Gyrokinetic simulations of turbulent transport: size scaling and chaotic behaviour

2010 Plasma Phys. Control. Fusion 52 124038

(http://iopscience.iop.org/0741-3335/52/12/124038)

IOP PUBLISHING PLASMA PHYSICS AND CONTROLLED FUSION

Plasma Phys. Control. Fusion **52** (2010) 124038 (18pp) doi:10.1088/0741-3335/52/12/124038

Gyrokinetic simulations of turbulent transport: size scaling and chaotic behaviour

L Villard[1], A Bottino[2], S Brunner[1], A Casati[1], J Chowdhury[3], T Dannert[2], R Ganesh[3], X Garbet[4], T Görler[2], V Grandgirard[4], R Hatzky[2], Y Idomura[5], F Jenko[2], S Jolliet[5], S Khosh Aghdam[1], X Lapillonne[1], G Latu[4], B F McMillan[1], F Merz[2], Y Sarazin[4], T M Tran[1] and T Vernay[1]

[1] Ecole Polytechnique Fédérale de Lausanne (EPFL), Centre de Recherches en Physique des Plasmas, Association Euratom-Switzerland, 1015 Lausanne, Switzerland
[2] Max-Planck Institut für Plasmaphysik, Association Euratom, Garching, Germany
[3] Institute for Plasma Research, Bhat, Ghandinagar, India
[4] CEA, IRFM, Association Euratom, F-13108 Saint Paul Lez Durance, France
[5] Japan Atomic Energy Agency, Higashi-Ueno 6-9-3, Taitou, Tokyo 110-0015, Japan

Received 18 June 2010, in final form 5 August 2010
Published 15 November 2010
Online at stacks.iop.org/PPCF/52/124038

Abstract

Important steps towards the understanding of turbulent transport have been made with the development of the gyrokinetic framework for describing turbulence and with the emergence of numerical codes able to solve the set of gyrokinetic equations. This paper presents some of the main recent advances in gyrokinetic theory and computing of turbulence. Solving 5D gyrokinetic equations requires state-of-the-art high performance computing techniques involving massively parallel computers and parallel scalable algorithms. The various numerical schemes that have been explored until now, Lagrangian, Eulerian and semi-Lagrangian, each have their advantages and drawbacks. A past controversy regarding the finite size effect (finite ρ_*) in ITG turbulence has now been resolved. It has triggered an intensive benchmarking effort and careful examination of the convergence properties of the different numerical approaches. Now, both Eulerian and Lagrangian global codes are shown to agree and to converge to the flux-tube result in the $\rho_* \rightarrow 0$ limit. It is found, however, that an appropriate treatment of geometrical terms is necessary: inconsistent approximations that are sometimes used can lead to important discrepancies. Turbulent processes are characterized by a chaotic behaviour, often accompanied by bursts and avalanches. Performing ensemble averages of statistically independent simulations, starting from different initial conditions, is presented as a way to assess the intrinsic variability of turbulent fluxes and obtain reliable estimates of the standard deviation. Further developments

concerning non-adiabatic electron dynamics around mode-rational surfaces and electromagnetic effects are discussed.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

It is now well established that turbulent behaviour is responsible for the presence of anomalous transport of heat, particles and momentum in magnetically confined fusion plasmas. The typical frequency range of turbulence is experimentally observed much lower than the cyclotron frequency of the plasma species. This is at the basis of the development of gyrokinetic theory, which is considered as the model of choice for its description. In parallel to these analytical developments, a tremendous effort has been devoted to the numerical resolution of the pertinent equations, with the result that nowadays direct numerical simulations of turbulence are feasible. For a comprehensive review see [1].

This paper focuses on crucial aspects that gyrokinetic simulations of turbulence have faced in the recent past and how previous controversies have been resolved.

Solving the set of gyrokinetic equations together with the appropriate field (Maxwell) equations is a rather formidable task. The main difficulties come from the disparate range of space and time scales of the plasma species (ions and electrons), the extreme anisotropy created by the background magnetic field and the various nonlinearities in the system. This has prompted the development of advanced numerical techniques that take advantage of the massively parallel high performance computing (HPC) platforms that have reached Petaflops range (i.e. 10^{15} floating point operations per second). Nevertheless, the full radius direct numerical simulation of an ITER-size plasma with gyrokinetic codes including all the relevant physics is still a long way ahead. Moreover, tapping the high-end HPC power requires a continuing adaptation of the numerical algorithms, which poses a non-trivial challenge.

The paper is structured as follows. First, a brief summary of the gyrokinetic equations and numerical methods is given in section 2. In section 3, the question of finite size (ρ_*) scaling of ITG turbulence [2, 3] is revisited, where $\rho_* = \rho_s/a$, ρ_s is the ion sound Larmor radius and a is the minor radius. Discrepancies between the results reported in [2, 3] had tentatively been ascribed to numerical noise inherent to the Lagrange-PIC method and finite aspect ratio effects. These issues are examined here. An often used geometrical model (s–α), containing inconsistencies at first order in inverse aspect ratio, is shown to be responsible for most of the difference. On the other hand, true finite aspect ratio effects, such as the zero β Shafranov shift, are only responsible for a very small difference in the results. It is also shown that numerical noise problems have now been solved and that when geometrical inconsistencies are removed, both Eulerian and Lagrangian-PIC codes are found to agree on the ρ_* scaling of ITG turbulence. A by-product of these developments is that another past controversy, concerning ETG turbulent transport level, has also been resolved: low-noise Lagrange-PIC simulations agree with Eulerian ones [4].

Another important finding is that ITG turbulence is bursty: a measurable fraction of the heat flux is carried by large, avalanche-like events [5–8]. There are signatures of chaos in the sequence of bursts. This poses difficulties as to how long a simulation should be carried out in order to reach statistical convergence, about the reproducibility of simulations using parallel algorithms and also implies that, strictly speaking, uniform numerical convergence cannot be achieved. A correct way to obtain statistically meaningful estimates, presented in section 4,

is to perform *ensemble averages* over a set of independent simulations differing only in their initial conditions.

The question of non-adiabatic electron response, in particular in the vicinity of mode-rational surfaces, and its consequence on the mode structure and on the turbulent heat transport are discussed in section 5.

2. Gyrokinetic equations and numerical models

Experimental observations in the core plasmas of magnetic confinement fusion systems suggest that small scale turbulence, responsible for anomalous transport, obeys the following ordering in a small parameter ϵ_g:

$$\frac{\omega}{\Omega_s} \sim k_\parallel \rho_s \sim \frac{v_E}{v_{ts}} \sim \frac{\delta n_s}{n_0} \sim \frac{B_1}{B_0} \sim \frac{\rho_s}{L_n} \sim \mathcal{O}(\epsilon_g), \tag{1}$$

where ω is a characteristic frequency of micro-turbulence, $\boldsymbol{B}_0$ is the equilibrium magnetic field, $\boldsymbol{B}_1$ is the perturbed magnetic field, $\boldsymbol{b} = \boldsymbol{B}_0/B_0$, $k_\parallel = \boldsymbol{k} \cdot \boldsymbol{b}$, $\boldsymbol{k}$ is the perturbation wavevector, v_E is the perturbed $\boldsymbol{E} \times \boldsymbol{B}$ drift velocity, v_{ts} is the thermal velocity, n_0 is the equilibrium density, δn_s is the perturbed density, $\rho_s = v_\perp/\Omega_s$ is the Larmor radius, $L_n = |\nabla \ln n_0|^{-1}$ is the characteristic scale length of equilibrium density, and the subscript s indicates the species type. Another parameter is $\epsilon_B = \rho_s/L_B$, where $L_B = |\nabla \log B_0|^{-1}$ is the characteristic scale length of equilibrium magnetic field. In most cases the gyrokinetic equations are obtained at first order in ϵ_B. It must be noted that while $k_\parallel \rho_s$ is ordered as a small quantity, the perpendicular wavenumber is not assumed small, $k_\perp \rho_s \sim \mathcal{O}(1)$. This reflects the very strong anisotropy of the perturbations, which typically have parallel wavelengths of the order of the system size whereas perpendicular wavelengths range down to the Larmor radius scale.

The essential steps to obtain the gyrokinetic equations are as follows (for more details see [1, 9] and references cited therein). The starting point is the Fokker–Planck equation (or its collisionless version, the Vlasov equation), describing the time evolution of the plasma species distribution function f_s in the 6D phase space, together with Maxwell's equations. Non-canonical guiding centre coordinates are used. Then the Lie-type gyro-centre transform is applied to find coordinates for which the gyro-angle remains a cyclic variable even in the presence of fluctuating fields, therefore reducing the dimensionality from 6D to 5D.

