A global collisionless PIC code in magnetic coordinates

S. Jolliet [a,*], A. Bottino [b], P. Angelino [a], R. Hatzky [c], T. M. Tran [a], B. F. Mcmillan [a], O. Sauter [a], K. Appert [a], Y. Idomura [d], L. Villard [a]

[a] *Ecole Polytechnique Fédérale de Lausanne Centre de Recherches en Physique des Plasmas Association Euratom-Confédération Suisse CH-1015 Lausanne, Switzerland*

[b] *Max Planck Institut für Plasmaphysik IPP-EURATOM Association, Garching, Germany*

[c] *Rechenzentrum der Max-Planck-Gesellschaft und des Max-Planck-Institutes für Plasmaphysik, EURATOM-Association, D-85748, Germany*

[d] *Japan Atomic Energy Agency, Higashi-Ueno 6-9-3, Taitou, Tokyo 110-0015, Japan*

Abstract

A global plasma turbulence simulation code, ORB5, is presented. It solves the gyrokinetic electrostatic equations including zonal flows in axisymmetric magnetic geometry. The present version of the code assumes a Boltzmann electron response on magnetic surfaces. It uses a Particle-In-Cell (PIC), δf scheme, 3D cubic B-splines finite elements for the field solver and several numerical noise reduction techniques. A particular feature is the use of straight-field-line magnetic coordinates and a field-aligned Fourier filtering technique that dramatically improves the performance of the code in terms of both the numerical noise reduction and the maximum time step allowed. Another feature is the capability to treat arbitrary axisymmetric ideal MHD equilibrium configurations. The code is heavily parallelized, with scalability demonstrated up to 4096 processors and 10^9 marker particles. Various numerical convergence tests are performed. The code is validated against an analytical theory of zonal flow residual, geodesic acoustic oscillations and damping, and against other codes for a selection of linear and nonlinear tests.

Key words: gyrokinetics, PIC, micro-instabilities, ITG, zonal flows
PACS: 52.30.Gz, 52.35.Qz, 52.35.Ra

* Corresponding author.
Email address: sebastien.jolliet@epfl.ch (S. Jolliet).

1 Introduction

Anomalous transport is currently one of the most active fields of research in magnetic confinement fusion. The cross-field transport measured in Tokamak experiments exceeds the neoclassical predictions by up to two orders of magnitude for the low confinement regime (L-mode) [1]. Even in the high confinement regime (H-mode) where transport is reduced near the plasma edge, anomalous transport still persists in the core. This phenomenon is attributed to plasma turbulence [2], which is driven by micro-instabilities [3]. The latter are mainly generated by free sources of energy in the plasma such as density and temperature gradients. It is now widely believed that self-organization of turbulence plays a crucial role. It refers to a process in which the internal organization of a system increases automatically without being guided or managed by an outside source. In magnetized plasmas, self-organization mainly appears through two different structures: zonal flows [4] and large-scale transport, such as avalanches [5] and streamers [6, 7].
In general, a full kinetic treatment of micro-instabilities is needed. In this framework, the gyrokinetic model [8] is usually sufficient for drift wave turbulence at low frequency. It removes the gyorangle dependance in the original equations, thus reducing the phase space from 6D to 5D. Various gyrokinetic equations can be found for example in [9–12]. Among the different approaches used to solve the gyrokinetic equations, the Particle-In-Cell (PIC) method [8, 13–18] is one of the most promising schemes. The distribution function is sampled along trajectories with numerical particles (markers). The main advantage is that the distribution function is not discretized on a 5D grid as in Eulerian [19–21] and semi-Lagrangian [22–24] codes, the latter being more difficult to handle numerically. Moreover, it is conceptually simple, easily generalized to multi-dimensional simulations and more adapted for complex geometries such as stellerators [25, 26] than Eulerian codes. However, the PIC method unavoidably gives rise to statistical noise which can lead to an unphysical behaviour in the nonlinear phase of the simulation, but significant progresses have been made to limit this effect [27, 28]. In addition to the solving methods, another important distinction between the different models is the simulated domain. In local simulations, turbulence is studied on a single magnetic surface. A less restrictive method is to simulate a flux-tube following a given magnetic-field line. This approach reduces the computer ressources needed, but imposes inconsistent T and ∇T profiles (i.e. $T = \text{const}, \nabla T = \text{const}$) and generally employs unphysical radially periodic boundary conditions. The global approach is certainly the most realistic model, because it contains the whole radial domain and therefore the effects of profile variation.
The aim of this paper is to present the ORB5 code, originally written by Parker [18] and further developed by Tran [13]. ORB5 is a nonlinear gyrokinetic global code which solves the Vlasov-Poisson system in the electrostatic

and collisionless limit, and has the unique capability of handling true MHD equilibria [29, 30].
The paper is organized as follows. In Sec. 2, the gyrokinetic equations are derived. Sec. 3 presents the implementation of these equations and the associated numerical algorithms. Sec. 4 shows the parallel performance of the code. Results related to the numerical schemes and benchmarks are shown in Sec. 5. Finally, conclusion and future works are exposed in Sec. 6.

2 Gyrokinetic model

2.1 Definitions and normalization

The code ORB5 solves the Vlasov-Poisson system in the gyrokinetic limit for an axisymetric toroidal plasma. Magnetic surfaces are labelled with the poloidal flux ψ, or by the radial coordinate $s = \sqrt{\psi/\psi_{\text{edge}}} = \sqrt{\tilde{\psi}}$. The geometrical radial coordinate will be noted ρ and the poloidal angle θ. The magnetic field is defined as $\vec{B} = F(\psi)\nabla\varphi + \nabla\psi \times \nabla\varphi$, where $F(\psi)$ is the poloidal current flux function. Two different kinds of magnetic equilibria are implemented: circular concentric magnetic surfaces, referred to as *ad hoc* equilibrium, and true MHD equilibria. For the latter case, ORB5 is coupled with the CHEASE code [31], which solves the Grad-Shafranov equation with a fixed plasma boundary. The poloidal coordinate used is either the geometrical angle θ or the straight-field-line coordinate θ_* defined by

$$\theta_* = \frac{1}{q(s)} \int_0^\theta \frac{\vec{B} \cdot \nabla\varphi}{\vec{B} \cdot \nabla\theta'} \mathrm{d}\theta', \tag{1}$$

where $q(s)$ is the safety factor. In this paper, χ represents a general poloidal coordinate. In ORB5, the options $\chi = \theta$ and $\chi = \theta_*$ have been implemented. All symbols with the subscript i will be used for ion quantities, while the subscript e will be used for electrons. It is assumed that the plasma contains electrons and an ion species with a mass m_i and a charge $q_i = eZ_i$. Four normalization quantities are used: q_i, m_i, B_0 and $\rho_s = c_s/\Omega_i$, where B_0 is the magnetic field at geometrical axis, $T_e(s_0)$ is the electron temperature in eV at a given reference magnetic surface s_0, and $c_s = \sqrt{eT_e(s_0)/m_i}$ is the ion sound speed.
The gyrokinetic equations solved in ORB5 are derived from Ref. [9] which describe the evolution of the plasma in an inhomogeneous static equilibrium magnetic field. Only the electrostatic component of the perturbation is considered, and magnetic perturbations are neglected.

2.2 Gyrokinetic Vlasov equation

The usual gyrokinetic ordering is assumed:

$$\begin{aligned} \frac{\omega}{\Omega_i} \sim \frac{k_\parallel}{k_\perp} \sim \frac{e\phi}{T_e} \sim \frac{\rho_{L,i}}{L_n} \sim \frac{\rho_{L,i}}{L_{T,i}} \sim \frac{\rho_{L,i}}{L_{T,e}} \sim \epsilon_g, \\ \frac{\rho_{L,i}}{L_B} \sim \epsilon_B, \end{aligned} \tag{2}$$

where ω is the characteristic fluctuation frequency, $\Omega_i = q_i B_0/m_i$ is the cyclotron frequency $k_\parallel$ and $k_\perp$ are the parallel and perpendicular components of the wave vector with respect to the magnetic field, ϕ is the fluctuating electrostatic potential, $\rho_{L,i}$ is the ion gyroradius, and L_n, $L_{T,i}$, $L_{T,e}$, L_B are the characteristic lengths associated with the density, the ion temperature, the electron temperature and the magnetic field profiles. ϵ_g and ϵ_B are small parameters, $\epsilon_g \sim \rho^*$, $\epsilon_B \sim \epsilon_a \epsilon_g$, with $\rho^* = \rho_s/a \ll 1$ and ϵ_a is the inverse aspect ratio. ORB5 solves the equations of motion in a 5D phase space $(\vec{R}, v_\parallel, \mu)$, where $\vec{R}$ is the position of the guiding center, $\mu = v_\perp^2/(2B)$ is the magnetic moment, with the conservation property $\mathrm{d}\mu/\mathrm{d}t = 0$, $v_\parallel$ and $v_\perp$ are the parallel and perpendicular components of the velocity with respect to the magnetic field. The renormalized potential $\Psi(\vec{R}, v_\parallel, \mu)$ in Ref. [9] has been approximated here by the gyroaveraged electrostatic potential

$$\Psi(\vec{R}, v_\parallel, \mu, t) \cong \langle\phi\rangle(\vec{R}, \mu, t) = \frac{1}{2\pi}\int_0^{2\pi} \phi(\vec{x}, t)\, \delta(\vec{R} + \vec{\rho}_{L,i} - \vec{x})\, \mathrm{d}\alpha, \tag{3}$$

thus neglecting a term of order $\mathcal{O}(\epsilon_g^2)$. Here α is the gyroangle. The gyroaveraged electric field is approximated by

$$\langle\vec{E}\rangle(\vec{R}, \mu, t) \cong -\langle\nabla_{\vec{x}}\phi(\vec{x})\rangle(\vec{R}, \mu, t) = -\frac{1}{2\pi}\int_0^{2\pi} \mathrm{d}\alpha \nabla_{\tilde{\mathrm{x}}}\, \phi(\tilde{\mathrm{x}}), \tag{4}$$

where a term of order $\mathcal{O}(\epsilon_g \epsilon_b)$ has been neglected. The polarization density is [9]:

$$\begin{aligned} n_{i,\mathrm{pol}} = \int \Bigg\{ \frac{q_i}{m_i B}\Big(\phi - \langle\phi\rangle\Big)\frac{\partial f}{\partial \mu} + \frac{q_i}{m_i \Omega_i^2}\nabla\Big[\int \mathrm{d}\alpha\Big(\phi - \langle\phi\rangle\Big)\Big] \times \frac{\vec{B}}{B}\cdot\nabla f \Bigg\} \\ \delta(\vec{R} + \vec{\rho}_{L,i} + \vec{x}) B_\parallel^*\, \mathrm{d}\vec{R}\, \mathrm{d}v_\parallel\, \mathrm{d}\mu\, \mathrm{d}\alpha, \end{aligned} \tag{5}$$

where $f = f(\vec{R}, v_\parallel, \mu)$ is the guiding center distribution function of the ion species, $B_\parallel^* = \vec{B}_\parallel^* \cdot \vec{B}/B$, and $\vec{B}_\parallel^* = \vec{B} + m_i v_\parallel/q_i \left(\nabla \times \vec{h}\right)$. The second term of Eq. (5) has been neglected, although it is of order $\mathcal{O}(\epsilon_g)$: ORB5 solves the linearized quasi-neutrality equation, so when f is replaced by the f_0 in the second term, it becomes smaller than the first one by a factor ϵ_g due to the density and temperature gradients appearing in ∇f_0. Sources, sinks and collisions are neglected, so the gyrokinetic Vlasov equation is:

$$\frac{\partial f}{\partial t} + \frac{\partial f}{\partial \vec{R}}\cdot\frac{\mathrm{d}\vec{R}}{\mathrm{d}t} + \frac{\partial f}{\partial v_\parallel}\frac{\mathrm{d}v_\parallel}{\mathrm{d}t} = 0. \tag{6}$$

The equations of motion are given by

$$\frac{\mathrm{d}\vec{R}}{\mathrm{d}t} = v_\parallel \vec{h} + \frac{1}{\Omega_i B_\parallel^*}\left(v_\parallel^2 + \frac{v_\perp^2}{2}\right)\left(\vec{h}\times\nabla B\right) - \frac{v_\parallel^2}{\Omega_i B_\parallel^*}\vec{h}\times\left[\vec{h}\times(\nabla\times\vec{B})\right] + \frac{\langle\vec{E}\rangle\times\vec{B}}{B_\parallel^* B}, \tag{7}$$

$$\frac{\mathrm{d}v_\parallel}{\mathrm{d}t} = \frac{1}{2}v_\perp^2\nabla\cdot\vec{h} + \frac{v_\perp^2 v_\parallel}{2\Omega_i B B_\parallel^*}\left\{\vec{h}\times\left[\vec{h}\times(\nabla\times\vec{B})\right]\right\}\cdot\nabla B + \langle\vec{E}\rangle\cdot\left\{\frac{q_i}{m_i}\vec{h} + \frac{v_\parallel}{B B_\parallel^*}\left(\vec{h}\times\nabla B\right) + \frac{v_\parallel}{B B_\parallel^*}\left(\vec{h}\times\frac{\mu_0\nabla p}{B}\right)\right\}, \tag{8}$$

where $\vec{h} = \vec{B}/B$, μ_0 is the permittivity of free space and $\nabla p = p'(\psi)\nabla\psi$ is the pressure gradient. The second term on the right hand side of Eq. (8) is of order $\mathcal{O}(\epsilon_B^2)$, does not influence neither energy conservation (see Sec. 2.5) nor the physical results and will be neglected in all the simulations. These equations of motion for a particle include the parallel motion, the drifts due to the curvature and the gradient of the magnetic field, the diamagnetic drift and a mirror term. For ad hoc equilibria, $\mu_0\nabla p/B$ is replaced by $-\vec{h}\times(\nabla\times\vec{B})$. Nonlinear terms are the $\vec{E}\times\vec{B}$ drift and the $v_\parallel$ nonlinearity, which is mainly an acceleration term due to the parallel electric field.

