

# Performance Evaluation of Kinetic Code on Scalar Processors

Takayuki Umeda

Nagoya University, Institute for Space-Earth Environmental Research

## Eulerian Kinetic (Vlasov) Simulations for Space Plasma Studies

Basic equations for collisionless space plasma:

- Maxwell equations (for electromagnetic wave propagations)

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

Computational load less than 0.1%

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

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

$f(x, y, z, vx, vy, vz)$   
6D!  $\Rightarrow$  5D

Operator splitting into three equations [Umeda et al. 2012a]

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \frac{\partial f_s}{\partial \mathbf{x}} = 0$$


$$\frac{\partial f_s}{\partial t} + \frac{q_s}{m_s} \mathbf{E} \cdot \frac{\partial f_s}{\partial \mathbf{v}} = 0$$

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

(advection in position by  $\mathbf{v}$ ) (advection in velocity by  $\mathbf{E}$ ) (rotation by  $\mathbf{B}$ )

Non-oscillatory, positive and conservative interpolation is used

[Umeda 2009; Umeda et al. 2012]



[Umeda et al. 2012b]

$40^5 \sim 4\text{GB}$

$40^6 \sim 160\text{GB}$

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

[Umeda et al. 2016]

## Performance Tuning

```

1 do n=0, nvz+1
2 do m=0, nvx+1
3 do l=0, nvx+1
4 ffi(1)=1.0d0/ff(1,m,n,i,j)
5 end do
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30 do l=0, nvx+1
31 gfyx = abs((gfy(1)*ffi(1))* gfy(1)) *0.5d0
32 gfyx = abs((gfy(1)*ffi(1))* gfy(1)) *0.5d0
33 gfyx = abs((gfy(1)*ffi(1))* gfy(1)) *0.5d0
34 gfyx = abs((gfy(1)*ffi(1))* gfy(1))*inv3
35 dfx(1+l0,m ,n ) = dfx(1+l0,m ,n ) + (gfy(1)+(gfyx-gfyx)*sx)
36 dfx(1+l0,m ,n+nv) = dfx(1+l0,m ,n+nv) - (gfyx-gfyx) *sx
37 dfx(1+l0,m+mv,n ) = dfx(1+l0,m+mv,n ) - (gfyx-gfyx) *sx
38 dfx(1+l0,m+mv,n+nv) = dfx(1+l0,m+mv,n+nv) + gfyx
39 dfy(1 ,m+m0,n ) = dfy(1 ,m+m0,n ) + (gfy(1)+(gfyx-gfyx-gfyx)*sy)
40 dfy(1+l0,m+m0,n ) = dfy(1+l0,m+m0,n ) - (gfyx-gfyx) *sy
41 dfy(1 ,m+m0,n+nv) = dfy(1 ,m+m0,n+nv) - (gfyx-gfyx) *sy
42 dfy(1+l0,m+m0,n+nv) = dfy(1+l0,m+m0,n+nv) + gfyx
43 dfz(1 ,m ,n+n0) = dfz(1 ,m ,n+n0) + (gfy(1)+(gfyx-gfyx-gfyx)*sz)
44 dfz(1 ,m+mv,n+n0) = dfz(1 ,m+mv,n+n0) - (gfyx-gfyx) *sz
45 dfz(1+l0,m ,n+n0) = dfz(1+l0,m ,n+n0) - (gfyx-gfyx) *sz
46 dfz(1+l0,m+mv,n+n0) = dfz(1+l0,m+mv,n+n0) + gfyx
47 end do
48 end do
49 end do

