Gyrokinetic Particle Simulation of Fusion Plasmas: Path to Petascale Computing*

W. W. Lee, S. Ethier, W. X. Wang, and W. M. Tang

Plasma Physics Laboratory, Princeton University, Princeton, NJ 08543

S. Klasky

Oak Ridge National Laboratory, Oak Ridge, TN

Gyrokinetic particle simulation of fusion plasmas for studying turbulent transport on state-of-the-art computers has a long history of important scientific discoveries. The primary examples are: (i) the identification of ion temperature gradient (ITG) drift turbulence as the most plausible process responsible for the thermal transport observed in tokamak experiments; (ii) the reduction of such transport due to the presence of zonal flows; (iii) the confinement scaling trends associated with size of the plasma and also with the ionic isotope species. With the availability of terascale computers in recent years, we have also been able to carry out simulations with improved physics fidelity using experimentally relevant parameters. Computationally, we have demonstrated that our lead Particle-in-Cell (PIC) code, the Gyrokinetic Turbulence Code (GTC), is portable, efficient, and scalable on various MPP platforms. Convergence studies with unprecedented phase-space resolution have also been carried out. Since petascale resources are expected to be available in the near future, we have also engaged in developing better physics models and more efficient numerical algorithms to take advantage of this exciting opportunity. For the near term, we are interested in understanding some basic physics issues related to burning plasmas experiments in International Thermonuclear Experimental Reactor (ITER) – a multi-billion dollar device to be constructed over the next decade. Our long range goal is to carry out integrated simulations for ITER plasmas for a wide range of temporal and spatial scales, including high-frequency short-wavelength wave heating, low-frequency meso-scale transport, and low-frequency large scale magnetohydrodynamic (MHD) physics on these computers.

I. INTRODUCTION

Particle-in-Cell (PIC) simulation, which began in the late sixties, took advantage of the physics of Debye shielding through the use of finite size particles on a grid to dramatically reduce the numerical noise associated with close encounters between the particles, while leaving intact their long range interactions outside the grid. This approximation also reduced the number of calculations for particle interactions from N^2 to $NlogN$, and greatly reduced the computational time. For simulations of magnetic fusion plasmas, further improvements came in the eighties and nineties with the development of the gyrokinetic particle simulation and perturbatve particle simulation methods. Briefly, under the gyrokinetic approximation the spiral motion of a charged particle is represented as a charged ring centered around its gyrocenter, and perturative methods are used to greatly reduce the discrete particle noise. These improvements along with the first simulation of plasma turbulence in global toroial geometry have driven modern research efforts on microturblence in fusion plasmas with the goal of: (i) understanding the particle and energy transport observed in the tokamak fusion experiments; and (ii) assisting in designing the next generation of fusion reactors.

In this paper, we will outline the fundamental equations associated with these simulation efforts and describe the discoveries made possible through the use of leadership computing facilities for the past dozen years in plasma turbulence research. These include: (i) the identification of ion temperature gradient (ITG) drift instabilities as the primary cause for the observed turbulence spectra in tokamak experiments; (ii) the observation of isotope scaling of confinement; (iii) the importance of the nonlinearly-generated zonal flows in global ITG turbulence simulations; (iv) the favorable transition from Bohm to gyroBohm scaling as a function of the plasma size; and (v) the the role of parallel acceleration of the particles in steady state turbulence. Moreover, we will describe the use of terascale computing capabilities to conduct convergence studies with unprecedented resolution to help resolve the controversial issue of discrete particle noise in PIC codes. Motivated by the promise of using petascale computing in the near future for studying ITER-type plasmas, we will also discuss our effort in developing the capabilities for understanding confinement scaling with size and with isotope species for burning plasma experiments. Finally, the proposed integrated multi-scale gyrokinetic model with disparate frequencies and wavelengths, which can cover a very wide range of physical processes, will be briefly described.

II. GOVERNING GYROKINETIC-POISSON EQUATIONS

The governing equations for the GTC code [1] are based on the gyrokinetic Vlasov-Poisson equations [2]. The corresponding equations of motion for pushing particles using perturbative (δf) method [3] in terms of the gyrokinetic units of Ω_i, ρ_s for time and space, and $e\phi/T_e$ for the perturbed potential can be written as

$$\frac{d\mathbf{R}}{dt} = v_{\|}\hat{\mathbf{b}} + \mathbf{v}_d - \frac{\partial\bar{\phi}}{\partial\mathbf{R}} \times \hat{\mathbf{b}} \tag{1}$$

