

# Electromagnetic Turbulence Simulations of the Tokamak Scrape-Off Layer

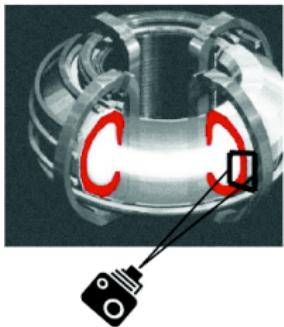
F.D. Halpern, T.-M. Tran, F. Musil,  
P. Ricci, F. Riva, C. Wersal

École Polytechnique Fédérale de Lausanne  
Centre de Recherches en Physique des Plasmas  
CH-1015 Lausanne, Suisse

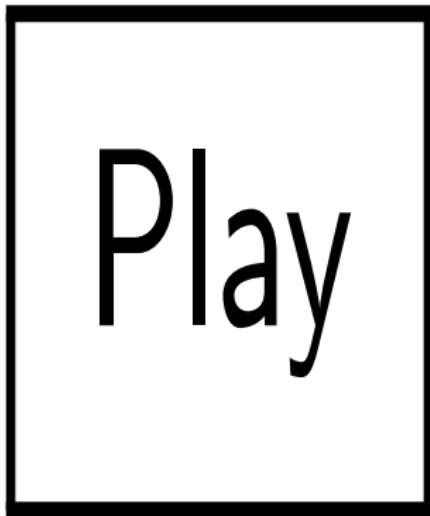
June 3rd, 2015

# A peek at the very edge of a tokamak plasma...

Visible light fluctuations at 400'000fps [S.Zweben (Princeton), J.Terry (MIT)]

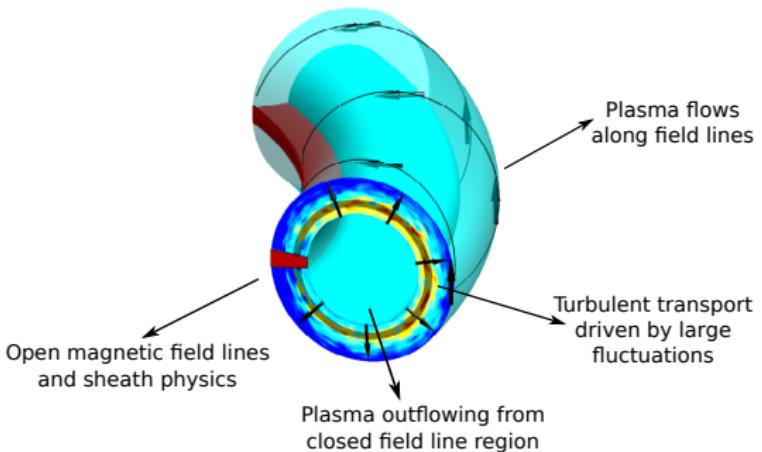


- ▶ Low temperature ( $\sim 10\text{eV}$ ) magnetized plasma
- ▶ Open  $B$  field lines  $\rightarrow$  plasma not confined
- ▶ Low frequency modes  $\omega \ll \omega_{ci}$
- ▶  $L_{fluc} \sim L_{eq}$
- ▶  $n_{fluc} \sim n_{eq}$



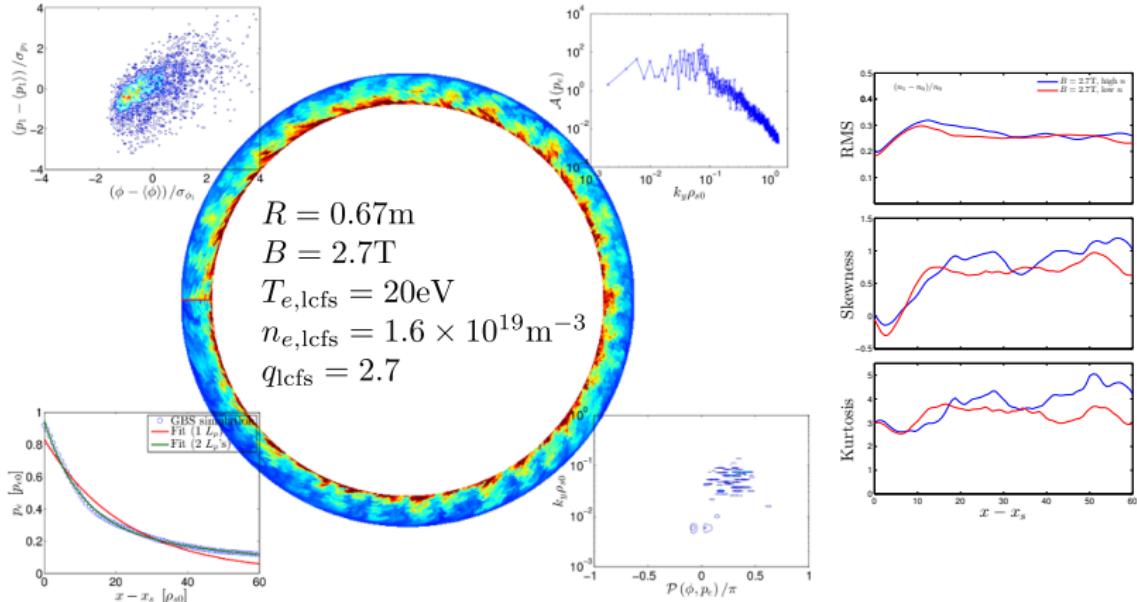
# How do we develop an understanding of SOL dynamics?