```

As-is code

- There dependences on array access in the loop iteration at Lines 35—46

```

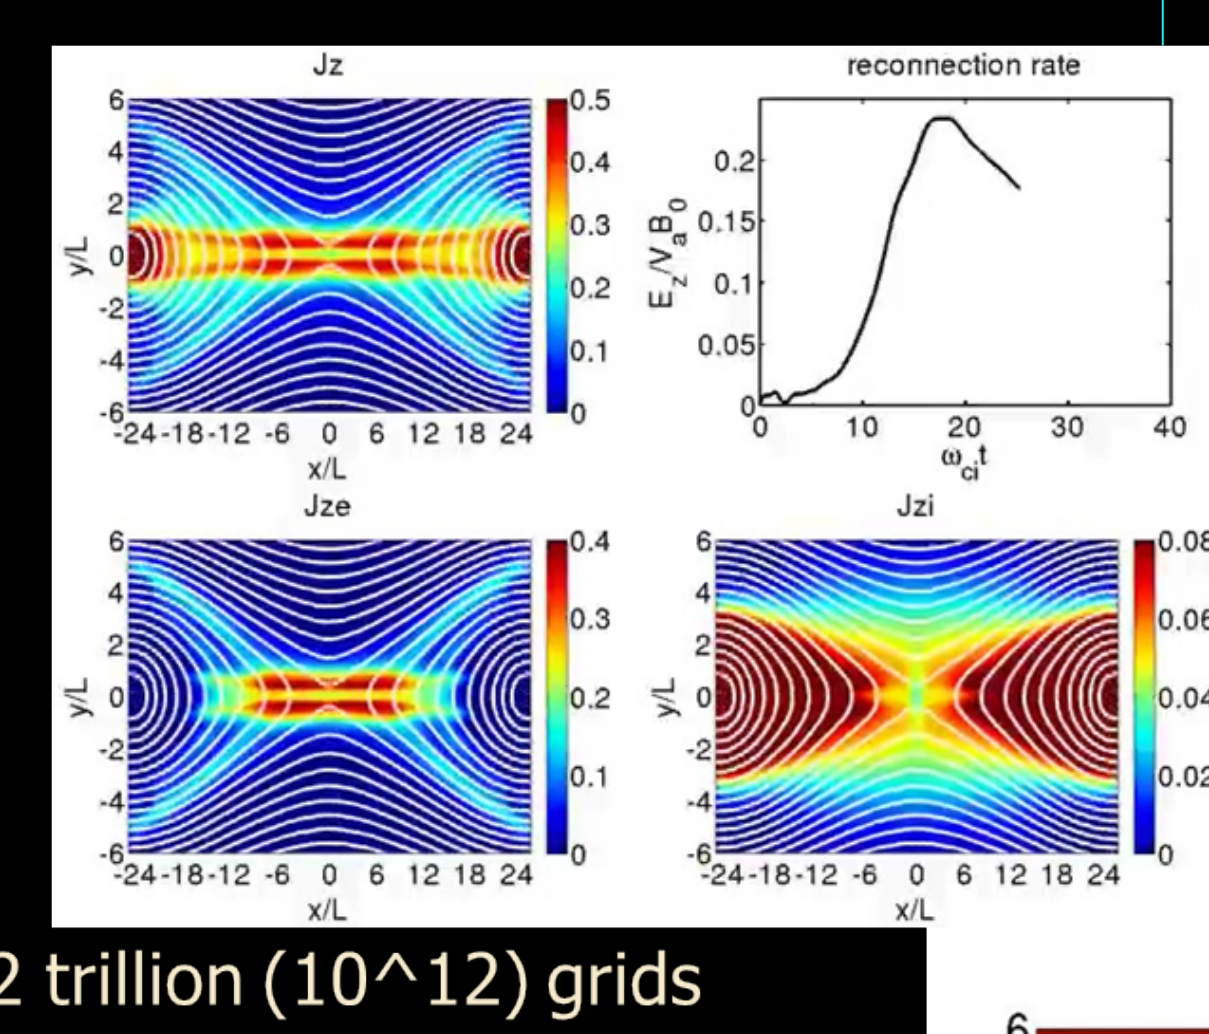
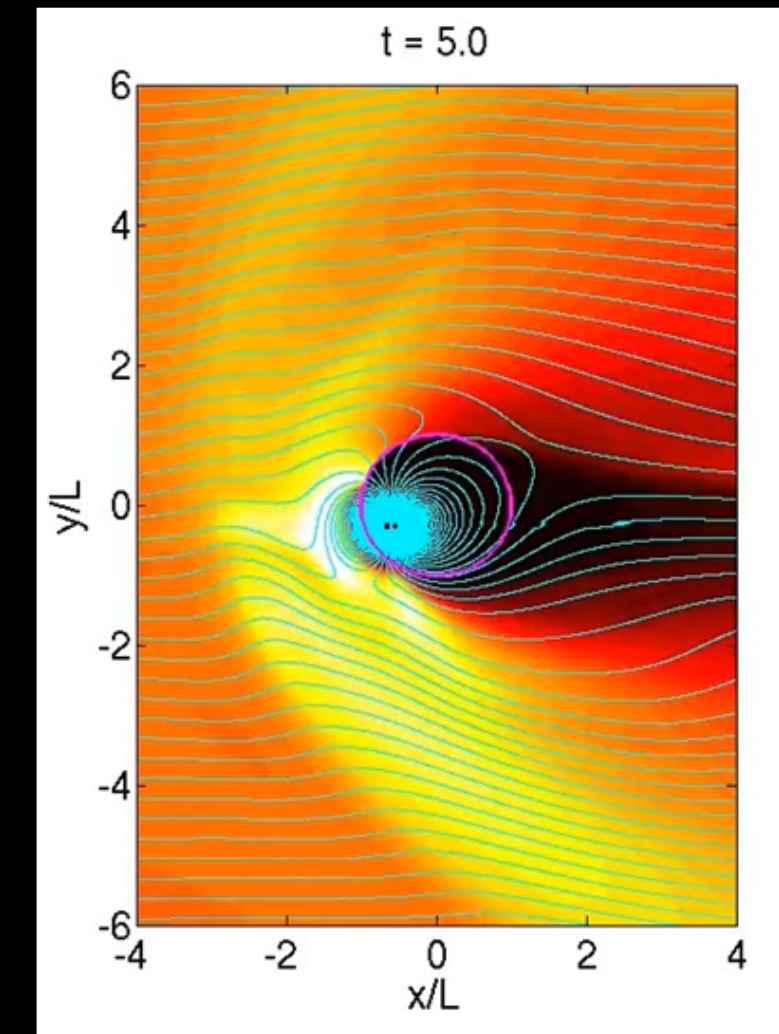
30 do ll=nvxs-1,nvxe+1
31 ffi = 1.0d0/ff(ll,mm,nn,ii,jj)
32 gfy(11) = abs(gfy(11)*ffi)
33 gfy(11) = abs(gfy(11)*ffi)
34 gfy(11) = abs(gfy(11)*ffi)
35 end do
36 do ll=nvxs-1,nvxe+1
37 dfx(11+l0,mm ,nn ) = dfx(11+l0,mm ,nn ) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)+(1.0d0-gfy(11)*0.5d0)*sx
38 dfx(11 ,mm+m0,nn ) = dfx(11 ,mm+m0,nn ) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)+(1.0d0-gfy(11)*0.5d0)*sy
39 dfz(11 ,mm ,nn+n0) = dfz(11 ,mm ,nn+n0) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)+(1.0d0-gfy(11)*0.5d0)*sz
40 end do
41 do ll=nvxs-1,nvxe+1
42 dfx(11+l0,mm ,nn+nv) = dfx(11+l0,mm ,nn+nv) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sx