$$\frac{dv_{\|}}{dt} = -\hat{\mathbf{b}}^* \cdot \left(\frac{v_{\perp}^2}{2}\frac{\partial}{\partial\mathbf{R}} lnB + \frac{\partial\bar{\phi}}{\partial\mathbf{R}}\right), \tag{2}$$

$$\frac{dw}{dt} = -(1-w)\left(\kappa\frac{\partial\bar{\phi}}{\partial\mathbf{R}} \times \hat{\mathbf{b}} \cdot \hat{\mathbf{r}} + \frac{T_e}{T_i}(v_{\|}\hat{\mathbf{b}} + \mathbf{v}_d) \cdot \frac{\partial\bar{\phi}}{\partial\mathbf{R}}\right), \tag{3}$$

and

$$\mu_B \equiv \frac{v_{\perp}^2}{2B}(1 - v_{\|}\hat{\mathbf{b}} \cdot \frac{\partial}{\partial\mathbf{R}} \times \hat{\mathbf{b}}) \approx const., \tag{4}$$

where

$$\hat{\mathbf{b}}^* = \hat{\mathbf{b}} + v_{\|}\hat{\mathbf{b}} \times (\hat{\mathbf{b}} \cdot \frac{\partial}{\partial\mathbf{R}})\hat{\mathbf{b}},$$

$$\mathbf{v}_d = v_{\|}^2\hat{\mathbf{b}} \times (\hat{\mathbf{b}} \cdot \frac{\partial}{\partial\mathbf{R}})\hat{\mathbf{b}} + \frac{v_{\perp}^2}{2}\hat{\mathbf{b}} \times \frac{\partial}{\partial\mathbf{R}} lnB,$$

$\kappa = \kappa_n - (3/2 - v^2/2v_{ti}^2)\kappa_{Ti}$ is the background inhomogeneity with $\kappa_n \equiv 1/L_n$ and $\kappa_T \equiv 1/L_T$. [Note that the ion acoustic speed c_s is unity in these units.] The perturbed distribution is defined as $\delta f = \sum_{j=1}^{N} w_j \delta(\mathbf{R} - \mathbf{R}_{\alpha j})\delta(\mu - \mu_{\alpha j})\delta(v_{\|} - v_{\|\alpha j})$, where N is the total number of particle ions in the simulation, $F = F_0 + \delta f$, F_0 is the background Maxwellian with $\int F_0 d\mathbf{x} = 1$,

$$w \equiv \delta f/F \tag{5}$$

and $F \equiv \delta f(w_j = 1)$. The transformation between the gyrocenter coordinates $\mathbf{R}$ and the particle coordinates $\mathbf{x}$ are $\bar{\phi}(\mathbf{R}) = \langle \int \phi(\mathbf{x})\delta(\mathbf{x} - \mathbf{R} - \rho)d\mathbf{x}\rangle_{\varphi}$, where $\langle\cdots\rangle_{\varphi}$ is the average over the gyro-angle φ and ρ is the particle gyroradius. The gyrokinetic Poisson's equation can be written as

$$\tau[\phi(\mathbf{x}) - \tilde{\phi}(\mathbf{x})] = -\bar{n}_i(\mathbf{x}) + n_e(\mathbf{x}), \tag{6}$$

where $\tilde{\phi}(\mathbf{x}) \equiv < \int \bar{\phi}(\mathbf{R})\delta f_i(\mathbf{R}, \mu, v_{\|})\delta(\mathbf{R} - \mathbf{x} + \rho)d\mathbf{R}d\mu dv_{\|} >_{\varphi}$, $\bar{n}_i(\mathbf{x}) = e\langle \int \delta f_i(\mathbf{R})\delta(\mathbf{R} - \mathbf{x} + \rho)d\mathbf{R}dv_{\|}d\mu\rangle_{\varphi}$,

$$n_e(\mathbf{x}) = \begin{cases} \phi(\mathbf{x}), & (m,n) \neq (0,0), \\ 0, & (m,n) = (0,0), \end{cases}$$

