Data communications and analysis in simulation of plasma turbulence on a five-dimensional phase space

Tomo-Hiko Watanabe and Shinya Maeyama
Department of Physics, Nagoya University, Japan

watanabe.tomohiko@nagoya-u.jp
www.p.phys.nagoya-u.ac.jp

Outline

- Introduction
  - What is plasma turbulence?
  - Why do we meet big data in plasma physics?
  - Basic equations and simulation model
- Multiscale turbulence simulation in plasma
  - Peta-scale simulation of fusion plasma turbulence
  - Data communication optimized on peta-scale computer
  - Nonlinear interactions in kinetic plasma turbulence
- Summary

What is plasma turbulence?

Biggest plasma in the universe

- Virgo clusters

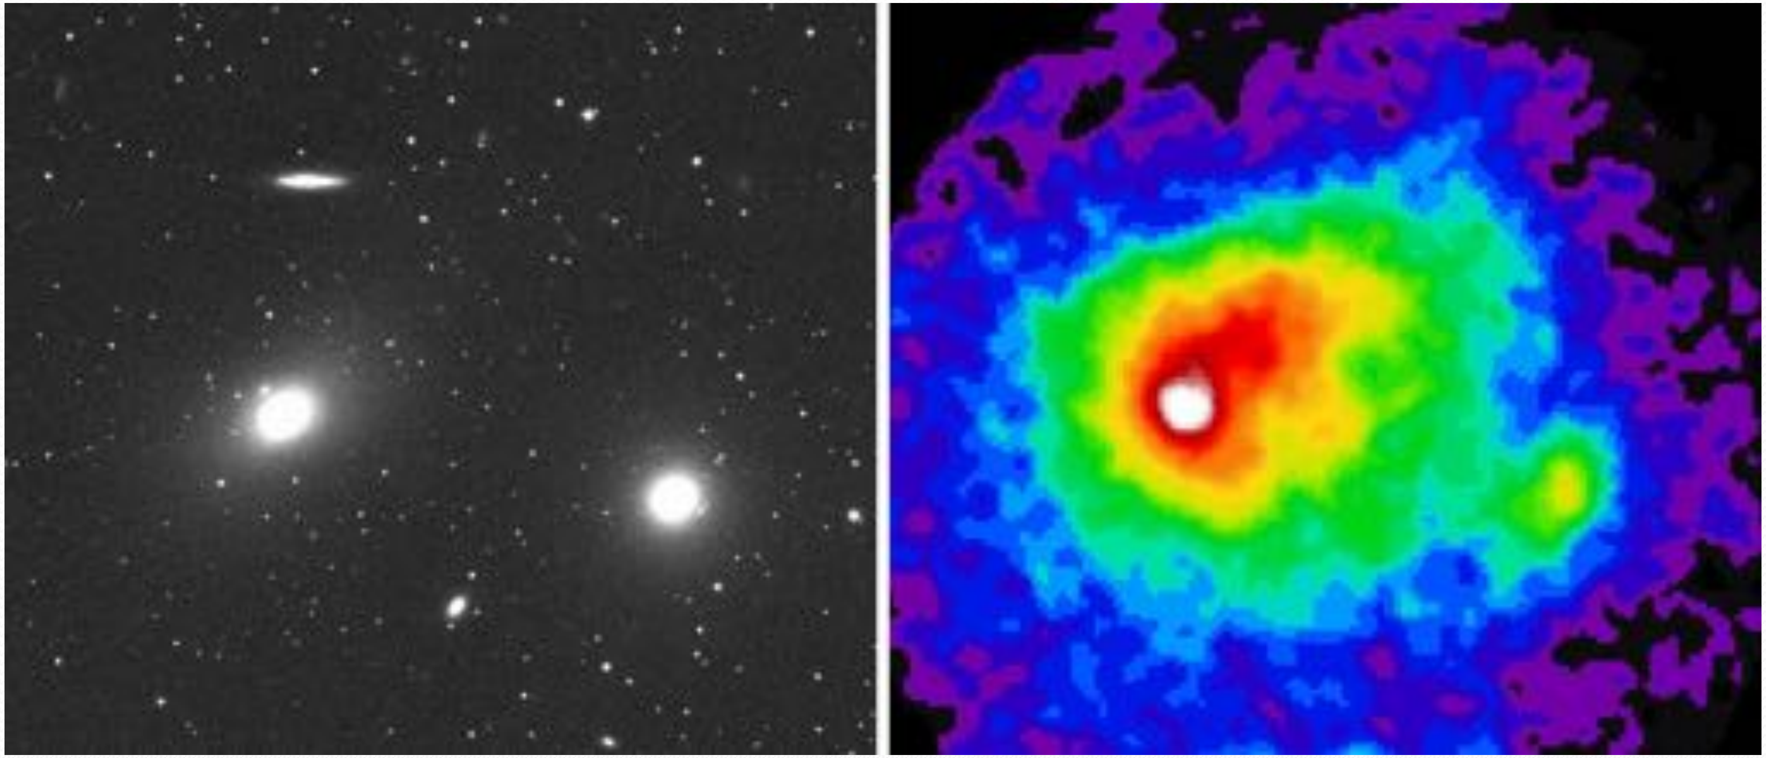

Visible (left) and X-ray (right) images

http://www.astro.isas.jaxa.jp/xjapan/asca/5/cggas/

Plasma turbulence in the interplanetary space

Bale+ PRL 2005

- Solar wind observation by spacecraft confirms the power law scaling of fluctuations

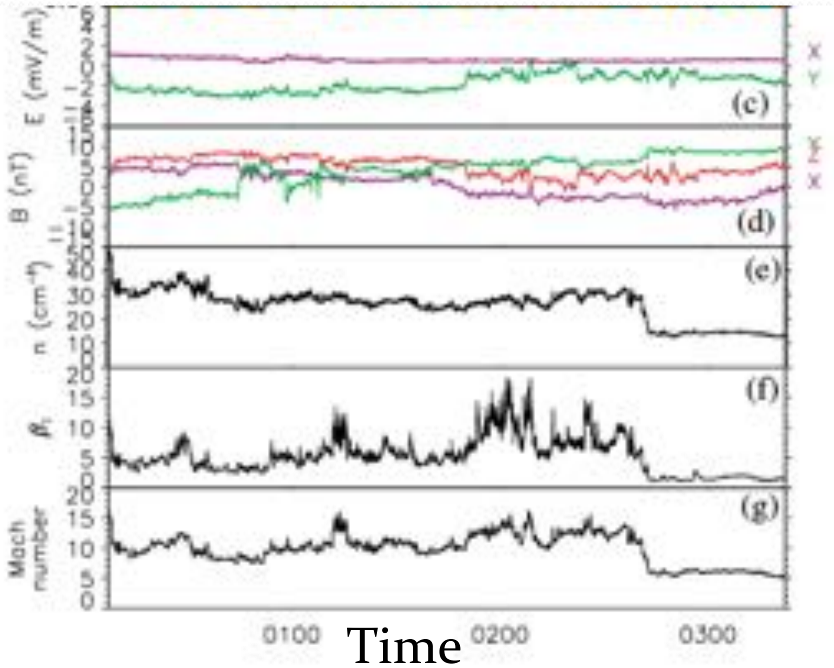

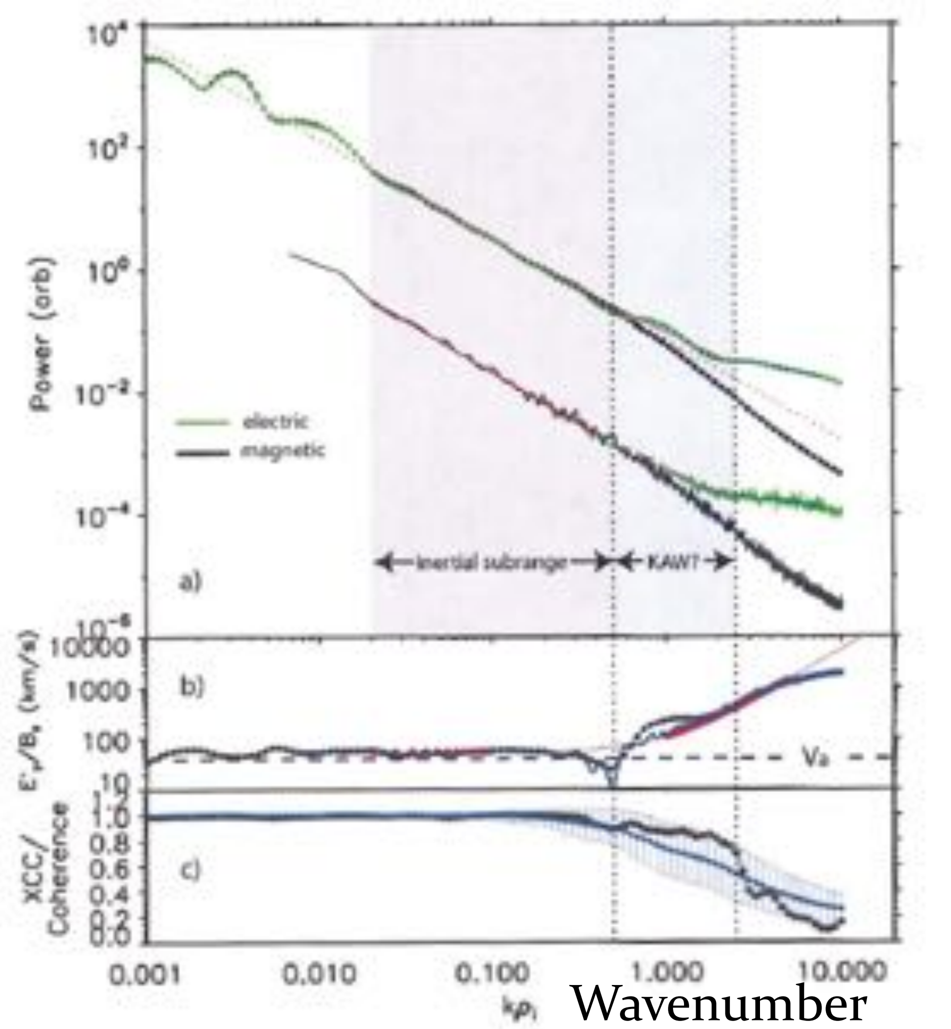

Strong turbulence and transport in magnetic fusion plasma

- Strong turbulence drives the particle and heat transport, if mean gradients of n and/or T exist in magnetic fusion plasma.
- TFTR experiments [Fonck+ PRL (1993)]