43 dfy(11+l0,mm+m0,nn ) = dfy(11+l0,mm+m0,nn ) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sy
44 dfz(11 ,mm+mv,nn+n0) = dfz(11 ,mm+mv,nn+n0) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sz
45 end do
46 do ll=nvxs-1,nvxe+1
47 dfx(11+l0,mm+mv,nn ) = dfx(11+l0,mm+mv,nn ) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sx
48 dfy(11 ,mm+m0,nn+nv) = dfy(11 ,mm+m0,nn+nv) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sy
49 dfz(11+l0,mm ,nn+n0) = dfz(11+l0,mm ,nn+n0) - ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*(gfy(11)*inv3-0.5d0)*sz
50 end do
51 do ll=nvxs-1,nvxe+1
52 dfx(11+l0,mm+mv,nn+nv) = dfx(11+l0,mm+mv,nn+nv) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*gfy(11)*inv3*sx
53 dfy(11+l0,mm+m0,nn+nv) = dfy(11+l0,mm+m0,nn+nv) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*gfy(11)*inv3*sy
54 dfz(11+l0,mm+mv,nn+n0) = dfz(11+l0,mm+mv,nn+n0) + ffi(ll,mm,nn,ii,jj)*gfy(11)*gfy(11)*gfy(11)*inv3*sz
55 end do