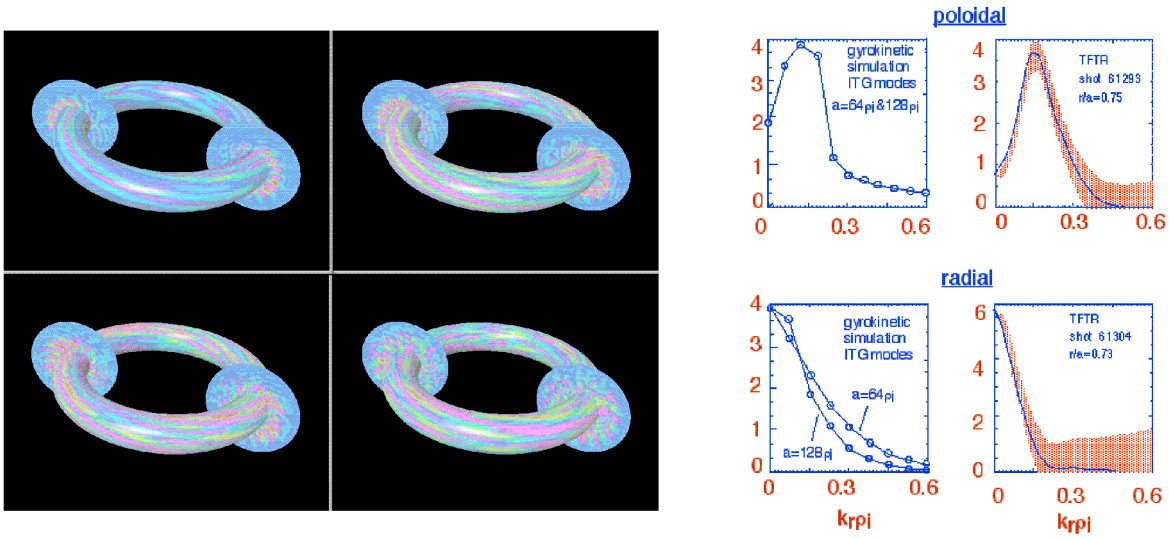

FIG. 1: *First simulations of ion temperature gradient (ITG) drift instabilities on Cray C90 at NERSC and the comparisons between the simulations results and the experimental measurements in the Toroidal Fusion Test Reactor (TFTR).*

(m, n) are the poloidal and toroial mode numbers, respectively, and $\mu \equiv v_\perp^2/2$. Numerically, the transformation between $\mathbf{R}$ and $\mathbf{x}$ can be accomplished through a 4-point average process valid for $k_\perp \rho_i \leq 2$ [4]. The approximation used for n_e is the so-called adiabatic electron model, which is adequate for the studies described in the present paper. We should remark here that the parallel nonlinearity, which is taken into account in the present studies, is the last term on the right hand side of Eq. (2). This term has mostly been ignored in the microturbulence community. However, without this term, the energy conservation cannot be satisfied in the simulation [2], whereas the only other nonlinearity in the simulation, the last term on the right hand side of Eq. (1), called the $\mathbf{E} \times \mathbf{B}$ nonlinearity, is unrelated to energy conservation.

III. DISCOVERIES THROUGH LEADERSHIP COMPUTING

The original work by Parker et al. [5], that initiated modern gyrokinetic simulation of microturbulence in realistic tokamak geometry, was carried out on the Cray C90 at National Energy Research NERSC. As shown in Fig. 1, finger-like streamers associated with the perturbed potentials in the weak field side of a tokamak were observed for the first time. These streamers were eventually modified by the turbulence into large eddies in the nonlinear stage.

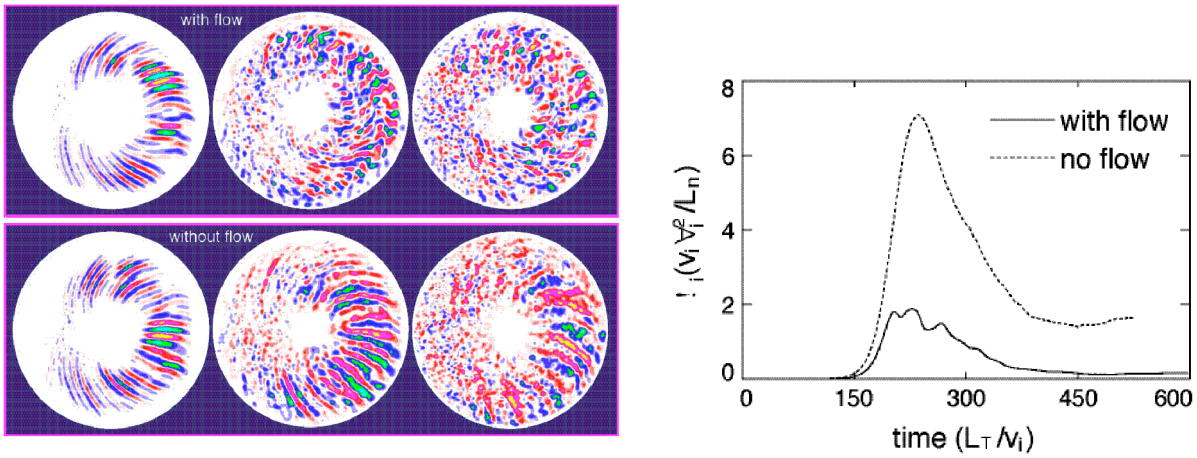

FIG. 2: *The identification of global zonal flows due to ITG turbulence as the cause for the observed reduction of turbulent transport. The simulations were carried out using Cray T3E at NERSC in 1998.*