Heat load to PFCs, rotation, impurities, L-H transition...



- ▶ **GBS code:** drift-reduced Braginskii eqns., flux-driven
- ▶ Full power balance → profile formation,  $\lambda_q \rightarrow$  heat-flux length
- ▶ Good scaling up to medium size tokamak [Halpern, NF/PPCF (2013-14)]

# Towards the next generation SOL turbulence code...



- ▶ Larger simulations required for experimental comparison
- ▶ Main roadblock to larger plasma sizes is Poisson solver

# Drift-reduced Braginskii equations to describe the SOL

$$\begin{aligned}
 \frac{\partial n}{\partial t} &= -\frac{\rho_\star^{-1}}{B} [\phi, n] + \frac{2}{B} [nC(T_e) + T_e C(n) - nC(\phi)] - n\nabla_{||} v_{||e} - v_{||e}\nabla_{||} n \\
 \frac{\partial \tilde{\omega}}{\partial t} &= -\frac{\rho_\star^{-1}}{B} [\phi, \tilde{\omega}] - v_{||i}\nabla_{||}\tilde{\omega} + \frac{B^2}{n}\nabla_{||}j_{||} + \frac{2B}{n}C(p) + \frac{B}{3n}C(G_i), \quad \tilde{\omega} = \nabla_{\perp}^2(\phi + \tau T_i) \\
 \frac{\partial}{\partial t} \left( v_{||e} + \frac{m_i}{m_e} \frac{\beta_e}{2} \psi \right) &= -\frac{\rho_\star^{-1}}{B} [\phi, v_{||e}] - v_{||e}\nabla_{||}v_{||e} + \frac{m_i}{m_e} \left[ \nu j_{||}/n + \nabla_{||}\phi - \frac{\nabla_{||}p_e}{n} - 0.71\nabla_{||}T_e - \frac{2}{3n}\nabla_{||}G_e \right] \\
 \frac{\partial v_{||i}}{\partial t} &= -\frac{\rho_\star^{-1}}{B} [\phi, v_{||i}] - v_{||i}\nabla_{||}v_{||i} - \frac{2}{3}\nabla_{||}G_i - \frac{1}{n}\nabla_{||}p \\
 \frac{\partial T_e}{\partial t} &= -\frac{\rho_\star^{-1}}{B} [\phi, T_e] - v_{||e}\nabla_{||}T_e + \frac{4}{3}\frac{T_e}{B} \left[ \frac{7}{2}C(T_e) + \frac{T_e}{n}C(n) - C(\phi) \right] + \\
 &\quad + \frac{2}{3} \left\{ T_e \left[ 0.71\nabla_{||}v_{||i} - 1.71\nabla_{||}v_{||e} \right] + 0.71T_e(v_{||i} - v_{||e})\frac{\nabla_{||}n}{n} \right\} + \mathcal{D}_{T_e}^{\parallel}(T_e) \\
 \frac{\partial T_i}{\partial t} &= -\frac{\rho_\star^{-1}}{B} [\phi, T_i] - v_{||i}\nabla_{||}T_i + \frac{4}{3}\frac{T_i}{B} \left[ C(T_e) + \frac{T_e}{n}C(n) - C(\phi) \right] + \\
 &\quad + \frac{2}{3}T_i(v_{||i} - v_{||e})\frac{\nabla_{||}n}{n} - \frac{2}{3}T_i\nabla_{||}v_{||e} - \frac{10}{3}\frac{T_i}{B}C(T_i) + \mathcal{D}_{T_i}^{\parallel}(T_i)
 \end{aligned}$$