Three classes of numerical methods have been applied to solve the set of gyrokinetic equations: Lagrangian, Eulerian and semi-Lagrangian. The Lagrangian approach, often referred to as 'PIC' (particle-in-cell), has a Monte Carlo aspect in that quasi-random positions in phase space, called markers, are used to obtain moments of the distribution function. The marker orbits are followed in time according to the characteristic equations. The Eulerian approach (sometimes referred to as 'continuum method') uses fixed grids in phase space and discretizes the differential operators by means of finite difference, finite element, finite volume and/or spectral methods. The semi-Lagrangian approach also uses fixed grids in phase space but integrates the orbits from each mesh point back in time to obtain the value of f at the previous time step by interpolation. For a more detailed description of these algorithms and their relative merits, see [1]. Given the high complexity and computational demands of nonlinear gyrokinetic simulations, numerical convergence is often difficult to achieve and cross-code comparisons involving these three different approaches are crucial to verify key results.

Let us give a quantitative estimate of the computing resources required and their scaling with physical system size, considering the simplest gyrokinetic global model, i.e. gyrokinetic

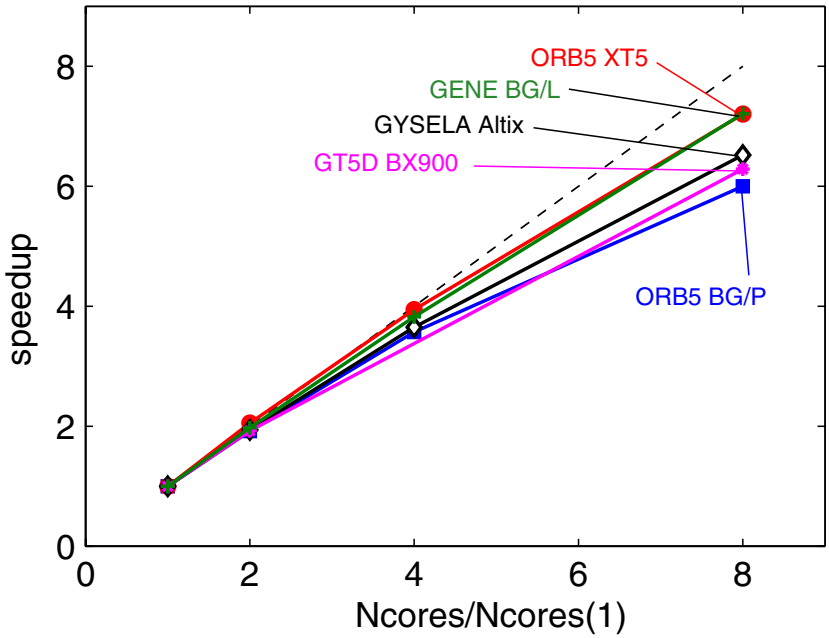

Figure 1. Strong scaling speedup achieved with the Lagrange-PIC code ORB5 [10], the Eulerian codes GT5D [6] and GENE [11], and the semi-Lagrangian code GYSELA [12] on a Cray XT5 (circles) from 1024 to 8192 cores, on an IBM BG/P from 4096 to 32768 cores (squares), on a Fujitsu BX900 from 1024 to 8192 cores (stars), on an IBM BG/L from 1024 to 8192 cores (plus) and on a SGI Altix ICE from 1024 to 8192 cores (diamonds). Parameters are for ORB5, 2×10^9 markers, 3D field solver grid $128 \times 512 \times 256$; for GT5D, phase space grid $240 \times 240 \times 256 \times 128 \times 32$ (6×10^{10} grid points); for GENE, $64 \times 32 \times 128 \times 64 \times 32$ (5×10^8 grid points); for GYSELA, phase space grid $512 \times 256 \times 256 \times 47 \times 32$ (5×10^{10} grid points).

ions and adiabatic electrons. For an ITER-size plasma, we have $1/\rho_* \approx 1000$. Turbulence spatial scales extend from the full machine size down to the gyroradius $k_\perp \rho_s \sim 1$. Considering that a minimum of four grid points per wavelength are minimal in each spatial direction for the field solver, this gives $\sim 7 \times 10^9$ grid cells in configuration space. The minimal velocity space resolution required is $\sim$300 per grid cell, per Fourier mode (see next section), resulting in $\sim 2 \times 10^{12}$ phase space cells (or markers). The necessary time to reach statistical steady-state (see section 4) is of the order of $2 \times 10^3 a/c_s$, corresponding to $\sim 10^5$–10^6 time steps depending on the time integration scheme. The computing cost scales as $(1/\rho_*)^4$ and such a simulation for an ITER-size plasma would need a prohibitive time even on the world's best performing computer ($\sim$2 Petaflops) at the time of writing. Hence, better formulations and algorithms have been devised: it is for example most efficient to use field-aligned coordinates, which take advantage of the anisotropy of the perturbations (parallel versus perpendicular to the background magnetic field). The parallel wavelength scales with system size, thus the grid requirements are reduced by a factor $\sim \rho_*$ and moreover the time step can be increased by a factor $\propto 1/\rho_*$: the scaling of computing cost is then ideally $\sim (1/\rho_*)^2$ and a global simulation of ITG turbulence in the ITER core plasma appears feasible.

Still, the gyrokinetic codes must be run on massively parallel platforms in order to get results in a reasonable amount of time. Thus great efforts have been made to obtain highly scalable algorithms. Examples in figure 1 show that gyrokinetic codes based on all three approaches (Lagrange-PIC, Euler and semi-Lagrange) perform very well in that respect.

3. Finite size effects in ITG turbulence

With the advent of global gyrokinetic codes, an obvious question that has been addressed concerns finite size effects. The finite size parameter is usually defined as $\rho_* = \rho_s/a$, with a the minor radius of the configuration. This is an important question, since present day tokamaks cannot achieve the value of $\rho_* \approx 10^{-3}$ expected in ITER. Flux-tube simulations rely on an approximation that is valid in the infinite size limit $\rho_* \to 0$, and the thermal conductivity naturally scales as

$$\chi \propto \chi_{\mathrm{GB}} = \rho_s^2 c_s/a = \rho_* \chi_{\mathrm{B}}, \tag{2}$$

where χ_{GB} is known as the gyro-Bohm diffusion and χ_{B} is the Bohm diffusion. Global gyrokinetic simulations [2, 3] found a deviation from gyro-Bohm scaling at large values of ρ_*. There was, however, discrepancy as to the absolute value of the heat transport obtained in the limit $\rho_* \to 0$ in the two above cited works. Various tentative explanations were put forward to identify the source of the discrepancy: (1) geometrical effects (finite aspect ratio); (2) the statistical numerical noise problem inherent to the Lagrange-PIC approach. Differences were also exposed as to the interpretation of the broken gyro-Bohm scaling through linear (e.g. through profile shearing [3, 13, 14]) or nonlinear effects (e.g. through turbulence spreading [15]). The gradient profile shape was shown to have an effect on the value of ρ_* above which gyro-Bohm scaling is broken, but in all cases the finite ρ_* heat diffusivity (expressed in gyro-Bohm units) always converges from below to the flux-tube limit for $\rho_* \to 0$: no possible reason was found to have finite ρ_* effects giving a heat transport above the flux-tube limit.

Let us first focus on the numerical noise issue in PIC codes. The Lagrange-PIC algorithm can be viewed as a statistical method to obtain, via Monte Carlo integration, estimates of the moments of the distribution function [16, 17]. The standard deviation of this estimator is known to scale as $1/\sqrt{N}$, with N the number of markers. This statistical error is called 'numerical noise'. The main problem for time-evolution nonlinear gyrokinetic simulations is that the noise level can accumulate indefinitely [18] unless special countermeasures are enforced. Another problem, known as the *entropy paradox*, is that the numerical fluctuation entropy can increase indefinitely, even if lower order moments such as the energy or the heat flux may *seem* to have converged [19]. In PIC codes, this entropy increase is accompanied by an increase in noise. The solution to this problem is demonstrated in [20].

The necessary steps to overcome the difficulty are described below.