The other interesting results of the paper were the fluctuation wavenumber spectra in the radial (k_r) and poloidal (k_θ) directions, which correlated well with the experimental measurements [6] in Tokamak Fusion Test Reactor (TFTR) at Princeton Plasma Physisc Laboratory (PPPL). The simulation used $a/\rho_s = 64-128$ with 1 million particles, where a is the minor radius. Although the simulation volume is considerably smaller than that of TFTR ($a \sim 400\rho_s$), this paper convinced the fusion community that ITG turbulence was the most likely cause for the anomalous transport of ions in tokamaks.

In 1998, using the T3E computer at NERSC with 100 million particles for $a/\rho_s = 160$, Lin et al [1] has found that the zonal flows generated by ITG turbulence can breakup the finger-like streamers into even smaller eddies and, in turn, can considerably reduce the ensuing steady state thermal transport. This is clearly a favorable trend for magnetic fusion. The comparisons between those with and without the nonlinearly-driven zonal flows are shown in Fig. 2. This is again an example of the impact of leadership computing on fusion research.

In order to design future fusion experiments, fusion researchers need a scaling law which can provide reliable estimates for how large a machine is needed in order to confine the plasma for sufficient time to achieve fusion. Two possible scaling laws derivable from the random

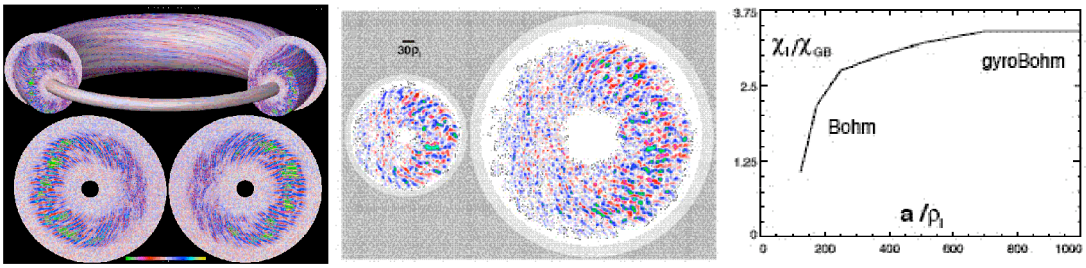

FIG. 3: *The transition from Bohm diffusion to GyroBohm diffusion as the size of simulation volume increases. The simulations were carried out on IBM S3P at NERSC in 2002.*

walk model due to the ITG turbulence can be expressed as

$$\chi_i \propto \begin{cases} \sqrt{M}/B^2, \texttt{GyroBohm} \\ 1/B, \texttt{Bohm} \end{cases} \tag{7}$$

where χ_i is the ion diffusivity, M is the ion mass and B is the magnitude of the external magnertic field. As we can see, the GyroBohm scaling is more desirable than the Bohm scaling, since it reduces the ion energy diffusion faster by using a stronger B field. On the other hand, GyroBohm scaling implies that a mixture of deterium-tritium plasma will have a larger x_i and, therefore, a negative impact on confimenment. However, most of the experimental evidence in the '90's showed that the scaling trend was Bohm-like, but with a scaling of $\chi_i \propto 1/\sqrt{M}$ with respect to the ion mass. To shed some light on these seemingly contradictory experimental trends, simulations were carried out on the C90 at NERSC with 4 million particles and $a/\rho_s = 128$. The results confirmed this experimental trend on the isotope effects [7]. Thus, the simple random walk model, Eq. (7), can be misleading.

In 2002, Lin et al. [8] used the IBM SP3 at NERSC for ITER-size plasmas for $a/\rho_s = 1000$ with 1 billion particles and discovered the transition from Bohm to gyroBohm diffusion as a function of the plasma size. The elusive gyroBohm scaling was finally found by this global simulation study, which is not based on the local approximation. This is a very encouraging trend because, with the gyroBohm scaling, the confinement improves with $1/B^2$, instead of $1/B$ for the Bohm scaling. If the GyroBohm scaling is valid for large tokamaks, should we expect a degradation in confinement with the introduction of tritium in the experiments? To address this important question, we will need more realistic physics models and, of course, more powerful computers.

