Numerics and computation in gyrokinetic simulations of electromagnetic turbulence with global particle-in-cell codes

A. Mishchenko, A. Bottino, A. Biancalani, T. Hayward-Schneider,
A. Koenies, R. Kleiber, M. Borchardt,
M. Cole, and L. Villard

Introduction

Goal: EM simulations of burning plasmas

- Burning plasmas are complex systems with multiple spatial and temporal scales
- A substaintial energetic-particle minority couples electromagnetic turbulence, global Alfvénic and MHD modes, zonal flows
- A single framework is needed which includes all these parts of the problem
- Preparation to future exascale systems

Tool: ORB5&EUTERPE

- Use the gyrokinetic PIC codes ORB5 and EUTERPE for this purpose (proposed for EUROfusion's TSVV Task 10)
- Refactor ORB5 and EUTERPE aiming at a single framework for global gyrokinetic PIC simulations

- Gyrokinetic Vlasov equation: method of characteristics

$$\frac{\partial f_{1s}}{\partial t} + \dot{\boldsymbol{R}} \cdot \frac{\partial f_{1s}}{\partial \boldsymbol{R}} + \dot{v}_\| \frac{\partial f_{1s}}{\partial v_\|} = -\dot{\boldsymbol{R}}^{(1)} \cdot \frac{\partial F_{0s}}{\partial \boldsymbol{R}} - \dot{v}_\|^{(1)} \frac{\partial F_{0s}}{\partial v_\|}.$$

- Gyrocenter trajectories: $\partial \langle A_\| \rangle / \partial t$ does not appear in $p_\|$-GK!

$$\dot{R} = \left(v_\| - \frac{q}{m} \langle A_\| \rangle \right) \boldsymbol{b}^* + \frac{1}{q B_\|^*} \boldsymbol{b} \times \left[\mu \nabla B + q \left(\nabla \langle \phi \rangle - v_\| \nabla \langle A_\| \rangle \right) \right]$$

$$\dot{v}_\| = -\frac{1}{m} \left[\mu \nabla B + q \left(\nabla \langle \phi \rangle - v_\| \nabla \langle A_\| \rangle \right) \right] \cdot \boldsymbol{b}^*$$

- Gyrokinetic field equations:

$$\int \frac{q_i F_{0i}}{T_i} \left(\phi - \langle \phi \rangle \right) \delta(\boldsymbol{R} + \boldsymbol{\rho} - \boldsymbol{x}) \, \mathrm{d}^6 Z = \bar{n}_i - \bar{n}_e$$

$$\frac{\beta_i}{\rho_i^2} \langle \overline{A}_\| \rangle_i + \frac{\beta_e}{\rho_e^2} A_\| - \nabla_\perp^2 A_\| = \mu_0 \left(\bar{j}_{\| i} + \bar{j}_{\| e} \right)$$

Discretization

- "Klimontovich" representation for perturbed distribution function:

$$\delta f_s(\boldsymbol{R}, v_\parallel, \mu, t) = \sum_{\nu=1}^{N_p} w_{s\nu}(t)\delta(\boldsymbol{R} - \boldsymbol{R}_\nu)\delta(v_\parallel - v_{\nu\parallel})\delta(\mu - \mu_\nu) \ ,$$

- Maxwellian distribution for all species:

$$F_{0s} = n_0 \left(\frac{m}{2\pi T_s}\right)^{3/2} \exp\left[-\frac{m_s v_\parallel^2}{2T_s}\right] \exp\left[-\frac{m_s v_\perp^2}{2T_s}\right]$$

- Finite-element discretization for fields:

$$\phi(\boldsymbol{x}) = \sum_{l=1}^{N_s} \phi_l(t)\Lambda_l(\boldsymbol{x}) \ , \quad A_\parallel(\boldsymbol{x}) = \sum_{l=1}^{N_s} a_l(t)\Lambda_l(\boldsymbol{x}) \ ,$$

