

Parallelization on a Hybrid Architecture of GBS, a Simulation Code for Plasma Turbulence at the Edge of Fusion Devices

C. Wersal, T.M. Tran, P. Emonts, F.D. Halpern, R. Jorge, J. Morales, P. Paruta, P. Ricci, F. Riva
École Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC), CH-1015 Lausanne, Switzerland

Introduction




Figure 1: Plasma turbulence simulation performed on Piz Dora (project s549)

- Magnetic fusion research intends to create a star on Earth and to harvest the energy it releases.
- The fusion fuel, heated to 100 million degrees, is in the plasma state and it is confined in a magnetic cage.
- A Tokamak is an axisymmetrical torus-shaped device that creates such a magnetic cage to confine the hot plasma.
- At the Tokamak edge magnetic field lines intercept the wall:
 - Exhaust heat
 - Control impurity transport, fusion ashes removal, and plasma fueling (recycling and gas puffing)

- GBS is a simulation code to evolve plasma turbulence in the edge of fusion devices. [Halpern *et al.*, JCP 2016], [Ricci *et al.*, PPCF 2012]
- GBS solves 3D fluid equations for electrons and ions, Poisson's and Ampere's equations, and a kinetic equation for neutral atoms.

The GBS code

Two fluid drift-reduced Braginskii equations, $k_{\perp}^2 \gg k_{\parallel}^2$, $d/dt \ll \omega_{ci}$

$$\frac{\partial n}{\partial t} = -\frac{1}{B}[\phi, n] + \frac{2}{eB} [C(p_e) - enC(\phi)] - \nabla_{\parallel}(nv_{\parallel e}) + D_n(n) + S_n + n_n \nu_{iz} - n \nu_{rec} \quad (1)$$

$$\frac{\partial \tilde{\omega}}{\partial t} = -\frac{1}{B}[\phi, \tilde{\omega}] - v_{\parallel} \nabla_{\parallel} \tilde{\omega} + \frac{B^2}{m_n n} \nabla_{\perp}^2 \tilde{\omega} + \frac{2B}{m_n n} C(p) + D_{\tilde{\omega}}(\tilde{\omega}) - \frac{n_n}{n} \nu_{cx} \tilde{\omega} \quad (2)$$

$$\frac{\partial v_{\parallel e}}{\partial t} + \frac{e}{m_e} \frac{\partial \Psi}{\partial t} = -\frac{1}{B}[\phi, v_{\parallel e}] - v_{\parallel e} \nabla_{\parallel} v_{\parallel e} + \frac{e}{\sigma_{\parallel} m_e} j_{\parallel} + \frac{e}{m_e} \nabla_{\parallel} \phi - \frac{T_e}{m_e n} \nabla_{\parallel} n - \frac{1.71}{m_e n} \nabla_{\parallel} T_e + D_{v_{\parallel e}}(v_{\parallel e}) + \frac{n_n}{n} (\nu_{en} + 2\nu_{iz})(v_{\parallel n} - v_{\parallel e}) \quad (3)$$

$$\frac{\partial v_{\parallel i}}{\partial t} = -\frac{1}{B}[\phi, v_{\parallel i}] - v_{\parallel i} \nabla_{\parallel} v_{\parallel i} - \frac{1}{m_i n} \nabla_{\parallel} p + D_{v_{\parallel i}}(v_{\parallel i}) + \frac{n_n}{n} (\nu_{iz} + \nu_{cx})(v_{\parallel n} - v_{\parallel i}) \quad (4)$$