(i) Numerical convergence of the results with increasing N should be demonstrated. While this may seem a trivial requirement, there are two main difficulties. First, the convergence of the error is rather slow ($1/N^{1/2}$) and the computing cost becomes prohibitive if no other measures are taken. Second (and this problem is faced by *all* numerical methods), the turbulence behaviour is most often accompanied by bursts and avalanches, which are of chaotic nature. A signature of chaos is found when infinitesimal differences between two states lead, after a finite time, to markedly different evolutions, in particular a different sequence of bursts. Strict numerical convergence, i.e. obtaining the same time evolution in the limit of increasing numerical resolution, is impossible to achieve after a finite time in chaotic systems. Figure 2 shows such a study performed with the ORB5 code [10] for Cyclone-base case parameters and without source terms. Clearly, whereas strict numerical convergence of heat flux can be shown until $t \approx 0.7 \times 10^4 \Omega_{\mathrm{i}}^{-1}$, it is not the case at later times. Note that agreement, in a time-average sense, of the results in the last half of the simulation neither implies strict numerical convergence nor that the noise problem is solved: this is simply due to the fact that, in the absence of sources, the profiles relax to

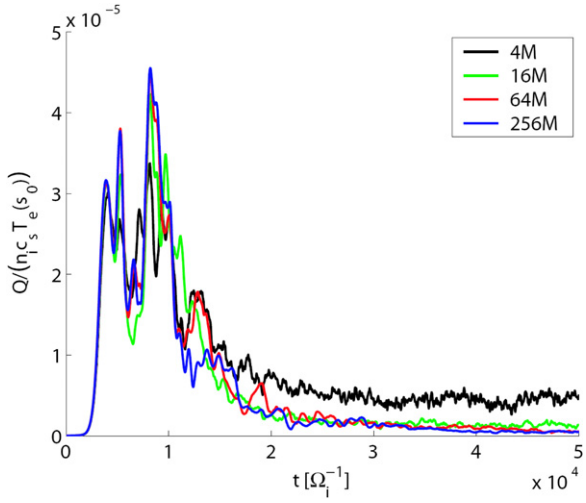

Figure 2. Radial heat flux versus time for different numbers of markers, in global simulations without sources using the ORB5 code. More details in [10].

their nonlinear marginal state with a decaying turbulence level. We have found that such 'decaying' simulations sooner or later always end up noise dominated.

(ii) Reducing the error on estimates can be achieved by variance reduction techniques. For example, *importance sampling* consists of finding an appropriate distribution of the marker positions. This 'optimized loading' scheme is exposed in [17]. In particular, it is shown [21] that it leads to a large improvement over the simple 'proportional loading'. Another way is to apply *control variate* techniques. The idea is to represent as little as possible of the full distribution function f with markers and integrate the rest analytically. The δf scheme consists in writing $f = f_0 + \delta f$, with f_0 a time independent given function of phase space coordinates and δf sampled with markers. The sampling noise error is thus reduced by the factor $|\delta f|/f$, which is usually much smaller than unity for core plasma turbulence. An improved variation of the δf scheme, called 'adjustable control variates', has been devised and applied to electromagnetic gyrokinetic simulations and resulted in a reduction of the required number of markers by orders of magnitude [22].

(iii) Knowledge of the underlying physics can inspire other major improvements. For example, microturbulent modes are characterized by very small parallel wavenumbers, $|k_\parallel|\rho_s \lesssim \rho_*$, as a consequence of gyrokinetic ordering. Modes that may be present in the numerical representation but do not satisfy this ordering can be considered as unphysical. Thus, it is better to remove them from the numerical representation before accumulating significant levels of numerical noise. A very efficient way to do so is to apply a *field-aligned Fourier filter*: it is shown in [10] that orders of magnitude improvement can be gained, because the signal to noise ratio depends on the number of markers per Fourier mode retained in the filter (instead of per grid cell). This distinction between *a priori* unphysical modes and modes with physical content is at the basis of a diagnostic of the signal to noise ratio which is very practical to implement and very instructive as to the quality of the simulation. It is found empirically (see, e.g. [4]) that *a signal to noise ratio larger than 10* is necessary. It is to be noted that unfiltered (in the parallel direction), noisy simulations show a *larger* transport level than filtered, low-noise simulations.

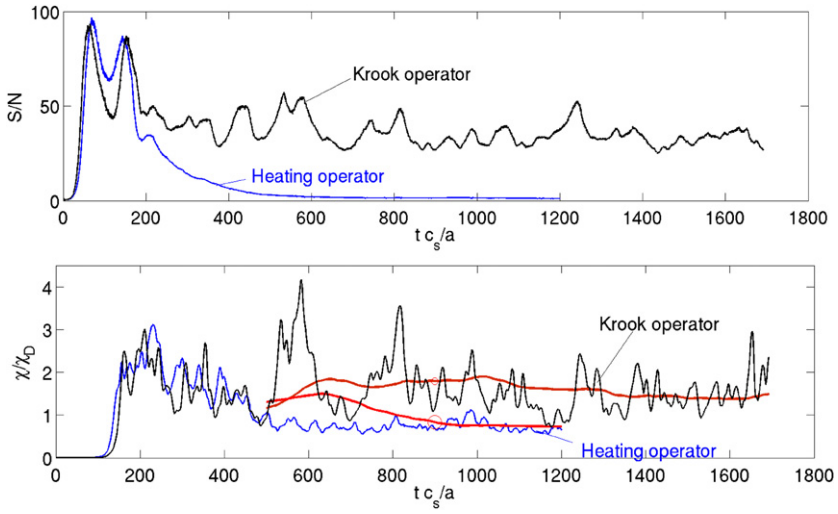

Figure 3. Signal to noise ratio (top) and heat diffusivity (bottom) versus time. One simulation with noise control modified Krook operator, the other with heating operator but without noise control. The thick red lines in the bottom plot are moving time-averaged values over a time window of 500 a/cs. For more details see [23].

(iv) In order to obtain a turbulent steady-state solution sufficiently long simulations must be performed. Thus there is the necessity to *control* the noise level over long time scales. Moreover, in global simulations profile relaxation occurs and sources are necessary to maintain the system above its marginal nonlinear gradient. A noise-control algorithm using modified Krook operators [23] was successfully applied. Figure 3 shows two simulations with the ORB5 code, one using a Krook-based noise control operator, the other using a heating operator without noise control. In the latter case, when the signal/noise ratio drops well below $\sim$10, i.e. for $t \gtrsim \sim 500\ a/cs$, the bursty nature of the transport is lost and the time-averaged heat flux gets *lower* than for the noise-controlled case.

(v) It has been demonstrated that a finite amount of dissipation is required in order to reach a statistically converged state [24]. Hence, strictly speaking, collisionless gyrokinetics cannot reach such a state. In grid-based codes (Eulerian or semi-Lagrangian), numerical dissipation can be introduced by (1) using dissipative discretization or interpolation schemes [12, 25–28]; (2) adding hyperdiffusion terms [29]. Finite dissipation can be introduced in Lagrangian-PIC codes by (1) application of modified Krook operators [23]; (2) coarse graining techniques where smoothing of the marker weights is applied in small portions of phase space [30]. It is to be noted, however, that Lagrange-PIC methods have finite numerical collisionality. Even in a δf model, the system can be collisionless only in the zero weight limit [31]. A good way to check the simulation results, besides the monitoring of the signal/noise ratio, is a verification of the entropy production rate [24, 32, 33], which can be shown to vanish on the long time scales [20].

Even with all these noise control measures taken, it is empirically found that 100–300 markers per radial cell per Fourier mode are typically required to obtain sufficiently accurate results. This is comparable in order of magnitude to the requirement in the number of velocity space grid points for Eulerian and semi-Lagrangian schemes: 8 to 16 μ values and 32 to 128 $v_\parallel$ values.