2.3 Background equilibrium and δf method

The δf method [27] is used, i.e. the distribution function f is split between a time independent part f_0 and a time dependent part δf:

$$f(\vec{R}, v_\parallel, \mu, t) = f_0(\vec{R}, v_\parallel, \mu) + \delta f(\vec{R}, v_\parallel, \mu, t). \tag{9}$$

A Maxwellian is chosen for f_0, i.e.

$$f_0(\vec{R}, v_\parallel, \mu) = n_0(\Upsilon)(2\pi)^{-3/2} v_{\mathrm{th},i}^{-3}(\Upsilon)\exp\left(-E/T_i(\Upsilon)\right), \tag{10}$$

where $E = 1/2 m_i v^2$ is the kinetic energy of a single ion, and is a constant of the unperturbed motion, and $v_{\mathrm{th},i} = \sqrt{eT_i/m_i}$ is the ion thermal velocity. Three different kinds of Maxwellian can be used, depending on the choice of the variable Υ as either $\psi, \psi_0, \hat{\psi}$, where ψ_0 is the canonical toroidal momentum, $\psi_0 = \psi + v_\parallel \Omega_i F(\psi)/B^2$ and $\hat{\psi}$ is a function of (E, μ, ψ_0) defined below. The axisymmetry of the tokamak implies that ψ_0 is a constant of the unperturbed motion, i.e. $\left.\frac{\mathrm{d}\psi_0}{\mathrm{d}t}\right|_0 = 0$. The quantity $\hat{\psi} = \psi_0 + \psi_{0,\mathrm{corr}}(E, \mu)$ is another constant

of motion, with

$$\psi_{0,\mathrm{corr}}(E,\mu) = -\mathrm{sign}(v_\parallel(t_0))\frac{q_i}{m_i}R_0\sqrt{2(E-\mu B_0)}\mathcal{H}(E-\mu B_0), \tag{11}$$

where $\mathcal{H}(x)$ is the Heavyside function, therefore $\psi_{0,\mathrm{corr}}$ is defined only for passing particles. $\hat{\psi}$ can be seen as the closest constant of motion to ψ. In the following, $f_0(\psi)$ is referred to as a *local* Maxwellian, $f_0(\psi_0)$ as a *canonical* Maxwellian and $f_0(\hat{\psi})$ as a *canonical* Mawellian *with correction*. The Vlasov equation is then $\frac{\mathrm{d}\delta f(\vec{R},v_\parallel,\mu)}{\mathrm{d}t} = -\frac{\mathrm{d}f_0(\Upsilon,v_\parallel,\mu)}{\mathrm{d}t}$ which can be written $\frac{\mathrm{d}\delta f(\vec{R},v_\parallel,\mu)}{\mathrm{d}t} = \tau(\vec{E})$, with

$$\tau(\vec{E}) = -f_0(\Upsilon,v_\parallel,\mu)\left(\kappa(\Upsilon)\frac{\mathrm{d}\Upsilon}{\mathrm{d}t}\bigg|_0 + \kappa(\Upsilon)\left.\frac{\mathrm{d}\Upsilon}{\mathrm{d}t}\right|_1 - \frac{q_i}{T(\Upsilon)}\langle\vec{E}\rangle\cdot\left.\frac{\mathrm{d}\vec{R}}{\mathrm{d}t}\right|_0\right), \tag{12}$$

where $\kappa(\Upsilon) = \mathrm{d}\ln f_0(\Upsilon,v_\parallel,\mu)/\mathrm{d}\Upsilon$. The subscript 0 refers to the unperturbed orbits and the subscript 1 refers to the terms that depend on the perturbed electric field. In most other gyrokinetic codes, with the exception of the GT3D code [16], a local Maxwellian is chosen, i.e $\Upsilon = \psi$, and the first term on the r.h.s. of Eq. (12) is neglected. A local Maxwellian can lead to spurious zonal flow oscillations [16], since it is not a true equilibrium distribution function, as $\left.\frac{\mathrm{d}f_0(\psi)}{\mathrm{d}t}\right|_0 \propto \left.\frac{\mathrm{d}\psi}{\mathrm{d}t}\right|_0 \neq 0$. When using a canonical Maxwellian, the quasi-neutrality equation is no longer satisfied as electron and ion equilibrium densities are different. In order to enforce quasi-neutrality, a radial electric field quickly develops and possibly suppresses instability. To eliminate this spurious field generation, the electron equilibrium density is further integrated from the ion distribution function after the particle loading and averaged over the poloidal angle,

$$n_{e0}(\psi) = \frac{1}{2\pi}\int_0^{2\pi}\left(\int f_0(\psi_0,\epsilon,\mu)\delta(\vec{R}+\vec{\rho}_{L,i}-\vec{x})B_\parallel^*\mathrm{d}\vec{R}\,\mathrm{d}v_\parallel\,\mathrm{d}\mu\,\mathrm{d}\alpha\right)\mathrm{d}\theta, \tag{13}$$

which minimizes the difference between n_{0e} and n_{0i}. Note that for small ρ^* plasmas, there is little difference between ψ and ψ_0 and the local Maxwellian becomes close to the canonical Maxwellian. Issues related to this choice are discussed in details in Ref. [32]. When $\Upsilon = \hat{\psi}$, the approximation $\left.\frac{\mathrm{d}\hat{\psi}}{\mathrm{d}t}\right|_1 \cong \left.\frac{\mathrm{d}\psi_0}{\mathrm{d}t}\right|_1$ is done.

As another alternative, since f is constant along the trajectories, δf can be simply obtained by:

$$\delta f(\vec{R},v_\parallel,\mu,t) = f\Big(\vec{R}(t_0),v_\parallel(t_0),\mu(t_0)\Big) - f_0\Big(\Upsilon(t),v_\parallel(t),\mu(t_0)\Big). \tag{14}$$

Details of this scheme, called *direct* δf, are given in Refs. [33] and [32]. ORB5 can be used with the standard or the direct δf scheme.

2.4 Quasi-neutrality equation

Using the usual quasi-neutrality constraint, the Poisson equation is $n_e = Z_i n_i$. The polarization density, Eq. (5), contained in $n_i(\vec{x},t)$ is simplified by using a long wavelength approximation $(k_\perp \rho_{L,i})^2 \ll 1$. The electron density $n_e(\vec{x},t)$ is evaluated by assuming adiabatic (or Boltzmann) electrons along the magnetic field lines. After linearization of the ion polarization density the quasi-neutrality equation becomes

$$\begin{aligned} n_{e0}(\psi) + \frac{e n_{e0}(\psi)}{T_e(\psi)} \Big(\phi(\vec{x},t) - \bar{\phi}(\psi,t)\Big) \\ = Z_i \langle n_{i0} \rangle (\vec{x}) + Z_i \nabla_\perp \cdot \left(\frac{n_{i0}(\psi)}{B\Omega_i} \nabla_\perp \phi(\vec{x},t) \right) + Z_i \delta n_i, \end{aligned} \tag{15}$$

where $\langle ... \rangle$ is the average over the gyroangle,

$$\begin{aligned} \langle n_{i0} \rangle (\vec{x}) = \int f_0 \Big(\Gamma\Big(\psi(\vec{x}), \epsilon(v_\parallel, \mu, \vec{x}), \mu\Big), \epsilon\Big(v_\parallel, \mu, \vec{x}\Big), \mu \Big) \\ \delta(\vec{R} + \vec{\rho}_{L,i} - \vec{x}) B_\parallel^* \mathrm{d}\vec{R}\, \mathrm{d}v_\parallel\, \mathrm{d}\mu\, \mathrm{d}\alpha, \end{aligned} \tag{16}$$

and

$$\delta n_i = \int B_\parallel^* \mathrm{d}\vec{R}\, \mathrm{d}v_\parallel\, \mathrm{d}\mu\, \mathrm{d}\alpha \delta f(\vec{R}, v_\parallel, \mu, t) \delta(\vec{R} + \vec{\rho}_{L,i} - \vec{x}). \tag{17}$$

At this point, it is assumed that $\langle n_{i0} \rangle (\vec{x}) = n_{i0}(\vec{x})$ and $n_{i0}(\vec{x}) = n_{e0}(\psi)/Z_i$ for a local Maxwellian only. Nevertheless, in this paper it is assumed that this relation holds for any type of equilibrium distribution function, so $n_{i0}(\vec{x}) = n_{e0}(\psi)/Z_i \equiv n_0(\psi)$. The final quasi-neutrality equation is therefore:

$$\frac{e Z_i n_0(\psi)}{T_e(\psi)} \Big(\phi(\vec{x},t) - \bar{\phi}(\psi,t)\Big) - \nabla_\perp \cdot \left(\frac{Z_i n_0(\vec{x})}{B\Omega_i} \nabla_\perp \phi(\vec{x},t) \right) = Z_i \delta n_i. \tag{18}$$

In Eq. (18), $\bar{\phi}$ is the flux-surface averaged potential. This term is nonzero only for axisymmetric perturbations: only the toroidal Fourier component $n = 0$ gives a contribution to this term. The $n = 0, m = 0$ mode is commonly called the *zonal flow*.

2.5 Particle and energy conservation

Despite all the approximations made in the previous sections, a particle number and an energy invariant can be derived (see Ref. [28]). The kinetic energy

of the plasma is

$$E_{\mathrm{kin}} = \int m_i \left(\mu B + \frac{v_\parallel^2}{2}\right) f B_\parallel^* \,\mathrm{d}\vec{R}\,\mathrm{d}v_\parallel\,\mathrm{d}\mu\,\mathrm{d}\alpha. \tag{19}$$

Its time derivative is

$$\frac{\mathrm{d}E_{\mathrm{kin}}}{\mathrm{d}t} = q_i \int \frac{\mathrm{d}\vec{R}}{\mathrm{d}t} \cdot \langle \vec{E} \rangle f B_\parallel^* \,\mathrm{d}\vec{R}\,\mathrm{d}v_\parallel\,\mathrm{d}\mu\,\mathrm{d}\alpha. \tag{20}$$

In this model, the electrostatic energy is defined as

$$E_{\mathrm{f}} = \frac{q_i}{2} \int \mathrm{d}\vec{x} \Big(\langle n_i \rangle (\vec{x}, t) - n_0(\vec{x}) \Big) \phi(\vec{x}, t). \tag{21}$$

The energy and particle number conservation are written:

$$\frac{\mathrm{d}E_{\mathrm{kin}}}{\mathrm{d}t} = -\frac{\mathrm{d}E_{\mathrm{f}}}{\mathrm{d}t}, \tag{22}$$

$$\frac{\mathrm{d}N_{\mathrm{ph}}}{\mathrm{d}t} = \frac{\mathrm{d}}{\mathrm{d}t} \int f B_\parallel^* \,\mathrm{d}\vec{R}\,\mathrm{d}v_\parallel\,\mathrm{d}\mu\,\mathrm{d}\alpha = 0. \tag{23}$$

3 Numerical implementation

This section describes the implementation of the gyrokinetic model exposed in Sec. 2. The general scheme of the code is shown in Fig. 1.

3.1 δf discretization: the Particle-In-Cell method

The Particle-In-Cell (PIC) method [34] is commonly employed in gyrokinetic simulations [8,16,17,28,35]. The perturbed distribution function is discretized in the 5D phase space along trajectories with N markers (also called tracers), carrying a weight w. In ORB5, δf is discretized as:

$$\delta f = \frac{N_{\rm ph}}{N} \sum_{p=1}^{N} \frac{1}{2\pi B_{\parallel}^{*}} w_p(t) \delta\Big(\vec{R} - \vec{R}_p(t)\Big) \delta\Big(v_{\parallel} - v_{\parallel p}(t)\Big) \delta\Big(\mu - \mu_p(t_0)\Big), \quad (24)$$

where $N_{\rm ph}$ is the number of physical particles. Each marker is characterized by its weight $w_p(t)$ and by its location $\Big(\vec{R}_p(t), v_{\parallel p}(t), \mu_p(t_0)\Big)$ in phase space. Integrating Eq. (24) over a small volume Ω_p on which δf is assumed to be constant, the temporal evolution of δf is obtained by solving the Vlasov equation for each weight $w_p(t)$:

$$\frac{{\rm d}w_p(t)}{{\rm d}t} = \frac{N}{N_{\rm ph}} \tau(\vec{E})_p \Omega_p, \quad (25)$$

where $\tau(\vec{E})$ is given by Eq. (12), $\Omega_p = B_{\parallel}^{*}\, {\rm d}\vec{R}\, {\rm d}v_{\parallel}\, {\rm d}\mu\, {\rm d}\alpha\, / {\rm d}N$ represents the volume of a marker in phase space centered around its location $(\vec{R}_p, v_{\parallel,p}, \mu_p)$ and ${\rm d}N$ is the number of markers in an infinitesimal volume of phase space:

$${\rm d}N = \frac{N}{N_{\rm ph}} f_L(\vec{R}, v_{\parallel}, v_{\perp}) J(\chi, s, \varphi)\, {\rm d}s\, {\rm d}\chi\, {\rm d}\varphi\, v_{\perp} {\rm d}v_{\parallel}\, {\rm d}v_{\perp}\, {\rm d}\alpha. \quad (26)$$

The loading of the markers is specified by the probability density function $f_L(\vec{R}, v_{\parallel}, v_{\perp})$. The conventional loading is such that markers have a probability function $p(s) = 1 - f_g + f_g \exp\Big[(s - s_{0,L})^2 / \Delta s_L^2\Big]$ in space, where $f_g \in [0:1]$, $s_{0,L}$ and Δs_L are input parameters. The markers are uniformly loaded in $(v_{\parallel}, v_{\perp})$, with a cut-off is applied at $(v_{\parallel}, v_{\perp}) = \kappa_v v_{{\rm th},i}(s)$, where κ_v is specified on input. It has been checked that all the simulations in this paper are converged for $\kappa_v = 5$. A disadvantage of this loading is that it is not specifically optimized for δf in velocity space. In the nonlinear phase, the size of the weights may vary significantly. During the charge assignment, noise will be produced where large weights (in absolute value) are present, which, due to its cumulative character, will alter the quality of the simulation. A way to overcome this difficulty is to apply an *optimized loading* [28], based on the statistical method

of importance sampling [36] of $|\delta f|$. The basic idea is to load markers with a probability function proportional to $|\delta f|$ at a time t_{NL} in the nonlinear phase. A first simulation is done with a uniform loading. Information about f_L is obtained by binning the markers at their initial position at a given time t_{NL} in the nonlinear phase. A second simulation is then performed using this new $f_L(\vec{R}, v_\parallel, v_\perp)$, which determines initial phase space coordinates $(\vec{R}_p, v_{\parallel p}, \mu_p)$ and initial phase space volume Ω_p of the markers. In the nonlinear phase of this restarted simulation, the statistical variance of the weights will be smaller, hence the numerical noise will have a lower level (see Eq. (1) of Ref. [38]. The reason is that the optimized loading avoids the appearance of large weights in the nonlinear phase. This technique has been successfully applied in cylindrical [28] and Tokamak [37] geometry.
Two different ways of initializing the weights (or equivalently the perturbation δf) have been implemented. The *white noise initialization* is defined by

$$\delta f_p(t_0) = A_p f_0\Big(\vec{R}_p(t_0), v_{\parallel p}(t_0), v_{\perp,p}(t_0)\Big)\Omega_p. \tag{27}$$

A_p is a pseudo-random number obtained with a Hammersley sequence, whose maximum amplitude is given on input. This scheme has the disadvantage that the initial perturbation is inversely proportional to the number of markers in the simulation. The simulation takes an increasing time, roughly proportional to the number of markers, until the physical modes emerge of the initial noise. Instead, the idea is to build a physical initialization, called *mode initialization*, independent of the number of markers:

$$\begin{aligned}\delta f_p(t_0) = \Omega_p \frac{A_0 f_0\Big(\vec{R}_p(t_0), v_{\parallel p}(t_0), \mu_p(t_0)\Big)}{(m_2 - m_1 + 1)(n_2 - n_1 + 1)} \left|\frac{T\Big(s_0\Big)}{\nabla T\Big(s_0\Big)}\right| \left|\frac{\nabla T\Big(s_p(t_0)\Big)}{T\Big(s_p(t_0)\Big)}\right| \\ \times \sum_{m=m_1}^{m_2} \sum_{n=n_1}^{n_2} \cos\Big(m\chi_p(t_0) - n\varphi_p(t_0)\Big).\end{aligned} \tag{28}$$

where $m_0, m_1, m_2, n_0, n_1, n_2$ are input parameters. This initialization will be stronger by choosing $m_0 = [-nq(s_0)]$, where [...] denotes the integer part, as ITG modes are aligned with the field lines. It has the advantage that the initial perturbation converges with the number of markers and the initial phase of the simulation, until the physical modes develop, is independent of the number of markers.