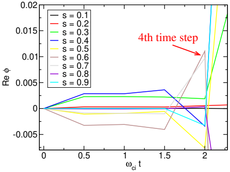

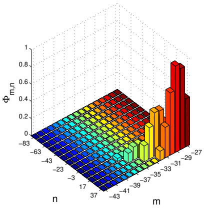

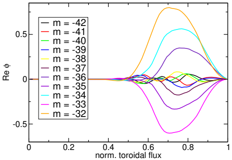

LHD-like geometry, electromagnetic ITG mode

Severe numerical instability at the very beginning of simulation

Small unavoidable inconsistencies: imbalance of side bands, small distortions of equilibrium at the axis, markers leaving and re-entering simulation domain

Cancellation problem can strongly magnify this small numerical issues

- Split the magnetic potential into the 'symplectic' and 'hamiltonian' parts:

$$A_\parallel = A_\parallel^{(\mathrm{s})} + A_\parallel^{(\mathrm{h})}$$

- The perturbed guiding-center phase-space Lagrangian

$$\gamma = q\boldsymbol{A}^* \cdot \mathrm{dR} + \frac{\mathrm{m}}{\mathrm{q}}\,\mu\,\mathrm{d}\theta + \mathrm{q}\,\mathrm{A}_\parallel^{(\mathrm{s})}\mathrm{b}\cdot\mathrm{dx} + \mathrm{q}\,\mathrm{A}_\parallel^{(\mathrm{h})}\mathrm{b}\cdot\mathrm{dx} - \left[\frac{\mathrm{mv}_\parallel^2}{2} + \mu\mathrm{B} + \mathrm{q}\phi\right]\mathrm{dt}$$

- "Mixed" Lie transform: $A_\parallel^{(\mathrm{h})} \to$ Hamiltonian, $A_\parallel^{(\mathrm{s})} \to$ symplectic structure

$$\Gamma = q\boldsymbol{A}^* \cdot \mathrm{dR} + \frac{\mathrm{m}}{\mathrm{q}}\mu\,\mathrm{d}\theta + \mathrm{q}\left\langle \mathrm{A}_\parallel^{(\mathrm{s})}\right\rangle \cdot \mathrm{dR} - \left[\frac{\mathrm{mv}_\parallel^2}{2} + \mu\mathrm{B} + \mathrm{q}\left\langle \phi - \mathrm{v}_\parallel \mathrm{A}_\parallel^{(\mathrm{h})}\right\rangle\right]\mathrm{dt}$$

- The corresponding perturbed equations of motion are

$$\dot{\boldsymbol{R}}^{(1)} = \frac{\boldsymbol{b}}{B_\parallel^*} \times \nabla \left\langle \phi - v_\parallel A_\parallel^{(\mathrm{s})} - v_\parallel A_\parallel^{(\mathrm{h})} \right\rangle - \frac{q}{m} \langle A_\parallel^{(\mathrm{h})} \rangle \, \boldsymbol{b}^*$$

$$\dot{v}_\parallel^{(1)} = -\frac{q}{m} \left[\boldsymbol{b}^* \cdot \nabla \left\langle \phi - v_\parallel A_\parallel^{(\mathrm{h})} \right\rangle + \frac{\partial}{\partial t} \left\langle A_\parallel^{(\mathrm{s})} \right\rangle \right] - \frac{\mu}{m} \frac{\boldsymbol{b} \times \nabla B}{B_\parallel^*} \cdot \nabla \left\langle A_\parallel^{(\mathrm{s})} \right\rangle$$

- An equation for $\partial A_\parallel^{(\mathrm{s})} / \partial t$ is needed

$$\frac{\partial}{\partial t} A_\parallel^{(\mathrm{s})} + \boldsymbol{b} \cdot \nabla \phi = 0$$