Turbulence with ion gyro-radius scales

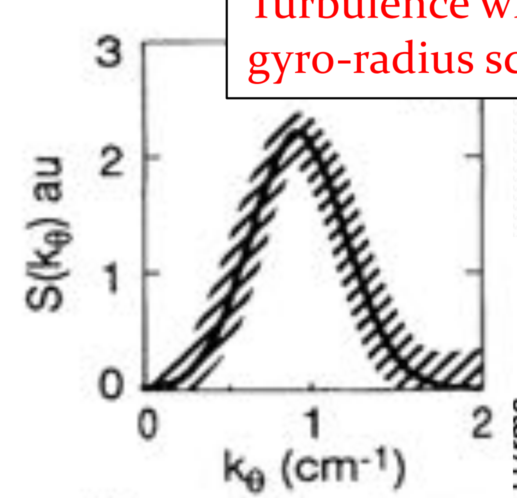

Poloidal wavenumber

Turbulent density fluctuations observed in a whole torus

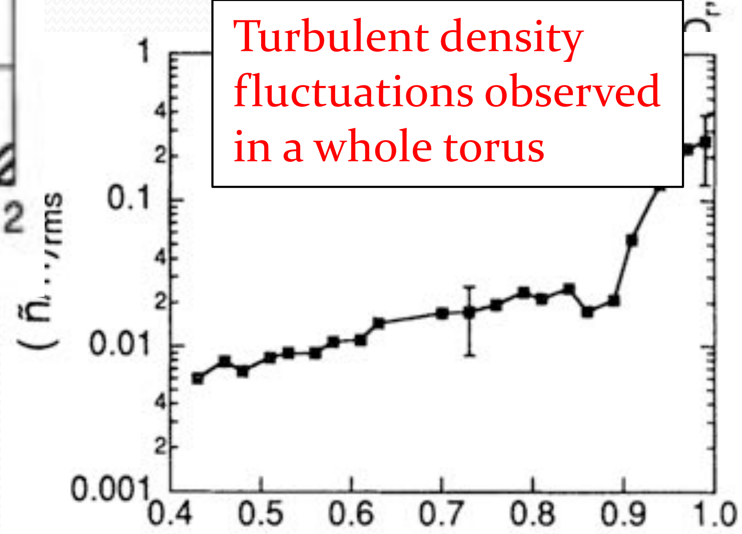

Normalized minor radius

Turbulent transport coefficients

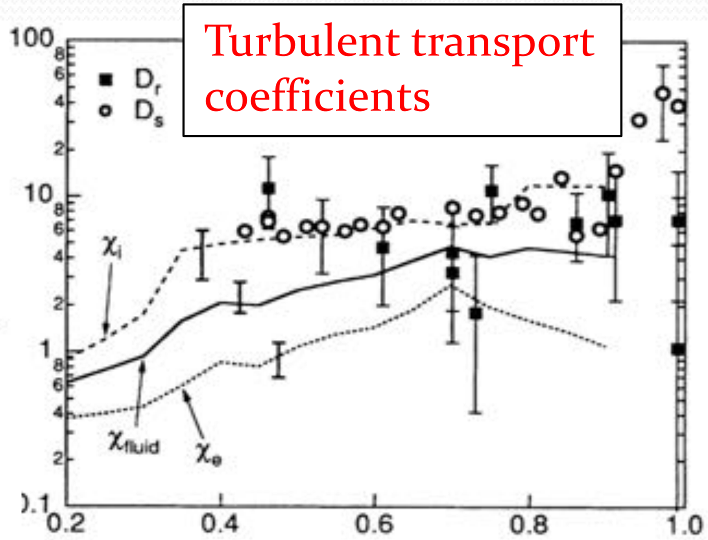

Normalized minor radius

Fluid approximation may break down in high temperature plasmas

- Fluid approximation can be valid for L

  $\gg \lambda_{ii}$: mean-free-path (// to $\boldsymbol{B}$)

  $\gg \rho_i$: gyro-radius (perpendicular to $\boldsymbol{B}$)

- In fusion plasmas of T_i~10 keV, n~10^{14}/cc,

  $\nu_{ii} \sim 10^2\ s^{-1}$, $\lambda_{ii} \sim 10^4$ m, a~1m, qR_0~10m

  Thu, the Knudsen number $\lambda_{ii} / qR_0 \sim 10^3$!!

- How large is λ_{ii} in the Earth's magnetosphere?

  $\lambda_{ii} \sim O(10^8$ km) !! (for T_i~10 eV, n~5/cc)

Why do we meet big data in plasma physics?

Properties of plasma turbulence

- Plasma turbulence consists of fluctuations of particle and velocity distributions and electromagnetic fields.
  - $f = f_0 + \delta f, \quad \boldsymbol{E} = \boldsymbol{E}_0 + \delta \boldsymbol{E}, \quad \boldsymbol{B} = \boldsymbol{B}_0 + \delta \boldsymbol{B}$
- Due to the high temperature of $T >$ keV, the one-body distribution function, $f(\boldsymbol{x}, \boldsymbol{v}, t)$, may deviate from the equilibrium with the Maxwellian F_M.
  - $\delta f \neq F_M$
- Due to the magnetic field $\boldsymbol{B}_0$, the charged particle motions, and thus, the turbulence is anisotropic.
  - $\delta f = \delta f(x_\parallel, \boldsymbol{x}_\perp, v_\parallel, \boldsymbol{v}_\perp)$

The Vlasov equation

- Advection of f along particle trajectories in the phase space is describe by the Vlasov equation,

$$\frac{\partial f}{\partial t} + \boldsymbol{v} \cdot \nabla f + \frac{q}{m}(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \cdot \frac{\partial f}{\partial \boldsymbol{v}} = 0$$

or

$$\frac{\partial f}{\partial t} + \{H, f\} = 0$$

which involves a variety of kinetic effects, i.e., Landau damping, particle trapping, finite gyroradius effects, …

- Fine structures are generated by the advection terms on the phase space, that is, shearing of f by the Hamiltonian flow,

$$u(x_i) \frac{\partial}{\partial x_j} f(x_i, x_j, \dots)$$

=> Generation of "Big Data" on the phase space

How does the distribution function develop on the phase space?

- Collisionless damping in 1-D Vlasov-Poisson system is shown below, where $f(x, v, t = 0) = F_M(1 - A\cos kx)$
- Fine structures of f continuously develop
  - Ballistic modes with scale-lengths of $1/kt$ in v-space
  - Stretching of f due to shear of the Hamiltonian flow

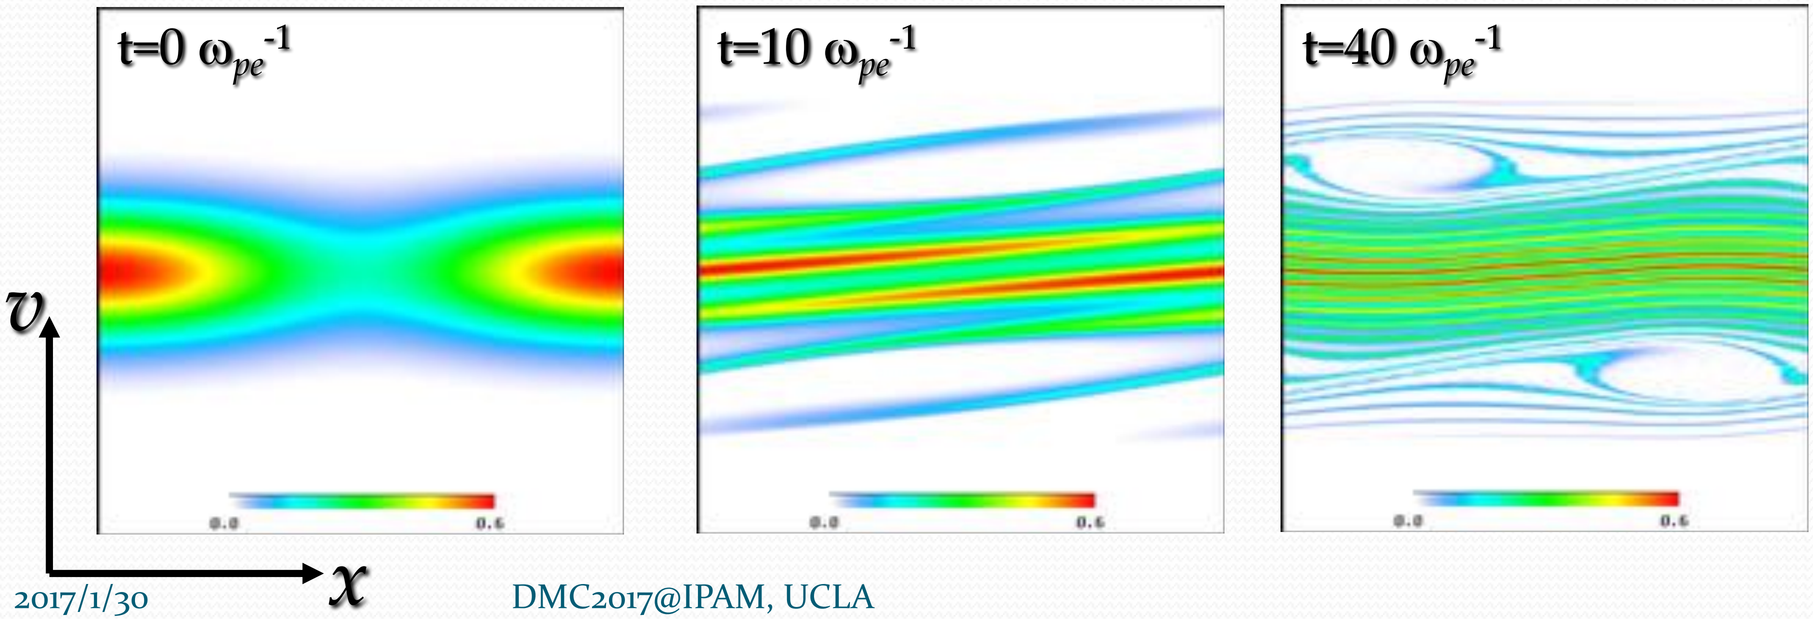

Basic equations and simulation model

Kinetic model simplified for low-frequency phenomena

- Although the Vlasov equation is "the first principle" for describing collisionless plasma behaviors, it involves short time scale of $\Omega_i^{-1}, \Omega_e^{-1}, \omega_p^{-1} \ldots$
  - In a magnetic fusion plasma with $B = 1\text{T}$,

$$\Omega_i = \frac{eB}{m_i} \sim 1\times 10^8 \ \ [\text{rad} \cdot \text{sec}^{-1}]$$

- We need reduced kinetic equations to eliminate the fast gyro-motion as well as ω_p, while keeping finite gyro-radius and other kinetic effects.

=> Gyrokinetic equations

From Vlasov to gyrokinetic eqs.

- To deal with fluctuations slower than the gyro-motion, reduce the Vlasov equation to a gyro-averaged form:

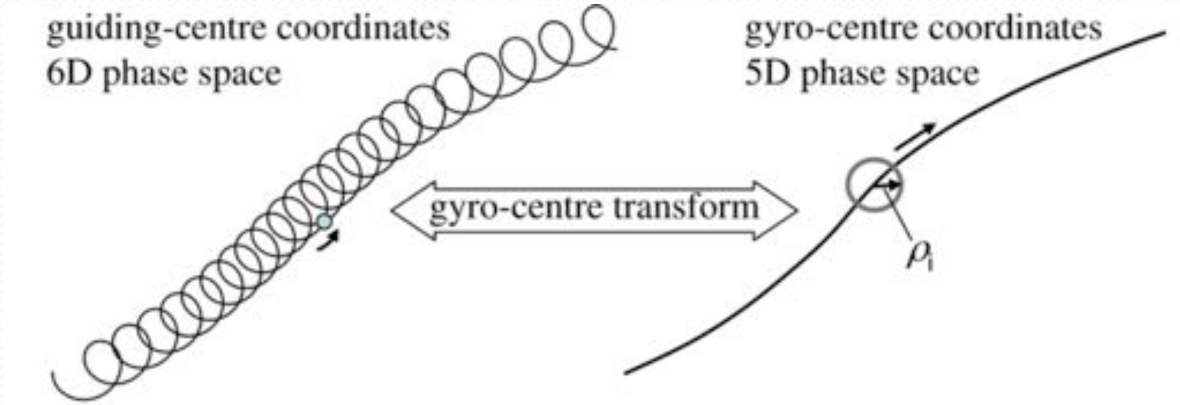

- Gyrokinetic ordering and perturbation expansion

$$\varepsilon \sim \frac{\omega}{\Omega} \sim \frac{\rho}{L} \sim \frac{k_\parallel}{k_\perp} \sim \frac{\delta f}{f_0} \sim \frac{e\phi}{T} \sim \frac{\delta B}{B_0}, \quad f = f_0 + \delta f$$