3.2 Equations of motion

Tracers can be pushed either in cylindrical coordinates (r, φ, z) or in magnetic coordinates (θ_*, s, φ). In order to avoid the singularity at the magnetic axis that would appear in $\mathrm{d}\theta_*/\mathrm{d}t$, it is more adequate to use $(\xi, \eta, \varphi) =$

$(s\cos\theta_*, s\sin\theta_*, \varphi)$. Among the numerous advantages of using magnetic coordinates, this set of coordinates avoids numerical interpolations during the construction of the perturbed density. Depending on the computer architecture, the simulation can speed up by 40%. For both ad hoc and MHD equilibria, equilibrium coefficients needed for the pushing are obtained with linear interpolations from a (r, z) or a (s, θ_*) grid.
For the integration of Eqs. (7), (8) and (25), a Runge Kutta integrator of order four is implemented in ORB5.
When a particle leaves the plasma, or equivalently when $s > 1$, it is reflected: $\chi \to -\chi$. This scheme may lead to a small violation of unperturbed conserved quantities for equilibria that are not up-down symmetric.
The particle pushing is speeded up (up to 40 %) by using a cache sorting algorithm [39]: particles are sorted in the poloidal plane every n_{CS} time steps, where n_{CS} is given on input.

3.3 Discretisation of quasi-neutrality equation

The quasi-neutrality Eq. (18) is solved with linear, quadratic or cubic B-splines finite elements [40]. The perturbed potential is discretized as:

$$\phi(\vec{x}, t) = \sum_{\mu} \phi_\mu(t)\Lambda_\mu(\vec{x}), \tag{29}$$

where μ stands for (j, k, l), $\{\phi_\mu(t)\}$ are coefficients and $\{\Lambda_\mu(\vec{x})\}$ are tensor products of $1D$ B-splines of order p, $\Lambda_\mu(\vec{x}) = \Lambda_j^p(s)\Lambda_k^p(\chi)\Lambda_l^p(\varphi)$. The B-splines are defined on a (N_s, N_χ, N_φ) grid. To get a linear system for $\{\phi_\mu(t)\}$, the Galerkin method [41] is used. The system is then:

$$\sum_{\mu} A_{\mu\nu}\phi_\mu(t) = b_\nu(t), \tag{30}$$

with:

$$\begin{aligned} A_{\mu\nu} &= \int \mathrm{d}\vec{x} \frac{n_0(\psi)}{Z_i T_e(\psi)} \Big(\Lambda_\mu(\vec{x})\Lambda_\nu(\vec{x}) - \bar{\Lambda}_\mu(s)\bar{\Lambda}_\nu(s)\Big) + \frac{n_0(\psi)}{B\Omega_i}\nabla_\perp\Lambda_\mu(\vec{x})\cdot\nabla_\perp\Lambda_\nu(\vec{x}), \\ b_\nu(t) &= \frac{N_{ph}}{N}\sum_{p=1}^{N}\frac{w_p(t)}{2\pi}\int_0^{2\pi}\mathrm{d}\alpha\Lambda_\nu\Big(\vec{R}_p + \vec{\rho}_{L,i,p}(\alpha)\Big). \end{aligned} \tag{31}$$

Note that the matrix is real, symmetric and positive definite. The building of $b_\nu(\vec{x}, t)$ is called the *charge assignment*. This projection onto the finite element basis is the main source of numerical noise, due to the particle discretization of δf. It becomes obvious that higher order splines are beneficial for the noise reduction, as they have a more extended shape. The perpendicular gradients are approximated to lie in the poloidal plane, $\nabla_\perp \cong \nabla_{\mathrm{pol}} = \nabla s\frac{\partial}{\partial s} + \nabla\chi\frac{\partial}{\partial\chi}$.

Due to the axisymmetry in the toroidal direction, the system can furthermore be decoupled in φ. Applying a discrete Fourier transform on ϕ_μ and b_ν, and inserting these relations in Eq. (30) yields, in Fourier space:

$$\sum_\mu A_{\mu\nu}\hat{\phi}_\mu^{(n)}(t) = \frac{\hat{b}_\nu^{(n)}(t)}{M^{(n),p}}, \tag{32}$$

where (μ,ν) now stand for $2D$ indices, $\hat{\phi}_\mu^{(n)}$, $\hat{b}_\nu^{(n)}(t)$ are Fourier coefficients of ϕ_μ and b_ν, and $M^{(n),p}$ is defined by:

$$\sum_{l'=1}^{N_\varphi} \int \mathrm{d}\varphi \Lambda_{l'}^p(\varphi)\Lambda_l^p(\varphi) \exp\left(\frac{2\pi i}{N_\varphi} n(l'-l)\right) = M^{(n),p}. \tag{33}$$

$M^{(n),p}$ can be computed analytically for any spline of order p (see Ref. [13]). The matrix $A_{\mu\nu}$ and the right hand side b_ν defined in Eq. (31) are modified with the Dirichlet boundary conditions $\phi(s=1,\chi,\varphi,t)=0$ and the regularity condition $\phi(s=0,\chi,\varphi,t) = \phi(s=0,\chi=0,\varphi,t)$. The matrix defined in Eq. (31) can be decomposed as $A_{\mu\nu}(\vec{x}) = A_{\mu\nu}^{\mathrm{noZF}} + A_{\mu\nu}^{\mathrm{ZF}}$, where $A_{\mu\nu}^{\mathrm{ZF}}$ contains the flux surface averaged terms only and $A_{\mu\nu}^{\mathrm{noZF}}$ contains all the other terms. The linear system associated with $A_{\mu\nu}^{\mathrm{noZF}}$ can be solved with iterative solvers, such as PETSc [42] or WSMP [43], or with a direct solver using LAPACK [44] routines. Although faster than iterative solvers, a direct solver requires more memory. The memory needed to store the matrix becomes rapidly prohibitive as the grid number is increased. To reduce the memory storage, a parallel solver using ScaLAPACK [45] routines has been implemented in ORB5, giving a reduction factor up to $N_s/8$ for the memory of the matrix. Iterative solvers are not adequate for $A_{\mu\nu}^{\mathrm{ZF}}$ because its band is full, therefore the associated linear system is only solved with direct solvers.

3.4 Gyro-averaging

The integral over the gyration angle for both electric field and perturbed density is computed with a $N_g = \min\left(32, \max(4, 4\rho_{L,i,p}/\rho_{L,i})\right)$ discrete sum, where $\rho_{L,i,p} = v_{\perp p}/\Omega_{i,p}$ is the marker gyroradius and $\Omega_{i,p}$ is the cyclotron frequency of the marker. Although it has been shown that a 4-points discretization is sufficient to describe perturbations up to $k_\perp \rho_{L,i} \sim 1$ [35], a gyro-adaptative method reduces noise since it acts like a Bessel filter [28]. The position (r_α, z_α) of the marker on the Larmor ring is assumed to be in the poloidal plane and is simply:

$$r_\alpha = r_{GC} + \rho_{L,i,p}\cos\alpha \tag{34}$$
$$z_\alpha = z_{GC} + \rho_{L,i,p}\sin\alpha \tag{35}$$

When orbits are evolved in magnetic coordinates, the position of a marker is obtained by a Taylor expansion:

$$\xi_\alpha \cong \xi_{GC} + \vec{\rho}_{L,i,p} \cdot \nabla\xi|_{GC}\,, \tag{36}$$
$$\eta_\alpha \cong \eta_{GC} + \vec{\rho}_{L,i,p} \cdot \nabla\eta|_{GC}\,, \tag{37}$$

where:

$$\vec{\rho}_{L,i,p} = \rho_{L,i,p}\left(\frac{\nabla s}{|\nabla s|}\cos\alpha + \frac{\vec{B}\times\nabla s}{|\vec{B}\times\nabla s|}\sin\alpha\right). \tag{38}$$

This implementation has been successfully applied in [46] for the linear gyrokinetic code LORB5 [47].
Using Eqs. (4) and (29), the gyro-averaged electric field is given by

$$\langle\vec{E}\rangle = -\frac{1}{2\pi}\sum_\mu \phi_\mu(t)\int \mathrm{d}\alpha\nabla\Lambda_\mu\Big(\vec{R} + \vec{\rho}_{L,i}(\alpha)\Big). \tag{39}$$

To achieve the best energy conservation, it is important to employ exactly the same gyro-averaging procedure for the charge assignment and the electric field [28].

3.5 Fourier filter

In order to reduce the noise, we apply a Fourier filter on the discretized perturbed density:

$$Fb_{i,j,k} = \sum_{m,n} f_{i,m,n}\widehat{b}_{i,m,n}(t)\ e^{im\chi_j}\ e^{in\varphi_k}. \tag{40}$$

A Fourier filter is of great interest since physically non relevant modes can be removed from the simulation. The poloidal Fourier filter should be chosen wide enough so as to include all modes relevant to the physics problem at hand, although a formal proof of energy conservation including Fourier filters cannot be done in toroidal geometry. One of the simplest filter, called *square* filter, is obtained by suppressing modes outside a window $[n_{\min} : n_{\max}]\times[m_{\min} : m_{\max}]$ specified on input. However, this kind of filter is inconsistent with the gyrokinetic ordering, as it retains modes with $k_\parallel/k_\perp$ much bigger than ρ^* [49]. Indeed, ITG modes tend to align with the field lines. Mathematically, it means that these modes are such that $k_\parallel = (q(s)r)^{-1}\Big(m + nq(s)\Big) \to \rho^*$, consistent with the gyrokinetic ordering Eq. (2). So m should be near $-nq(s)$. Toroidal-ITG modes are composed of a superposition of a few m components around $m = -nq(s)$, resulting in $k_\parallel \approx (q(s)r)^{-1}$. In both cases, it is enough to keep a narrow window of modes in order to describe all physically relevant modes.

The idea here is to define a surface-dependent filter that suppresses high $k_\parallel$ modes. For each $n \in [n_{\min} : n_{\max}]$, the poloidal modes $[-nq(s) \pm \Delta m] \bigcap [m_{\min} : m_{\max}]$, where Δm is an input parameter, are retained. It is very useful to use θ_* as the poloidal coordinate: the poloidal width of the spectrum for a toroidal mode is narrower with the straight-field-line coordinate than with the geometrical poloidal angle θ. In addition, when θ differs significantly from θ_* (in case of low aspect ratio, non concentric or shaped magnetic surfaces), the mode is generally not peaked around $m = nq(s)$. Finally, the width of a toroidal mode spectrum increases with the plasma size when θ is used, whereas it is in principle independent of plasma size with θ_*. The beneficial influence of a small Δm will be explained in Sec. 5.

3.6 Code Parallelization

Gyrokinetic simulations are extremely CPU time and memory demanding. Therefore, ORB5 is massively parallelized with Message Passing Interface (MPI) [48] routines. A decomposition concept called *domain cloning* [50,51] is applied: $P = P_C P_\varphi$, where P is the total number of processors, P_φ is the number of sub-domains in the φ direction and each of these sub-domain is cloned P_C times. After each charge assignment, the perturbed density is summed over the clones. After each step in the particle pushing, the markers have moved also in the φ direction and are therefore communicated if necessary to their appropriate φ sub-domain. This parallelization scheme offers great flexibility because it can be tuned to different types of parallel platform architectures. The performance of this decomposition will be shown at Sec. 4.

3.7 Relations between physical and numerical parameters and $\rho*$ scaling

Depending on the physical case under study, the ORB5 numerical parameters should be set up as follows. The radial mesh resolution depends on the maximum $k_s \rho_{L,i}$ that should be resolved, where k_s is the radial component of the wave number. With cubic B-spline finite elements there should be at least 3 points per wavelength. Thus

$$N_s > \frac{3}{2\pi} \left(k_s \rho_{L,i}\right)_{\max} \frac{a}{\rho_{L,i}} \,. \tag{41}$$

Similarly, for a maximum $k_\chi \rho_{L,i}$ that should be resolved up to the magnetic surface $s_{\max}$, where k_χ is the poloidal component of the wave number, the poloidal mesh should be set to

$$N_\chi > 3 s_{\max} \left(k_\chi \rho_{L,i}\right)_{\max} \frac{a}{\rho_{L,i}} \,. \tag{42}$$

Since the perturbations tend to be aligned with magnetic field lines the toroidal mesh should be chosen as

$$N_\varphi \approx N_\chi / q(s) \ . \tag{43}$$

The time step should satisfy the following three requirements:

$$\omega_\parallel \Delta t \ll 1, \ \omega_* \Delta t \ll 1, \ \omega_{\mathrm{ExB}} \Delta t \ll 1, \tag{44}$$

with $\omega_\parallel = k_\parallel v_\parallel$, $k_\parallel = (q(s)r)^{-1}\Big(m + nq(s)\Big)$, $\omega_* = v_{\mathrm{th},i}(d \ln T / d\rho) k_\chi \rho_{L,i}$ and $\omega_{\mathrm{ExB}} = k_\perp v_{\mathrm{ExB}}$. There is a limiting value of $k_\parallel$ beyond which the ITG modes are strongly Landau damped. An estimate can be obtained from a dispersion relation [52], and taking the limit $\eta_i \to \infty$:

$$|k_\parallel^{\mathrm{lim}}| r = (1/2)(r/L_{T,i}) k_\chi \rho_{Li} \ . \tag{45}$$

The field-aligned Fourier filter of width $\pm \Delta m$ gives $|k_\parallel|_{\max} = \Delta m / \Big(rq(s)\Big)$. In order to resolve the physically relevant modes up to $|k_\parallel|^{\mathrm{lim}}$ the width of the field-aligned Fourier filter should be

$$|\Delta m| > \frac{q(s)}{2} \frac{r}{L_{T,i}} k_\chi \rho_{Li} \ . \tag{46}$$

The r.h.s of Eq. (46) is typically between 5 and 10. It is worth mentioning that it does not scale with $a/\rho_{L,i}$, since the typical values of $k_\chi \rho_{L,i}$ present in ITG turbulence do not depend on $a/\rho_{L,i}$. When quasi-neutrality equation is solved with the geometrical angle θ instead of the straight-field-line coordinate θ_*, Eq. (46) is no longer a good estimate, because the relation $|k_\parallel|_{\max} = \Delta m / \Big(rq(s)\Big)$ does not hold. Note also that Δm should be large enough so that the filtered perturbed density does not have discontinuities across magnetic surfaces. Mathematically, it is expressed as $\Delta m > nq'(s)/N_s$. Using Eqs. (41) and (42) with $(k_s \rho_{L,i})_{\max} = (k_\chi \rho_{L,i})_{\max} = s_{\max} = 1$, the last condition can be written:

$$\Delta m \gtrsim 2\hat{s}, \tag{47}$$

where $\hat{s}$ is the magnetic shear. The first criterion of Eq. (44) gives

$$\Omega_i \Delta t \ll \frac{r}{a} \frac{a}{\rho_{L,i}} \frac{q(s)}{|\Delta m|}, \tag{48}$$

or, if $|\Delta m|$ is chosen as in Eq.(46),

$$\Omega_i \Delta t \ll \frac{2(a/\rho_{L,i})(r/a)}{(r/L_{T,i}) k_\chi \rho_{L,i}} \ . \tag{49}$$

The second criterion gives

$$\Omega_i \Delta t \ll \frac{(a/\rho_{L,i})(r/a)}{(r/L_{T,i})k_\chi \rho_{L,i}} \ . \tag{50}$$

The third criterion can be expressed in terms of the Mach number $M = v_{\mathrm{ExB}}/v_{\mathrm{th},i}$ as

$$\Omega_i \Delta t \ll \frac{1}{M\, k_\perp \rho_{L,i}} \ . \tag{51}$$

For most cases of interest for ITG driven turbulence M is found to be of the order of 10^{-2} and $k_\perp \rho_{L,i}$ is at maximum of order unity. Using local estimates for the time step, it appears that $\omega_\parallel$ is the fastest frequency of the system with a square filter, independently of the plasma size. With a field-aligned filter, $k_\parallel$ scales with ρ^* and so for large plasmas $\omega_{E\times B}$ becomes the fastest frequency. The time step can be increased by a factor 10, hence the simulations are strongly shortened. More details can be found in Ref. [49].

