Performance Measurement of Eulerian Kinetic Code on the Xeon Phi KNL

Takayuki Umeda[1] & Keiichiro Fukazawa[2]

1.Nagoya University, Institute for Space-Earth Environmental Research

2.Kyoto University, Academic Center for Computing and Media Studies

Eulerian Kinetic (Vlasov) Simulations for Space Plasma Studies

Basic equations for collisionless space plasma:

- Maxwell equations (for electromagnetic wave propagations)

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \qquad \nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{1}{c^2}\frac{\partial \mathbf{E}}{\partial t}$$

Computational load less than 0.1%

- Collisionless Boltzmann equation with electromagnetic field (known as Vlasov equation for charged particle motions)

$$\frac{\partial f_s}{\partial t} + \boxed{\mathbf{v}\cdot\frac{\partial f_s}{\partial \mathbf{x}}} + \boxed{\frac{q_s}{m_s}(\mathbf{E}+\mathbf{v}\times\mathbf{B})\cdot\frac{\partial f_s}{\partial \mathbf{v}}} = 0$$

f(x, y, ~~z~~, vx, vy, vz)

6D! ⇒ 5D

Operator splitting into three equations

$$\frac{\partial f_s}{\partial t} + \mathbf{v}\cdot\frac{\partial f_s}{\partial \mathbf{x}} = 0 \qquad \frac{\partial f_s}{\partial t} + \frac{q_s}{m_s}\mathbf{E}\cdot\frac{\partial f_s}{\partial \mathbf{v}} = 0 \qquad \frac{\partial f_s}{\partial t} + \frac{q_s}{m_s}(\mathbf{v}\times\mathbf{B})\cdot\frac{\partial f_s}{\partial \mathbf{v}} = 0$$

(advection in position by v) (advection in velocity by E) (rotation by B)

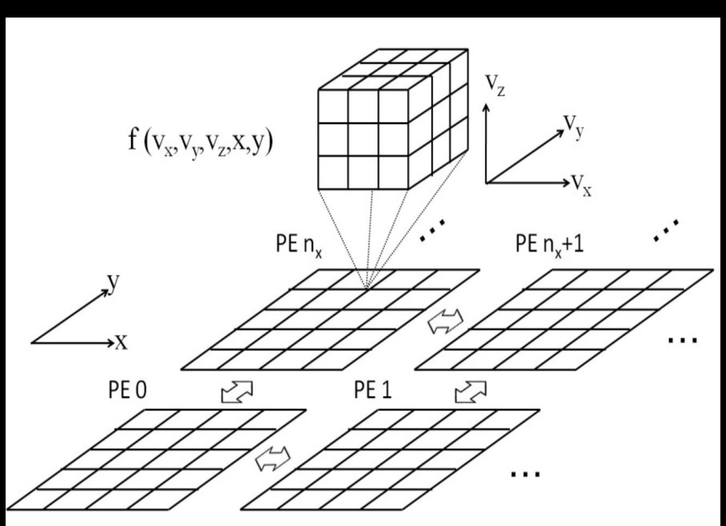

40^5 ~ 4GB

40^6 ~ 160GB

- Hybrid parallelism is adopted to reduce number of processes
- Large number of dimensions (up to 6)
⇒ Requires huge memory
- Length of each loop is short: 20-40
⇒ number of threads > loop length in many core environments.
- Multiple loops are thread-parallelized by loop collapsing of OpenMP

Examples

2D position & 3D velocity (5D)