  - Recursive formulation of linear gyrokinetic equations [Rutherford & Frieman (1968); Antonsen & Lane (1980)]

Perturbed gyrokinetic equation

- Gyrocenter coordinates $\left(X^{(g)}, v_{\parallel}, \mu, \xi\right)$
  - μ: magnetic moment, ξ: gyrophase

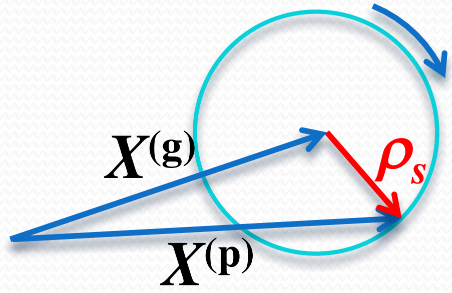

- Nonlinear gyrokinetic equation for $\delta f_s^{(g)}$

$$\left[\frac{\partial}{\partial t} + v_{\parallel}\mathbf{b}\cdot\nabla + \mathbf{v}_{ds}\cdot\nabla - \frac{\mu_s}{m_s}\mathbf{b}\cdot\nabla B\frac{\partial}{\partial v_{\parallel}}\right]\delta f_s^{(g)} + \frac{c}{B_0}\left\{\Phi - \frac{v_{\parallel}}{c}\Psi, \delta f_s^{(g)} + \frac{e_s\varphi}{T_s}\right\}$$

Magnetic drift　　Mirror force　　ExB drift and advection along $\tilde{B}$

$$= -v_{\parallel}F_{0s}\frac{e_s}{T_s}\left(\mathbf{b}\cdot\nabla\Phi + \frac{1}{c}\frac{\partial\Psi}{\partial t}\right) + F_{0s}\frac{e_s}{T_s}\left[\mathbf{v}_{*s}\cdot\nabla\left(\Phi - \frac{v_{\parallel}}{c}\Psi\right) - \mathbf{v}_{ds}\cdot\nabla\Phi\right]$$

Parallel electric field　　Diamagnetic drift

[Friemann & Chen, '82]

- Potential Φ and Ψ act on the gyrocenter (nonlinear term).

- Major "flow shear" terms in the GK equation generate fine structures of $\delta f_s^{(g)}$ on the phase space,

$$v_{\parallel}\boldsymbol{b}\cdot\nabla\delta f_s^{(g)}, \qquad \boldsymbol{v}_{ds}\cdot\nabla\delta f_s^{(g)}, \qquad \left\{\Phi - \frac{v_{\parallel}}{c}\Psi, \delta f_s^{(g)}\right\}.$$

Fluctuations of δf on (x,v)-space

- Gyrokinetic simulation of the ion temperature gradient driven turbulence causes energy transport and fluctuations of $\delta f_s^{(g)}$ on the velocity space. [Watanabe & Sugama, NF2006]

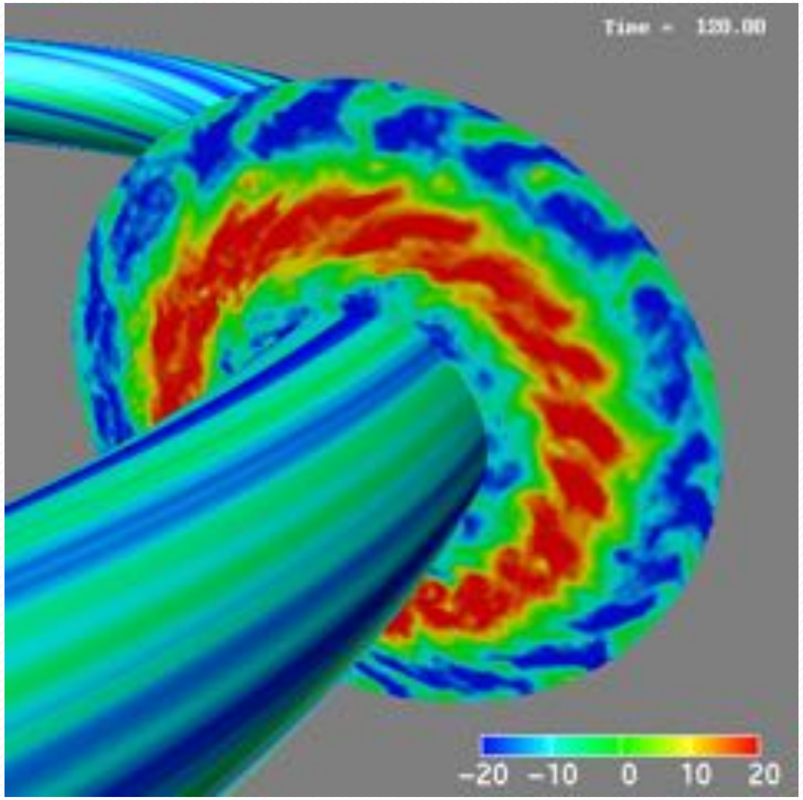

Snapshot of $\delta f_s^{(g)}$ on v-space

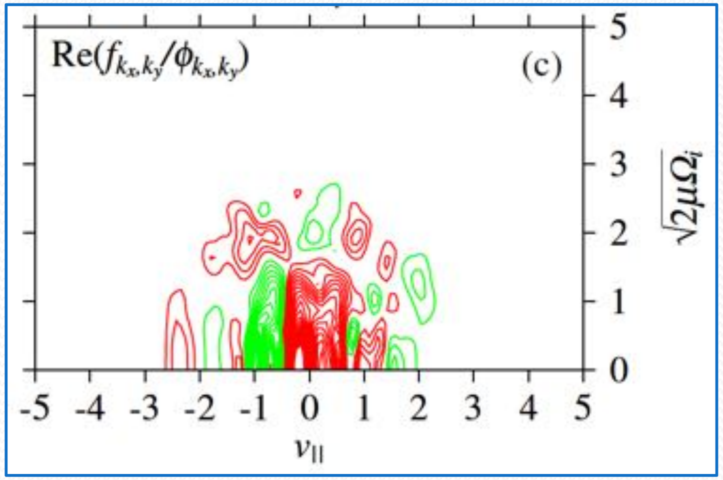

Grid points ~ 50 billion (5×10^{10})
Memory ~ 2.6TB
Computation ~ 5TFlops
on Earth Simulator 192 nodes
(peak 12TFlops) 24 hours.

(Nx, Ny, Nz, Nv, Nm) = (256, 256, 128, 128, 48)

Peta-scale simulation of fusion plasma turbulence

The K computer

http:www.aics.riken.jp/en/kcomputer/

"Kei" means 10^{16}.

- CPU: SPARC64 VIIIfx 2 GHz
  8 cores/processor
  16 GFlops/core
  Memory BW 8 GB/s/core
- Interconnect: Torus fusion (Tofu)
  6D mesh/torus topology
  Interconnect BW 5 GB/s × 4
  4 send + 4 recv. simultaneously
- 88128 nodes (705024 cores)
- 10.51 PFlops (No. 1 of Top500 in Nov 2011)
- Top 1 of Graph500 still in Nov 2016

<u>Tofu interconnect with</u>
<u>6D mesh/torus topology</u> [Ajima,2012]
(3D torus network as a user view)

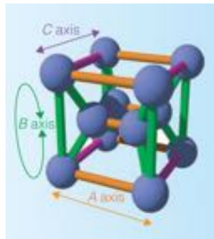

Gyrokinetic simulation of multiscale plasma turbulence

- The flux tube code, GKV, has been applied to the direct numerical simulation of the multiscale turbulence.
- Ion and electron scale turbulence are simultaneously computed with high spatial resolution.
- Employ the periodic boundary in $\boldsymbol{x}_\perp$
- $m_i/m_e = 1836$ leads to 43 times difference of the two scales.

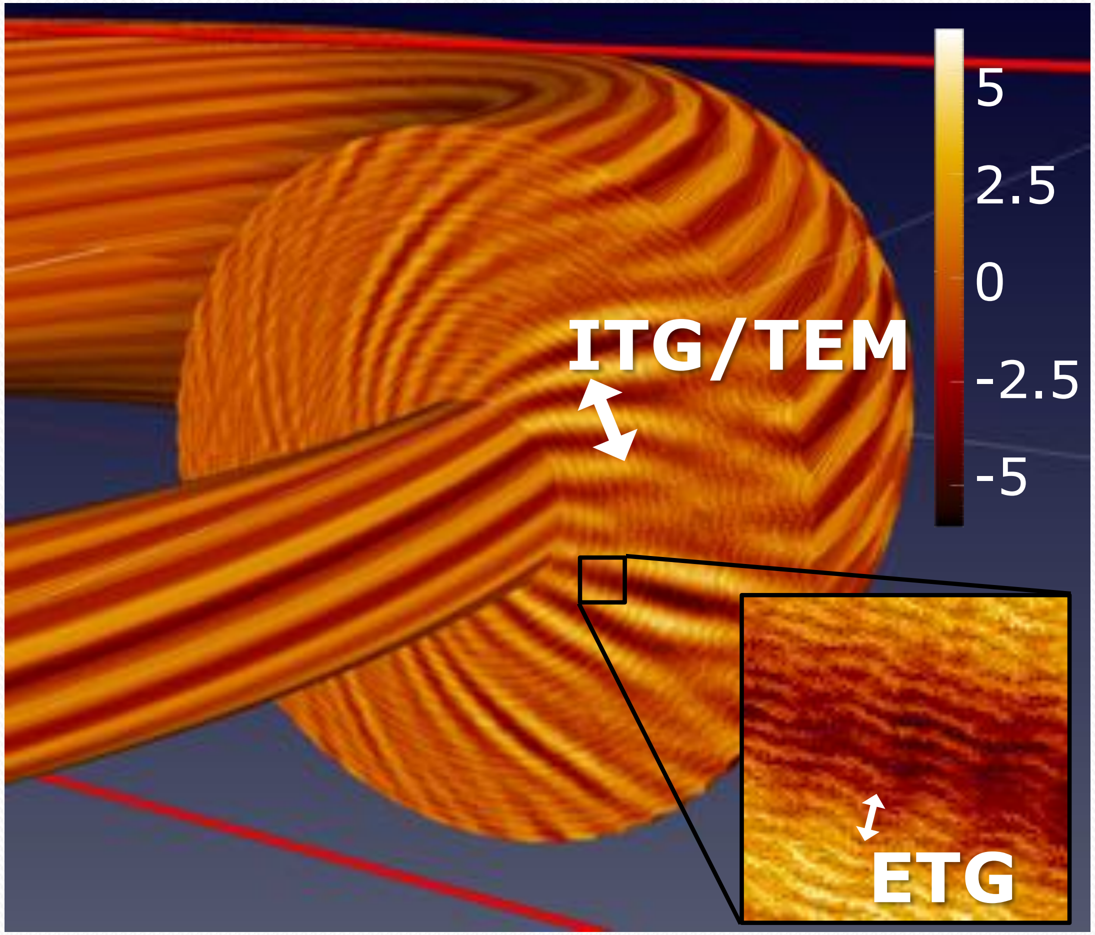

Maeyama+ PRL (2015)

Transport in multi-scale turbulence: "More is different"

- Transport in the multi-scale turbulence is characterized *neither* by the ITG and ETG transport in a single-scale.