- Field equations law takes the form

$$\sum_{s=i,f} \int \frac{q_s^2 F_{0s}}{T_s} \left(\phi - \langle \phi \rangle \right) \delta_{\mathrm{gy}} \, \mathrm{d}^6 \mathrm{Z} = \sum_{\mathrm{s=i,e,f}} \mathrm{q_s} \bar{\mathrm{n}}_\mathrm{s}$$

$$\sum_{s=i,e,f} \frac{\beta_s}{\rho_s^2} \left\langle \overline{A}_\parallel^{(\mathrm{h})} \right\rangle_s - \nabla_\perp^2 A_\parallel^{(\mathrm{h})} = \mu_0 \sum_{s=i,e,f} j_{\parallel 1s} + \nabla_\perp^2 A_\parallel^{(\mathrm{s})}$$

Mixed-variable formulation: nonlinear algorithm

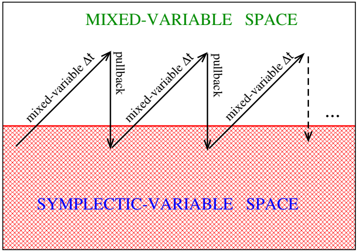

$$f_{1s}(Z_s, A_\parallel^{(s)}) = f_{1m}(Z_m, A_\parallel^{(s)}, A_\parallel^{(h)})$$
$$v_\parallel^{(\mathrm{s})} = v_\parallel^{(\mathrm{m})} - \frac{e}{m}\left\langle A_\parallel^{(\mathrm{h})}\right\rangle$$

Additional nonlinear terms appear in equations of motion [R. Kleiber et al, PoP 2016] (symplectic-hamiltonian equivalence at the 2nd order)

1. Push coordinates and weights along the nonlinear mixed-variable trajectories
2. Transform coordinates into symplectic space keeping weights constant
3. Set $A^{(\mathrm{s})}_{\parallel(\mathrm{new})}(t_i) = A_\parallel(t_i) = A^{(\mathrm{s})}_{\parallel(\mathrm{old})}(t_i) + A^{(\mathrm{h})}_{\parallel(\mathrm{old})}(t_i)$ and $A^{(\mathrm{h})}_{\parallel(\mathrm{new})}(t_i) = 0$.

EUTERPE: electromagnetic ITG instability in “LHD”

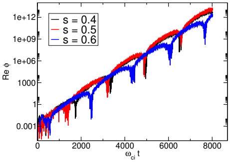

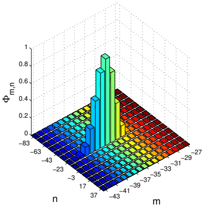

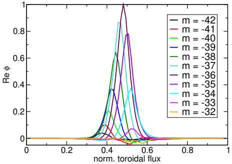

LHD-like geometry, electromagnetic ITG mode

Severe numerical instability at the very beginning of simulation: mitigated!

Clean modes is observed

EUTERPE: electromagnetic instabilities in W7-X

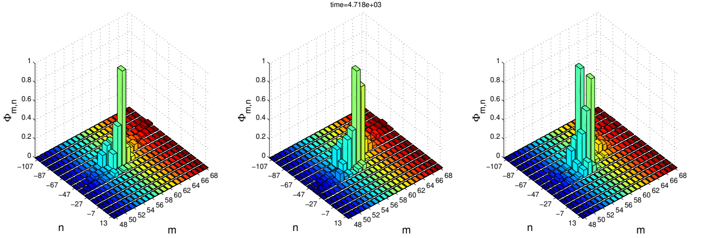

W7-X standard configuration, $\beta = 2\%$: electromagnetic ITG/TEM/KBM spectra

1. Flat electron temperature and density (only ion temperature gradient)
2. Flat density, ion and electron temperature gradients: mode structure changes
3. Flat electron temperature, gradient in ion temperature and densities

Further studies are needed; applications to "stability valley" in W7-X (global EM)