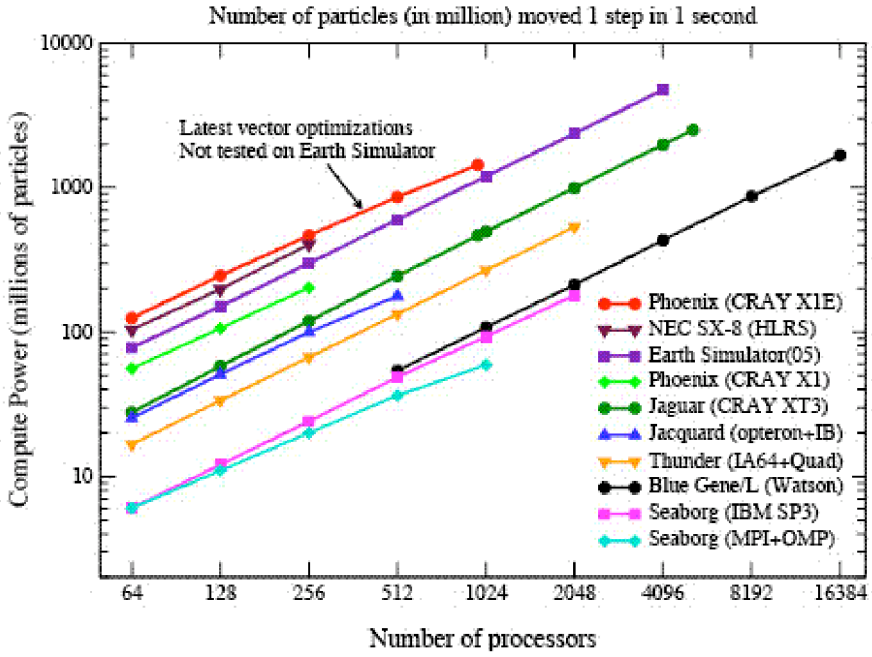

FIG. 4: *Performance of GTC on various MPP platforms.*

IV. RECENT TERASCALE COMPUTING

With the availability of terascale computing capabilities in recent years, we have ported GTC to various MPP platforms in the US and Japan. As shown in Fig. 4, the performance of GTC is very encouraging. It gives very good scaling on both vector-parallel and cache-based processors up to thousands of processors, and is very portable.

The demonstrated excellent performance on MPP platform has enabled us to run GTC with higher resolution and for longer duration. As such, we are now able to address the challenge of carrying out long-time steady-state simulations with more complete nonlinear effects taken into account. A particular nonlinearity of interest is associated with parallel acceleration and is the second term on the RHS of Eq. (2). Although this term has not attracted much attention in the fusion community, it is nevertheless important for properly maintaining energy conservation. With regard to relevant parameters, there are 64 toroidal grid with $a/\rho_i = 250$ on each poloidal plane. Based on the field line following coordinate system, the code uses an unstructured grid of the size ρ_s, (i.e., the ion thermal radius measured with the electron temperature), on each poloidal plane. The grid on each poloidal plane then rotates to the next plane by following the prescribed field lines. Thus, the shortest wavelength modes that can be