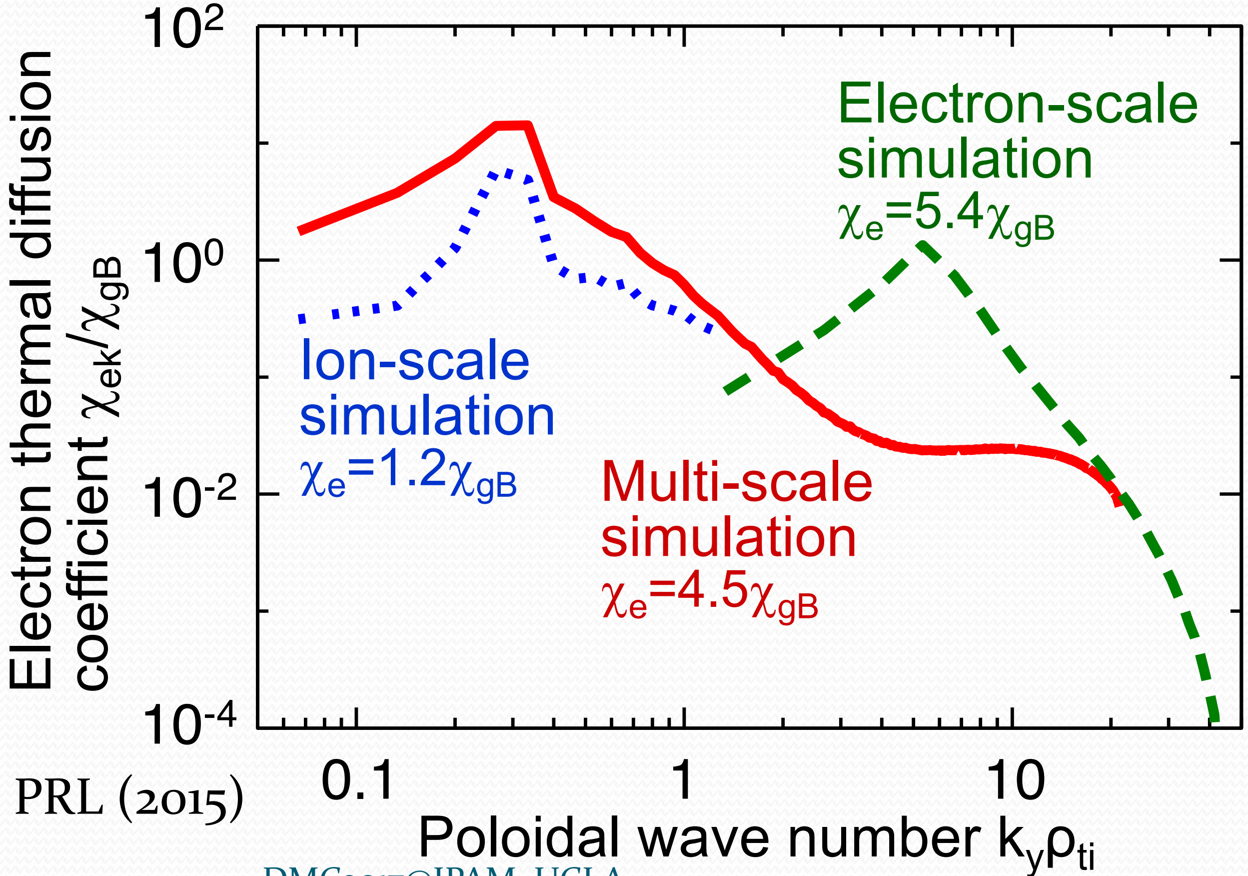

Maeyama+ PRL (2015)

Computation meeting big data needs optimization

- The peta-scale turbulence simulation needs large amount of data communications
  - ~100MB data transfer for each MPI process / time step
  - ~ 1TB data transfer / time step for 12,288 MPI processes
  - It costs ~ 0.9 sec / time step (best estimate after optimization)
- Computational cost for the same run
  - 190 TFLOP / time step
  - 140 TFLOPS achieved on 12,288 nodes (~9% to peak)
  - It costs ~ 1.36 sec / time step (memory bottle neck)
- Computation / communication ~ only 1.5
- How to achieve the high performance => Optimization

Data communication optimized on peta-scale computer

Requirements for multi-scale simulations

- Fine resolutions in x and y to resolve electron and ion scales.
- Small time step size to resolve rapid electron motions.

Resource: ~100 EFlop
Time steps: ~10^5 steps
Problem size: $1024 \times 1024 \times 96 \times 96 \times 32 \times 2 = 6 \times 10^{11}$ grids
FFT FD, Reduction

Parallelization: over 100 k cores

The improvement of the strong scaling (**reducing the cost of inter-node communications**) is critically important.

- ➢ Optimizations of MPI-rank mapping and communications
  → reduce the communication costs
- ➢ Computation-communication overlaps
  → mask the communication costs

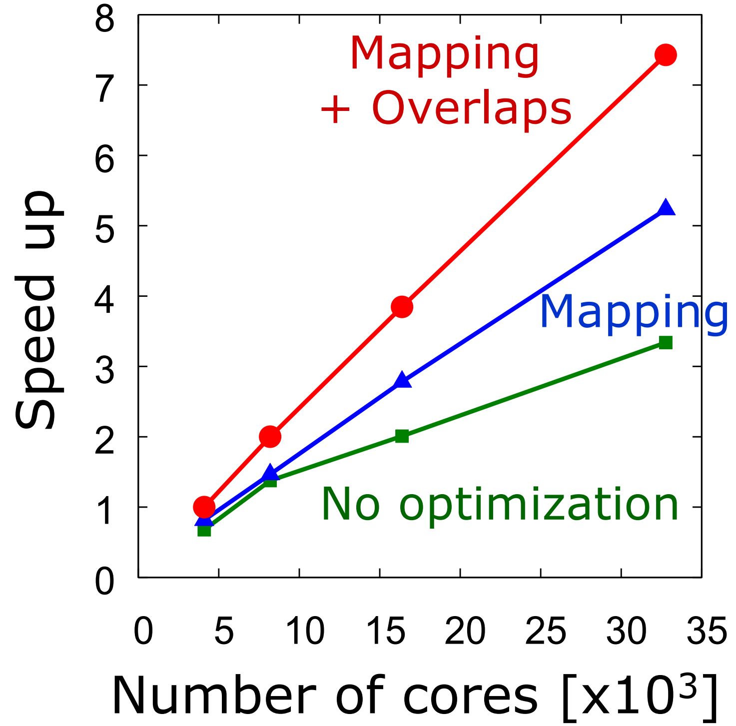

Segmented rank mapping on 3D torus network

①Arrange rank_xy:
Data transpose is
performed in a segment.

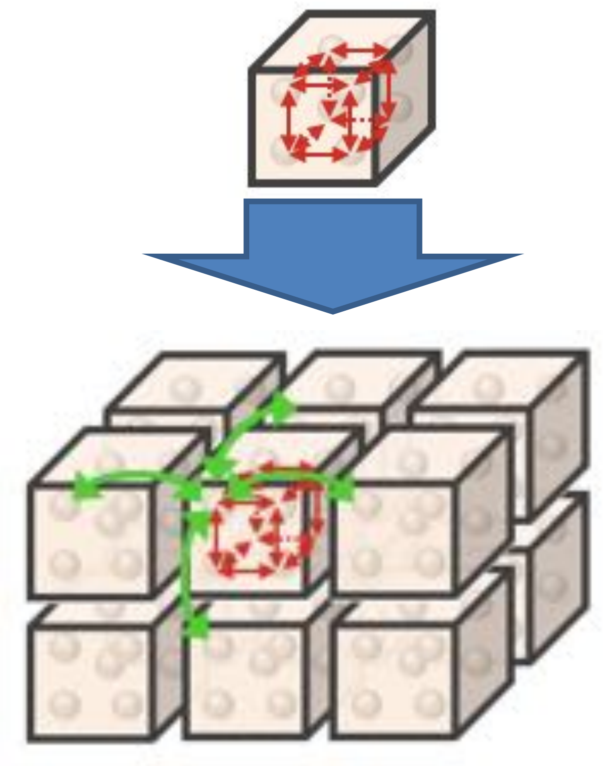

②Arrange rank_z, rank_v, rank_m:
Point-to-point communications are performed
between adjacent segments.

Segmented rank mapping on 3D torus network

①Arrange rank_xy:
Data transpose is
performed in a segment.