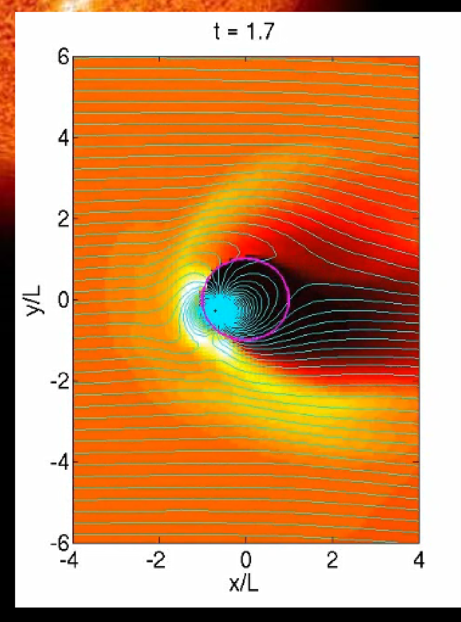

Magnetosphere *Umeda & Fukazawa EPS 2015*

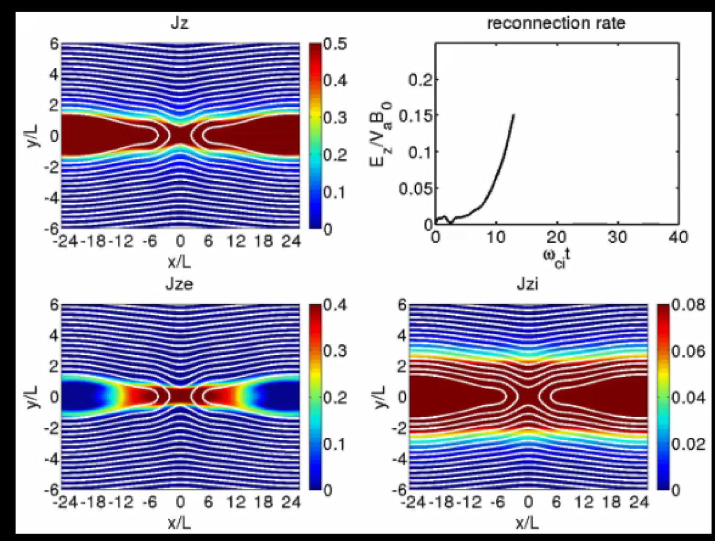

Current Layer *Umeda et al. PoP 2012*

2D Position & 2D Veclotiy (4D)

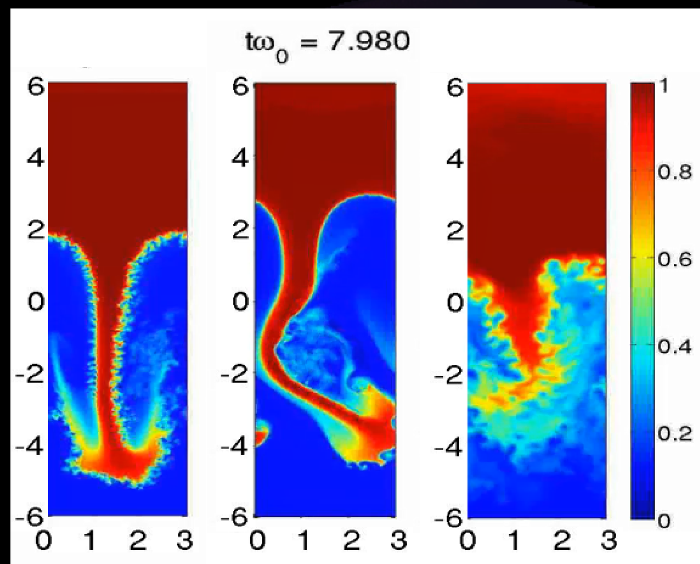

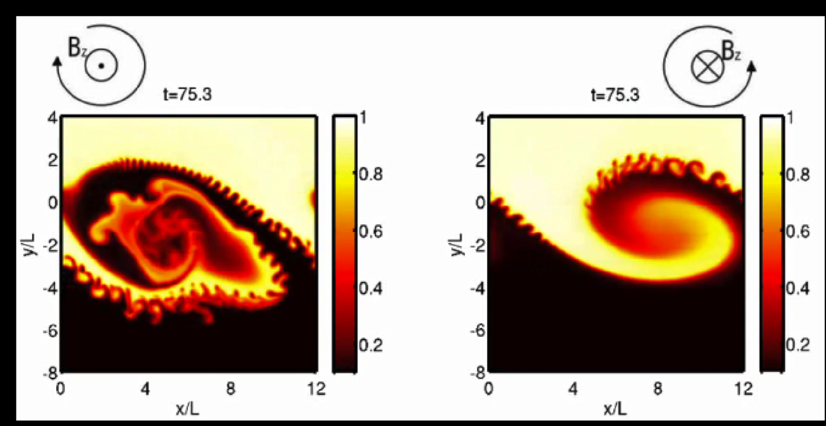

Velocity shear *Umeda et al. PPCF 2014* Density shear *Umeda et al. PoP 2017*

System Description

- Xeon Phi 7250 (68 cores, 16GB MCDRAM)
- 96GB DDR4
- Intel Compiler Ver.17.0.1
  – Option: -ipo -ip -O3 -xMIC-AVX512
- Number of grids: 40*40*40*128*64*2 (~28GB > MCDRAM)
- Elapsed time for 5 time steps is measured
- Memory mode: Flat, Cache
- Cluster mode: All2All, Hemisphere, Quadrant, SNC-2, SNC-4
⇒ Memory and Cluster modes are selected from BIOS
- Default environmental variables