For all profiles, numerically clean mode is observed

ORB5: saturation of EM turbulence in adhoc tokamak

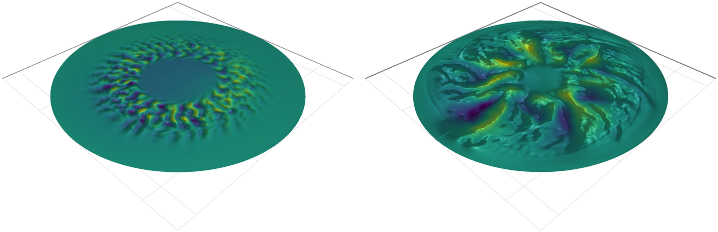

1. Low-beta EM-ITG turbulence ($\beta = 0.01\%$): zonal-flow saturation mechanism
2. Large beta case ($\beta = 1.6\%$): global eigenmode (BAE) dominates $\phi(\boldsymbol{x})$ including NL harmonics.

Physics changes at larger beta! To be studied with ORB5 in detail.

ORB5: saturation of EM turbulence in "ITER"

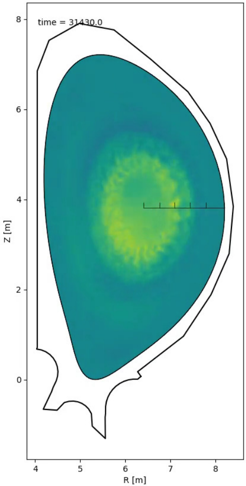

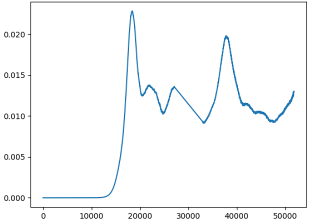

ITER geometry; plasma profiles and ρ_* similar to cyclone-base case, low β: saturated EM turbulence is observed (heat flux and $\phi(\boldsymbol{x})$ shown)

We acknowledge PRACE for awarding us access to Marconi100 (CINECA)

XGC: KBM instability

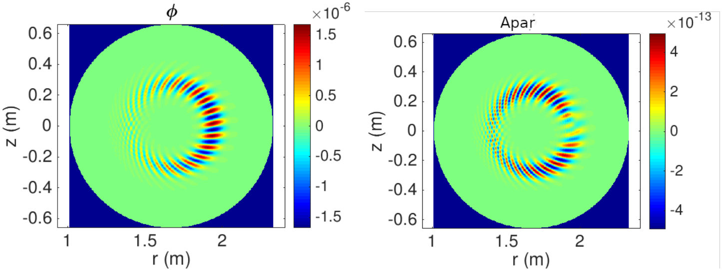

- KBM instability: $\beta = 2.5\%$ [M. Cole et al, submitted to Phys. Plasmas]
- Electrostatic and magnetic potential; pullback mitigation
- “Goerler benchmark” (ENR NumKiN)

Computation performed on Cori (NERSC)

GPU in HPC

- Fraction of computer time with mandatory GPUs is increasing
- Eventual future: no GPU-enabling ⇒ no computer time ⇒ no results/papers etc.
- Codes running on hetergeneous systems have competitive advantage

Example: HPC system in 21st PRACE call

- HAWK: 345 mln core hours (total)
- Joliot-Curie (KNL/Rome/SKL): 88/459/124 mln core hours (total)
- JUWELS (Booster/Cluster): 35.04/70 mln core hours (total)
- Marconi100: 660 mln core hours (total)
- MareNostrum4: TBD (30 mln core hours minimum)   has GPU partition
- Piz Daint: 510 mln core hours (total)
- SuperMUC-NG: 121 mln core hours (total)

1205.04 mln GPU core hours vs. 1207 non-GPU core hours (w/o MareNostrum4)
49.959% of all core hours available in the call are GPU-mandatory