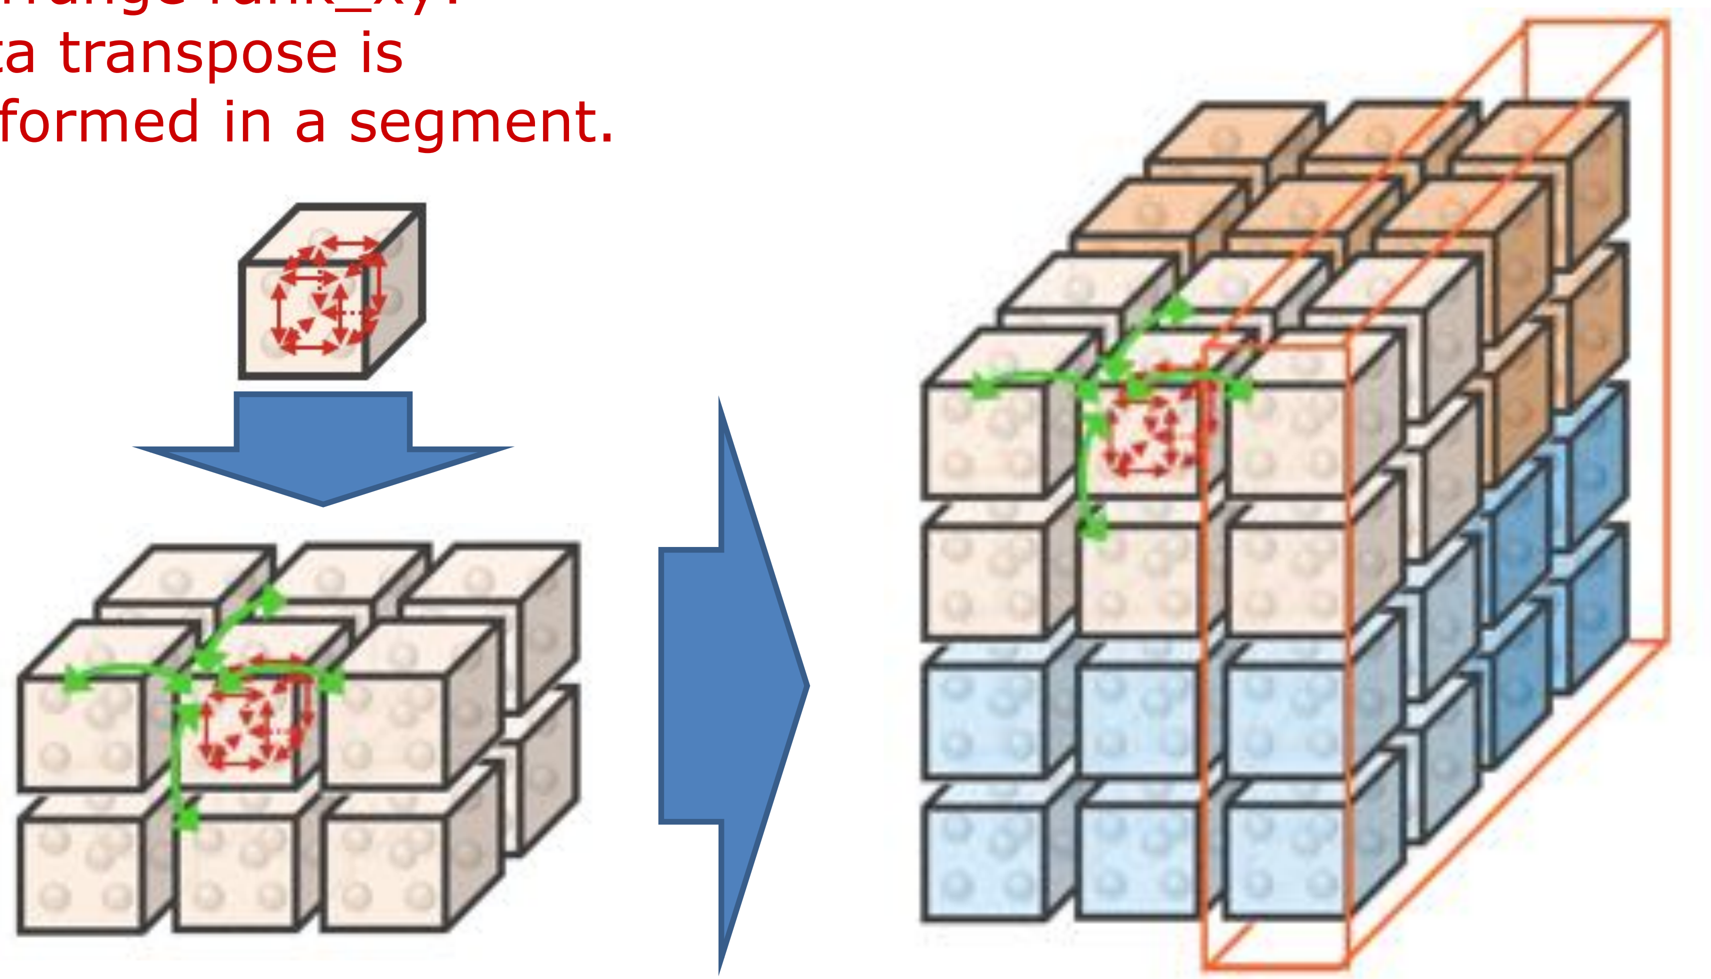

②Arrange rank_z, rank_v, rank_m:
Point-to-point communications are performed
between adjacent segments.

③Arrange rank_s:
Reduction is
performed in a
cross section.

Computation-communication overlaps

Overlap method for MPI/OpenMP hybrid parallelization [Idomura, 2013]

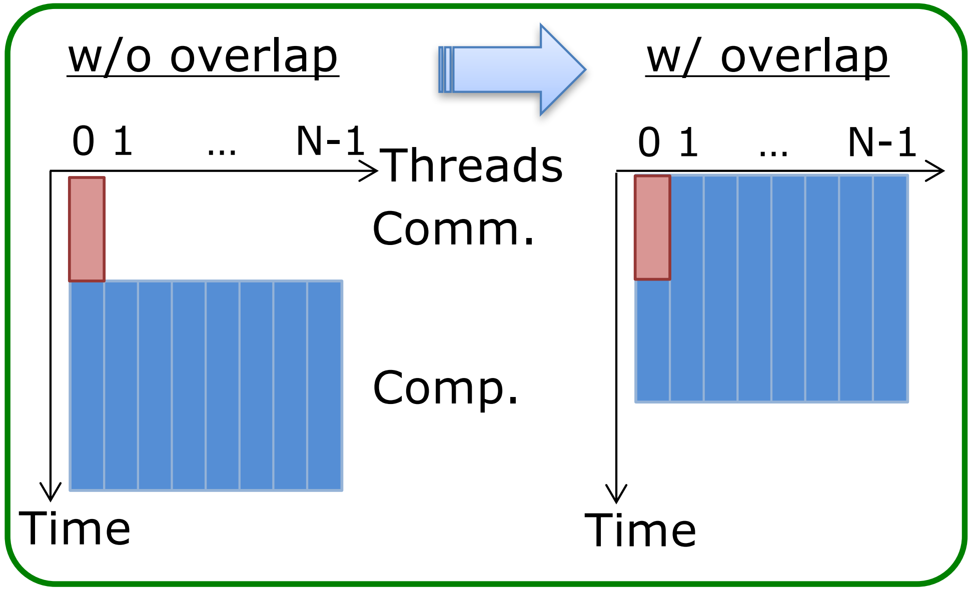

- The communication thread enables overlaps for All-to-All.
- The overlap techniques are applied for spectral and finite difference calculations.

Pipelined overlap of spectral calculation [Maeyama+, Parallel Comput.]

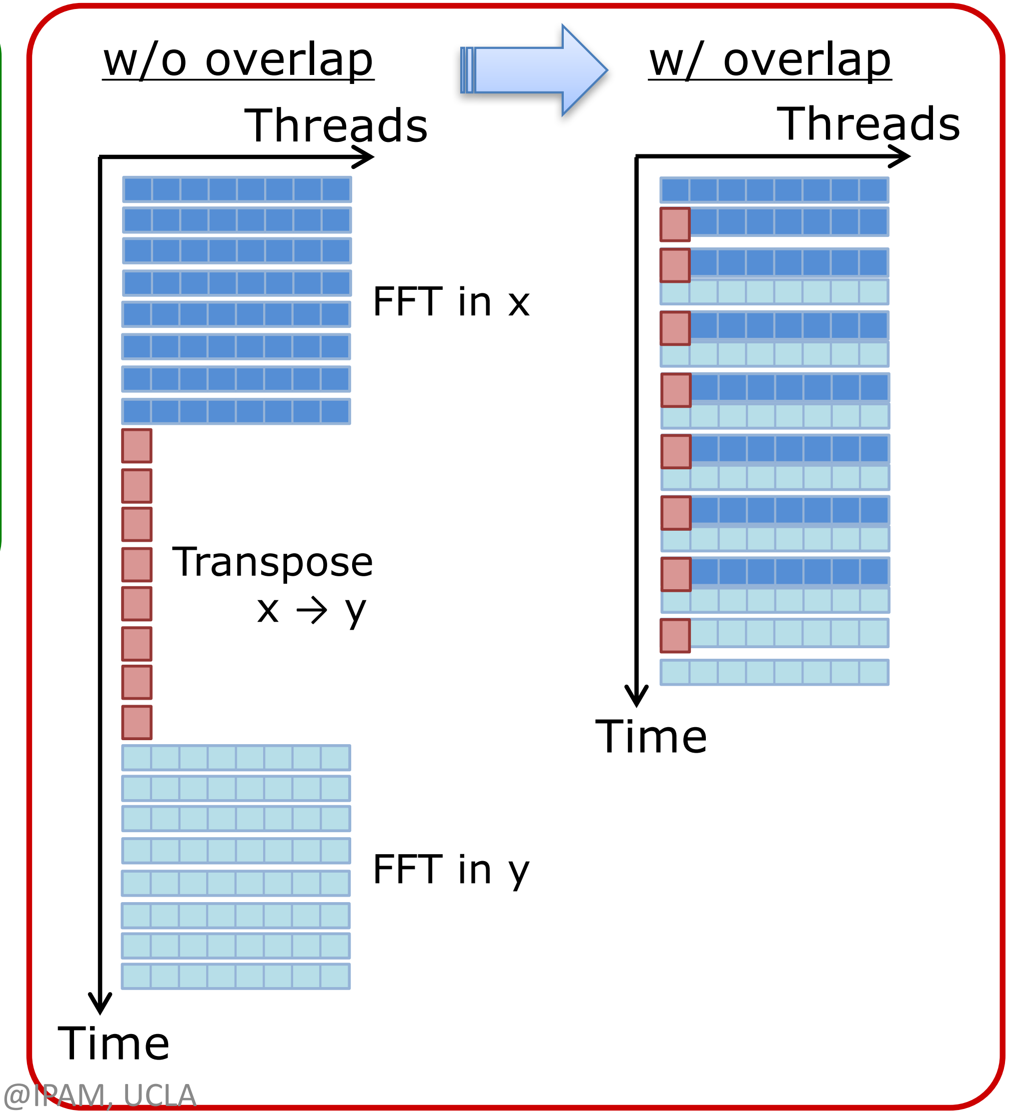

Effects of the optimizations

- The segmented rank mapping reduces comm. cost.
- The pipelined overlaps efficiently mask comm. cost.