- + Sheath BCs consistent with PIC simulations [Loizu, PoP (2012)]
- + Equations for neutral ion physics [Wersal, to be submitted]

## Status of GBS circa 2014 [Ricci, PPCF (2012)]

- ▶ 2nd order FD spatial discretization, RK4 time advance
- ▶ Arakawa scheme for  $\mathbf{E} \times \mathbf{B}$  n.l. advection terms
- ▶ Field-aligned parallel gradient operators
- ▶ MPI domain decomposition in  $(x, z)$  directions
- ▶ Sparse linear solver for elliptical operators  $\sim \nabla_{\perp}^2 \phi$
- ▶ Reasonable scalability up to medium size tokamak
  - ▶  $(ny, nx, nz) = (1384, 128, 256)$ ,  $\Delta t = 10^{-5} R/c_s \sim 1\text{sec}$

## Summary of new features in GBS

	Circa 2014	New GBS
Language	Fortran90	F2003/8
Parallelization	2-D MPI DD ( $z, x$ )	3-D MPI DD ( $z, y, x$ )
Elliptical solver	Parallel sparse	Parallel multigrid
Scalability	$\sim 1024$ cores	$\sim 8192$ cores
ES potential	$\nabla \cdot (n \nabla_{\perp} \phi) \sim \nabla_{\perp}^2 \phi / n$ (Boussinesq)	Retain full $\tilde{\phi} \Leftarrow \tilde{n}$ coupling
EM effects	Small plasma size	Arbitrary size

- ▶ Also: Ongoing effort to port GBS to manycore/GPU architecture using OpenMP/OpenACC/MPI (2016)

## Stencil based parallel multigrid implemented in GBS

- ▶ 2D Cartesian  $(x, y)$  grid topology mapped to a **2D domain decomposition**

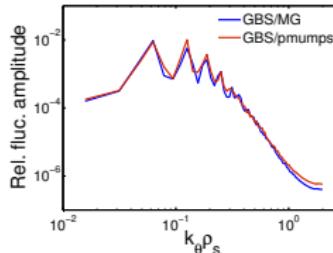
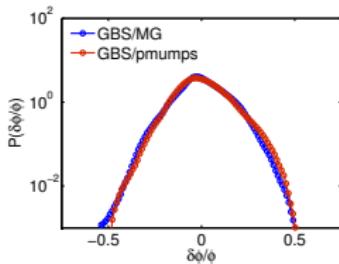
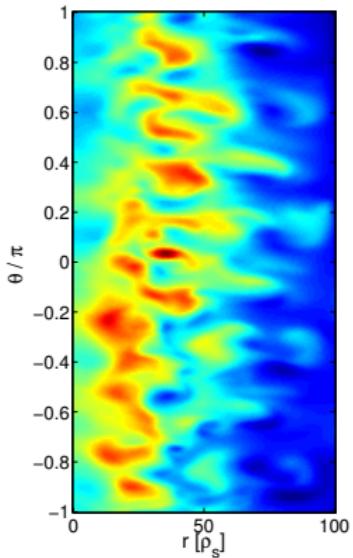
$$\mathbf{D}_i = \begin{pmatrix} \delta_{xy,(-1,1)} & \delta_{yy(0,1)} & \delta_{xy,(1,1)} \\ \delta_{xx,(-1,0)} & \delta_{xx}(0,0) + \delta_{yy}(0,0) & \delta_{xx,(1,0)} \\ \delta_{xy,(-1,-1)} & \delta_{yy}(0,-1) & \delta_{xy,(1,-1)} \end{pmatrix}$$

$$\mathbf{R}_i = \frac{1}{16} \begin{pmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{pmatrix}, \quad \mathbf{l}_{i,x} = \frac{1}{2} \begin{pmatrix} 1 & 1 \end{pmatrix}, \quad \mathbf{l}_{i,y} = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

- ▶ Use  $\delta_{\alpha\beta}$  to describe diagonally dominant 2-D elliptic operators
  - ▶ Including plasma shaping, generalized Poisson operators,...
- ▶ Damped Jacobi/**RB Gauss-Seidel**/SOR relaxation
- ▶ In GBS, we converge the residue to  $\varepsilon \sim 10^{-10}$  within 3-4 V(3,3)-cycles