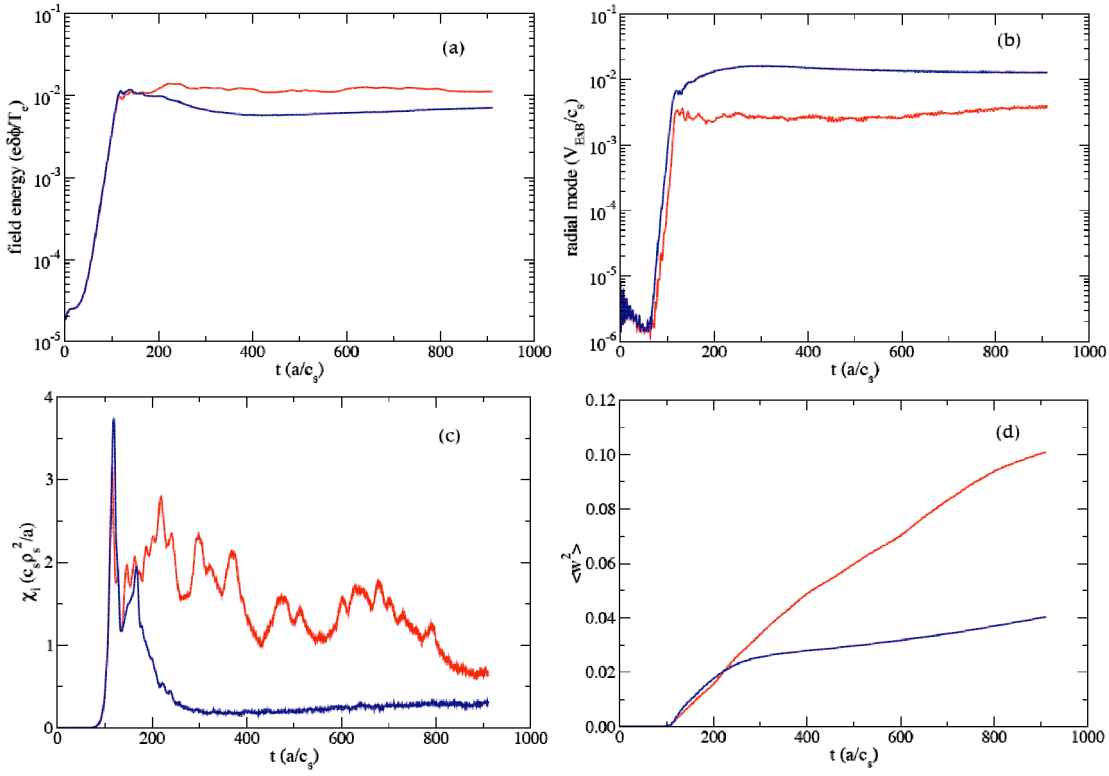

FIG. 5: *Time evolutions for the ITG turbulence for* $a/\rho_i = 250$ *with (blue or dark lines) and without (red or light lines) the velocity space nonlinearity for the (a) perturbed field energy, (d) radial mode amplitude, (c) ion thermal diffusivity, and (d) particle-weight-square.*

resolved in the code is $k_\perp \rho_s \approx 1$ on the poloidal plane, while allowing the physics with $k_{||} \ll k_\perp$. Here, $||$ and $\perp$ designate the directions parallel and perpendicular to the direction of the external magnetic field, respectively. The other parameters are: $n_0 = 20$ (number of particles per cell) , $R/L_T = 6.9$, $R/a = 2.79$, $L_n/L_{Ti} = 3.13$, $\Omega_i \Delta t = 15$ and $T_e/T_i = 1$. The radial profile of the inhomogeneity is given by $(1/L)e^{-[(r-r_c)/r_w]^6}$, where L represents either the temperaure scale length L_{Ti} or the density scale length L_n with $r_c/a = 0.5$ and $r_w/a = 0.35$. The simulation results for two runs, one without the parallel velocity space nonlinearity represented by red (or light) lines, and the other with the nonlinearity represented by blue (dark) lines are shown in Fig. 5. As we can see, there is no appreciable difference between the two cases in the linear stage of the development. However, the saturation amplitude of the spatially-averaged $e\phi/T_e$ at 1% for the case without this nonlinearity is about a factor of two higher than the case with the nonlinearity as indicated by Fig. 5(a). On the other hand, Fig. 5(b) shows that the zonal flow amplitude in term of v_{ExB}/c_s is about four times larger. Most

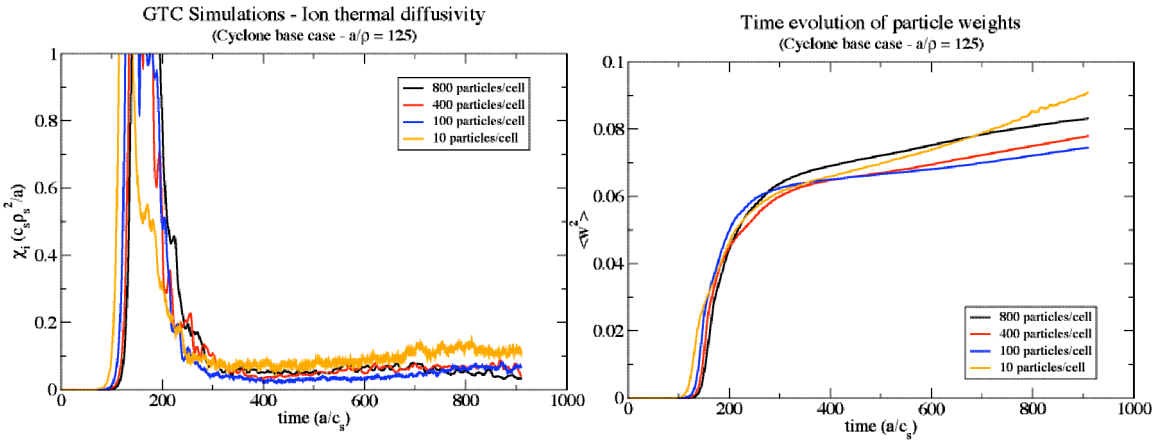

FIG. 6: *Particle number convergence studies for the ITG simulation: thermal diffusivty and the time rate of change of entropy for 10, 100, 400, and 800 particles per cell.*