Comparison of the elapsed time

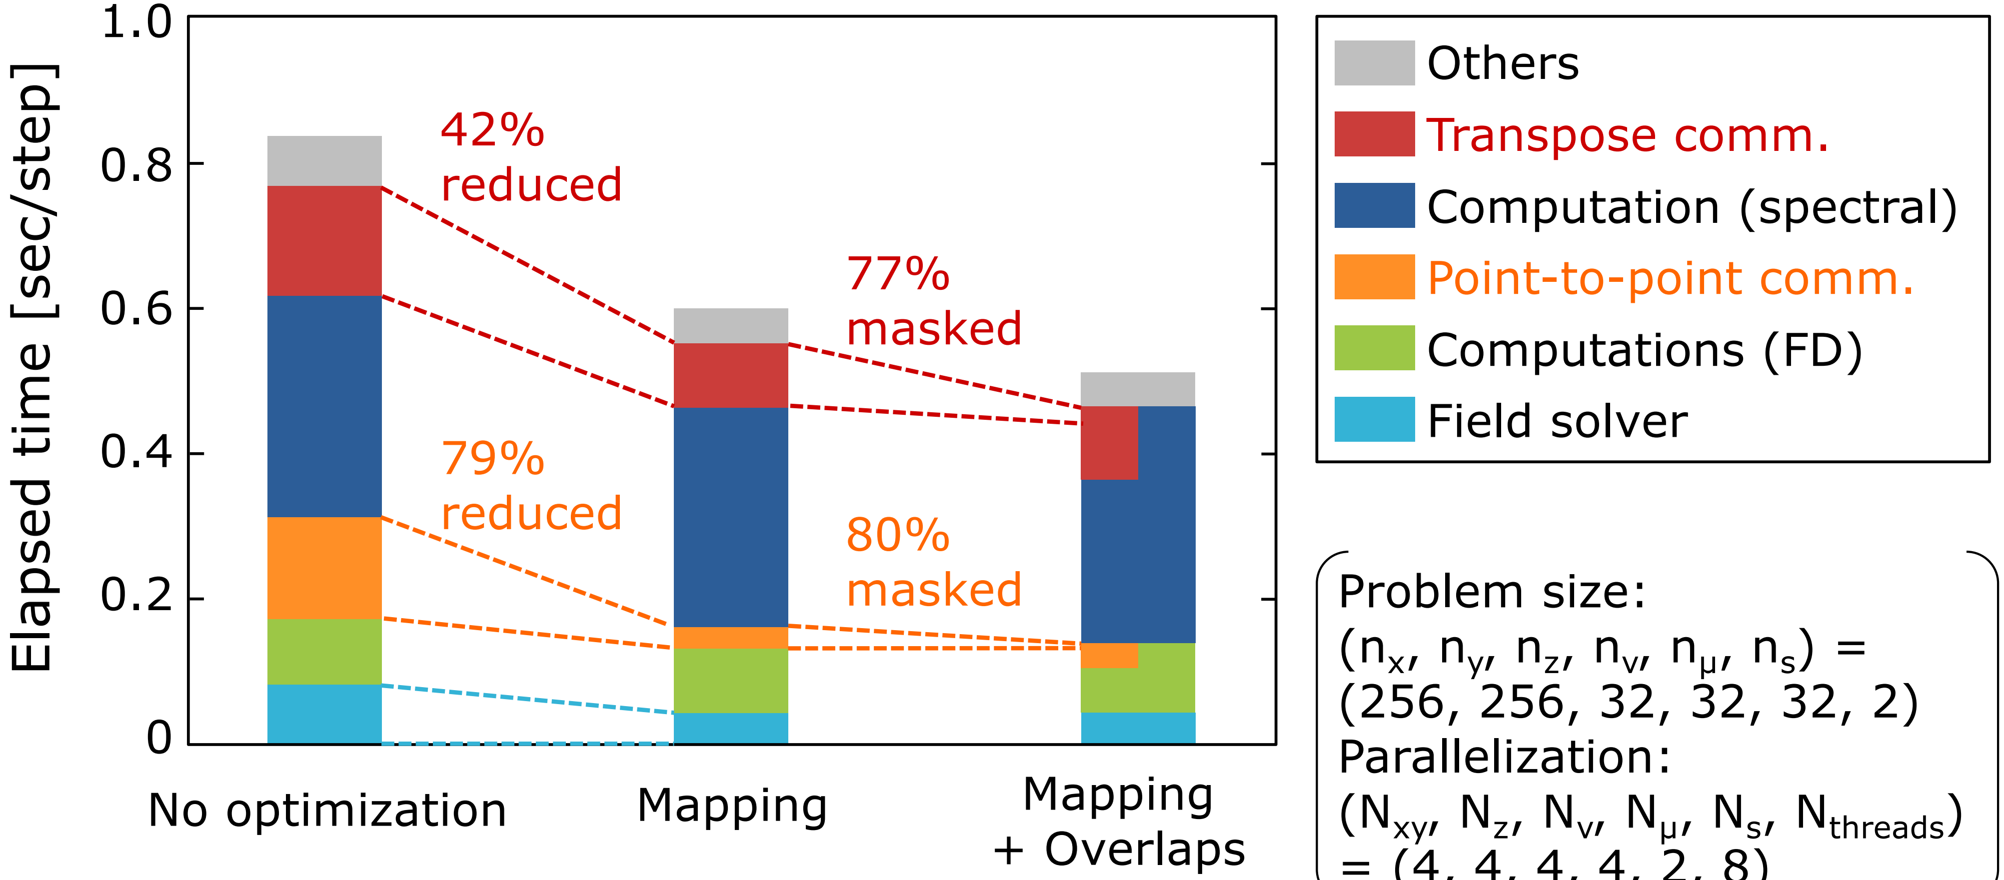

Problem size:
$(n_x, n_y, n_z, n_v, n_\mu, n_s)$ =
(256, 256, 32, 32, 32, 2)
Parallelization:
$(N_{xy}, N_z, N_v, N_\mu, N_s, N_{threads})$
= (4, 4, 4, 4, 2, 8)

* Pipelined overlaps ideally mask $(N_{threads}-1)/N_{threads}$ ~ 87.5%.

Strong scaling toward million cores

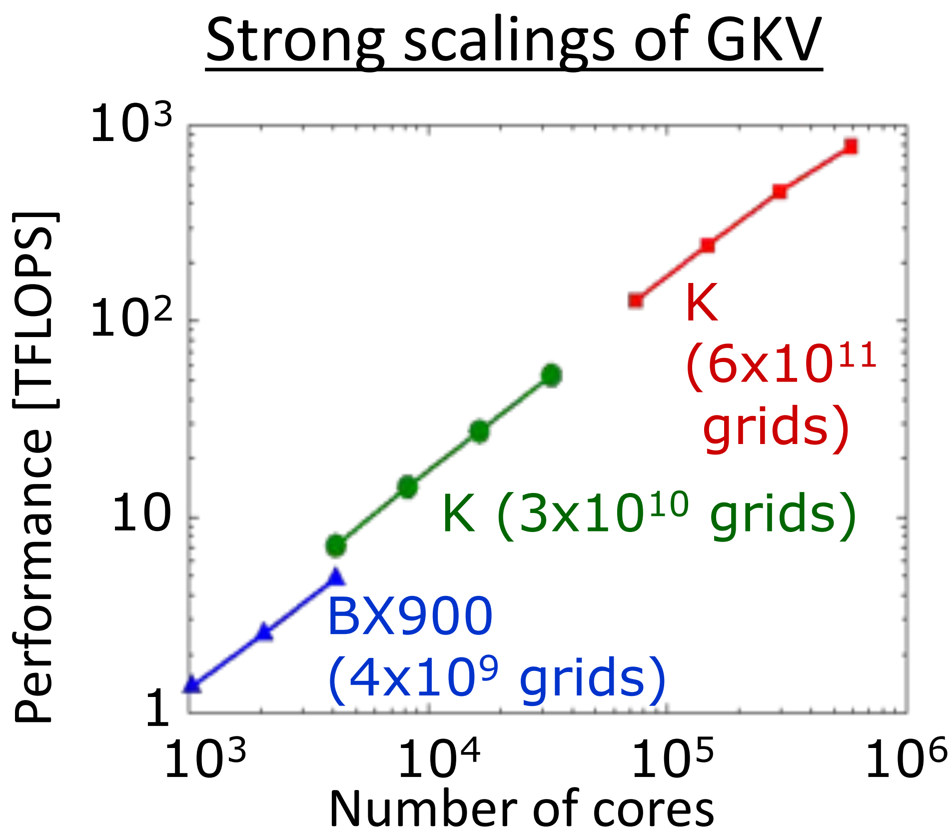

- Excellent strong scaling up to ~ **600k cores**.
- High parallel efficiency ~ **99.99994%**.
- Flops/Peak is 8.3~10.8%.

The highly-optimized code enables multi-scale turbulence simulations from electron to ion scales.

Problem size:
$(n_x, n_y, n_z, n_v, n_\mu, n_s) = (1024, 1024, 96, 96, 32, 2)$
Parallelization:
$(N_{xy}, N_z, N_v, N_\mu, N_s, N_{threads}) = (8\text{-}64, 12, 12, 4, 2, 8)$

Nonlinear interactions in kinetic plasma turbulence

Fluctuations of δf on (x,v)-space

- Analysis of the distribution function δf provides us fundamental information on plasma turbulence.
  - Anisotropic flow patterns (zonal flows and 2D turbulence)
  - Generation of smaller (x,v)-scales (cascading)

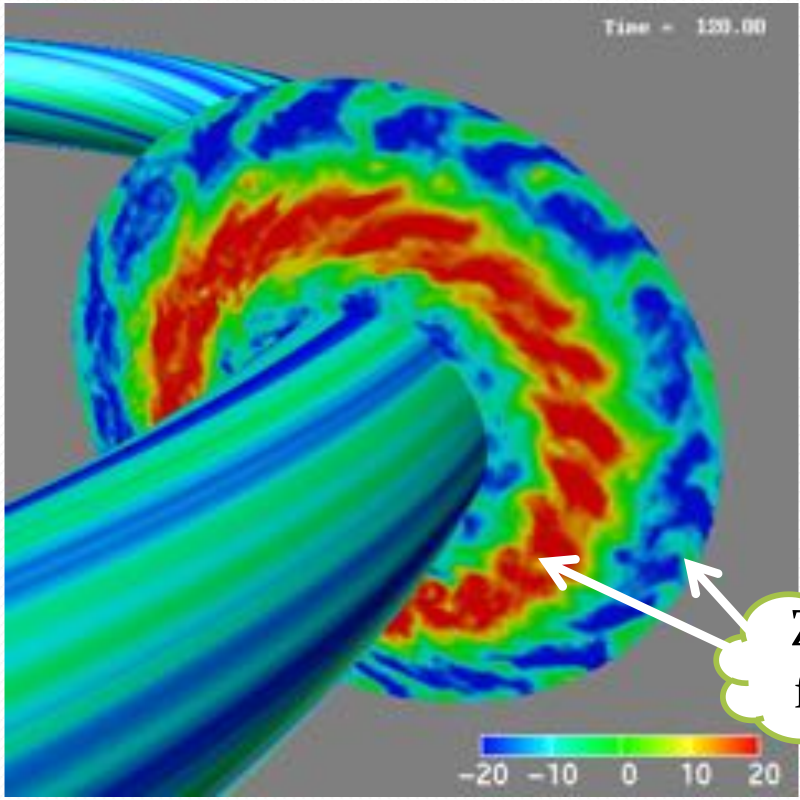

Zonal flows

Snapshot of $\delta f_s^{(g)}$ on v-space

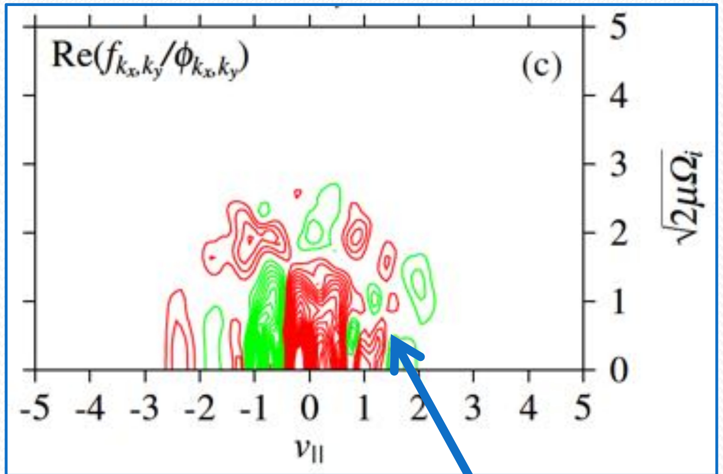

Fine-scale fluctuations

Entropy Balance and Transfer

- A quadratic functional of $\delta f_{i\boldsymbol{k}_\perp}^{(g)}$, that is, $\delta S_{i\boldsymbol{k}_\perp}$, is a measure of fluctuation, "entropy variable"
- Production rate of $\delta S_{i\boldsymbol{k}_\perp}$ balances with transport $Q_{i\boldsymbol{k}_\perp}$ and dissipation $D_{i\boldsymbol{k}_\perp}$
- In kinetic plasma turbulence, $\delta S_{i\boldsymbol{k}_\perp}$ is produced with fine velocity-space structures by $u(x_i)\dfrac{\partial}{\partial x_j} f(x_i, x_j, \ldots)$,
- and is transferred in the k space through interactions of turbulence and zonal flows

Under the periodic boundary condition in $\boldsymbol{x}_\perp$

$$\delta S_{\mathrm{i}\boldsymbol{k}_\perp} = \left\langle \int d\boldsymbol{v} \frac{|\delta f_{\mathrm{i}\boldsymbol{k}_\perp}^{(\mathrm{g})}|^2}{2F_\mathrm{M}} \right\rangle$$

$$\frac{\partial}{\partial t}(\delta S_{\mathrm{i}\boldsymbol{k}_\perp} + W_{\boldsymbol{k}_\perp}) = L_{T_\mathrm{i}}^{-1} Q_{\mathrm{i}\boldsymbol{k}_\perp} + \mathcal{T}_{\mathrm{i}\boldsymbol{k}_\perp} + D_{\mathrm{i}\boldsymbol{k}_\perp}$$

$$Q_{\mathrm{i}\boldsymbol{k}_\perp} = \mathrm{Re}\left\langle v_{\mathrm{ti}} \int d\boldsymbol{v} \delta f_{\mathrm{i}\boldsymbol{k}_\perp}^{(\mathrm{g})} \left(\frac{m_\mathrm{i} v_\parallel^2 + 2\mu B}{2T_\mathrm{i}} \right) i k_y \rho_{\mathrm{ti}} \frac{e\delta\psi_{\boldsymbol{k}_\perp}^*}{T_\mathrm{i}} \right\rangle$$

$$D_{\mathrm{i}\boldsymbol{k}_\perp} = \mathrm{Re}\left\langle \int d\boldsymbol{v}\, C[h_{\mathrm{i}\boldsymbol{k}_\perp}] \frac{h_{\mathrm{i}\boldsymbol{k}_\perp}^*}{F_\mathrm{M}} \right\rangle$$