```

Tuned code

- Decomposition of most-inner loops  $\Rightarrow$  no loop dependence
- Reduction of temporary work arrays

2D position & 3D velocity (5D)




## Examples

Magnetic shear (current layer)  
Umeda et al. PoP 2012

2D Position & 2D Velocity (4D)

2 trillion ( $10^{12}$ ) grids on K computer (6144 nodes)  
Umeda & Fukazawa EPS 2015



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

## System Description

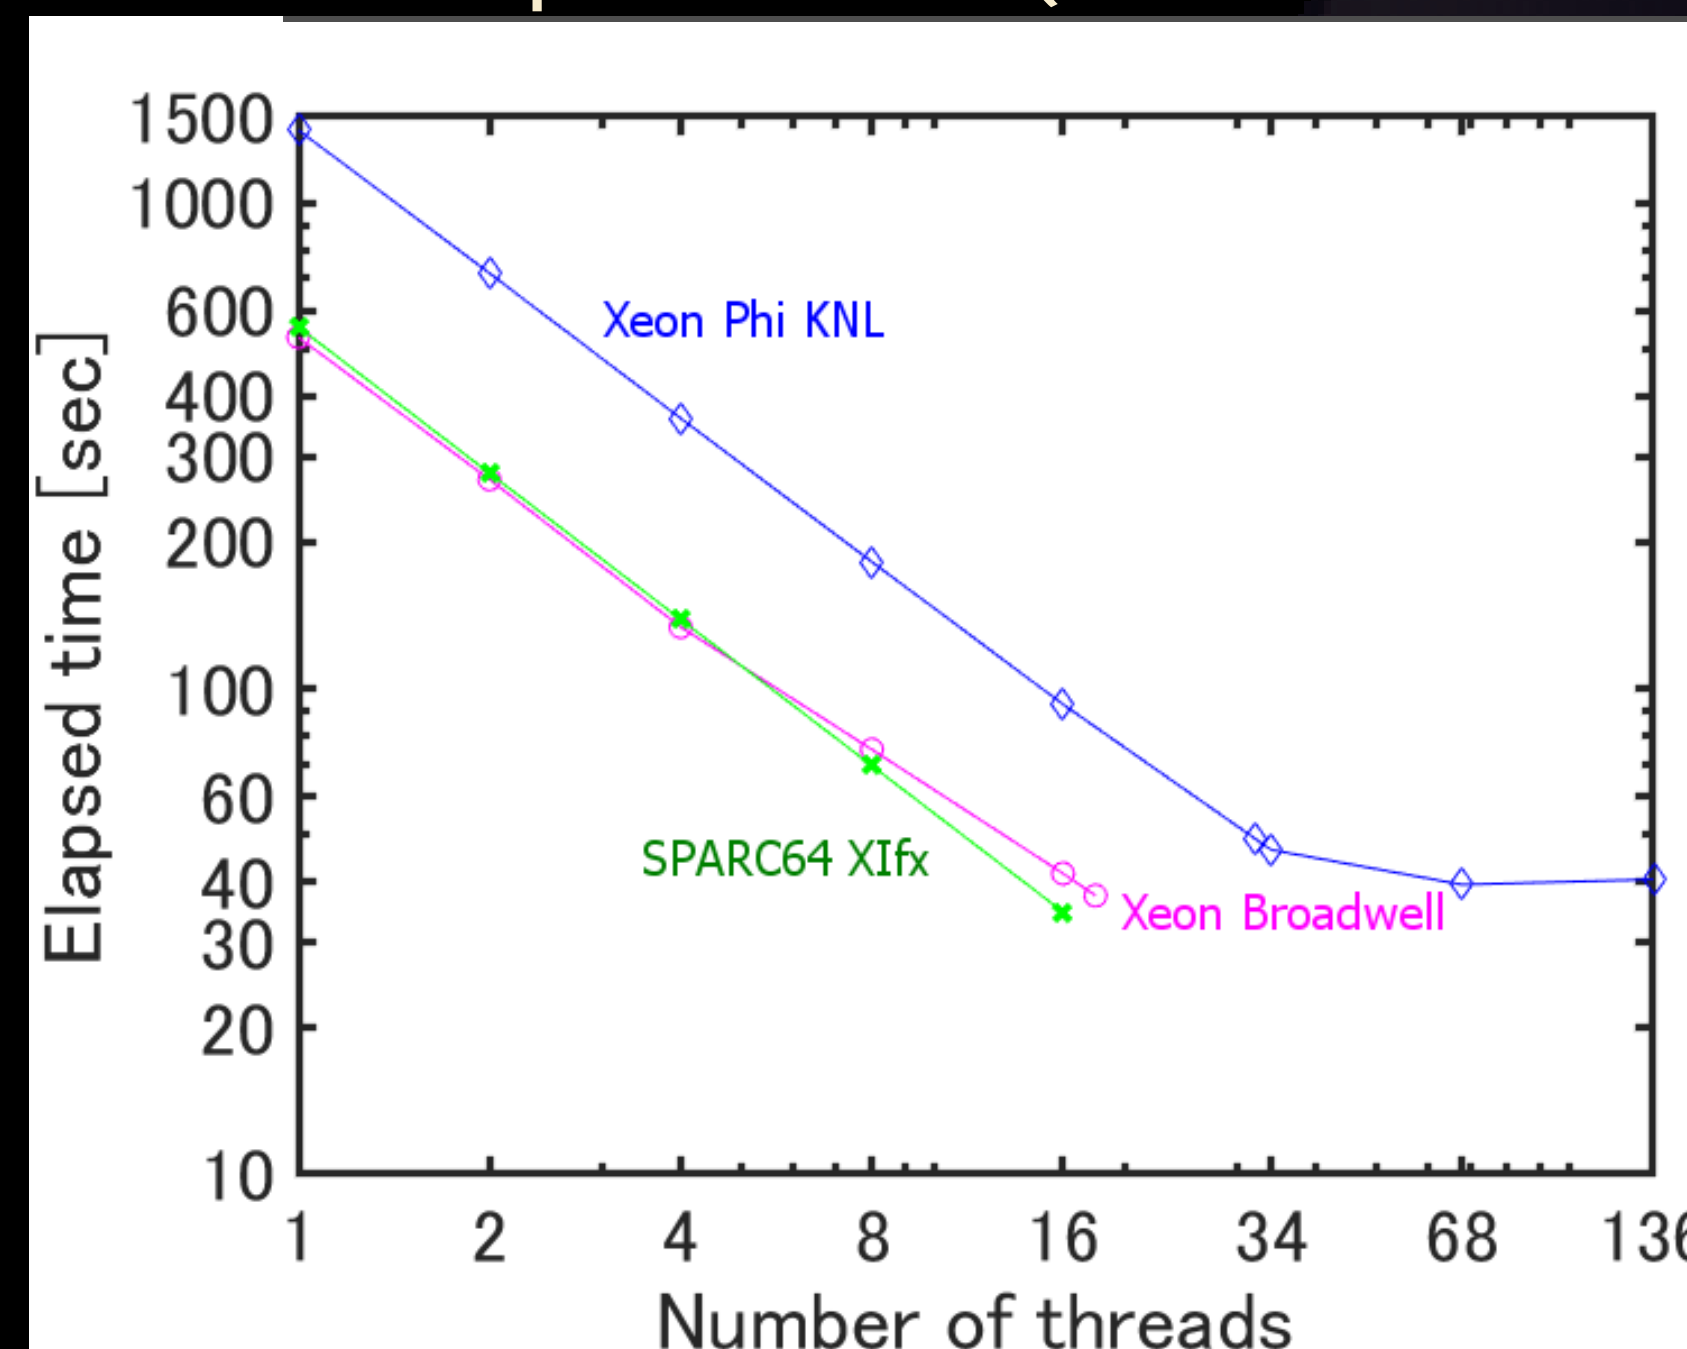
- Xeon Phi 7250 – 96GB DDR4  
– Intel Compiler Ver.17.0.1  
– Option: -ipo -ip -O3 -xMIC-AVX512
- Xeon E5-2697 v4 (dual) – 512GB DDR4  
– Intel Compiler Ver.17.0.1  
– Option: -ipo -ip -O3 -xCORE-AVX2
- SPARC64 (Fujitsu FX100) – 32GB HBM  
– Option: -Kfast,ocl,simd,openmp,parallel
- Number of grids:  $40 \times 40 \times 40 \times 128 \times 64 \times 2$  (~28GB)
- Elapsed time for 5 time steps is measured