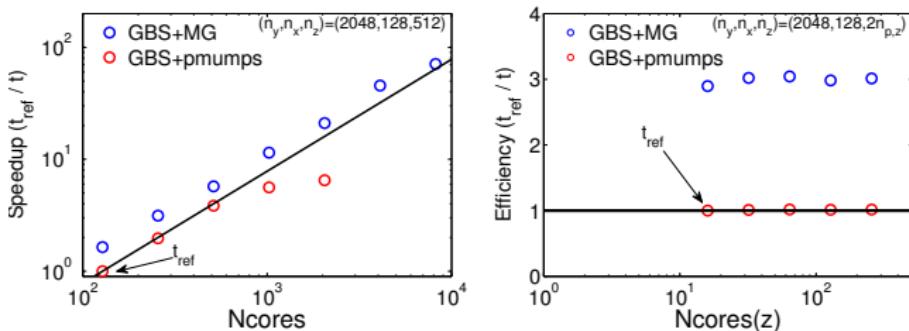
# Turbulent steady state successfully verified

- ▶ Compare steady state profiles, fluctuation moments...



## Parallel benchmarks show improved performance/scalability

- ▶ Compare GBS strong/weak scalings in PizDaint:  
**parallel multigrid vs pmumps backsolve**



- ▶ New GBS scales up to  $\sim 8192$  cores for large plasmas
  - ▶ Parallel multigrid solver yields better performance at **any size**

## Electromagnetic effects now possible at large size

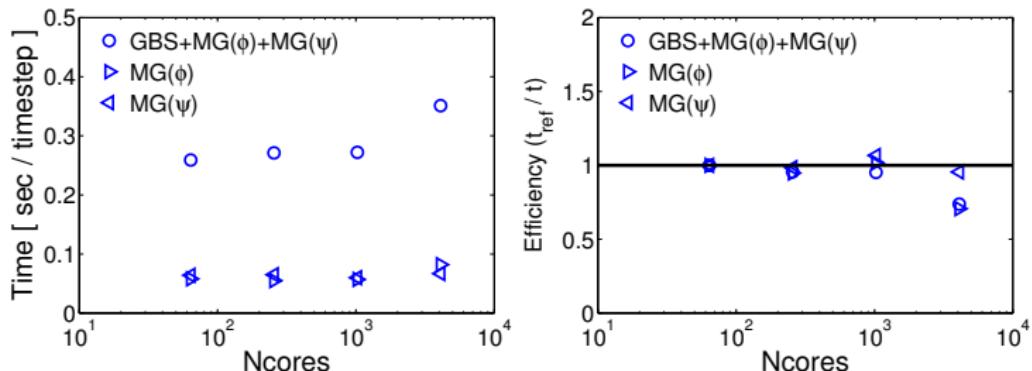
New multigrid solver allows efficient inversion of **time dependent** operators:

- ▶ Ampère's equation from Ohm's law including electron inertia

$$\begin{aligned}\frac{\partial}{\partial t} \left( v_{\parallel e} + \frac{\beta_e m_i}{2 m_e} \psi \right) &= \dots \\ \Rightarrow \left[ \nabla_{\perp}^2 - \frac{\beta_e m_i}{2 m_e} n(y, x, z, t) \right] v_{\parallel e} &= S_{v_{\parallel e}}\end{aligned}$$

- ▶ Severe disadvantages of direct inversion removed in new GBS
  - ▶ No need to LU decompose regularly (computational cost)
  - ▶ Memory limits plasma size because of matrix size

## Extended weak scaling including EM effects



- Here, the plasma volume was increased by a factor of 64
- Multigrid solution time does not degrade with core number
- Modest cost increase respect to ES simulations

## Boussinesq approximation removed from GBS

- ▶ The Boussinesq approximation

$$[\nabla \cdot (n(y, x, z, t) \nabla_{\perp})] \phi \sim \nabla_{\perp}^2 \phi / n(y, x, z, t)$$

is not really justifiable in the tokamak SOL  $\rightarrow \delta n/n \sim \mathcal{O}(1)$

- ▶ We have reformulated drift-reduced Braginskii eqns to obtain new vorticity equation retaining full  $(n, \phi)$  coupling