Entropy Transfer Function T_k

- Entropy transfer function describes nonlinear interactions in anisotropic turbulence including drift waves and zonal flows. [Sugama+ PoP 2009; Nakata+ PoP 2012]

$$\mathcal{T}_{\mathrm{i}\boldsymbol{k}_\perp} = \sum_{\boldsymbol{q}_\perp} \sum_{\boldsymbol{p}_\perp} \delta_{\boldsymbol{k}_\perp+\boldsymbol{p}_\perp+\boldsymbol{q}_\perp,0} \mathcal{J}_{\mathrm{i}}[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp]$$

$$\mathcal{J}_i[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp] = \left\langle \frac{c}{B}\boldsymbol{b}\cdot(\boldsymbol{p}_\perp\times\boldsymbol{q}_\perp)\int d\boldsymbol{v}\frac{1}{2F_{\mathrm{M}}}\mathrm{Re}[\delta\psi_{\boldsymbol{p}_\perp}h_{\mathrm{i}\boldsymbol{q}_\perp}h_{\mathrm{i}\boldsymbol{k}_\perp} - \delta\psi_{\boldsymbol{q}_\perp}h_{\mathrm{i}\boldsymbol{p}_\perp}h_{\mathrm{i}\boldsymbol{k}_\perp}] \right\rangle$$

h_k: non-adiabatic part of $f^{(g)}$

- Detailed balance relation for the triad transfer function $J[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp]$ holds for the triad interaction with $\boldsymbol{k}_\perp+\boldsymbol{p}_\perp+\boldsymbol{q}_\perp=0$

$$\mathcal{J}_{\mathrm{i}}[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp] + \mathcal{J}_{\mathrm{i}}[\boldsymbol{p}_\perp|\boldsymbol{q}_\perp,\boldsymbol{k}_\perp] + \mathcal{J}_{\mathrm{i}}[\boldsymbol{q}_\perp|\boldsymbol{k}_\perp,\boldsymbol{p}_\perp] = 0$$

Successive Entropy Transfer in Ion Scale Turbulence

[Sugama+ PoP 2009; Nakata+ PoP 2012]

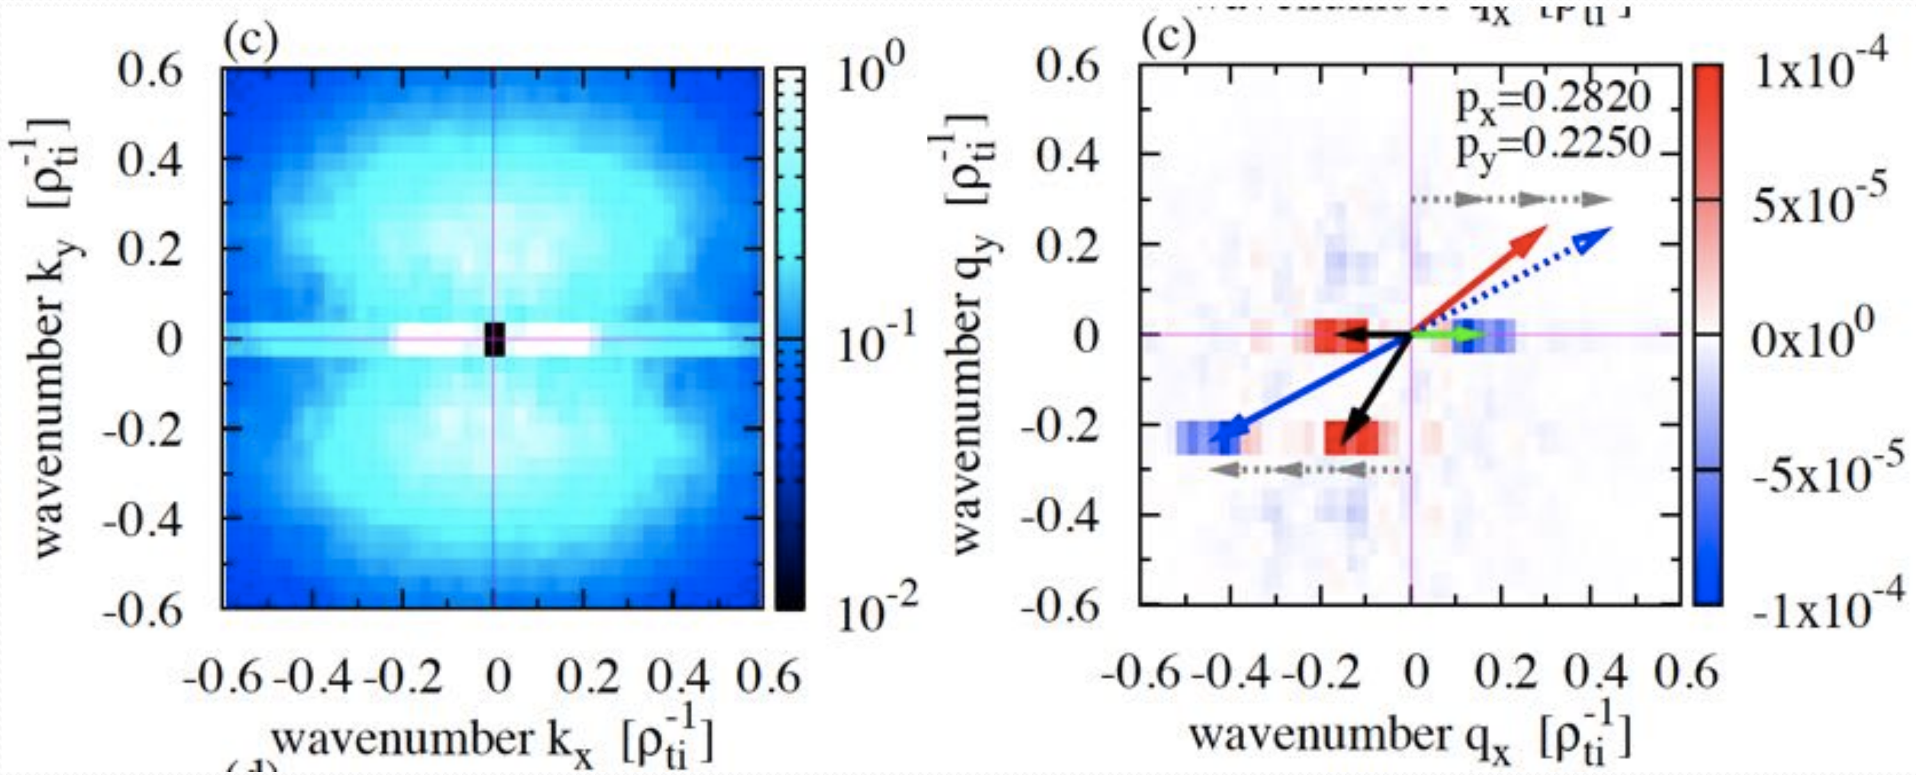

- Turbulence ($k_y \neq 0$) entropy is transferred to higher-k_r side via nonlinear interactions with zonal flows ($k_y = 0$)
  => Spectral broadening and transport reduction

Detailed Entropy Transfer Analysis Demands Computational Costs

- The triad transfer function $J[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp]$ represents an element of nonlinear interactions, and is useful when several players co-exist in the multi-scale turbulence.
- But, computation of $J[\boldsymbol{k}_\perp|\boldsymbol{p}_\perp,\boldsymbol{q}_\perp]$ for the whole $\boldsymbol{k}_\perp$-space demands huge computational costs.

  (nkx,nky,nz,nv,nm,ns)x(nkx,nky)
  = (320,640,64,96,16,2)x(320,640)= 8x10^{15} loops (!?)

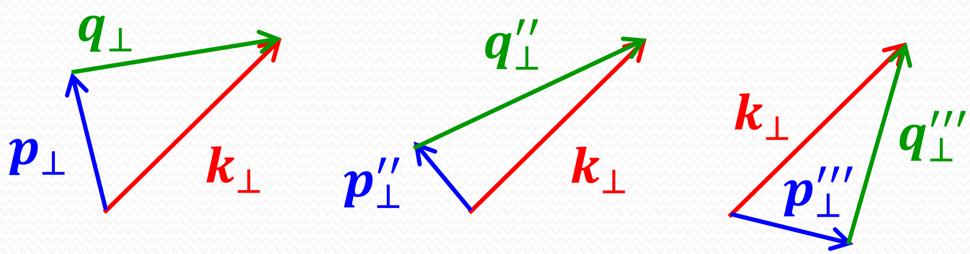

… $\boldsymbol{k}'_\perp$, $\boldsymbol{k}''_\perp$, $\boldsymbol{k}'''_\perp$, …

- A reduce model is necessary for the triad transfer analysis of the multiscale turbulence.

A reduced model for the triad entropy transfer

- Apply the Hermite-Laguerre polynomial expansion,

$$J_{0s\boldsymbol{k}} g_{s\boldsymbol{k}}(z, v_\parallel, \mu) = \sum_{l=0}^{\infty} \sum_{m=0}^{\infty} \frac{M_{s\boldsymbol{k}lm}(z)}{l!} H_l(v_\parallel) L_m(\mu) F_{Ms}$$

$$M_{s\boldsymbol{k}lm}(z) = \int dv^3 H_l(v_\parallel) L_m(\mu) J_{0s\boldsymbol{k}} g_{s\boldsymbol{k}}(z, v_\parallel, \mu)$$

- With an approximation of $J_{0s\boldsymbol{p}} \approx J_{0s\boldsymbol{q}} J_{0s\boldsymbol{k}}$ for $\boldsymbol{k} + \boldsymbol{p} + \boldsymbol{q} = 0$

$$J_{s\boldsymbol{k}}^{\boldsymbol{p},\boldsymbol{q}} \approx \delta_{\boldsymbol{k}+\boldsymbol{p}+\boldsymbol{q},0} \frac{n_s T_s}{2B} \boldsymbol{b} \cdot \boldsymbol{p} \times \boldsymbol{q} Re \sum_{l} \sum_{m} \left\langle \frac{M_{s\boldsymbol{k}lm}}{l!} \left(\phi_{\boldsymbol{p}} M_{s\boldsymbol{q}lm} - \phi_{\boldsymbol{q}} M_{s\boldsymbol{p}lm} \right) \right\rangle$$

  for the electrostatic part.

- In practice, sum over *l* and *m* are taken up to the third order, reducing the computational cost by a factor of $O(10^2)$.