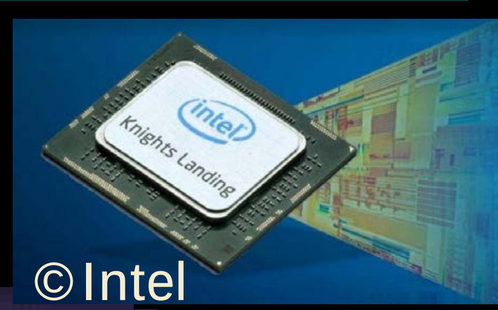

Performance Measurements

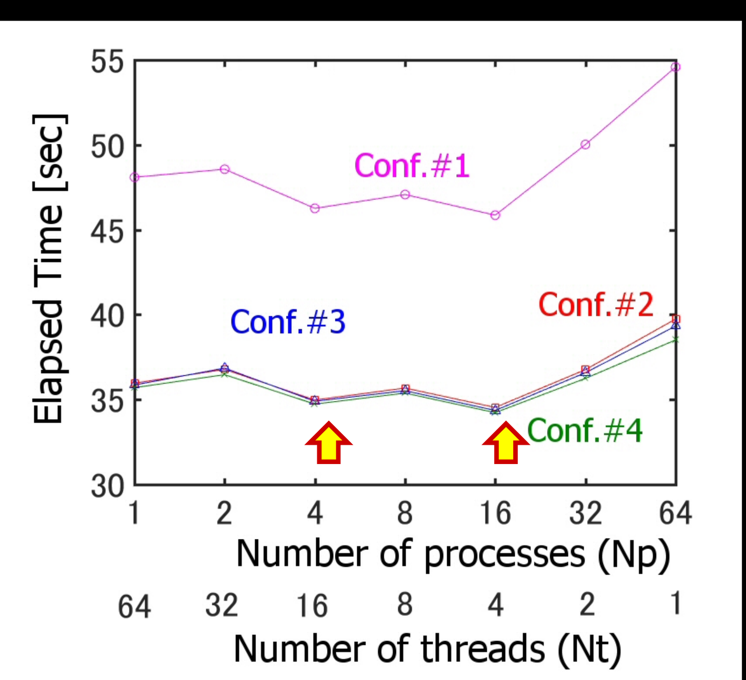

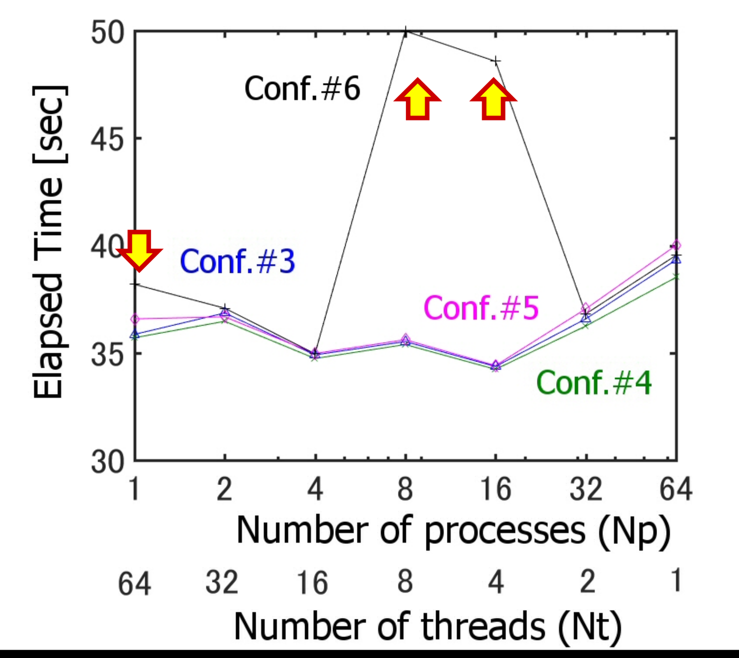

- #1 Flat – All2All
- #2 Cache – All2All
- #3 Cache – Hemisphere
- #4 Cache – Quadrant
- #5 Cache – SNC-2
- #6 Cache – SNC-4