Let us now focus on the geometrical issue. It was recognized early on that the geometry of the equilibrium magnetic configuration plays an important role to determine the behaviour of microinstabilities and turbulence. Thus, a tentative explanation for the discrepancy between the results of [2, 3] was that it was due to finite aspect ratio effects which were accounted for differently in the different codes. A recent work [34] has examined this issue in detail. The problem comes from the implementation of a simplified geometry using the so-called 's–α' model which, in the $\alpha = 0$ limit, reduces to circular concentric magnetic surfaces and in which the straight-field line coordinate, θ_*, is approximated by the geometrical poloidal angle θ, which is true only in the infinite aspect ratio (i.e. cylinder) limit. On the other hand, the magnetic field variation in the poloidal direction is retained so as to keep the magnetic mirror effect: $B = B_0/(1 + \varepsilon \cos\theta)$. This results in a magnetic configuration that is *inconsistent at first order in the inverse aspect ratio* $\varepsilon = a/R_0$. For example, the magnetic field strength computed from the metric coefficients of the s–α model ($\theta \approx \theta_*$) is $B = B_0$. For Cyclone-base case parameters [35] this results in significative differences between the s–α model (at $\alpha = 0$) and a model with circular concentric magnetic surfaces treated with consistent metric: the maximum ITG growth rate is lower by 30% (figure 4, left). Note that using a true ideal MHD equilibrium with circular boundary at $\beta \approx 0$ results in only a small difference (about 6%) with the circular concentric case with consistent metric. Thus, the main source of the difference is the inconsistency in the metric, and not a true geometrical effect, nor the fact that a circular concentric configuration is only an approximation to the true ideal MHD equilibrium. (Indeed, the MHD case does have, even at zero β, a finite Shafranov shift which, for these parameters, is slightly destabilizing.) These differences persist in the nonlinear regime, and are even larger: ion heat diffusivities are lower by a factor of 2 when using the inconsistent (s–α) model (figure 4, right), whereas the differences between a circular concentric model with consistent metric and an ideal MHD $\beta = 0$ equilibrium are not significant.

There is a fundamental problem with inconsistent models such as s–α. We have actually checked, using the GENE code [36–40] that the results obtained do depend on the detailed implementation of the coefficients of the gyrokinetic equations. More precisely, we have performed simulations in which some of the terms were expressed using the consistent metric, and some other terms used the approximated metric: the results were all different, depending on the choice of the terms. The largest differences were obtained, however, for terms related to the parallel direction of the equilibrium magnetic field; more precisely, the straight-field line coordinate approximation $\theta_* \approx \theta$ appears as being responsible for most of the difference between s–α and a correct consistent model.

In order to determine the part of the finite ρ_* effects due to linear physics, a comparison of flux-tube results is made with the global code GYGLES [41] using a circular concentric magnetic equilibrium and consistent metric (figure 4, left). We note that the growth rates from the global result at $\rho_* = 180$ are very close to those of the infinite size case using the inconsistent (s–α) metric. This relatively good apparent agreement is, however, fortuitous, and may have led to the wrong conclusion that finite ρ_* effects are not due to linear dynamics.

We now consider sequences of finite size cases with varying ρ_*. First, we try to identify linear global effects of profile shearing by studying a series of ion temperature gradient profiles of varying shape (figure 5, left):

$$\frac{R}{a}\frac{\mathrm{d}\log T}{\mathrm{d}\rho} = -\left(\frac{R}{L_T}\right)_{\max} \cosh^{-2}\left(\frac{\rho-\rho_0}{\Delta}\right), \tag{3}$$

with $\rho = r/a$ the normalized minor radius of the (circular concentric) magnetic surface, $(R/L_T)_{\max}$ is the maximum value of R/L_T, ρ_0 is the radial position of the maximum R/L_T and Δ is a normalized width parameter. The linear ITG growth rates are shown in figure 5,

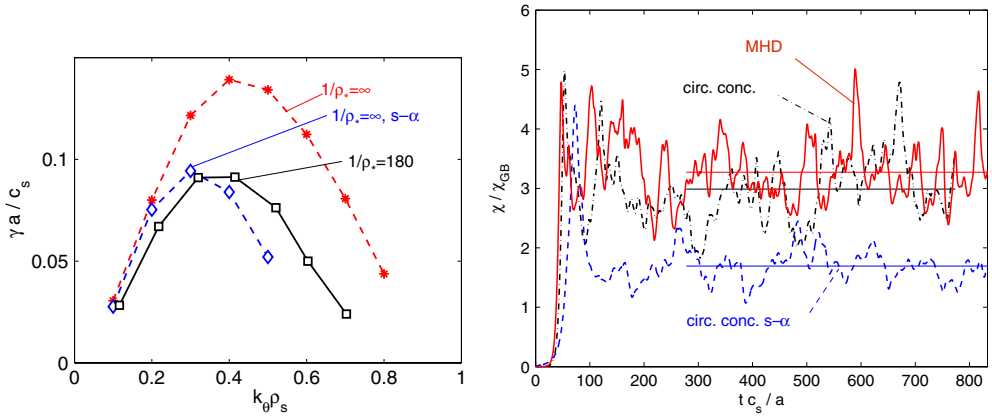

Figure 4. Left: linear growth rates of ITG modes computed using an equilibrium with circular concentric magnetic surfaces. Flux-tube limit ($1/\rho_* = \infty$) with consistent (stars) and inconsistent (s–α) metric (diamonds) computed with GENE. Global result ($1/\rho_* = 180$) computed with GYGLES. Right: ion heat diffusivity computed in the flux-tube limit with GENE using an MHD equilibrium (solid line), circular concentric with consistent metric (chain line) and inconsistent (s–α) metric (dashed). See [34] for more details.

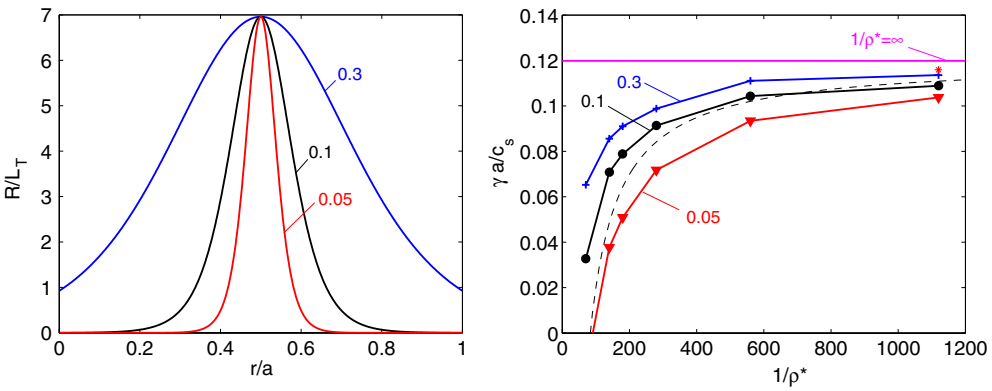

Figure 5. Left: radial profiles of normalized temperature gradient inverse scale length R/L_T for different values of the width parameter $\Delta = 0.05, 0.1, 0.3$. Right: growth rates computed with the global code GYGLES as a function of system size $1/\rho_*$ for different values of the temperature gradient profile parameter Δ, for $k_\theta \rho_s = 0.3$. The horizontal line $1/\rho_* = \infty$ was computed with the flux-tube version of the GENE code. The dashed line is an analytical estimate (see text).

right. This confirms that the global linear results converge to the infinite size (flux-tube) limit, and that the way they do so depends on the temperature gradient profile width. In order to understand this, we show in figure 6 (left) the characteristic radial mode number $k_r = 2\pi/\Delta r$, where Δr is the width of the mode amplitude envelope. Clearly, the finite system size has the effect of limiting the radial wavelength of the mode. As a matter of fact, this effect depends on the profile width parameter Δ. This increase in k_r with decreasing system size and profile width is partly responsible for the decrease in growth rate, see figure 6, right, where we compare the global results with flux-tube results at finite radial wavenumber. Another explanation is based on the works of [42, 43] using ballooning representation, which show that the radial extent of a global mode is a compromise between micro- and macroscopic scales. This leads

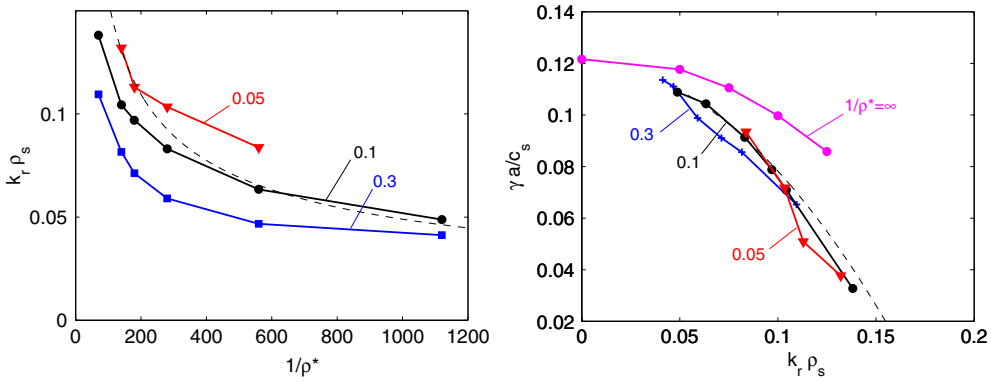

Figure 6. Left: characteristic radial wavenumber as a function of system size $1/\rho_*$. Right: growth rate as a function of the characteristic radial wavenumber for different temperature profiles labelled with the value of Δ. The curve labelled $1/\rho_* = \infty$ is computed using the flux-tube version of the GENE code at finite k_r. The dashed lines are analytical estimates (see text).