MareNostrum4 excluded (TBD); some multi-core CPUs (KNL etc) require OpenMP

ORB5: GPU vs. CPU speedup on DAINT

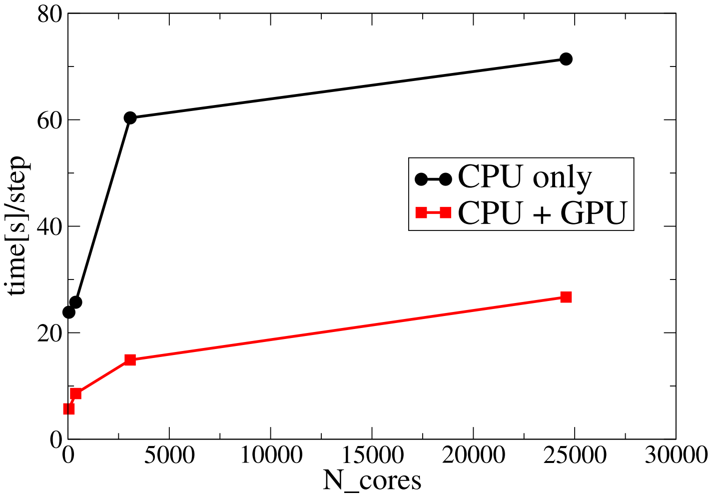

Result of this type can justify computer time on a GPU machine

EM ITG heat flux on GPUs: Marconi vs. M100

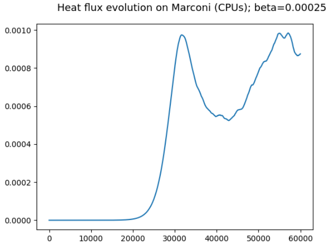

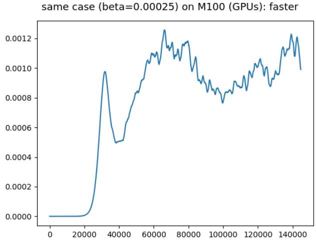

EM ITG + BAE case

- Large-aspect-ratio tokamak (physical $\beta = 0.01\%$) [Biancalani et al]
- GPU speedup: 48 hours on 24 Marconi nodes vs. 24 hours on 16 M100 nodes

Problems

- Number of the markers is limited by the number of GPUs (memory).
- Only 16 nodes were allowed originally: high-marker resolution runs impossible.
- At a larger node number, memory is volatile and tends to crash with out-of-memory: Details of OpenACC implementation in ORB5? Issues with PGI compiler enviroment? Configuration of M100 GPUs?

Status

Tokamak simulations

- Majority of results using ORB5
- EM turbulence in ad-hoc geometry including fast particles
- EM turbulence in "ITER" (down-scaled, small β); to be extended to real
- Alfvén Eigenmodes with fast particles in realistic ASDEX-U, ITER
- Runs on GPUs (M100, Dain, Summit): GPU memory limitations (many GPUs needed for many markers)

Stellarator simulations

- EUTERPE is needed
- Electromagnetic linear instabilities, electrostatic turbulence (W7-X)
- Memory requirements increase for turbulence with machine size (large matrices)
- Noise control in stellarators [E. Sanchez et al]
- CPU-only; push is similar to ORB5 (track for GPU-enabling)

Outlook: directions of the code development

Production code

Merging EUTERPE and ORB5; creating appropriate data structures and modules

Adaptation to available hardware

Heterogeneous Systems replacing conventional CPUs; pure MPI is not suffucient; solution algorithms must be designed with hardware properties in mind

Algorithms

Traditional: noise control, collisions, electromagnetics, electron time stepping
Novel: large perturbations (semi-lagrangian control variate), Maxwell solvers

Applications

Driven by experimental programs: ITER, W7-X, ASDEX-U, TCV, JET, JT60-SA
Global gyrokinetics, zonal flows, fast particles, and MHD; Tokamaks&Stellarators
Beyond gyrokinetics?: ion-cyclotron time scales, core-pedestal-edge modelling