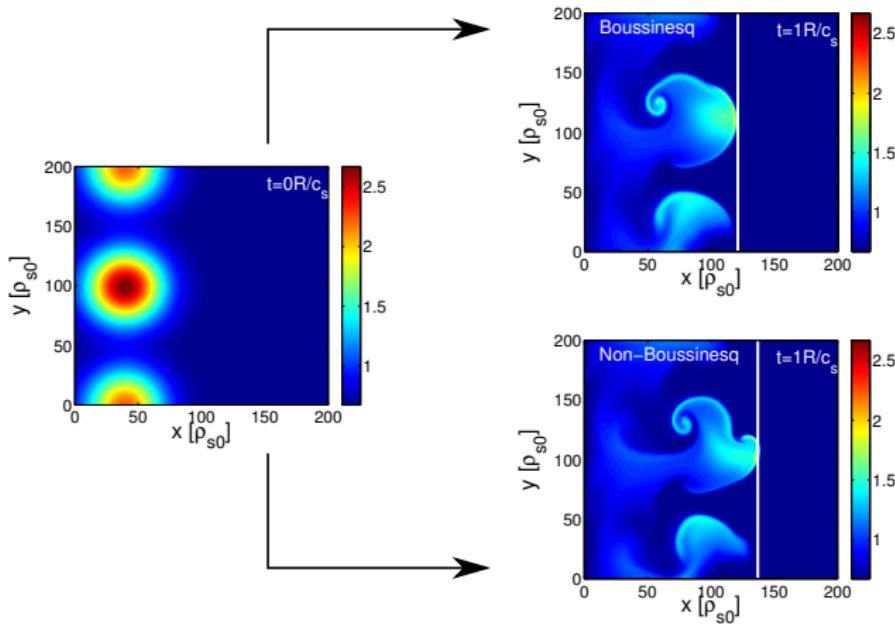
$$\begin{aligned} \frac{d\tilde{\omega}}{dt} &= \nabla \cdot (j_{\star} + j_{\parallel}) / n & \Rightarrow \frac{d\Omega}{dt} &= \nabla \cdot (j_{\star} + j_{\parallel}) \\ \tilde{\omega} &= \nabla \cdot (\nabla_{\perp} \phi + \tau \nabla_{\perp} T_i) & \Rightarrow \Omega &= \nabla \cdot (n \nabla_{\perp} \phi + \tau \nabla_{\perp} p_i) \end{aligned}$$

- ▶ Implemented in GBS using finite volume stencil for  $\phi$

$$\int \Omega dV = \int \nabla \cdot (n \nabla_{\perp} \phi) dV \rightarrow \int n (\nabla_{\perp} \phi \cdot \hat{n}) dl$$

# First Non-Boussinesq simulations now complete

We already start to see some differences in the plasma dynamics...



## Summary and conclusions

- ▶ First production runs with new GBS version now ongoing
- ▶ Code updated to F2003/8, pure MPI parallelism upgraded
- ▶ Poisson solver rewritten using multigrid techniques
- ▶ Main accomplishments:
  - ▶ EM runs at large size → size scaling of magnetic flutter effects
  - ▶ Non-Boussinesq equations → study flow/turbulence interaction
  - ▶ Typical simulation cost cut by half, walltime cut by factor of 4

## Outlook and future work

- ▶ Benchmarking of stencil operators for MG carried out using OpenMP/OpenACC/MPI
- ▶ Begin porting Poisson solver to MPI/OpenMP hybrid