Sub-space transfer analysis

Dividing the wave-number space into sub-spaces, we define the sub-space transfer by a sum over Ω_k

$$\frac{d}{dt}\left(S_{\Omega_k} + W_{\Omega_k}\right) = X_{\Omega_k} + D_{\Omega_k} + E_{\Omega_k} + I_{\Omega_k}$$

$$I_{\Omega_k} = \sum_{\Omega_p} \sum_{\Omega_q} J_{\Omega_k}^{\Omega_p,\Omega_q}$$

Use FFT with filters of Ω_p and Ω_q and compute in the real space

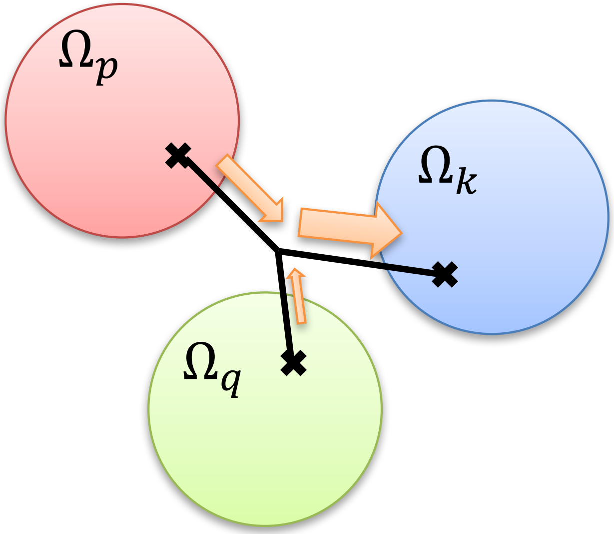

$$J_{\Omega_k}^{\Omega_p,\Omega_q} = \sum_{s=i,e} \sum_{\boldsymbol{k}\in\Omega_k} \sum_{\boldsymbol{p}\in\Omega_p} \sum_{\boldsymbol{q}\in\Omega_q} J_{\boldsymbol{sk}}^{\boldsymbol{p},\boldsymbol{q}}$$

which satisfies

- Symmetry $J_{\Omega_k}^{\Omega_p,\Omega_q} = J_{\Omega_k}^{\Omega_q,\Omega_p}$
- Detailed balance $J_{\Omega_k}^{\Omega_p,\Omega_q} + J_{\Omega_q}^{\Omega_k,\Omega_q} + J_{\Omega_p}^{\Omega_q,\Omega_k} = 0$
- $J_{\Omega_k}^{\Omega_p,\Omega_p} \neq 0 \ (if\ \Omega_p \neq \Omega_k)$

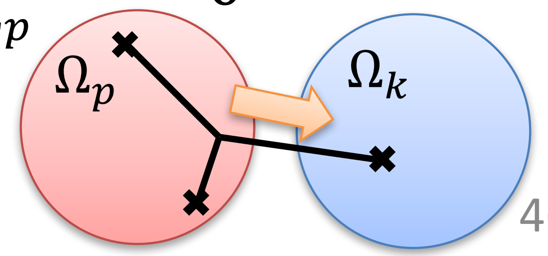

Analysis of the nonlinear mode coupling

Sub-space transfer is

- a generalization of the shell-to-shell transfer for the isotropic turbulence.
- is applied to anisotropic and multi-scale turbulence.

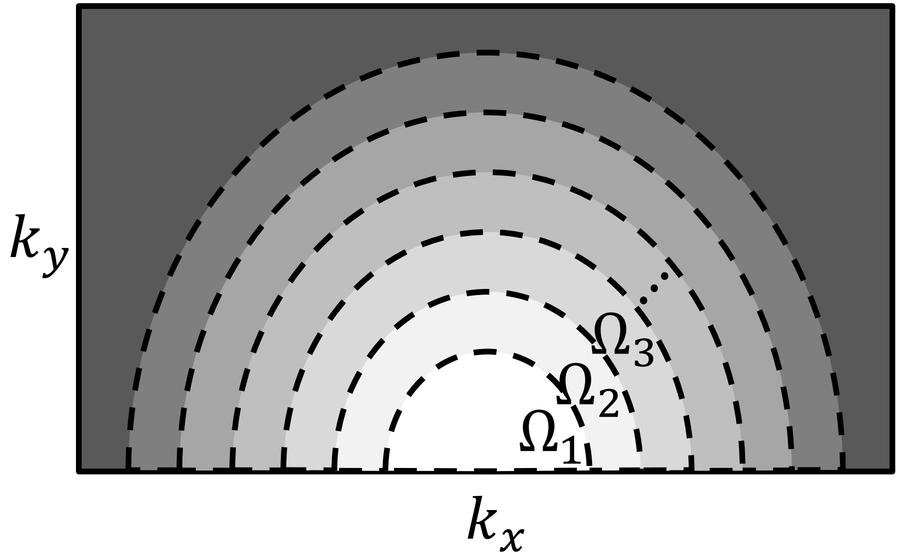

Zonal flows $\Omega_{ZF} = \{k_\theta = 0\}$
Ion-scale turbulence
$\Omega_i = \{k_\theta \neq 0 \cap k_\perp \rho_{ti} \leq 2\}$
Electron-scale turbulence
$\Omega_e = \{k_\theta \neq 0 \cap k_\perp \rho_{ti} > 2\}$

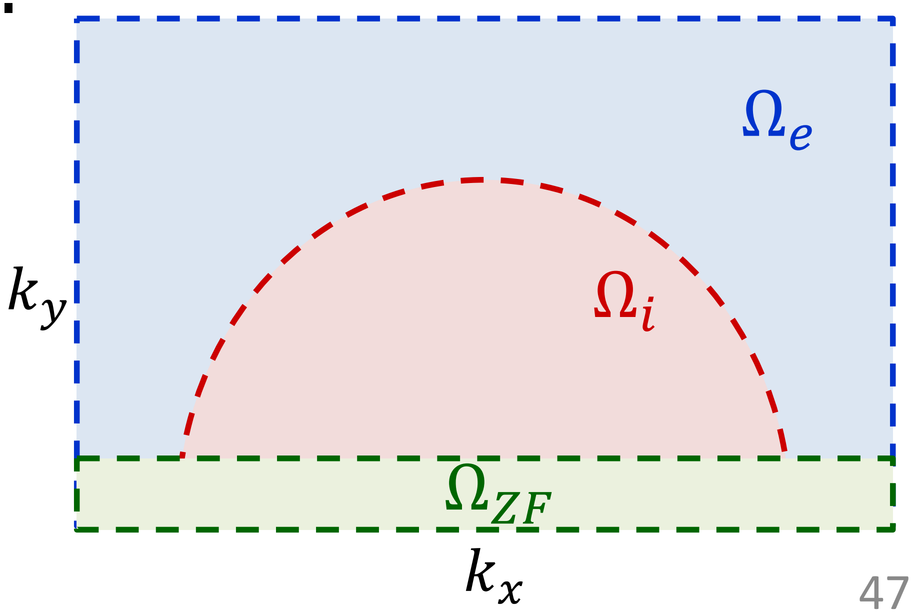

Enhancement mechanism of ITGs by ETGs

- Weaker zonal flow generation in the multi-scale run.
- Reduction of ZF enhances the ion-scale transport.
- Electron-scale turbulence has damping effects on short-wave-length zonal flows.

Ratio of zonal to non-zonal field energy

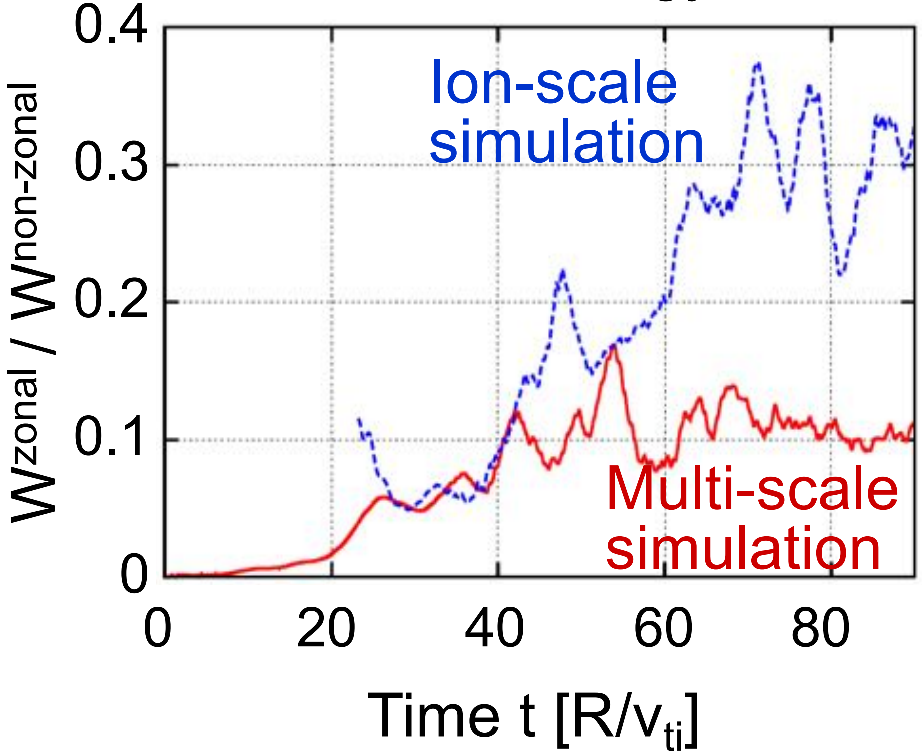

Normalized entropy transfer to zonal modes

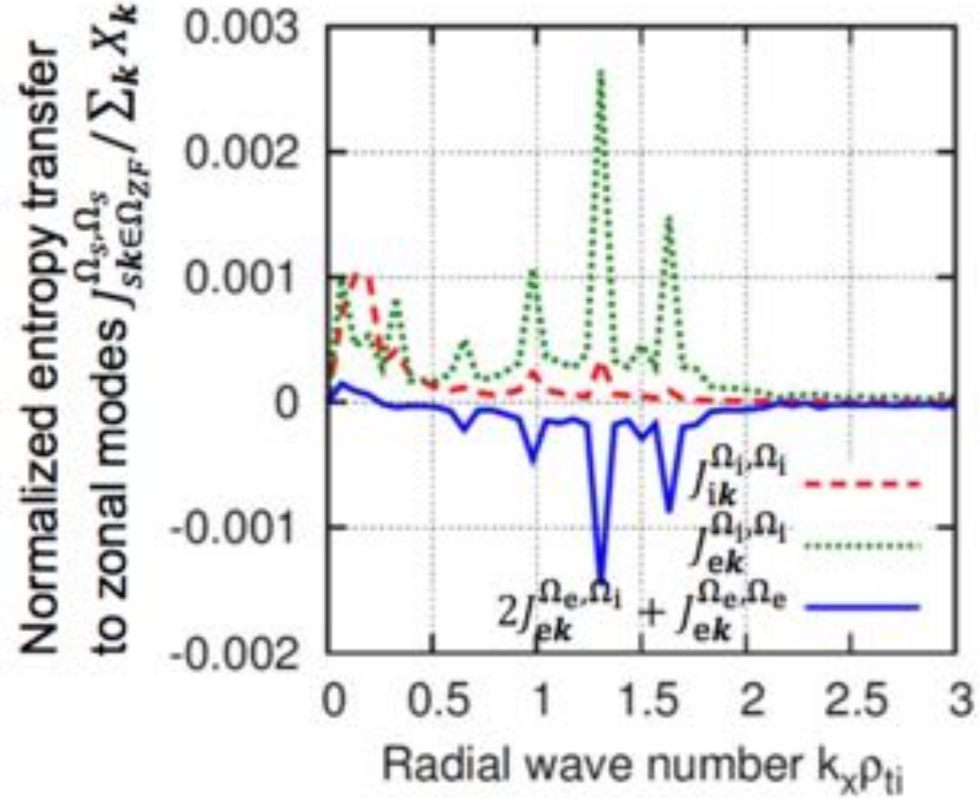

Summary

- Data communication in multiscale plasma turbulence simulation
  - Simulation data of ~5TB for a single variable distributed on 72 k nodes of the K computer are transferred through all-to-all, all-reduce, and one-to-one MPI communications.
  - The inter-node communications optimized for the network topology are efficiently overlapped with computations, achieving strong scaling to ~600 k cores
- Data analysis for the nonlinear turbulence interactions demands computational costs of $O(N^2)$. (N: # of Fourier modes)
  - Reduced model of triad transfer function is developed and applied to the multiscale turbulence results.
  - Sub-space transfer is useful for studying interactions among sub-groups of Fourier modes, such as ion- and electron-scale turbulence and zonal flows.