3.8 Density and temperature profiles

A given profile $A = T_i,\ T_e$ or n_0 is defined by:

$$\frac{1}{A}\frac{\mathrm{d}A}{\mathrm{d}\tilde{\psi}} = -\frac{a}{L_A}\frac{\left[1-\cosh^{-2}\left(\frac{s_0}{\Delta_A}\right)\right]}{\left[\cosh^{-2}\left(\frac{s-s_0}{\Delta_A}\right)-\cosh^{-2}\left(\frac{s_0}{\Delta_A}\right)\right]} \ . \tag{52}$$

$a/L_{T,i}$, $a/L_{T,e}$, a/L_n, $\Delta_{T,i}$, $\Delta_{T,e}$, Δ_n and s_0 are input parameters. All the gradients are peaked at s_0, therefore ITG modes should develop around that magnetic surface. Temperature profiles are normalized by their value at s_0. The density is normalized with the volume averaged density. Fig. 2 shows a typical profile used in ORB5.

4 Scalability

To measure the scalability, short simulations were performed at fixed grid size, varying the number of processors from 512 to 4096 and keeping the number of clones fixed. For each number of processors, several simulations with different number of markers per processor have been performed. Since the size of the matrix is fixed, the time dedicated to the solver (backsolve and Fourier transforms) should be proportional to the number of toroidal slices per processor. The time dedicated for the particles (pushing + charge assignment) is assumed to scale linearly with the number of markers per processor. Finally, the communication time is mainly due to the φ-partition of the markers. It is hard to give a simple estimate for the communication time, since it depends on the time step, the cluster architecture, and the ion temperature. However it represents a small fraction (maximum 10 %) of the total simulation time. Therefore, the time per iteration is modelled as:

$$t_{\mathrm{it}} = \underbrace{K_s \frac{N_\varphi P_c}{P}}_{\text{solver}} + \underbrace{K_m \frac{N}{P}}_{\text{markers}} + \underbrace{t_{\mathrm{comm}}}_{\text{communication}} , \tag{53}$$

where K_s and K_m are constants to be determined. Fig. 3 displays the time per iteration t_{it} as a function of the number of markers per processor N/P. The dependence is linear, as expected. The slope of these lines, namely K_m seems to slowly increase with P, which illustrates a slight deviation from an ideal scaling. In fact, a super-scaling can be observed from Fig. 3: for a fixed number of markers per processor, the time per iteration decreases when the number of processors increases. This is because P is increased by keeping P_C fixed and so according to Eq. (53) the time per iteration dedicated to the solver decreases. Fig. 4 shows the speedup of the code relative to $P = 1024$ for a fixed number of markers, which is very close to an ideal scaling. K_s and K_m have been obtained by a linear fit of t_{it} as a function of N/P at fixed $P = 512$. Therefore, the measured time per iteration can be further compared to the fitted time t_{it}^f. The maximum value of the relative degradation due to the increase of the number of processors is only 12 %. Effects of t_{comm} can be included by fitting $t_{\mathrm{it}} - t_{\mathrm{comm}}$ instead of t_{it}. In that case, the maximum relative degradation falls down to 10 %. Globally, these results show the excellent scaling properties of ORB5 up to 4096 processors.

5 Simulation results

5.1 Field-aligned filter

The field-aligned filter has been tested with the following input parameters: $m_i = 1, a = 40\rho^*, B_0 = 1\ [T], R_0 = 1\ [m], R_0/a = 5$. The density profile is flat, $T_i = T_e$, $R_0/L_{T,i} = 12$, $\Delta_{T,i} = 0.208$, $s_0 = 0.5$ and $\Delta t = 40\ \Omega_i^{-1}$. If not specified, $N = 2^{24} \cong 16$ M markers. The initial distribution f_0 is a canonical maxwellian with correction, (see Eq. (11)). δf is obtained through the conventional δf scheme and is initialized with white noise. The quasi-neutrality equation is solved with cubic B-splines on a $N_s = 128, N_{\theta_*} = 128, N_\varphi = 64$ grid with the parallel direct solver. The main parameter of the field-aligned filter is the width Δm. By keeping only small $k_\parallel$ modes, the number of Fourier modes kept in the filter is strongly reduced. A square filter has $N_s(2m_{\max}+1)(2n_{\max}+1)$ modes $\propto (\rho^*)^{-3}$ whereas a field-aligned filter has approximatively $N_s(2\Delta m+1)(2n_{\max}+1)$ modes $\propto (\rho^*)^{-2}$. The latter reduces the number of modes by about $(\rho^*)^{-1}$. In Ref. [53], it is shown that numerical noise, due to the projection of the charge density onto a finite number of markers, mainly depends on the square root of the number of markers per Fourier modes. In that sense, a field-aligned filter should improve the quality of a simulation without affecting CPU time, in contrary to an increase of the numbers of tracers.
Fig. 5 displays the volume-averaged radial heat flux Q, defined by

$$Q = \frac{1}{V}\sum_{p=1}^{N} w_p \frac{1}{2} m v_p^2 \left. \frac{\langle \vec{E} \rangle \times \vec{B}}{B_\parallel^* B} \cdot \frac{\nabla\psi}{|\nabla\psi|} \right|_{\vec{R}_p, v_{\parallel p}, \mu_p}, \tag{54}$$

where V is the volume of the torus, for field-aligned and square simulations. The simulation starts with the linear phase, in which the perturbation grows exponentially. Then nonlinear effects become important and saturation occurs. Since the gyrokinetic model applied in ORB5 contains no sources and no dissipation term, profiles are free to evolve. Therefore all the simulations presented in this paper exhibit profile relaxation. A transient phase consisting of avalanches and bursts occurs, while the turbulence further decays and finally the system evolves towards a quasi-steady state in which the gradients are totally relaxed and there is no perpendicular transport. In Fig. 5, for the simulation with the square filter, a numerical heating develops in the late nonlinear phase. This kind of phenomenon is typical when too few markers are used. A good indicator of the quality of a simulation is the energy conservation described in Sec. 2.5. Fig. 6 shows the relative energy conservation $\Delta E/E(t_0)$, where $\Delta E = E_{\rm kin}(t) + E_f(t) - E_{\rm kin}(t_0) - E_f(t_0)$ for simulations with square and field-aligned filters. In a noise-free simulation, or equivalently in the limit of an infinite number of markers and an infinitely small time step,

$\Delta E/E_f$ should be zero. By looking simultaneously at Figs. 5 and 6, one sees that the energy conservation for the square filter simulation starts to degrade after the top of the overshoot, when nonlinear effects become important. During the burst phase (from $t = 0.8 \cdot 10^4\ \Omega_i^{-1}$ to $t = 3 \cdot 10^4\ \Omega_i^{-1}$), the energy deviation increases and reaches 1% of the total plasma energy, which is an unacceptable value. From $t = 3 \cdot 10^4\ \Omega_i^{-1}$ to the end of the simulation, the system should reach a quasi-steady state but both the heat flux and the energy deviation continue to grow: the simulation would explode if it was continued. On the other hand, a gain of 2 orders of magnitude in the relative energy conservation is obtained as the filter goes from square to field-aligned: in this case, the relative energy conservation does not exceed 10^{-4}. However, the energy deviation is of the order of the field energy of the system so a good simulation in terms of energy conservation should have $|\Delta E(t)| < \epsilon_E E_f(t)$ in the nonlinear phase, with $\epsilon_E \ll 1$. Simulations presented in this section have a large ρ^* which implies a very fast relaxation. It is indeed difficult to obtain good energy conservation for these parameters. Simulations of realistic sized plasmas will be shown at Sec. 5.4.

The origin of the unphysical behaviour observed on Fig. 5 can be observed on Fig. 7: the energy of all toroidal modes is growing in time. However, for the field-aligned filter, the energy of $n \neq 0$ modes is constant for late times (see Fig. 8). On Fig. 9, the radially averaged energy spectrum of the $n = 6$ mode in the poloidal space, normalized to its maximum m component, is shown at different times for a square filter simulation. In the linear phase, the mode is peaked near $m = -12 = -nq(s_0)$ as $n = 6$ and $q(s_0) = 2$. After the saturation, the peak energy becomes smaller but the spectrum broadens: it is no longer peaked and m components which are far away from $nq(s_0)$ contain a very significant part of the total toroidal mode energy. This causes the field energy of all $n \neq 0$ modes to grow in time: it is a clear evidence that numerical noise is created because of high $k_\parallel$ modes. The bad quality of the square filter can also be observed in Fig. 10 (left), which shows the electric potential on a magnetic surface. The resulting structure is clearly a superposition of high $k_\parallel$ modes, whereas the field-aligned (right) filter naturally preserves the field-aligned structure of ITG modes.

These results show that a field-aligned filter instead of a square filter must be employed in ORB5, but it is still crucial to show how to fix the width Δm of the field-aligned filter, and more important to show that this filter will contain all the relevant modes of the system. As the field-aligned filter acts locally on a magnetic surface, a look on the local poloidal energy spectrum is necessary. It must be checked that no relevant poloidal harmonics is locally removed from the simulation: a proper determination of Δm is a crucial step. A too small value of Δm will obviously cut some relevant physics, whereas a too large Δm quickly introduces additional numerical noise. A necessary but not sufficient condition to fix Δm is to converge the growth rate of toroidal modes in the linear phase. Fig. 11 shows the evolution of the the field energy of the mode $n = 6$ for different values of Δm and a square filter simulation.

Small Δm cases yield lower growth rates and linear convergence is reached with $\Delta m = 5$. A careful look at the local energy spectrum during the non-linear phase is still needed, in order to account for a possible broadening of the spectrum. The problem is that for late times, numerical noise also grows inside the field-aligned filter and it becomes hard to separate the physical signal from the noise. The idea to overcome this difficulty is to identify the noisy components of the local energy spectrum by choosing a larger Δm and by varying the number of markers. Fig. 12 shows the local energy spectrum, averaged between $t = 1.5 \cdot 10^4\ \Omega_i^{-1}$ and $t = 2.4 \cdot 10^4\ \Omega_i^{-1}$, of the mode $n = 6$ for 2 simulations with $\Delta m = 15$ but with 16 M and 67 M markers. In both cases, 2 noise plateaux are visible near the edges of the filter, where high $k_\parallel$ modes are present. The plateau corresponding to the 67 M markers case is lower but the central parts of both spectrum are very similar. Therefore, from Fig. 12 one can conclude that $\Delta m = 9$ is a reasonable value. Fig. 13 shows that for a sufficiently high number of markers a field-aligned filter with $\Delta m = 9$ will contain all the locally relevant poloidal modes. The same conclusion can be drawn from Fig. 14 for the global energy spectrum. Note that it has also been checked that changing the averaging times of the spectrum does not modify this conclusion. Finally, Fig. 15 shows the radial heat flux as a function of Δm for simulations with constant noise, in the sense that the number of markers per Fourier mode is kept constant [38]. One sees that $\Delta m = 9$ is a reasonable value and that using a large Δm does not modify the final state of the system but requires a higher number of markers.
To summarize, smart Fourier filtering is a powerful numerical scheme to improve the quality of a PIC simulation: by relaxing the time step criterion and by decreasing the number of Fourier modes in the simulations, CPU time is reduced by 2 orders of magnitude. In addition, the field-aligned filter should even be more efficient in the limit of small ρ^* plasmas relative to the square filter. Its width Δm is set through the procedure applicable to any set of physical parameters described above.

5.2 Convergence with number of markers

In collisionless gyrokinetic simulations, convergence is a subtle notion. For PIC simulations, not only the time step and the grid resolution need to be carefully chosen, but the number of markers plays a crucial role as well. Indeed, numerical noise inherent to the PIC method may determine the level of transport in ETG simulations [54]. In ITG PIC simulations, due to the strong influence of the zonal flow, the situation is different, however the question of the required number of markers for convergence still remains. By measuring the level of numerical noise in ETG simulations, it has been established that the number of markers required is linked to the number of Fourier modes in the simulation [53]. A study of numerical noise in ITG simulations is beyond the

scope of this paper. Here, the question of convergence is approached by means of physics diagnostics. Fig. 16 shows the evolution of the volume averaged heat flux for different numbers of markers and the white noise initialization. The overshoot is shifted in time, as the initial level of the perturbation is inversely proportional to $\sqrt{N}$. The sole Fig. 16 is not sufficient to say if a reasonable convergence is reached or not. This difficulty can be overcome by using the mode initialization with a single mode (m_0, n_0) (see Sec. 3.1). Fig. 17 shows again Q for a marker scan performed with the mode initialization. As the number of tracers is increased, the different curves look more and more alike during the whole simulation and the top of the overshoot occurs always at the same time and converges to the same level. Fig. 18 shows the volume averaged heat flux Q, time-averaged between $t = 4 \cdot 10^4$ and $t = 5 \cdot 10^4$ Ω_i^{-1}, i.e. when the system has reached a quasi steady-state, for the 2 different initializations. For the white noise initialization, it is harder to say for which number of markers the simulation is converged. This is because the system relaxes differently when N is increased, which implies different series of bursts. When the mode initialization is used, the convergence curve is smoother because the different simulations have in this case a more similar nature, and the error bars representing the standard deviation over the averaging times tend to decrease as N is increased, which is not the case for the white noise initialization.
In addition, note that these 2 initializations give quite different overshoots and therefore quite different profile relaxations: for the mode initialization, a single toroidal mode strongly dominates since the beginning of the simulation, whereas all toroidal modes have approximatively the same initial energy when the white noise initialization is employed. In this context, a multiple mode initialization would be more appropriate, but in principle the quasi-stationary state of a simulation should not depend on the initialization. The important point is that by initializing the system independently of the number of markers, the notion of convergence with respect to the number of markers becomes easier to handle.