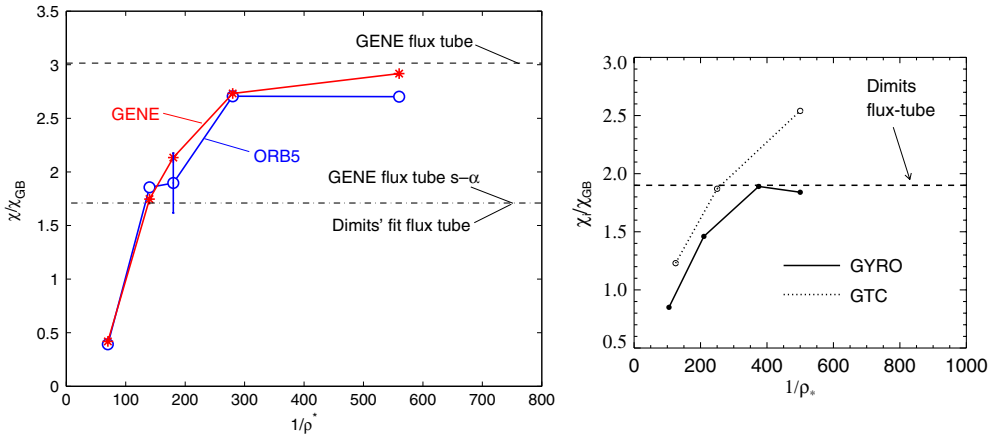

Figure 7. Ion heat diffusivity computed using the global codes GENE (stars, left) ORB5 (circles, left), GYRO (solid, right) and GTC (dotted, right) versus normalized system size. Flux-tube results computed with GENE with consistent (dashed, left) and inconsistent (s–α) metric (chain, left). The flux-tube results from Dimits' fit, equation (1) of [35], is also indicated. See [3, 40, 44] for more details. (Part (b) is reprinted with permission from [3]. Copyright 2004 American Institute of Physics.)

to $k_r\rho_s \sim C^{-1/4}\sqrt{k_\theta\rho_s\hat{s}}\sqrt{\rho_*/\Delta}$ and $\gamma \sim \gamma_{\max}[1 - (C^{1/2}/2k_\theta\rho_s\hat{s})(\rho_*/\Delta)]$, where $\hat{s}$ is the magnetic shear and C is a parameter describing the quadratic dependence of the growth rate on the ballooning angle. Analytical estimates for $\Delta = 0.05$, $\hat{s} = 0.8$, $C = 4$, are presented by the dashed lines in figures 5 and 6. Note that for small enough Δ it is equivalent to vary Δ at fixed ρ_* and to vary ρ_* at fixed Δ, but for too large Δ values the plasma and geometrical parameters vary substantially over the radial extent of the unstable region and global variations play a role. In summary, finite ρ_* 'confines' the radial extent of the mode, which implies finite k_r, which is stabilizing. Thus, ρ_* scaling is, at least in part, attributable to radial linear dispersive effects.

Finally, we show in figure 7 the results of nonlinear simulations using both the Eulerian code GENE [36–40] and the Lagrangian PIC code ORB5 [10, 23]. The initial temperature

profiles are specified as

$$\frac{R}{a}\frac{\mathrm{d}\log T}{\mathrm{d}\rho} = -\left(\frac{R}{L_T}\right)_{\max}\left[1-\cosh^{-2}\left(\frac{\rho-\rho_0-w/2}{\Delta}\right)-\cosh^{-2}\left(\frac{\rho-\rho_0+w/2}{\Delta}\right)\right] \quad (4)$$

for $|\rho-\rho_0| < w/2$, and 0 otherwise. The parameters used for this series of runs are $(R/L_T)_{\max} = 7.1$ and 7.5, $w = 0.8$, $\Delta = 0.04$, $\rho_0 = 0.5$, $T_\mathrm{e} = T_\mathrm{i}$, and similar density gradient profile with $(R/L_n)_{\max} = 2.2$. Source terms are included in these global simulations [23], allowing for some profile relaxation to occur but maintaining the averaged temperature gradient above the nonlinear threshold. (The error bar appearing in the figure will be explained later.) Because of profile relaxation, the time-averaged temperature gradient is different from the initial temperature gradient given as input. Therefore, for each ρ_* case, two simulations are made with slightly different values of initial $(R/L_T)_{\max}$, and then the time-averaged heat diffusivities are interpolated at $R/L_T = 6.96$ between these two runs using the values of the time-averaged temperature gradients and heat fluxes obtained in the late, quasi-steady phase of the simulations. A remarkable feature is the excellent agreement between both code results. Being based on entirely different numerical schemes, and potentially subject to numerical errors of a totally different nature, this gives confidence in the validity of the results.

The way the gyro-Bohm scaling is broken by finite size effects depends on the width of the temperature profile gradient. In a recent work [44], it was found that the finite size scaling is best expressed in terms of an 'effective ρ_*' defined as $\rho_*^{\mathrm{eff}} = \rho_s/w$, where w is the *width* of the temperature gradient profile. Thus, finite size effects may well be important even in ITER-size plasmas, for example in internal transport barriers.

To conclude this section, it is shown that Lagrangian-PIC and Eulerian global numerical schemes are both giving the same results, including for finite ρ_* effects, when both numerical (in particular the noise problem) and geometrical aspects are correctly treated. In particular, finite size simulations (both linear growth rates and nonlinear heat diffusivities) converge to the flux-tube result always from below (in gyro-Bohm units) in the limit $\rho_* \to 0$. Another consequence of this study is that code benchmark comparisons should not be made using models that have inconsistencies, in particular on the geometrical aspects. Note that our results (ORB5 and GENE) compare well with GTC results published in [2]: agreement is within 15%, which is the uncertainty we estimate due to the chaotic nature of the system (see next section). On the other hand, while there is quantitative disagreement with the GYRO results, which we attribute to the inconsistency in the $(s$–$\alpha)$ model used, qualitative agreement is obtained in that finite size heat transport reaches the infinite system size (flux-tube) result from below in the appropriate $1/\rho_* \to \infty$ limit. Thus, earlier controversies have been solved.

4. Bursty nature of turbulent transport

It is now well established that turbulent transport (at least in the ITG regime) is characterized by bursts and avalanches [5–8]. The presence of these is visible on time traces of heat diffusivity (see e.g. figure 3). As previously mentioned, a fundamental problem arises due to the chaotic nature of the sequence of bursts, which poses practical problems for the demonstration of numerical convergence (see figure 2). The exponential divergence, after a finite time, of any pair of simulations with infinitesimally small differences in the initial condition [45] has drastic consequences. It is sometimes observed that even tiny round-off errors are sufficient to cause problems: moreover, the parallelization of the algorithms implies that the order of the operations (e.g. for a global sum) is run-time dependent and thus produces different round-off errors that in some cases may lead to different results after a finite simulation time. Due to the intrinsic chaotic nature of the processes the reproducibility of the numerical simulations

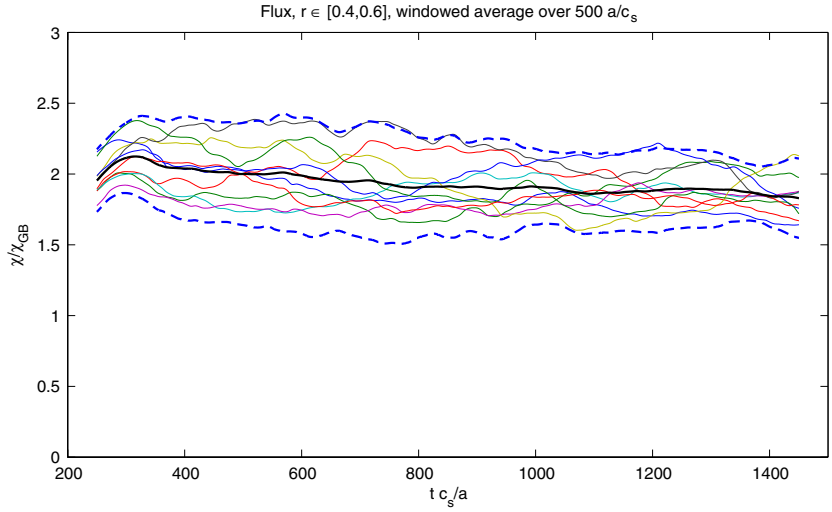

Figure 8. Heat diffusivities versus time, radially averaged over $r = 0.5 \pm 0.1a$, time averaged over $t \pm 250a/c_s$, for a set of 10 ITG Cyclone ORB5 global simulations differing only in their initial condition. The thick line is the ensemble average and the thick dashed lines are the average $\pm 2\sigma$ with σ the standard deviation of the ensemble.

cannot be guaranteed. The equivalence of two simulations can thus only be understood in a statistical sense.