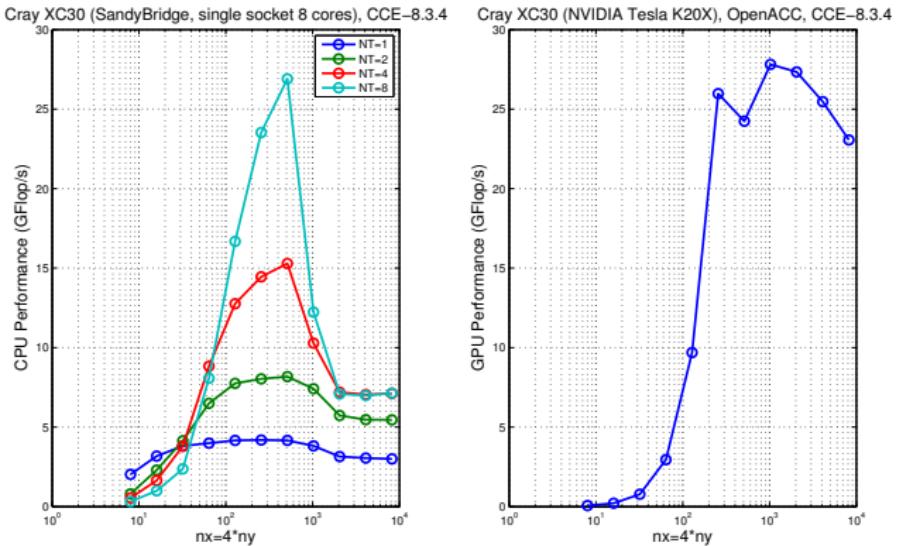
# Extra slides

# The Stencil Kernel

- ▶ Experiment with kernel computation (9-point stencil)
- ▶ OpenMP and OpenACC parallelization (single node)

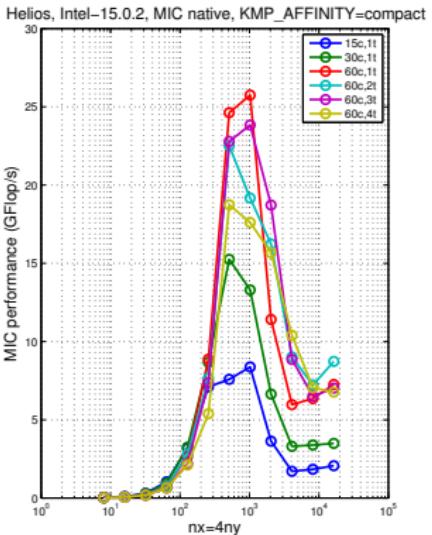
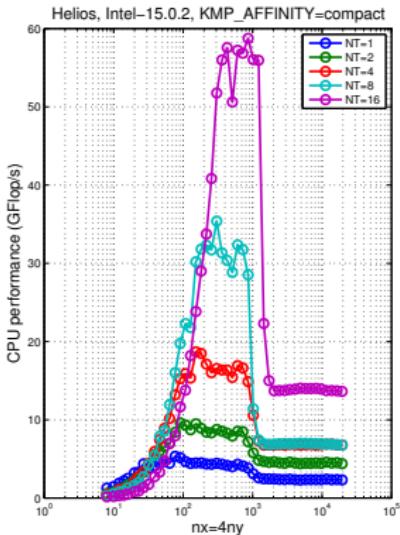
```
1 !$omp parallel do private(ix,iy)
2 !$acc parallel loop present(mat,x,y) private(ix,iy)
3   DO iy=0,ny
4     DO ix=0,nx
5       y(ix,iy) = mat(ix,iy,1)*x(ix-1,iy-1) &
6         & + mat(ix,iy,2)*x(ix,iy-1) &
7         & + mat(ix,iy,3)*x(ix+1,iy-1) &
8         & + mat(ix,iy,4)*x(ix-1,iy) &
9         & + mat(ix,iy,0)*x(ix,iy) &
10        & + mat(ix,iy,5)*x(ix+1,iy) &
11        & + mat(ix,iy,6)*x(ix-1,iy+1) &
12        & + mat(ix,iy,7)*x(ix,iy+1) &
13        & + mat(ix,iy,8)*x(ix+1,iy+1)
14   END DO
15 END DO
16 !$acc end parallel loop
17 !$omp end parallel do
```

# Performance on Piz Daint (Cray XC30)



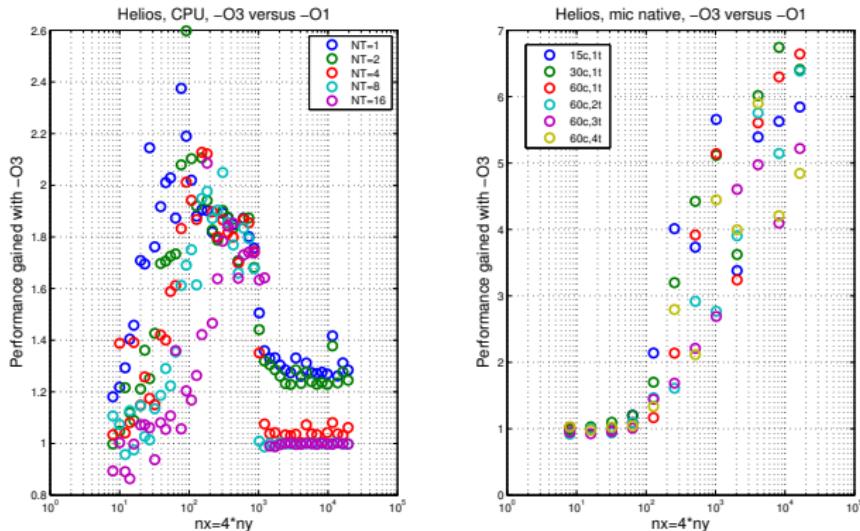
- ▶ NVIDIA GPU 3x faster than 8 CPU cores for  $(nx, ny) > 1024 \times 256$
- ▶ Small flop intensity per thread in smaller grids.