5.3 The Rosenbluth-Hinton test

The test consists in comparing the numerical calculation of the time evolution of the axisymmetric potential with the analytical result, valid for circular magnetic surfaces in the limit of large aspect ratio, obtained by Hinton and Rosenbluth [55]. The gyrokinetic Vlasov equation for the zonal flow component, i.e. $n = m = 0$, is analytically solved and the axisymmetric component of an electrostatic perturbation is found to linearly damp and a residual flow level is found. Therefore, the $E \times B$ velocity ($v_{\vec{E}\times\vec{B}}$, normalized to the initial value $v_{\vec{E}\times\vec{B}}(0)$), generated by a pure axisymmetric density perturbation, is

expected to behave as:

$$\frac{v_{\vec{E}\times\vec{B}}}{v_{\vec{E}\times\vec{B}}(0)} = (1 - A_{R_0})\, e^{-\gamma_G t} \cos(\omega_G t) + A_{R_0} \tag{55}$$

where A_{R_0} is the residual, at the radial position ρ:

$$A_{R_0} = \frac{1}{\left(1 + 1.6 q(s)^2 \left(\rho/R_0\right)^{-1/2}\right)} \tag{56}$$

and ω_G, γ_G are respectively the frequency and the decaying rate of the velocity perturbation [56]:

$$\omega_G = \left(1 + \frac{1}{2q(s)^2}\right)^{1/2} \frac{\sqrt{2} v_{\mathrm{th},i}}{R_0} \tag{57}$$

$$\gamma_G = \omega_G \exp\Big(-q(s)^2 - \frac{1}{2}\Big) \tag{58}$$

An underestimate or overestimate of the residual zonal flow would lead to an incorrect prediction of the radial transport. Therefore a numerical code must provide the correct residual A_{R_0} in order to produce reliable physical results. The zonal flow damping test has been performed in linear mode, which means that nonlinear terms in particle trajectory equations have been suppressed. In order to reproduce the results of Hinton and Rosenbluth, we solve only for a radial electrostatic potential, by initializing a radial perturbed density $\delta n_i(s,\theta,\phi) = \delta n_i(s)$ and by filtering density such that only modes $n = 0, m \in [-5:5]$ are kept. GAMS will be created through the poloidal dependence of the $\vec{E}\times\vec{B}$ velocity and, in addition, through the poloidal coupling generated by the magnetic field in the quasi-neutrality equation. Input parameters are $m_i = 1, a = 40\rho^*, B_0 = 1\ [T], R_0 = 5\ [m], R_0/a = 10$. The density and temperature profiles are flat, $T_i = T_e$. The safety profile is monotonic and two different magnetic surfaces have been studied, $s_0 = 0.5, q(s_0) = 1.15$ and $s_0 = 0.7, q(s_0) = 1.33$. The grid is $N_s = 64$, $N_\chi = 64, N_\varphi = 64$, the number of markers is $N = 16$ M and the time step is $\Delta t = 50\ \Omega_i^{-1}$. The initial condition has been prepared in order to obtain an axisymmetric ion density perturbation. The results must not depend on the initial conditions, therefore two different perturbations of the ion density have been tested, $\delta n_i(s) \cong \delta n_{i,0} \cos(\pi s)$ and $\delta n_i(s) \cong \delta n_{i,0} \sin(\pi s)$, where $\delta n_{i,0}$ is chosen so that $\left\langle v_{\vec{E}\times\vec{B}} \right\rangle_s (t=0) = 0.07\ v_{\mathrm{th},i}$. The results of the ORB5 simulations are plotted in Figs. 19 and 20. In these figures, the $E\times B$ velocity normalized at the initial value $v_{\vec{E}\times\vec{B}}/v_{\vec{E}\times\vec{B}}(0)$ is plotted as a function of time. As a reference, the residual evaluated from Eq. (56) and the exponential decay predicted by Eq. (55) are also plotted. In all cases the results are in good agreement with the residual predicted by Rosenbluth-Hinton theory. Table 1 gives a summary of the frequencies ω_G and decaying rates γ_G from the simulations, compared to the values predicted by the theory. We find an overall good agreement between

numerical results and theory predicted values. Finally, note that the inclusion of Finite-Orbit-Width effects (see Eq. (2.10) of [57]) would increase the theoretical damping rate by only 8% for these simulations and are therefore neglected.

5.4 CYCLONE benchmark

In Ref. [59], several fluid, gyrofluid and gyrokinetic codes are compared for the so called CYCLONE test case, which represents local parameters from an ITER-relevant DIII-D H-mode shot [58]. The physical parameters used here are $\rho^* = 1/175$, $a = 0.48$ $[m]$, $B_0 = 1.91$ $[T]$, $R_0 = 1.32$ $[m]$, $s_0 = 0.624$ (corresponds to $\rho_0 = 0.5a$), $q(s_0) = 1.4$, $T_i = T_e$, $R_0/L_{T,i} = 6.9$, $\eta_i = L_n/L_{T,i} = 3.12$, $\hat{s} = 0.78$, $\Delta_{T,i} = 0.3$. The value of ρ^* has been decreased by increasing the electronic temperature $T_e(s_0)$ to 2.52 keV, to avoid a too fast relaxation and to approach the flux-tube limit $\rho^* \to 0$ without going beyond the available computational resources. The numerical parameters are the following: $N = 2^{26} \cong 83$ M markers, $\Delta t = 40\ \Omega_i^{-1}$. The quasi-neutrality equation is solved with cubic B-splines on a $N_s = 128, N_\chi = 448, N_\varphi = 320$ grid and a field-aligned filter with $\Delta m = 5$, is applied. Benchmarking ORB5 for these parameters is crucial in order to have confidence in the code.
A first simple test is presented on Fig. 21, where the growth rates obtained with the ORB5 code run in the linear mode are compared with GT3D [16]. ρ^* has been changed to 1/140, a local maxwellian and the exact same equilibrium profiles have been employed in an effort to have similar parameters between the two codes, which show excellent agreement.
The numerical quality of CYCLONE nonlinear simulation is shown on Fig. 22 through $\epsilon_E(t) = \Delta E(t)/E_f(t)$ which has a meaning only in the nonlinear phase. For this CYCLONE simulation, energy is conserved up to very long times as $|\epsilon_E| < 0.3$ up to $t = 5 \cdot 10^4\ \Omega_i^{-1}$, a time at which the temperature profile is relaxed (see Fig. 23). This is a remarkable value for a global PIC code. This energy deviation represents 10^{-5} of the total initial energy of the system. As the quasi-equilibrium state establishes, numerical noise grows and slowly leads to the loss of energy conservation. However, for late times the system is close to marginal stability because of profile relaxation so the state of the system will not provide any new physical information. Hence it is useless to continue a PIC simulation up to very long times without sources and sinks. The situation could be different if collisional sources were added to the simulation.
Nonlinear benchmark is usually performed by plotting the ion diffusivity defined by $\chi_i \equiv -Q/(n_i \nabla T_i)$ versus $R_0/L_{T,i}$. Note that no assumption is done on Q, n_i and ∇T_i: these profiles are reconstructed with appropriate moments of the Vlasov equation and then smoothed using splines with tension interpolation [60]. In [59], Dimits proposed a fit to express χ_i as a function of $R_0/L_{T,i}$

when the system has reached (quasi-)steady state:

$$\frac{\chi_i}{\chi_{\text{Dimits}}} \cong 15.4\left(1 - 6\frac{L_{T,i}}{R_0}\right), \tag{59}$$

with $\chi_{\text{Dimits}} = \chi_{\text{GB}} a/L_n$ where $\chi_{\text{GB}} = (\rho^*)^2 c_s/a$ is the gyro-Bohm transport coefficient [61]. There are two difficulties in benchmarking ORB5 against the Dimits fit, which has been obtained with a flux-tube code. First, spatial averaging must be applied since ORB5 is a global code. Second, the temperature profile and consequently $R_0/L_{T,i}$ are not frozen as in flux-tube codes and relaxes during the simulation. It is therefore better to characterize the radial transport by a cloud of points $(\chi_i, R_0/L_{T,i})$ representing the time evolution of space averaged values rather than with a single point. Such a procedure has been applied in Ref. [62]. ORB5 results are displayed in Fig. 24. At the beginning of the simulation, the radial transport is null and the logarithmic gradient variation is very weak. Then the turbulence establishes, leading to profile relaxation. Finally, the system is in quasi-equilibrium state, meaning that the ion diffusivity decreases to 0 and the logarithmic gradient is close to marginal stability. The cloud of points taken from the relaxation phase is well located near the Dimits fit. The dispersion is more important for small averaging widths Δs. This quantity should be large enough to average bursts, but a too large value would move the temperature gradient too far away from the local value. Remark on Fig. 23 how fast the profile relaxes to a quasi-equilibrium state. A way to prevent this phenomenon would be to decrease ρ^*, thus reducing global effects. Unfortunately simulations at lower ρ^* were not possible because of the limited computational power.

6 Conclusion

The global collisionless PIC code ORB5 is a powerful tool for micro-instabilities studies. It relies on a gyrokinetic model which conserves energy and particle number, providing useful checks of the numerical solution at each time step. Based on the δf method, this code uses true equilibrium functions to prevent spurious zonal flows. Several noise reduction techniques are implemented, such as adaptative gyro-averaging, optimized loading and Fourier filtering. In this paper, it is shown that simulations based on field-aligned Fourier filtering combined with the use of straight-field line coordinates have lead to a speedup of 2 orders of magnitude. More precisely, the time step is relaxed by removing the high $k_{\|}$ modes from the turbulence spectrum, which are anyhow inconsistent with the gyrokinetic ordering, and the number of markers needed in the simulation is strongly reduced because the simulation contains much less Fourier modes. By using a straight-field-line poloidal coordinate instead of the poloidal geometrical angle, the required width of the field-aligned filter is minimal and most importantly becomes independent of the plasma size. The latter remark is of considerable interest as future tokamaks such as ITER will have extremely small ρ^*. Magnetic coordinates and the field-aligned filter are therefore important steps towards the simulation of realistic devices. ORB5 also differs from other gyrokinetic codes as it evolves particle orbits in magnetic coordinates, therefore avoiding costly interpolations during charge assignment. ORB5 shows excellent scalability properties: a proper parallelization is crucial with the development of large scale computers. The question of convergence with the number of markers has been studied by applying a physical initialization of the perturbation instead of random noise. Finally, the code has been successfully benchmarked against other gyrokinetic codes. The next step will be the implementation of several relevant effects missing in the actual model, such as kinetic electron dynamics, collisions, impurities, electromagnetic effects, sources and sinks, in order to reduce the gap between experimental and theoretical transport predictions.

Acknowledgements

This work was partly supported by the Swiss National Science Foundation. Computations were performed on the BG/L and Pleiades clusters of the Ecole Polytechnique Fédérale de Lausanne.

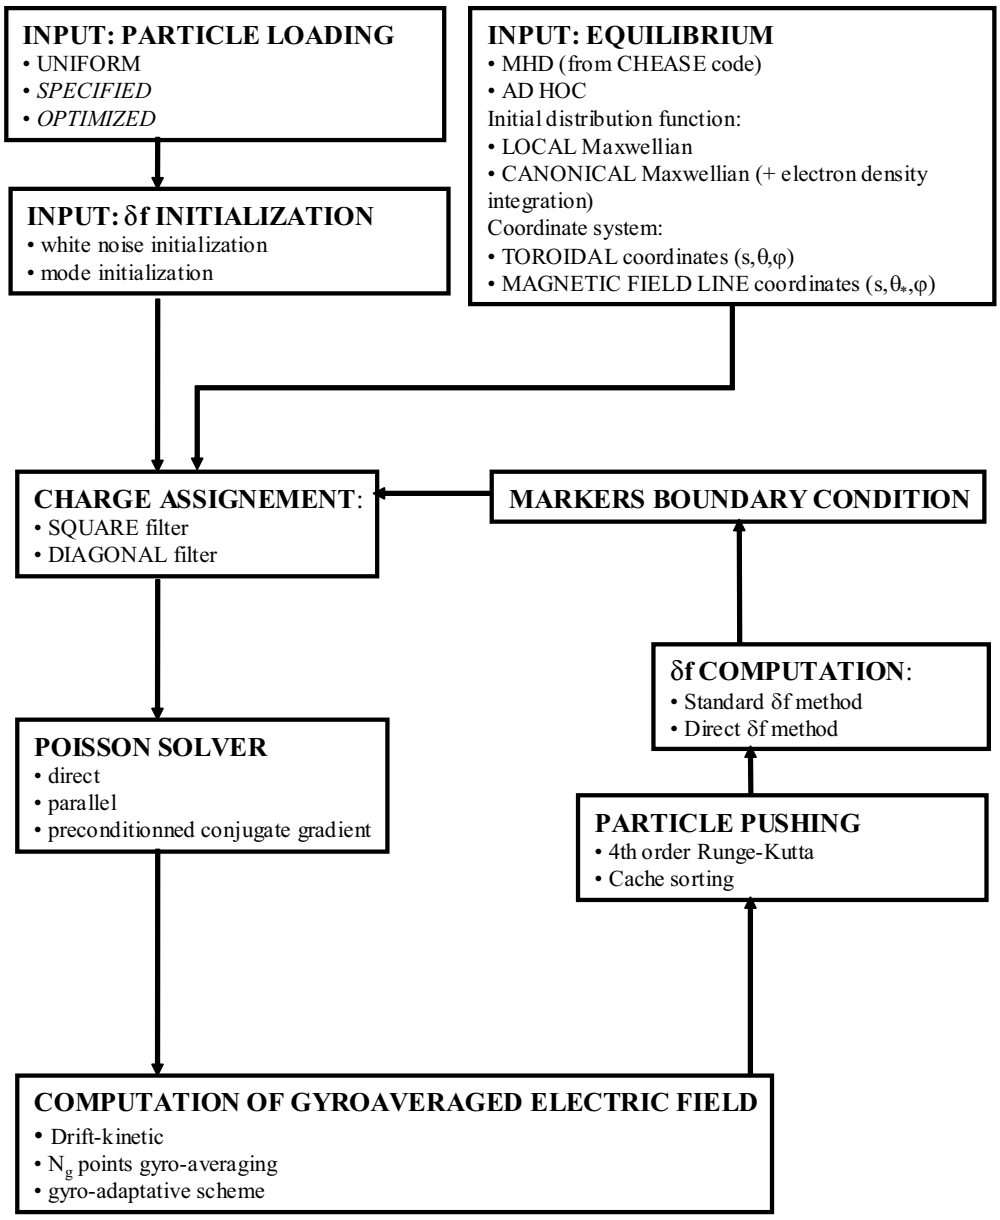

Figure 1. ORB5 scheme.