This immediately raises the following questions: (i) How long a numerical simulation must be in order to obtain a reliable estimate of transport coefficients? (ii) How reliable are the results of a given simulation, i.e. how to estimate a confidence interval? (iii) Given two different simulations (e.g. obtained by different codes based on the same physics, or obtained by the same code but including different physics), are the observed differences *significant*?

A systematic way to answer these questions is to perform an *ensemble of independent simulations* with the same code, the same physics, but starting from different initial conditions. Ensemble averages of these simulations and statistical analysis of this set of results should produce unbiased estimates of the average, standard deviation, etc. In particular, it is expected that the standard deviation thus obtained will give information on the intrinsic variability of the turbulent processes. In principle, a large number of such independent simulations should be performed. In practice, this is almost never done, due to the prohibitive computing resources required.

We have performed a set of ten ITG Cyclone-base case [35] simulations with the ORB5 code with the same parameters as for figure 7 at $\rho_* = 1/180$. All these simulations are noise controlled and run up to $t = 1700a/c_s$. The simulations differ only in the initialization: a small amplitude perturbation is applied at $t = 0$ consisting in a sum over a set of Fourier modes. The difference between these simulations is only in the values of the individual Fourier initial amplitudes. Figure 8 shows the time traces of heat diffusivity radially averaged at $r/a = 0.5 \pm 0.1$ and time window averaged over $500a/c_s$. The bold line is the time trace of the ensemble average heat diffusivity and the dashed lines are the average $\pm 2\sigma$, with σ the standard deviation over the ensemble. Clearly, while individual simulation results do differ from each other, the differences should not be significant. Thus $\pm 2\sigma$ is a reasonable estimate (i.e. 95% confidence interval if the distribution were Gaussian) of an 'error bar' due to the intrinsic variability caused by the underlying chaotic behaviour. Performing time averages

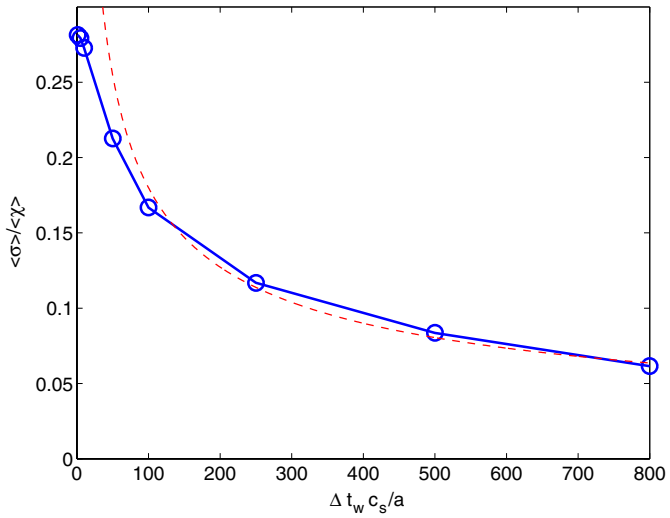

Figure 9. Standard deviation (normalized to the average heat diffusivity) versus time span Δt_w of the averaging window. Simulation data taken in the time interval $tc_s/a \in [600, 1700]$. The dashed line is the fit $1.8/\sqrt{\Delta t_w c_s/a}$.

over the last half of the simulation gives an average diffusivity $\langle\chi\rangle/\chi_{\mathrm{GB}} = 1.88$, an average standard deviation $\langle\sigma\rangle/\chi_{\mathrm{GB}} = 0.14$ (thus a relative error bar at 95% confidence of $\pm 15\%$).

The question of whether or not the simulations are long enough is delicate. The problem is that turbulent time scales extend from the frequencies of the highest mode number involved down to the zero frequency, meaning there is no upper limit on the time scale, and moreover the chaotic nature is reflected in an aperiodic behaviour: it is thus not easy, from first principles, to establish without any ambiguity a way of finding an appropriate time-averaging procedure. Thus in practice one still relies on an empirical approach. The last half of the simulations of figure 8 can be in a certain sense considered as having reached a steady-state: all curves lie within the range $\langle\chi\rangle \pm 2\langle\sigma\rangle$. Another way is to perform statistical tests: considering two separate time windows centred at $tc_s/a = 950$ and $tc_s/a = 1450$, regression analysis gives a relative drop in flux of 4% with a standard deviation of 8%: thus the drop is not significant. The averaged ensemble standard deviation $\langle\sigma\rangle$ does depend on the time span Δt_w of the moving time-average window. Figure 9 shows this quantity, using data over the interval $tc_s/a \in [600, 1700]$, versus Δt_w. An asymptotic dependence is found as $\langle\sigma\rangle/\langle\chi\rangle \approx 1.8/\sqrt{\Delta t_w\, c_s/a}$. This gives information as to the simulation time required to reach a given accuracy (to which a time of roughly $tc_s/a \approx 600$ must be added to pass the initial transient phase). Note that performing a single, 10 times longer simulation (instead of the ensemble of 10 simulations) would not have given this information. Moreover, this poses questions about the ergodicity assumption: *a priori*, it is not guaranteed that a single simulation would span enough of the possible states of the system.

More realistic global gyrokinetic simulations are performed in which a specified input power is given to the system. In contrast to fixed-gradient simulations, the heat flux is expected to reach a constant time-averaged value determined by the input power, while the plasma profiles are freely evolving. Ideally, all simulations with the same input power but starting from different initial conditions should reach a steady state with the same heat flux and a similar temperature profile. Also, the closer the system is to marginal state, the longer it takes to reach

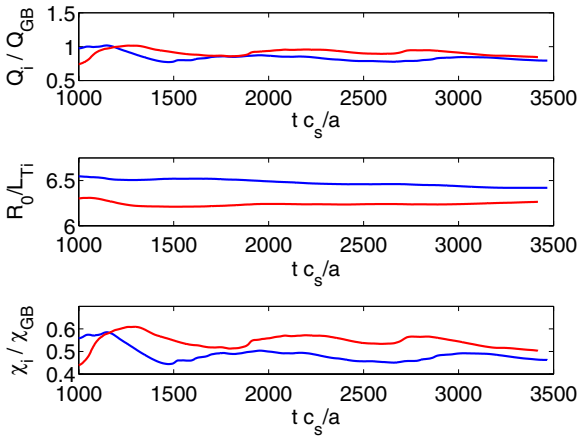

Figure 10. Time traces of the ion heat flux (top), normalized ion temperature gradient (middle) and ion heat diffusivity (bottom) for a pair of simulations using the GT5D code run at fixed input power. Quantities are time averaged between $t = 500c_s/a$ and time t and radially averaged between $r/a = 0.5$ and $r/a = 0.8$, which is in the source-free region.

such a steady-state. Examples of power-driven simulations using the GT5D code [6] are shown in figure 10 where time traces of the ion heat flux, ion temperature gradient scale length and ion heat diffusivity are represented. The quantities are time averaged between time $t_0 = 500\,a/c_s$ and time t, and radially averaged between $r/a = 0.5$ and $r/a = 0.8$, which is in the source-free region of the simulation. The parameters are standard Cyclone-base case except the system size $\rho_* = 1/135$. The input power is 2 MW and the initial temperature gradient is $R/L_T = 10$. The normalization for the heat flux is $Q_{\mathrm{GB}} = m_{\mathrm{i}} n_{\mathrm{i}0} v_{t\mathrm{i}}^3/\rho_*^2$. These simulations show that, in spite of differences of the order of 20% on the time-averaged heat diffusivity, the system is very close to having reached steady state. Note that the temperature gradient settles around $R/L_T = 6.35$, which is much closer to the marginal point $R/L_T \approx 6.0$ than the simulations of figure 8 at $R/L_T = 6.96$. The differences between the two simulations, obtained with two different initializations, are due to the intrinsic variability due to the presence of chaos in the system. Of course with only two simulations it does not make sense to make ensemble averages. On the other hand, we can perform a statistical analysis of these simulations, for example looking at the probability distribution functions (PDFs) of the local heat flux.