# Performance on Helios



- ▶ Good parallel speedup on host CPU
- ▶ MIC scales only up to 60 cores with 1 thread/core
- ▶ MIC performance not better than 8 CPU cores

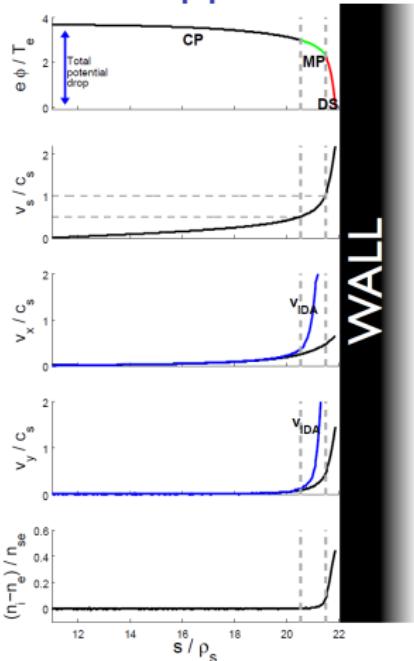
## Performance on Helios (cont.)



- Good **vector** speedup on both CPU and MIC
- Poor **parallel** performance on MIC: overhead of creation of large number of threads?

# Sheath boundary conditions from kinetic approach

- ▶ COLLISIONAL PRESHEATH (CP)
  - ▶ Quasi-neutral, IDA holds
  - ▶ Potential drop  $\sim 0.5 T_e$  over  $\sim L$
  - ▶ Ions accelerated to  $v_s = c_s \sin \alpha$
- ▶ MAGNETIC PRESHEATH (MP)
  - ▶ Quasi-neutral, IDA breaks
  - ▶ Potential drop  $\sim 0.5 T_e$  over  $\sim \rho$
  - ▶ Ions accelerated to  $v_s = c_s$
- ▶ DEBYE SHEATH (DS)
  - ▶ Non-neutral, IDA breaks
  - ▶ Potential drop  $\sim 3 T_e$  over  $\sim 10 \rho_s$
  - ▶ Ions accelerated to  $v_s > c_s$



## Fluid BCs at the Magnetic Pre-Sheath Entrance:

$$\begin{aligned}
 v_{||i} &= \pm c_s \left( 1 + \theta_n - \frac{1}{2} \theta_{T_e} - \frac{2\phi}{T_e} \theta_\phi \right) \\
 v_{||e} &= \pm \sqrt{T_e} \left( \exp(\Lambda - \eta_m) - \frac{2\phi}{T_e} \theta_\phi + 2(\theta_n + \theta_{T_e}) \right) \\
 \frac{\partial \phi}{\partial y} &= \mp c_s \left( 1 + \theta_n + \frac{1}{2} \theta_{T_e} \right) \frac{\partial v_{||i}}{\partial y} \\
 \frac{\partial n}{\partial y} &= \mp \frac{n}{c_s} \left( 1 + \theta_n + \frac{1}{2} \theta_{T_e} \right) \frac{\partial v_{||i}}{\partial y} \\
 \frac{\partial T_e}{\partial y} &= \frac{\partial T_i}{\partial y} \simeq 0 \\
 \tilde{\omega} &= -\cos^2 \alpha \left[ (1 + \theta_{T_e}) \left( \frac{\partial v_{||i}}{\partial y} \right)^2 + c_s (1 + \theta_n + \theta_{T_e}/2) \frac{\partial^2 v_{||i}}{\partial y^2} \right]
 \end{aligned}$$

where  $\theta_A = \frac{\rho_s}{2 \tan \alpha} \frac{\partial_x A}{A}$ , and  $\eta_m = e(\phi_{mpe} - \phi_{wall})/T_e$ . [Loizu et al PoP 2012]