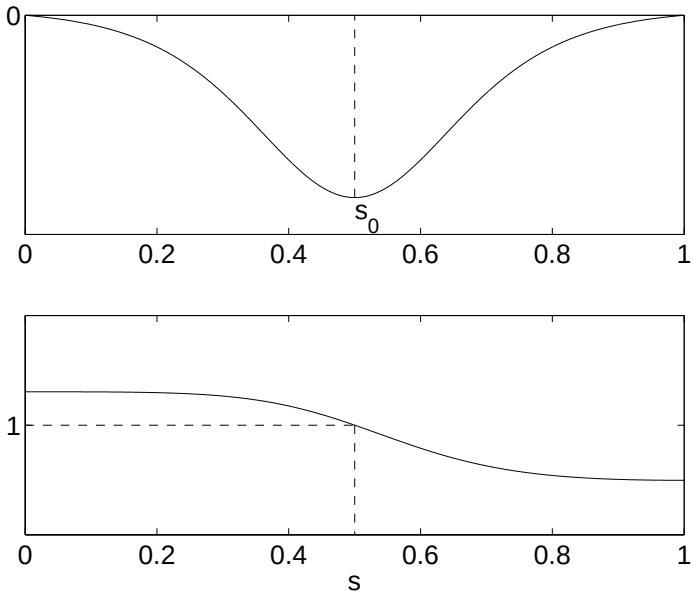

Figure 2. Logarithmic gradient (top) and profile (bottom) used in ORB5.

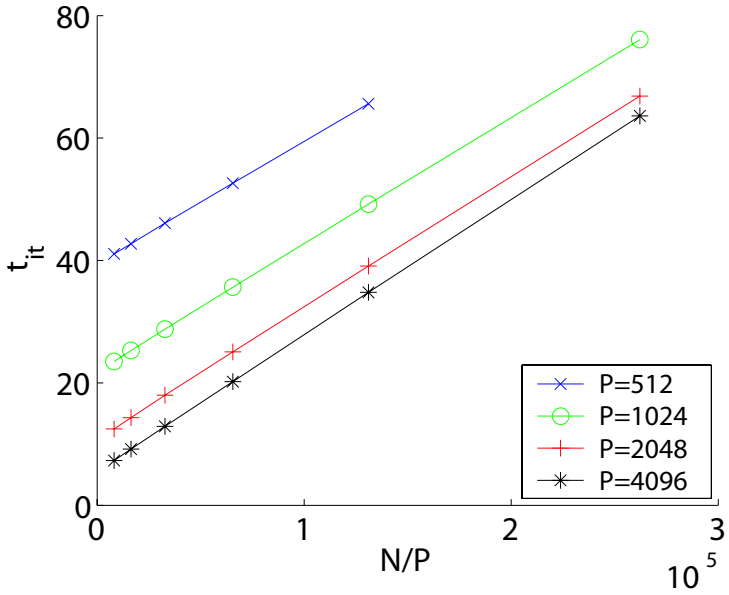

Figure 3. Time per iteration as a function of N/P.

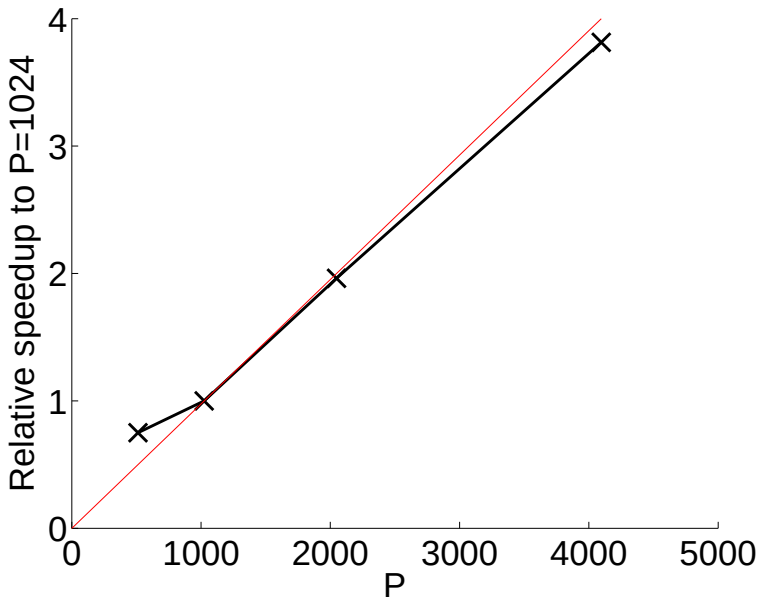

Figure 4. Speedup (relative to P=1024) as a function of P for a fixed number of markers (black, crosses) compared to an ideal scaling (red) .

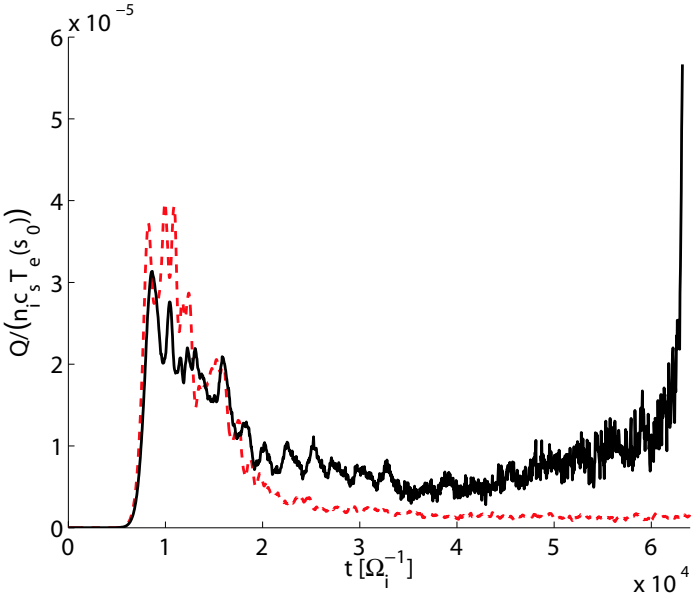

Figure 5. Volume averaged radial heat flux for a square filter simulation (black, solid) and a field-aligned simulation (red, dashed) with $\Delta m = 7$, keeping all other physical and numerical parameters fixed.

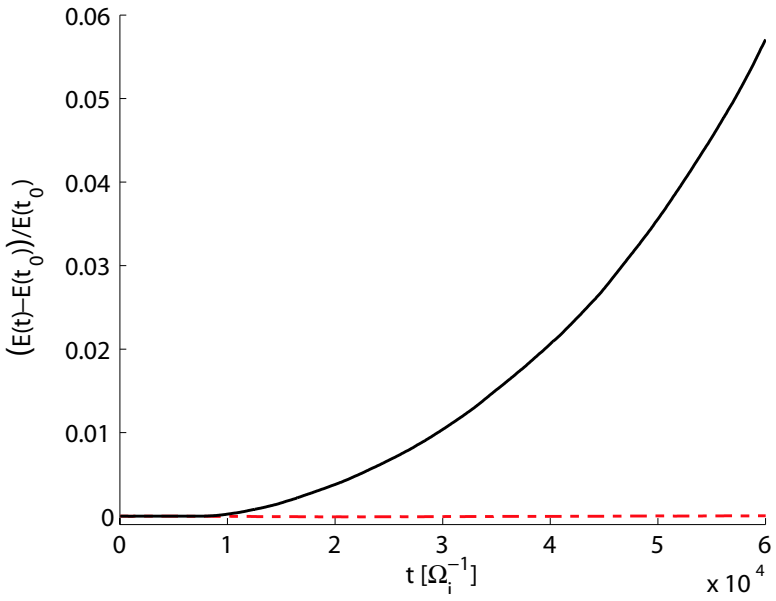

Figure 6. Relative energy conservation for a square filter simulation (black, solid) and a field-aligned simulation (red, dashed) with $\Delta m = 7$, keeping all other physical and numerical parameters fixed.

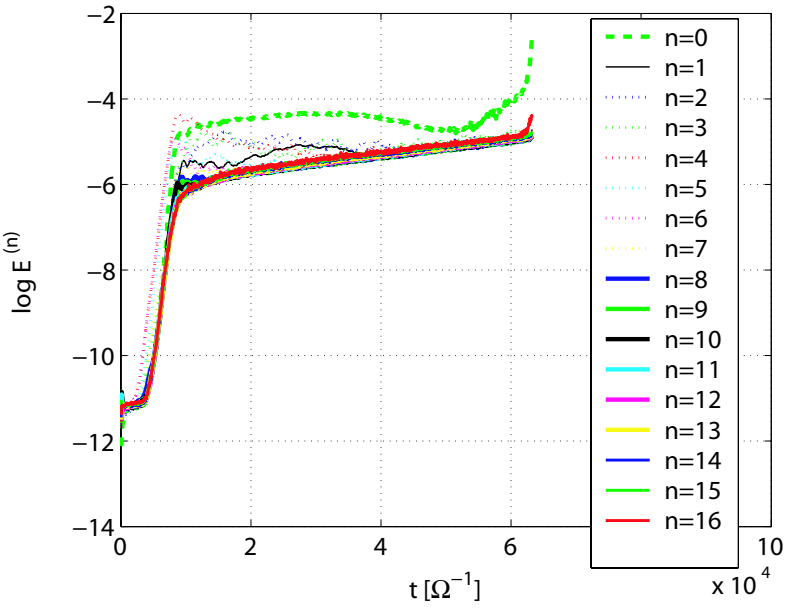

Figure 7. Energy of toroidal modes for a square filter simulation. Dashed line is the $n = 0$ mode.

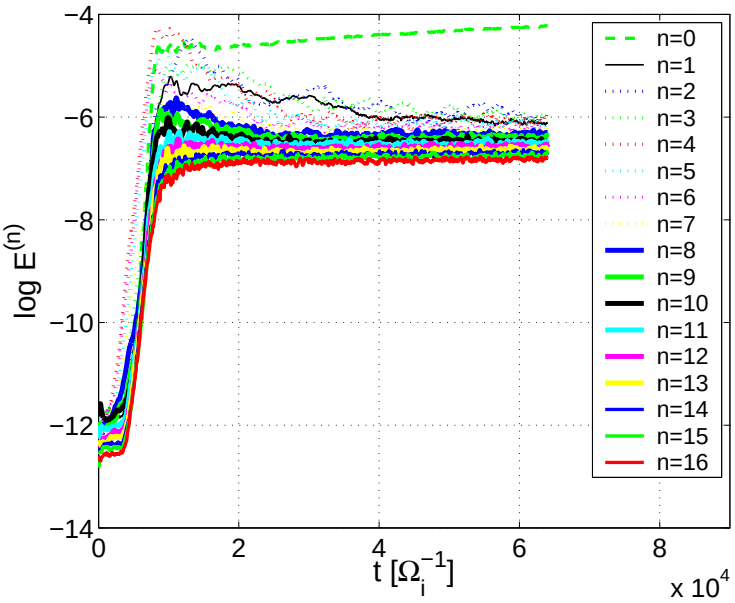

Figure 8. Energy of toroidal modes for a field-aligned filter simulation with $\Delta m = 7$. Dashed line is the $n = 0$ mode.

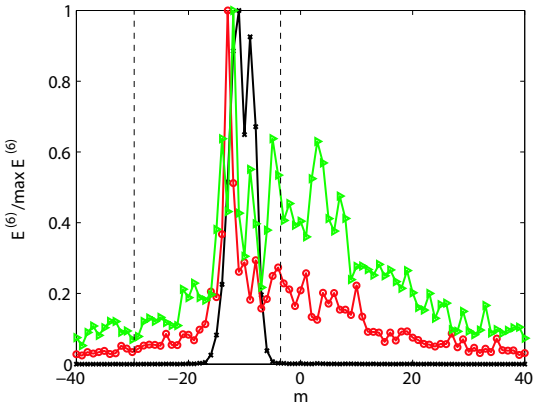

Figure 9. Radially averaged poloidal spectrum of energy for toroidal mode $n = 6$, normalized to the largest component, at time $t = 0.8 \cdot 10^4\ \Omega_i^{-1}$ (linear phase, black, crosses), $t = 2.4 \cdot 10^4\ \Omega_i^{-1}$ (red, circles) and $t = 4.8 \cdot 10^4\ \Omega_i^{-1}$ (green, triangles) for a square filter simulation. Dashed lines show $-nq_{\text{axis}} + \Delta m$ and $-nq_{\text{edge}} - \Delta m$.

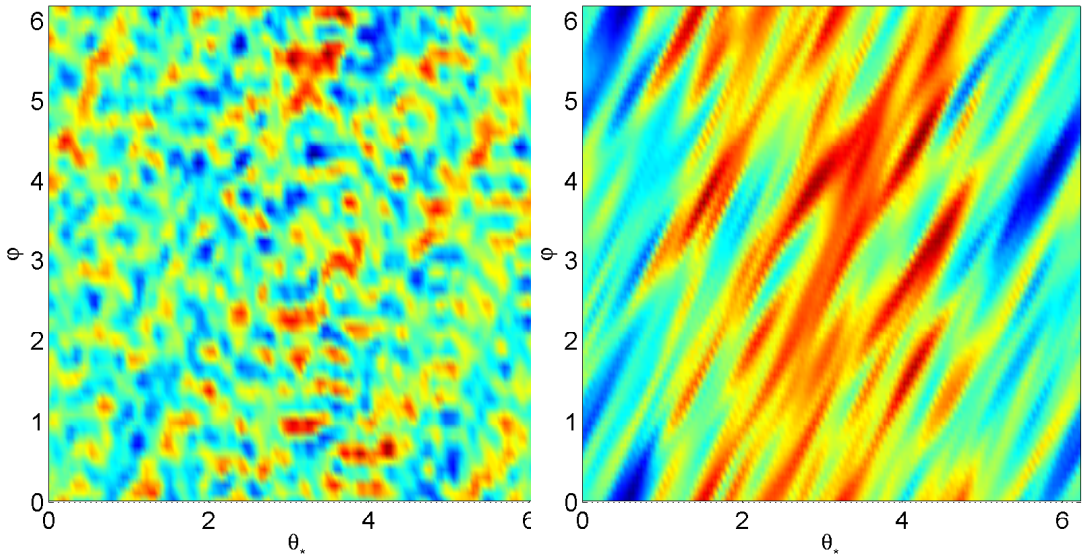

Figure 10. Perturbed electric potential along $s_0 = 0.5$ at $t = 5 \cdot 10^4\ \Omega_i^{-1}$ in the (θ_*, φ) plane for a square filter (left) and a diagonal filter with $\Delta m = 4$ (right). All other physical and numerical parameters are fixed.