These PDFs of the local instantaneous heat flux (figure 11) show a clear non-Gaussian behaviour with exponential tails, which reflects the presence of large, rare events. (Note that this does not mean that the density and temperature fluctuations themselves are non-Gaussian: large tails in the flux PDF can occur when these are cross-correlated [46].) The ORB5 PDF has been obtained using data from all 10 simulations of the ensemble in the time interval $tc_s/a = 500$–1700. The GT5D PDFs of two separate simulations are shown using data in the time interval $tc_s/a = 1000$–3400. In figure 11, the ORB5 simulations at $R/L_T \approx 6.96$ (shown in figure 8) and the GT5D simulations at $R/L_T \approx 6.35$ (shown in figure 10) show a similar behaviour but with some differences: we note that the GT5D simulations are closer to the critical gradient and have higher exponential tails compared with the ORB5 simulations at a higher temperature gradient. This indicates that the relative importance of large events is increasing as the critical gradient is approached. The two GT5D simulations, obtained with different initializations, exhibit almost the same PDF: this is a good indication that the simulations have converged in the statistical sense. Using the semi-Lagrangian code GYSELA, non-Gaussian PDFs of the

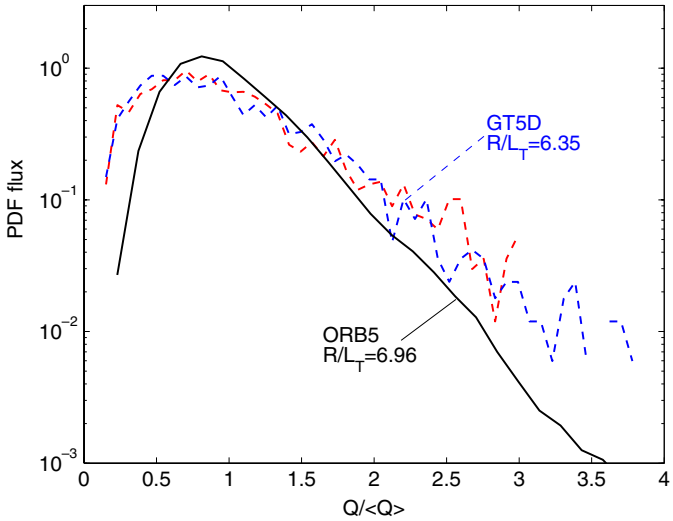

Figure 11. Probability distribution functions of the local heat flux for the simulations shown in figures 8–10 using ORB5 ($r/a = 0.5$, $R/L_T = 6.96$) and GT5D ($r/a = 0.7$, $R/L_T = 6.35$). The two GT5D curves are obtained with two different initializations. Exponential tails due to large events are clearly visible.

heat flux were also found, although the results seem to indicate a lower non-Gaussian content. The different regimes considered by these three codes so far prevent us to make a quantitative meaningful comparison.

5. Non-adiabatic electrons

Global gyrokinetic simulations of turbulence have now progressed to include more physical effects, such as non-adiabatic electron response and electromagnetic effects. For the ETG case, it is usually assumed that the ion response is treated adiabatically, although recent works have shown that this assumption has limitations: different results are obtained when both ions and electrons are treated kinetically [47–49].

The simultaneous kinetic treatment of both ion and electron species poses great challenges to the numerical simulations. The characteristic timesteps must be reduced by a factor of the order of the square root of the mass ratio if constrained by the parallel dynamics (though in practice the required factor is often less than that). Moreover, it has been found that, even in ITG-dominated cases, another radial space scale is involved: the electron response is strongly non-adiabatic in narrow regions centred around mode-rational $q = m/n$ surfaces. This is due to passing electron dynamics. The width of the non-adiabatic regions can be estimated from the regions for which $|\omega/k_\parallel v_{te}|$ is larger than unity for the corresponding (m, n) mode. This results in a width $\Delta r \approx 2Rq\,|\omega|/(|\hat{s}\,k_\theta|v_{te})$ which depends on the magnetic shear $\hat{s}$ and can be smaller than the ion Larmor radius. For $\omega \approx \omega_{*,T} = (c_s/R)(R/L_T)k_\theta\rho_s$ we can obtain an approximate expression for the width in units of the Larmor radius:

$$\frac{\Delta r}{\rho_s} \approx 2\frac{R}{L_T}\sqrt{\frac{m_\mathrm{e}}{m_\mathrm{i}}}\,\frac{q}{\hat{s}} \tag{5}$$

which gives e.g. for Cyclone parameters $\Delta r/\rho_s \approx 0.4$ in deuterium. Consequently, the radial grid resolution has to be increased typically by up to one order of magnitude as compared with

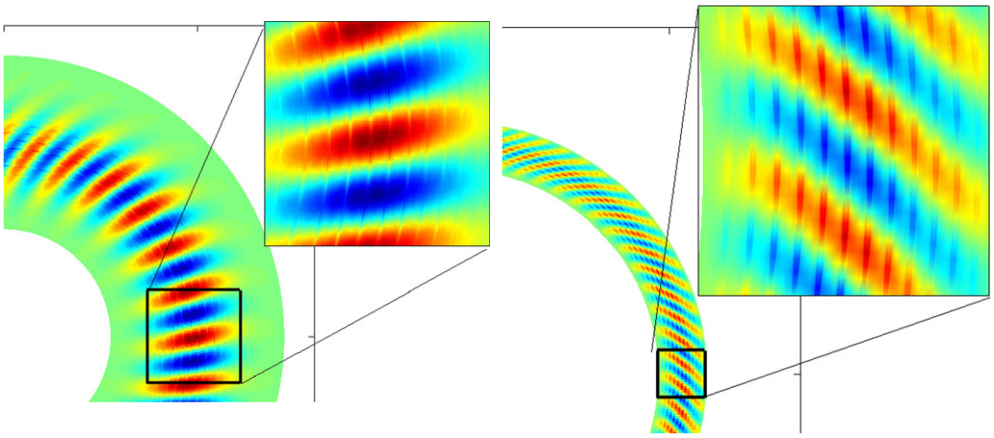

Figure 12. Contours of the perturbed electrostatic potential for global linear simulations with the GENE code including non-adiabatic electron response. The sharp radial features are numerically resolved and correspond to mode-rational surfaces. An artificially enhanced electron mass has been used, $m_i/m_e = 400$. The mode on the left has $k_\theta \rho_s = 0.345$ and is ITG-driven. The mode on the right has $k_\theta \rho_s = 0.776$ and is a TEM.

the adiabatic electron case. In order to save computational resources, simulations are often run using an artificially decreased mass ratio. An example of linear global calculation using the GENE code is shown in figure 12 for Cyclone parameters $(R/L_{T_e})_{\max} = (R/L_{T_i})_{\max} = 6.96$, $(R/L_n)_{\max} = 2.23$, profiles given by equation (3) with $\Delta = 0.3$ and $\rho_0 = 0.5$, $q = 0.85{+}2.2\rho^2$, $\beta_e = 10^{-3}$, $m_i/m_e = 400$. Both in ITG (left) and TEM (right) regimes, a resonant behaviour is observed at mode-rational surfaces. The width Δr is within 30% of the analytical estimate, equation (5).

It is not yet clear how far the sharp mode-rational behaviour of the electron non-adiabatic response affects the nonlinear turbulence results. Nonlinear electromagnetic global simulations with gyrokinetic ions and non-adiabatic (drift-kinetic) electrons with $m_i/m_e = 1000$ have been performed using the ORB5 code [50]. The parameters are $1/\rho_* = 160$, $R_0/a = 4$, $(R/L_{T_e})_{\max} = (R/L_{T_i})_{\max} = 10$, $\rho_0 = 0.6$, $\beta_e = 3 \times 10^{-3}$, $N_i = N_e = 512 \times 10^6$ markers. A timestep 20 times smaller than for the adiabatic electron case had to be taken and the radial resolution was increased from 128 to 512 grid points in order to avoid numerical instability. Figure 13 shows the comparison of three simulations using three different models: electrostatic with all adiabatic electrons, electrostatic with drift-kinetic trapped electrons and adiabatic passing electrons and electromagnetic with gyrokinetic ions and drift-kinetic electrons. For these parameters the dominant underlying instability is ITG driven. We can observe that the largest effect of the non-adiabatic electron response is due to trapped electrons, and that the effect of non-adiabatic passing electrons on the turbulent heat flux does not appear to be significant in this case. However, this conclusion is not necessarily applicable to all situations, e.g. at higher β. Also, the role of collisions should not be forgotten. Clearly, more work is required in this area.