- 64 cores are used (4 cores are free)
- Number of processes (Np) are changed by fixing the total load 64 = Np*Nt (number of threads)
- Flat-MPI (Np=64) is slowest
- Hybrid parallelism with 4 and 16 processes is fastest
- Cache mode is ~1.5 times faster than Flat mode
- Small difference among the performance of All2All, Hemisphere, and Quadrant
- SNC-2 and SNC-4 have a tendency of performance similar to Hemisphere/Quadrant
- However, performance becomes worse with some number of threads
  – 1 process – 64 threads for SNC-2(#5)
  – 1 process – 64 threads for SNC-4(#6)
  – 8 and 16 processes for SNC-4(#6)
  ⇒the performance becomes further worse with KMP_AFFINITY=compact environmental variable!!

Strong Scaling with Np = 4

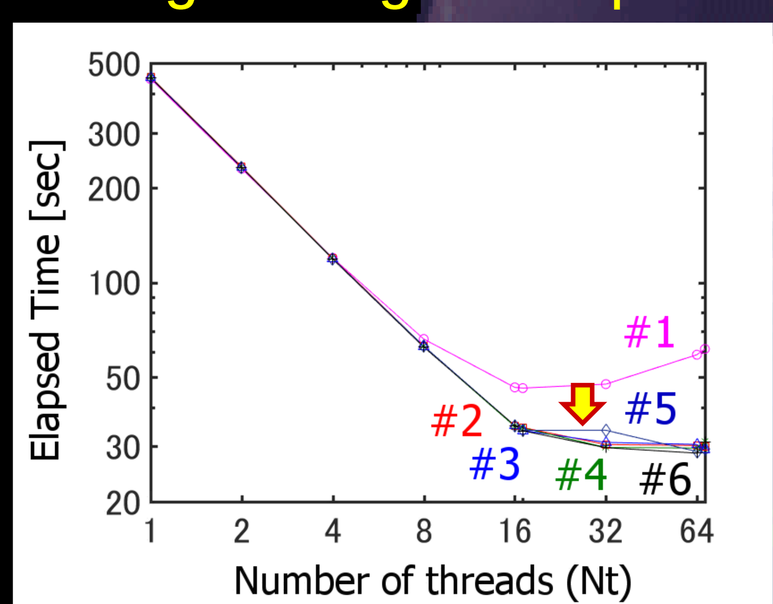

Strong Scaling with Np = 16

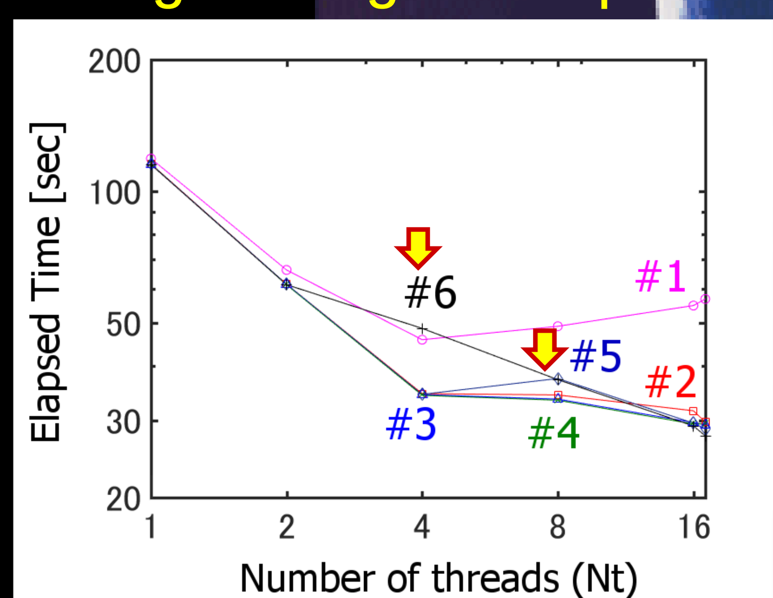

Np = 4

- Scales up to 17 threads
- Hyper Threads (HT) is effective with 32, 64, and 68 threads
- Performance loss with HT for Flat mode(#1)
- Performance loss with 32 threads for SNC-2(#5)

Np = 16

- Scales up to 4 threads
- HT is effective with 8, 16, and 17 threads
- Performance loss with HT for Flat mode(#1)
- Performance loss with 8 threads for SNC-2(#5)
- Performance loss with 4 and 8 threads for SNC-4(#6)