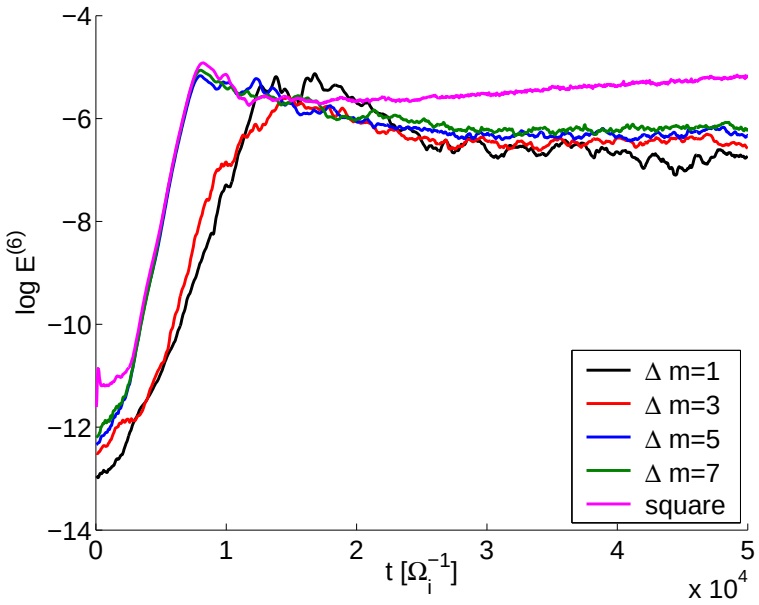

Figure 11. Field energy of the $n = 6$ mode for a field-aligned filter with different values of Δm and a square filter.

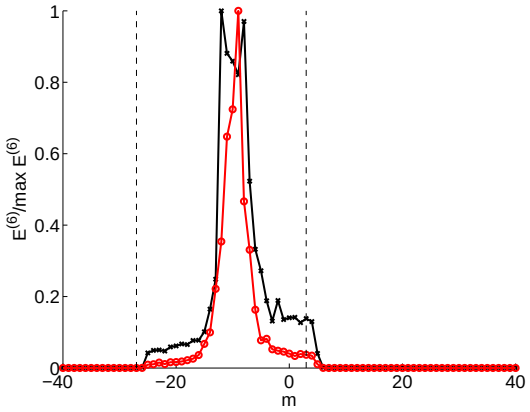

Figure 12. Local poloidal spectrum of energy for toroidal mode $n = 6$ at $s = s_0 = 0.5$, normalized to the largest component, for a simulation with 16 M markers (black, crosses) and 67 M markers (red, circles). Both simulations have a field-aligned filter with $\Delta m = 15$. The spectrum is averaged between $t = 1.5 \cdot 10^4\ \Omega_i^{-1}$ and $t = 2.4 \cdot 10^4\ \Omega_i^{-1}$. The vertical dashed lines show $-nq(s_0) \pm 15$.

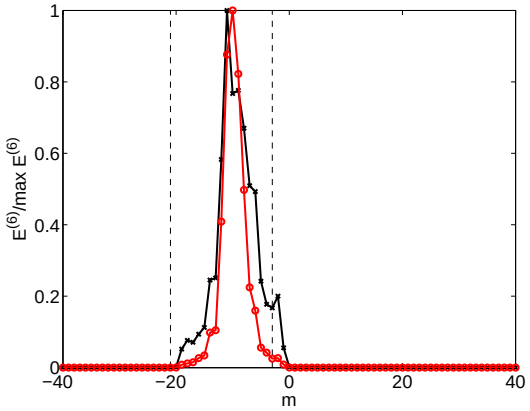

Figure 13. Local poloidal spectrum of energy for toroidal mode $n = 6$ at $s = s_0 = 0.5$, normalized to the largest component, for a simulation with 16 M markers (left, crosses) and 67 M markers (red, circles). Both simulations have a field-aligned filter with $\Delta m = 9$. The spectrum is averaged between $t = 1.5 \cdot 10^4\ \Omega_i^{-1}$ and $t = 2.4 \cdot 10^4\ \Omega_i^{-1}$. The vertical dashed lines show $-nq(s_0) \pm 9$.

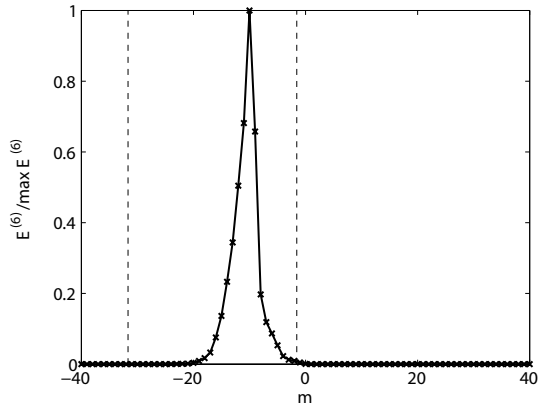

Figure 14. Global poloidal spectrum of energy for toroidal mode $n = 6$, normalized to the largest component, for a simulation with 67 M markers. and a field-aligned filter with $\Delta m = 9$. The spectrum is averaged between $t = 1.5 \cdot 10^4\ \Omega_i^{-1}$ and $t = 2.4 \cdot 10^4\ \Omega_i^{-1}$. The vertical dashed lines show $-nq_{\mathrm{axis}} + \Delta m$ and $-nq_{\mathrm{edge}} - \Delta m$.

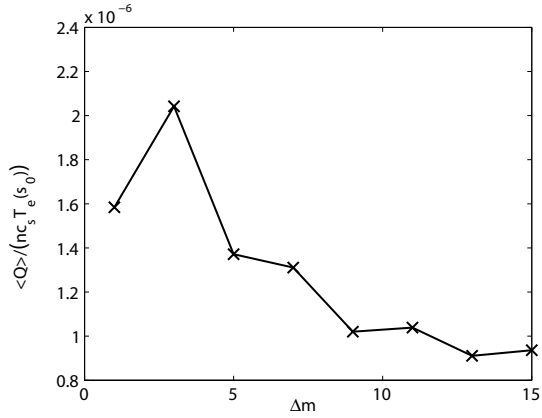

Figure 15. Volume averaged radial heat flux, time-averaged between $t = 4 \cdot 10^4$ and $t = 5 \cdot 10^4\ \Omega_i^{-1}$ as a function of Δm. The number of markers has been varied to keep the number of markers per Fourier mode constant.

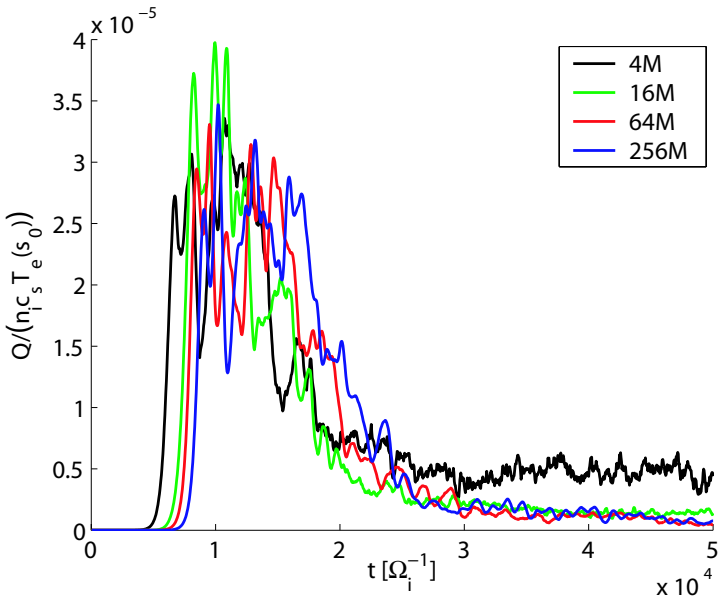

Figure 16. Time evolution of radial for different markers number with white noise initialization and a field-aligned filter.

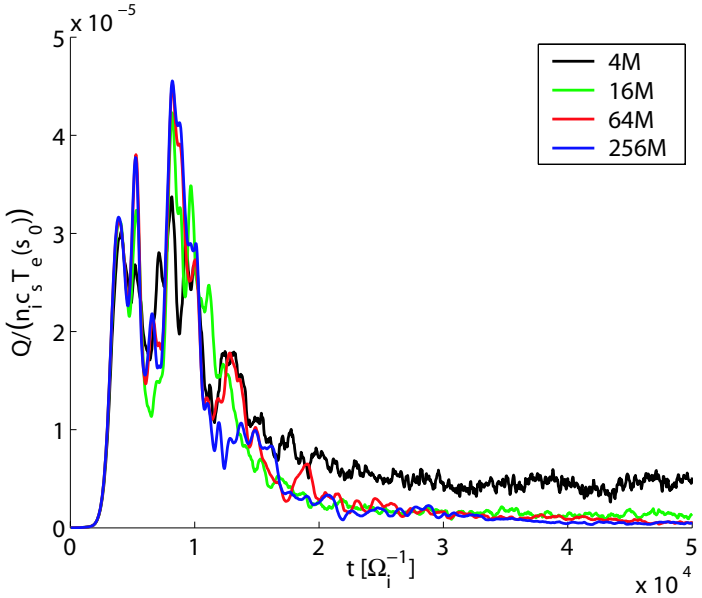

Figure 17. Time evolution of radial heat flux for different markers number with field-aligned filter and a single mode initialization.

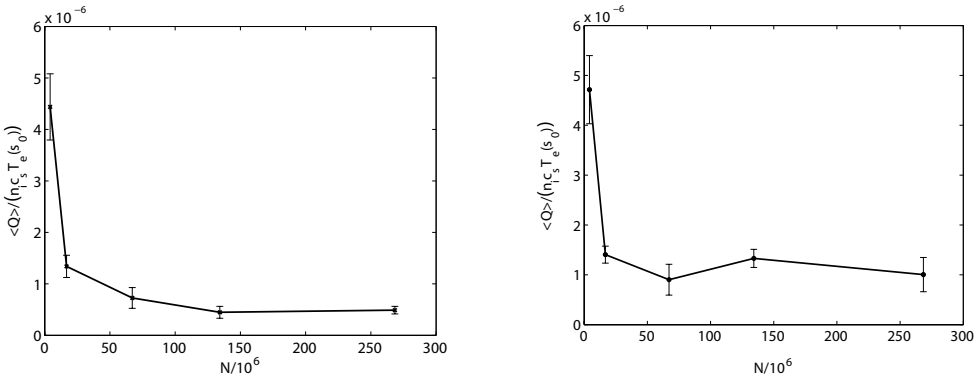

Figure 18. Radial heat flux as a function of the number of markers for the mode initialization (left) and for the white initialization (right). The heat flux is averaged between $t = 4 \cdot 10^4$ and $t = 5 \cdot 10^5 \Omega_i^{-1}$. Error bars represent the standard deviation.

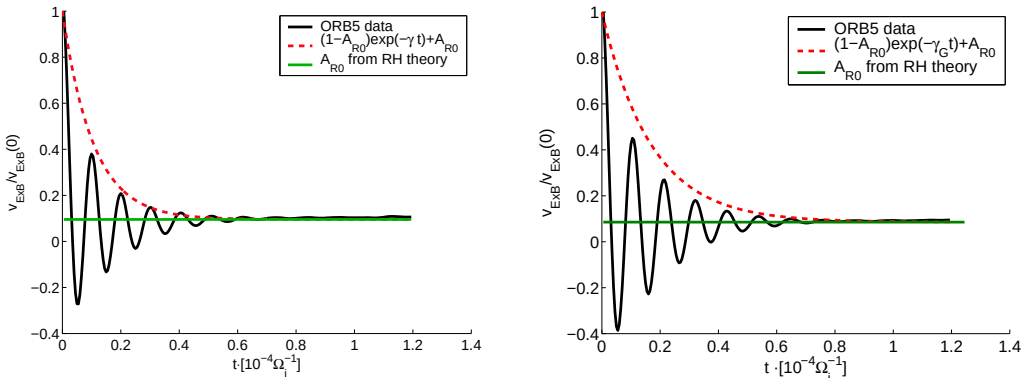

Figure 19. $\vec{E} \times \vec{B}$ velocity at $s_0 = 0.5$ (left) and $s_0 = 0.7$ (right) as a function of time. The solid line is the result of the ORB5 simulation in linear mode, with $\cos(\pi s)$ perturbation. The dotted line is the residual and the dashed line is the exponential decay from Rosenbluth-Hinton theory.

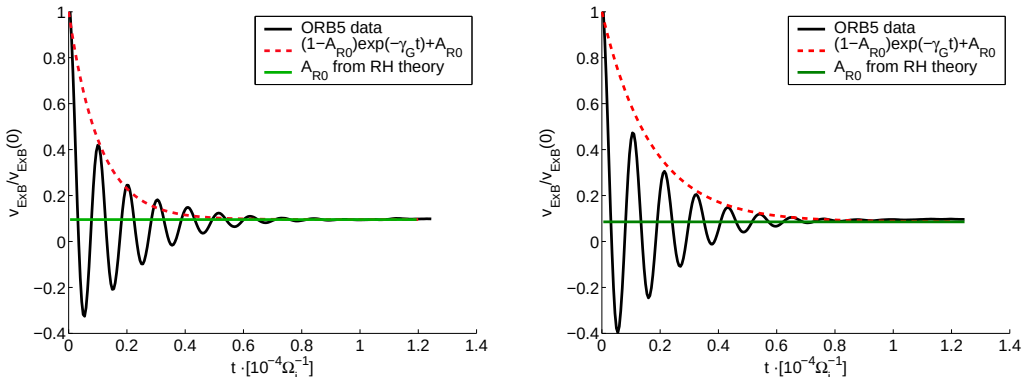

Figure 20. $\vec{E} \times \vec{B}$ velocity at $s_0 = 0.5$ (left) and $s_0 = 0.7$ (right) as a function of time. The solid line is the result of the ORB5 simulation in linear mode, with $\sin(\pi s)$ perturbation. The dotted line is the residual and the dashed line is the exponential decay from Rosenbluth-Hinton theory.