6. Conclusion

Gyrokinetic simulations have progressed to the stage where detailed quantitative comparisons between different physical models and different numerical approaches can be made. We have

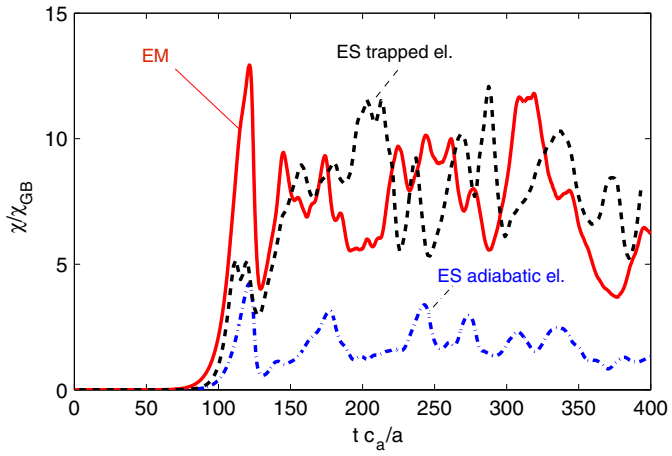

Figure 13. Time traces of heat diffusivity computed with the ORB5 code using three different models: electrostatic with all adiabatic electrons (chain line), electrostatic with drift-kinetic trapped electrons and adiabatic passing electrons (dashed line) and electromagnetic with gyrokinetic ions and drift-kinetic electrons (solid line). See [50]. (Reprinted with permission from [50]. Copyright ©2010 IEEE.)

shown in particular how the numerical noise problem in Lagrange-PIC simulations could be cured, and how important it is to use consistent geometric approximations. System size (ρ_*) scaling studies have confirmed that the heat diffusivity at finite size converges (always from below) towards the flux-tube limit, which is gyro-Bohm. A global linear analysis has shown that part of the finite ρ_* scaling can be attributed to linear dispersive effects.

The chaotic nature of turbulence underlying anomalous transport and its implication for numerical simulations have been studied. An ensemble of independent simulations has allowed us to obtain information on the intrinsic variability due to chaos and thus to obtain confidence intervals on the estimates of the average quantities. Global simulations at fixed input power require longer simulations than fixed-gradient simulations, especially close to the marginal stability, where it was shown that the transport is increasingly due to the presence of large events (bursts and avalanches), leading to PDFs with long exponential tails.

Issues related to non-adiabatic electrons, in particular in the vicinity of mode-rational surfaces, have also been addressed. The short, shear-dependent radial scale involved implies an increase in computational cost by at least an order of magnitude compared with the adiabatic electron case. More work is required, in particular concerning finite β cases.

Acknowledgments

Fruitful discussions with T H Watanabe are gratefully acknowledged. This work was partly supported by the Swiss National Science Foundation and by the HP2C project. Computations were performed on the BG/P at IDRIS in the framework of a DEISA-Extreme Computing Initiative, on the HPC-FF JUROPA server at the Jülich Rechenzentrum under EFDA, on the BG/P of EPFL under the CADMOS project, on the Cray XT5 at the CSCS and on the BX900 at the JAEA. The support of the EFDA-HLST is gratefully acknowledged. Some of the authors were partly supported by the European Communities under the contracts of Association between EURATOM and Switzerland, CEA and the Max-Planck Institut für Plasmaphysik, respectively, within the framework of the European Fusion Development

Agreement. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

[1] Garbet X, Idomura Y, Villard L and Watanabe T H 2010 *Nucl. Fusion* **50** 043002
[2] Lin Z, Ethier S, Hahm T S and Tang W M 2002 *Phys. Rev. Lett.* **88** 195004
[3] Candy J, Waltz R E and Dorland W 2004 *Phys. Plasmas* **11** L25
[4] Bottino A *et al* 2007 *Phys. Plasmas* **14** 010701
[5] McMillan B F *et al* 2009 *Phys. Plasmas* **16** 022310
[6] Idomura Y, Urano H, Aiba N and Tokuda S 2009 *Nucl. Fusion* **49** 065029
[7] Sarazin Y *et al* 2010 *Nucl. Fusion* **50** 054004
[8] Candy J and Waltz R E 2003 *Phys. Rev. Lett.* **91** 045001
[9] Brizard A J and Hahm T S 2007 *Rev. Mod. Phys.* **79** 421
[10] Jolliet S *et al* 2007 *Comput. Phys. Commun.* **177** 409
[11] Lederer H *et al* 2008 *Adv. Parallel Comput.* **15** 713
[12] Grandgirard V *et al* 2007 *Plasma Phys. Control. Fusion* **49** B173
[13] Waltz R E *et al* 2002 *Phys. Plasmas* **9** 1938
[14] Garbet X and Waltz R E 1996 *Phys. Plasmas* **3** 1898
[15] Lin Z and Hahm T S 2004 *Phys. Plasmas* **11** 1099
[16] Aydemir A Y 1994 *Phys. Plasmas* **1** 822
[17] Hatzky R, Tran T M, Könies A and Kleiber R 2002 *Phys. Plasmas* **9** 898
[18] Nevins W M *et al* 2005 *Phys. Plasmas* **12** 122305
[19] Watanabe T H and Sugama H 2002 *Phys. Plasmas* **9** 3659
[20] Jolliet S *et al* 2009 *Phys. Plasmas* **16** 052307
[21] Villard L *et al* 2004 *Nucl. Fusion* **44** 172
[22] Hatzky R, Könies A and Mishchenko A 2007 *J. Comput. Phys.* **225** 568
[23] McMillan B F *et al* 2008 *Phys. Plasmas* **15** 052308
[24] Krommes J A and Hu G 1994 *Phys. Plasmas* **1** 3211
[25] Candy J and Waltz R E 2006 *Phys. Plasmas* **13** 032310
[26] Grandgirard V *et al* 2008 *Commun. Nonlinear Sci. Numer. Simul.* **13** 81
[27] Dif-Pradalier G *et al* 2009 *Phys. Rev. Lett.* **103** 065002
[28] Garbet X *et al* 2009 *Phys. Plasmas* **16** 062503
[29] Pueschel M J, Dannert T and Jenko F 2010 *Comput. Phys. Commun.* **181** 1428
[30] Chen Y and Parker S E 2007 *Phys. Plasmas* **14** 082301
[31] Idomura Y, Watanabe T-H and Sugama H 2006 *C. R. Physique* **7** 650
[32] Lee W W and Tang W M 1988 *Phys. Fluids* **31** 612
[33] Sugama H *et al* 1996 *Phys. Plasmas* **3** 2379
[34] Lapillonne X *et al* 2009 *Phys. Plasmas* **16** 032308
[35] Dimits A M *et al* 2000 *Phys. Plasmas* **7** 969
[36] Jenko F, Dorland W, Kotschenreuther M and Rogers B N 2000 *Phys. Plasmas* **7** 1904
[37] Dannert T and Jenko F 2005 *Phys. Plasmas* **12** 072309
[38] Merz F 2008 Gyrokinetic simulation of multimode plasma turbulence *PhD Thesis* Technische Universität München
[39] Görler T 2009 Multiscale effects in plasma microturbulence *PhD Thesis* Universität Ulm
[40] Lapillonne X 2010 Local and global Eulerian gyrokinetic simulations of microturbulence in realistic geometry with applications to the TCV Tokamak *PhD Thesis* No 4684, Ecole Polytechnique Fédérale de Lausanne
[41] Fivaz M *et al* 1998 *Comput. Phys. Commun.* **111** 27
[42] Romanelli F and Zonca F 1993 *Phys. Fluids* B **5** 4081
[43] Kim J Y and Wakatani M 1994 *Phys. Rev. Lett.* **73** 2200
[44] McMillan B F *et al* 2010 *Phys. Rev. Lett.* at press
[45] Idomura Y, Tokuda S and Wakatani M 1997 *Comput. Phys. Commun.* **102** 68
[46] Carreras B A *et al* 1996 *Phys. Plasmas* **3** 2664
[47] Nevins W M *et al* 2006 *Phys. Plasmas* **13** 122306
[48] Candy J *et al* 2007 *Plasma Phys. Control. Fusion* **49** 1209
[49] Görler T and Jenko F 2008 *Phys. Rev. Lett.* **100** 185002
[50] Bottino A *et al* 2010 *IEEE Trans. Plasma Science* **99** at press, doi:10.1109/TPS.2010.2055583