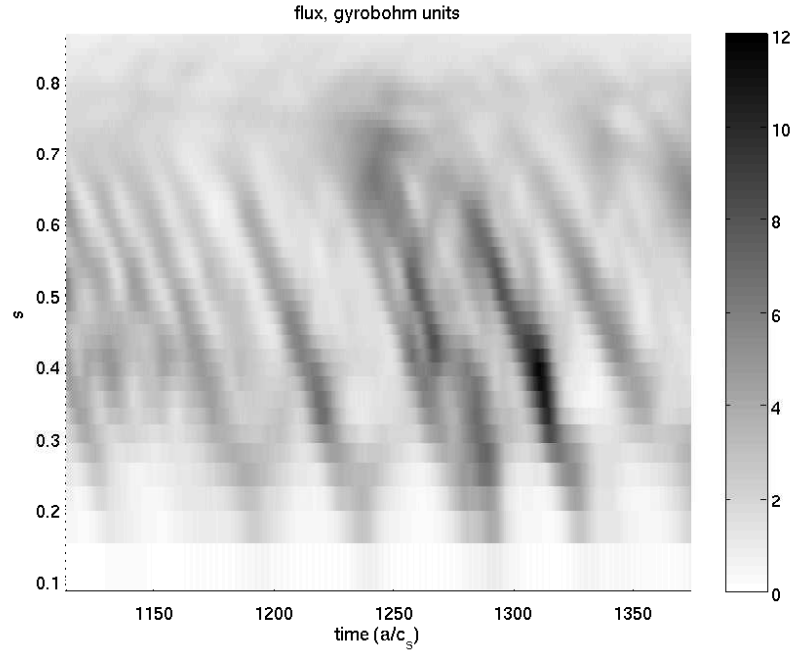


FIG. 13: Flux χ/χ_D as a function of time and radial position, during the late nonlinear phase, for the simulation with $R/L_{t0} = 7.4$.




FIG. 14: Temporal moving average (thin solid lines) over $500(a/c_s)$ of χ/χ_D on the radial range $s = [0.524, 0.724]$ for four simulations with $R/L_{T0} = 7.2$ with different numbers of markers and a naive Krook operator. The upper thick solid line shows a quadratic fit to these temporal moving averages. The lower thick solid line shows the fit from the simulations using a corrected Krook operator, and the dotted lines are the error bounds around this fit.




FIG. 15: Moving time average (over $500(a/c_s)$) of χ/χ_D on the radial range $s = [0.524, 0.724]$ for three simulations with $R/L_{T0} = 7.2$ but different values of the relaxation coefficient.