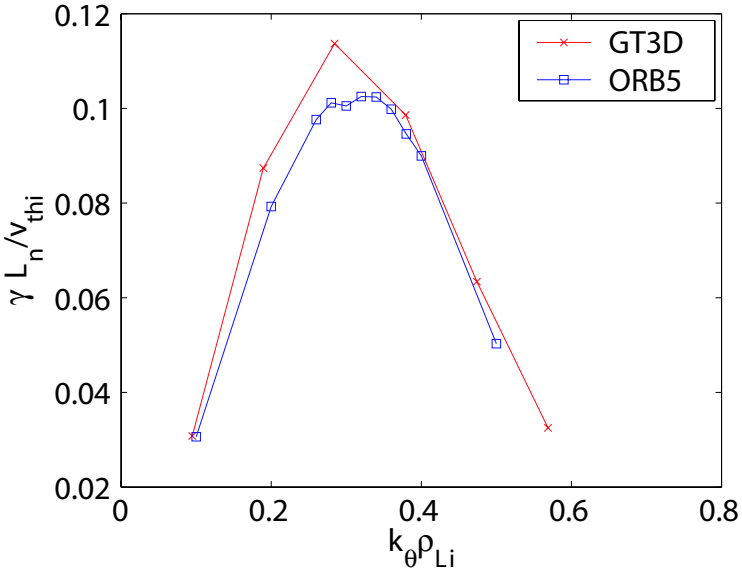

Figure 21. Linear growth rates for GT3D and ORB5.

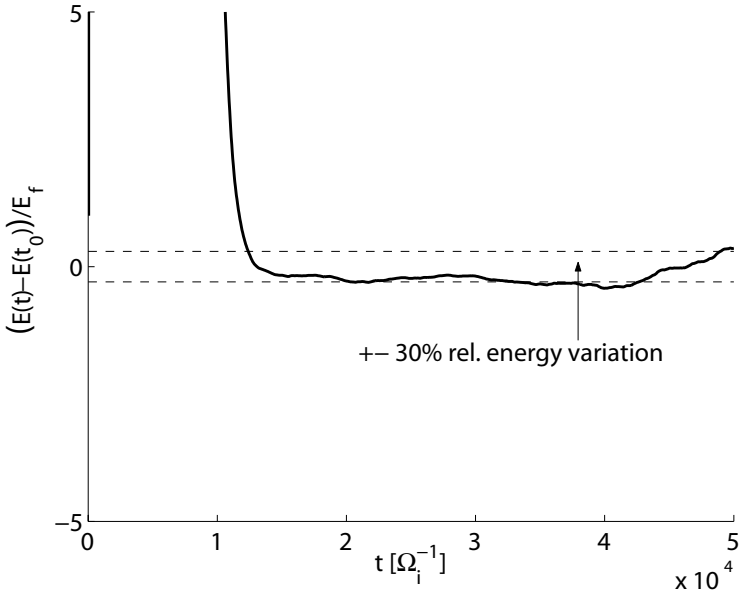

Figure 22. Relative energy conservation for a typical CYCLONE simulation.

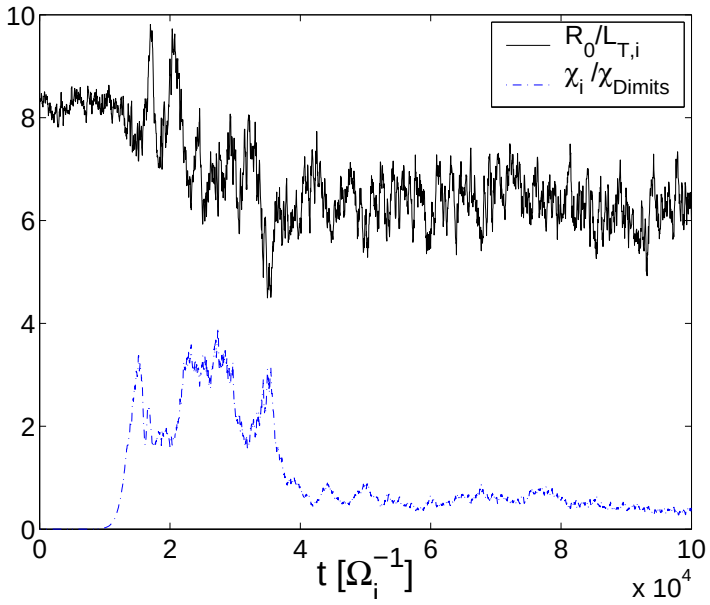

Figure 23. Time evolution of $R_0/L_{T,i}$ (blue) and $\chi_i/\chi_{\text{Dimits}}$ at $s = s_0$.

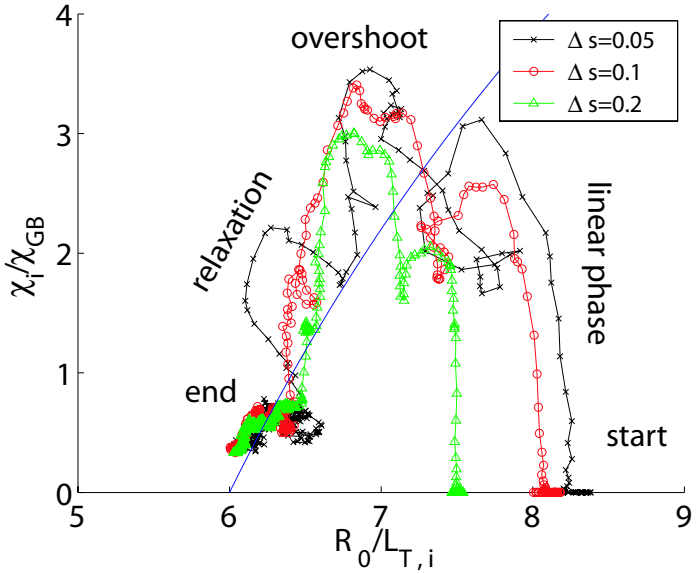

Figure 24. Evolution of χ_i vs $R_0/L_{T,i}$. Each point is averaged between $s_0 \pm \Delta s$. A point is computed every 320 Ω_i^{-1}. The simulation goes from $t = 0$ up to $t = 10 \cdot 10^4$ Ω_i^{-1}. The solid blue curve is the Dimits fit.

pert	s_0	ω_G (Eq. (57))	ω_G (ORB5)	γ_G (Eq. (58))	γ_G (ORB5)
cos	0.5	0.0059	0.0062 ± 0.0002	-0.0009	-0.0009 ± 0.0001
cos	0.7	0.0057	0.0058 ± 0.0002	-0.0006	-0.0007 ± 0.0001
sin	0.5	0.0059	0.0062 ± 0.0002	-0.0009	-0.0008 ± 0.0001
sin	0.7	0.0057	0.0058 ± 0.0002	-0.0006	-0.0006 ± 0.0001

Table 1
Comparison between analytical and numerical values for ω_G and γ_G.

References

[1] ITER Physics Expert Group on Confinement and Transport, ITER Physics Expert Group on Confinement Modelling and Database, Iter Physics Basis Editors. Nucl. Fusion 39:2175 (1999).

[2] X. Garbet. Plasma. Phys. Control. Fusion 43:A251 (2001).

[3] W. Horton. Rev. Modern Phys. 71:735 (1999).

[4] P. H. Diamond, S. I. Itoh, K. Itoh and T. S. Hahm. Plasma Phys. Control. Fusion 47:R35 (2005).

[5] L. Garcia and B. A. Carreras. Phys. Plasmas 12:092305 (2005).

[6] J. F. Drake, P. N. Guzdar and A. B. Hassam. Phys. Rev. Lett. 61:2205 (1988).

[7] F. Jenko, W. Dorland, M. Kotschenreuther and B. N. Rogers. Phys. Plasmas 7:1904 (2000).

[8] W. W. Lee. Phys. Fluids, 26:256 (1983).

[9] T. S. Hahm. Phys. Fluids. 31:2670 (1988).

[10] E. A. Frieman and L. Chen. Phys. Fluids 25: 502 (1982).

[11] A. J. Brizard. Phys. Plasmas 7:4816 (2000).

[12] R. G. Littlejohn J. Plasma Physics, 29:111 (1983).

[13] T. M. Tran, K. Appert, M. Fivaz, G. Jost, J. Vaclavik and L. Villard. Theory of fusion Plasmas, Int. Workshop, Editrice Compositori, Societá italiana di Fisica, Bologna, 1999, p. 45.

[14] A. M. Dimits, T. J. Williams, J. A. Byers and B. I. Cohen. Phys. Rev. Lett., 77:71 (1996).

[15] Z. Lin, T. S. Hahm, W. W. Lee, W. M. Tang and R. B. White. Science, 281, 1835 (1998).

[16] Y. Idomura, S. Tokuda and Y. Kishimoto. Nucl. Fusion 43:234 (2003).

[17] R. Sydora, V. Becyk and J. Bawson. Plasma Phys. Control. Fusion, 38:A281 (1996).

[18] S. E. Parker, C. Kim and Y. Chen. Phys. Plasmas 6:1709 (1999).

[19] M. Kotschenreuther, G. Rewoldt and W. M. Tang. Comp. Phys. Commun. 88:128 (1995).

[20] F. Jenko. Comp. Phys. Commun. 125:196 (2000).

[21] J. Candy and R.E. Waltz. Journ. Comp. Phys., 186:545 (2003).

[22] C. Z. Cheng and G. Knorr. Journ. Comput. Phys. 22:330 (1976).

[23] M. Brunetti, V. Grandgirard, O. Sauter, J. Vaclavik and L. Villard. Comput. Phys. Commun. 163:1 (2004).

[24] V. Grandgirard, M. Brunetti, P. Bertrand, N. Besse, X. Garbet, P. Gendrih, G. Manfredi, Y. Sarazin, O. Sauter, E. Sonnendrücker, J. Vaclavik and L. Villard.*A drift-kinetic Semi-Lagrangian 4D code for ion turbulence simulation.* Journ. Comp. Phys., 217:395 (2006)

[25] G. Jost, T. M. Tran, W. A. Cooper, L. Villard and K. Appert. Phys. Plasmas 8, 3321 (2001).

[26] V. Kornilov, R. Kleiber, R. Hatzky, L. Villard and G. Jost. Phys. Plasmas 11:3196 (2004).

[27] M. Kotschenreuther. Bull. Am. Phys. Soc., 33:2107 (1988).

[28] R. Hatzky, T. M. Tran, A. Könies, R. Kleiber and S. J. Allfrey. Phys. Plasmas, 9:898 (2002).

[29] P. Angelino, A. Bottino, R. Hatzky, S. Jolliet, T. M. Tran and L. Villard. Plasma Phys. Control. Fusion, 48:557 (2006).

[30] A. Bottino, P. Angelino, S. J. Allfrey, S. Brunner, R. Hatzky, Y. Idomura, S. Jolliet, O. Sauter, T. M. Tran and L. Villard. Theory of fusion Plasmas, Int. Workshop, Editrice Compositori, Societá italiana di Fisica, Bologna, 2005, p. 75.

[31] H. Lütjens, A. Bondeson and O. Sauter. Comp. Phys. Comm. 97:219 (1996).

[32] P. Angelino, A. Bottino, R. Hatzky, S. Jolliet, O. Sauter, T. M. Tran and L. Villard. Phys. Plasmas 13, 052304 (2006).

[33] S. J. Allfrey and R. Hatzky. Comp. Phys. Comm. 154:98 (2003).

[34] C. K. Birdsall and A. B. Langdon. Adam Hilger, Bristol (1989).

[35] W. W. Lee. Journ. Comp. Phys., 72(1):243 (1987).

[36] A. Y. Aydemir. Phys. Plasmas 1:822 (1994).

[37] A. Bottino. PhD Thesis No.2938, Ecole Polytechnique Fédérale de Lausanne, (2004).

[38] A. Bottino, A. G. Peeters, R. Hatzky, S. Jolliet, B. F. McMillan, T. M. Tran and L. Villard. Phys. Plasmas 14:010701 (2007).

[39] K. J. Bowers. Journ. Comp. Phys., 173(2):393 (2001).

[40] M. Fivaz, S. Brunner, G. de Ridder, O. Sauter, T. M. Tran, J. Vaclavik, L. Villard and K. Appert. Comp. Phys. Commun., 111:27 (1998).

[41] K. Höllig. *Finite Element Methods with B-Splines*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia (2003).

[42] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith and H. Zhang. *PETSc Web page* (2001) URL: http://www.mcs.anl.gov/petsc

[43] *WSMP Web Page*.URL:http: www-users.cs.umn.edu/∼agupta/wsmp.html

[44] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney and D. Sorensen.*LAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, third ed., 1999.

[45] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker and R. C. Whaley. *ScaLAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.

[46] S. Jolliet, P. Angelino, A. Bottino, Y. Idomura and L. Villard. Theory of fusion Plasmas, Int. Workshop, Editrice Compositori, Societá italiana di Fisica, Bologna, 2005, p. 345.

[47] A. Bottino, A. G. Peeters, O. Sauter, J. Vaclavik and L. Villard. Phys. Plasmas 11:198 (2004).

[48] Message Passing Interface Forum. *MPI: A Message-Passing Interface standard (version 1.1)*. Technical report, 1995. URL: http://www.mpi-forum.org

[49] S. Jolliet, A. Bottino, P. Angelino, T. M. Tran, B. F. Mcmillan, R. Hatzky, A. G. Peeters, E. Poli, O. Sauter and L. Villard. Proceedings of Theory of fusion Plasmas, 2006, to be published.

[50] C. C. Kim and S. E. Parker. Journ. Comp. Phys., 161(2):589 (2000).

[51] R. Hatzky. Parallel Computing 32:325 (2006).

[52] S. Brunner. PhD Thesis No.1701, Ecole Polytechnique Fédérale de Lausanne, (1997).

[53] A. Bottino, P. Angelino, R. Hatzky, S. Jolliet, A. G. Peeters, E. Poli, O. Sauter, T. M. Tran and L. Villard. *Proc.* 33^{rd} *EPS Conf. on Plasma Physics and Controlled Fusion (Roma, Italy, June 2006), to be published*.

[54] W. M. Nevins, G. W. Hammett, A. M. Dimits, W. Dorland and D. E. Shumaker. Phys. Plasmas 122305 (2005).

[55] M. N. Rosenbluth and F. L. Hinton. Phys. Rev. Lett., 80:724 (1998).

[56] N. Winsor, J. L. Johnson and J. M. Dawson. Phys. Fluids, 11:2448 (1968).

[57] H. Sugama and T. H. Watanabe. J. Plasma Physics 72:825 (2006).

[58] C. M. Greenfield, J. C. DeBoo, T. H. Osborne, F. W. Perkins, M. N. Rosenbluth and D. Boucher. Nucl. Fusion, 37:1215 (1997).

[59] A. M. Dimits, G. Bateman, M. A. Beer, B. I. COhen , W. Dorland, G. W. Hammet, C. Kim, J. E. Kinsey , M. Kotschenreuter, 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.Phys. Plasmas 7:969 (2000).

[60] S. P. Hirshman, D. K. Lee, F M. Levinton, S. H. Batha, M. Okabayashi and R. M. Wieland. Phys. Plasmas 1, 2277 (1994).

[61] R. E. Waltz, J. C. DeBoo and M. N. Rosenbluth. Phys. Rev. Let. 65: 2390 (1990).

[62] B. D. Scott. Phys. Plasmas 12: 102307 (2005).