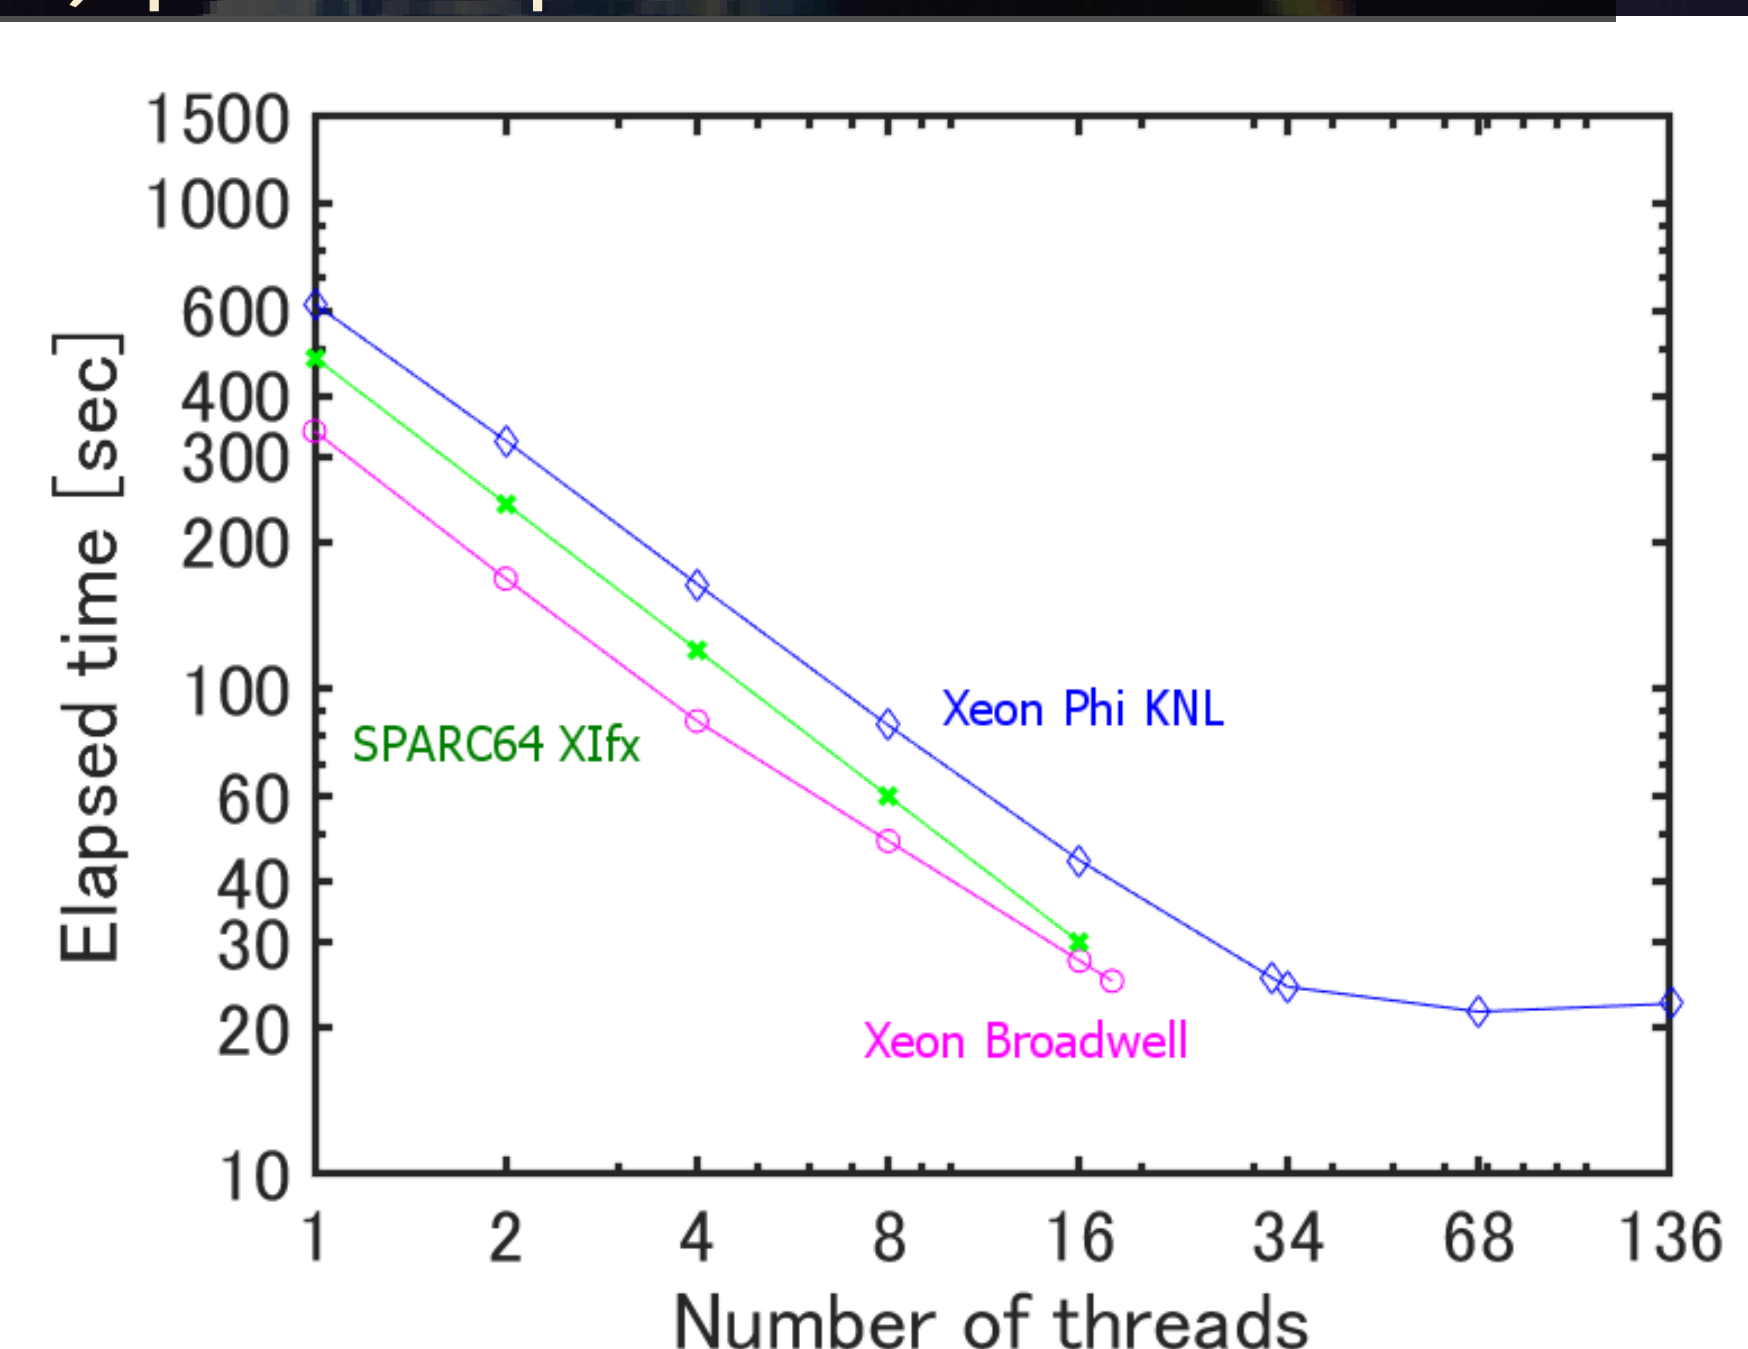
©Nagoya Univ.

## Performance Measurements

Strong scaling with  $N_x \times N_y \times N_{vx} \times N_{vy} \times N_{vz} = 134 \times 70 \times 40^3$   
Two processes (ion and electron) per compute node



As-is code



Tuned code

## Summary

- Scales very well by thread parallelism
- Speedup by performance tuning:
  - Xeon Phi 7250 (KNL) : 2.0 times
  - Dual Xeon E5-2697 v4 (Broadwell) : 1.5 times
  - SPARC64 (Fujitsu FX100) : 1.15 times
- Performance improved by reduction of temporary work arrays on Broadwell and KNL
- Bottlenecks in scheduling of instructions on SPARC64 Xifx