striking is the evolution of the ion thermal diffusivity shown in Fig, 5(c), where the case with the velocity-space nonlinearity reaches its steady state value at a much faster pace and at a level of 0.3 in GyroBohm units, which is a factor two smaller than the case without this term. Another interesting aspect of these simulation is the the time evolution of the particle weights defined in Eq. (5). As shown in Fig. 5(d), the weight gain between the two cases are quite different with $\sqrt{\langle w^2 \rangle} \approx 0.2$ at the end of the simulation for the case with the nonlinearity. The time derivatives of $\langle w^2 \rangle$ is related to the entropy production as first pointed out in Ref. [9]. With the presence of the velocity space nonlinearity, it can be written as

$$\frac{\partial}{\partial t} \sum_{j=1}^{N} (1 - \alpha/4) w_j^2 = \kappa_{Ti} \langle Q_{ir} \rangle, \tag{8}$$

where $\alpha \approx 1$ is related to the velocity space nonlinearity and κ_{Ti} denotes ion temperature inhomogeneity and

$$Q_{ir} \equiv \frac{1}{N} \sum_{j=1}^{N} w_j v_j^2 \mathbf{v}_{ExB} \cdot \hat{\mathbf{r}}$$

is the ion thermal flux, the corresponding ion thermal diffusivity is χ_i / κ_{Ti} and κ_{Ti} is the inverse of the ion temperature scale length. The noise level for the simulation plasma can be estimated by $\phi \approx \sqrt{\langle w_j^2 \rangle / N}$, [11] which is related to the high frequency wiggles depicted in Fig. 5(c).

With the availability of Cray X1E and XT3 at ORNL, we have been able to conduct high resolution convergence studies using GTC for addressing the discrete particle noise issue in

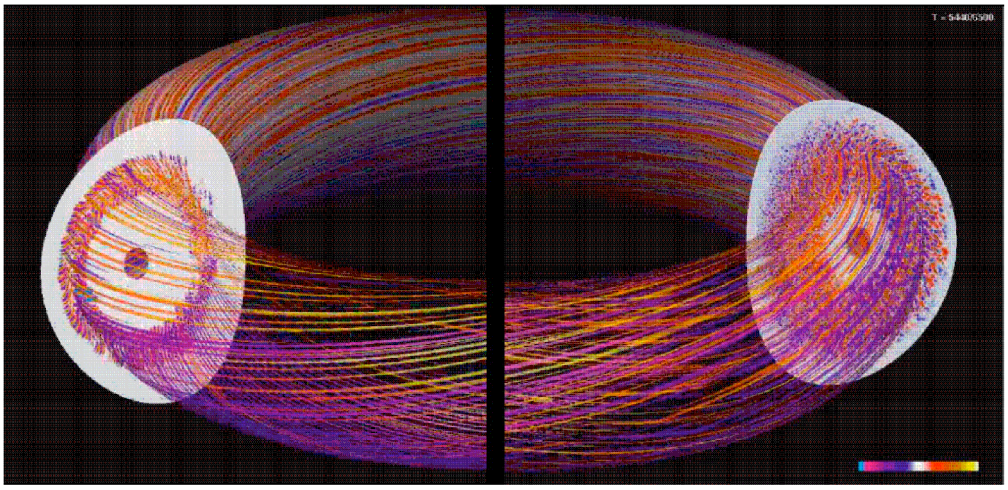

FIG. 7: Turbulence spreading (left to right) as depicted by the perturbed potentials of ITG turbulence on the poloidal plane as they follow the magnetic field lines around the torus.

steady state ITG simulations. For example, by using $10, 100, 400$,and 800 particles per cell in the simulation for parameters similar to those associated with Fig. 5 (except for $a/\rho_s = 125$), we can demonstrate that numerical noise has little effect on the resulting energy transport when a reasonable number of particles is used. As shown in Fig. 6, both the ion thermal diffusivities on the left as well as the time rate of change in entropy [see, Eq. (8)] the only numerical noise one can detect comes from the case for 10 particles/cell. The signature of the noise is the enhancement of the thermal flux and the high frequency noise associated with the high frequency damped modes, when 10 particles per cell is used. We should remark that, in PIC simulations, the relevant measure for the noise is the number of particles in the wavelengths of interest, but it is not really the number of particles per cell.

Finally, a new version of the GTC code has been generalized to more realistically represent experimentally relevant geometry. It has been developed and applied to the investigation of ITG turbulence in actual experiments. Considering a shaped geometry similar to the ITER plasma, we have found very interesting new phenomena related to the non-local nature of the associated turbulence [10]. As illustrated in Fig. 7, the turbulence spreads radially from a localized region at an earlier time (left) in the simulation to eventually cover most of the poloidal plane at a later time (right).

V. PETASCALE COMPUTING IN THE FUTURE

Based on our very positive experiences using GTC for turbulence research in recent years, we are very pleased by the successful implementation of PIC codes on MPP platforms. This also holds great promise for our future utilization of petascale computing. With two-dimensional (toroidal and radial) domain decomposition in place for both particles and fields, we expect to be able to use thousands of processors without significant degradation in performance. It is, therefore, possible to envision a scenario in which one trillion particle is used on a $10,000 \times 10,000 \times 100$ grid with 100 particles per cell on the 50,000 XT3 quad-core processors. This is an ITER-size plasmas with electron skin depth as the grid size. With this kind of resolution, we should be able to address the physics of the electron energy transport associated with shear-Alfven dynamics. At present, the cause of electron thermal transport remains an outstanding unresolved question in fusion research, and it is suspected that much higher resolution in the simulations will be needed to resolve this issue. Petascale computing capability is clearly needed to deal with this formidable challenge. Such an ambitious research activity is also a natural ground for cross-field fertilization. For example, the corresponding elliptic solver in GTC using such a high resolution needs to be able to handle 10^8 elements per poloidal plane. Thus, we will depend on close collaborations with researchers from the Terascale Optimal PDE Simulations (TOPS) Center [12] to help provide efficient petaflop-ready routines. Furthermore, such high resolution runs will generate many terabytes of data both for the particles and fields. Close collaborations with Scientific Data Management (SDM) Center [13] along with advanced visualization capabilities to help analyze and interpret such data will be essential. Important physics questions that will be addressed in this context will include the scaling of confinement with plasma size and with ionic isotopic species burning plasmas.

Beyond turbulent transport physics, we are also looking forward to simulating a wide range of tokamak physics using the gyrokinetic PIC approach. Research activities under the OASCR/MICS Mutli-Scale Gyrokinetics (MSG) project have focused on developing efficient PIC algorithms that can bridge the gap between magnetohydrodynamic (MHD) physics, turbulent transport and resonant wave heating [14, 15]. The ultimate goal is the development of an effective integrated modeling capability for fusion plasmas with high physics fidelity on petaflop computers.

* The work is supported by DoE contract DE-AC02-76-CHO3037 and the SciDAC Gyrokinetic Particle Simulation Center for Turbulent Transport in Burning Plasmas.

---

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

[2] T. S. Hahm, Phys. Fluid **31**, 2670 (1988).

[3] S. E. Parker, and W. W. Lee, Phys. Fluids B **5**, 77 (1993).

[4] W. W. Lee, J. Comp. Phys. **72**, 243 (1987).

[5] S. E. Parker, W. W. Lee, and R. A. Santoro, Phys. Rev. Lett. **71**, 2042(1993).

[6] R. Fonck, Crosby, Durst, S. Paul, N. Bretz, S. Scott, E. Synakowski, G. Taylor, Phys. Rev. Lett. **70**, 3736 (1993).

[7] W. W. Lee and R. A. Santoro, Phys. Plasmas **4**, 169 (1997).

[8] Z. Lin, S. Ethier, T. S. Hahm, W M. Tang, Phys. Rev. Lett. **88**, 195004 (2002).

[9] W. W. Lee and W. M. Tang, Phys. Fluids **31**, 612 (1988).

[10] W. X. Wang et al., PPPL-4142 Report, submitted to Phys. Plasmas (2006).

[11] W. W. Lee, J. Lewandowski, T. S. Hahm, and Z. Lin, Phys. Plasmas **8**, 4435 (2001).

[12] http://www-unix.mcs.anl.gov/scidac-tops/

[13] http://sdm.lbl.gov/sdmcenter/

[14] W. W. Lee and H. Qin, Phys. Plasmas **10**, 3196 (2003).

[15] R. Kolesnikov, W. W. Lee, H. Qin and E. Startsev, in preparation (2006).