$$\frac{\partial T_e}{\partial t} = -\frac{1}{B}[\phi, T_e] - v_{\parallel e} \nabla_{\parallel} T_e + \frac{4T_e}{3eB} \left[\frac{T_e}{n} C(n) + \frac{7}{2} C(T_e) - eC(\phi) \right] + \frac{2T_e}{3n} \left[\frac{0.71}{e} \nabla_{\perp}^2 j_{\parallel} - n \nabla_{\parallel} v_{\parallel e} \right] + D_{T_e}(T_e) + D_{T_e}^{\parallel}(T_e) + S_{T_e} + \frac{n_n}{n} \nu_{iz} \left[-\frac{2}{3} E_{iz} - T_e + m_e v_{\parallel e} \left(v_{\parallel e} - \frac{4}{3} v_{\parallel n} \right) \right] - \frac{n_n}{n} \nu_{en} m_e \frac{2}{3} v_{\parallel e} (v_{\parallel n} - v_{\parallel e}) \quad (5)$$

$$\frac{\partial T_i}{\partial t} = -\frac{1}{B}[\phi, T_i] - v_{\parallel i} \nabla_{\parallel} T_i + \frac{4T_i}{3eB} \left[C(T_e) + \frac{T_e}{n} C(n) - \frac{5}{3} C(T_i) - eC(\phi) \right] + \frac{2T_i}{3n} \left[\frac{1}{e} \nabla_{\perp}^2 j_{\parallel} - n \nabla_{\parallel} v_{\parallel i} \right] + D_{T_i}(T_i) + D_{T_i}^{\parallel}(T_i) + S_{T_i} + \frac{n_n}{n} (\nu_{iz} + \nu_{cx}) \left[T_n - T_i + \frac{1}{3} (v_{\parallel n} - v_{\parallel i})^2 \right] \quad (6)$$

$$\rho_x = \rho_s/R, \quad \nabla_{\parallel} f = \mathbf{b}_0 \cdot \nabla f, \quad \tilde{\omega} = \omega + \tau \nabla_{\perp}^2 T_i, \quad p = n(T_e + \tau T_i)$$

- A set of fluid boundary conditions where the magnetic field lines intersect the vessel: $\frac{\partial v_{\parallel i}}{\partial y} = -\frac{n}{\sqrt{T_e + T_i}} \frac{\partial v_{\parallel i}}{\partial y}$, $v_{\parallel e} = \sqrt{T_e} \exp(\Lambda - \frac{\phi}{\sqrt{T_e}})$, ... [Loizu *et al.*, PoP 2012]
- Gradients and curvature terms discretized using finite differences
- Poisson brackets, $[a, b] = \mathbf{b}_0 \cdot (\nabla a \times \nabla b)$, discretized using Arakawa scheme
- Time evolution using the classic Runge Kutta method
- GBS uses a 3D Cartesian MPI communicator decomposing the computational 3D domain. OpenMP directives have been included recently.

The Poisson and Ampere equations

- Poisson equation with Boussinesq approximation, $\nabla_{\perp}^2 \phi = \omega$, or without, $\nabla \cdot (n \nabla_{\perp} \phi) = \Omega - \tau \nabla p_i$
- Ampere's equation from Ohm's law, $(\nabla_{\perp}^2 - \frac{\beta_{e0} m_i n}{2 m_e n}) v_{\parallel e} = S_{v_{\parallel e}}$
- Stencil based **parallel multigrid** implemented in GBS
- The elliptic equations are separable in the parallel direction leading to independent 2D solutions for each x-y plane
- 2D Cartesian (x, y) grid topology mapped to a **2D domain decomposition**

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

$$R_i = \frac{1}{16} \begin{pmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{pmatrix}, \quad I_{i,x} = \frac{1}{2} \begin{pmatrix} 1 & 1 \end{pmatrix}, \quad I_{i,y} = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

- $\delta_{\alpha\beta}$ describe diagonally dominant 2-D elliptic operators
- Damped Jacobi/**RB Gauss-Seidel**/SOR relaxation
- In GBS, the residue converges to $\varepsilon \sim 10^{-10}$ within 3-4 V(3,3)-cycles

The kinetic equation for neutral atoms


$$\frac{\partial f_n}{\partial t} + \mathbf{v} \cdot \frac{\partial f_n}{\partial \mathbf{x}} = -\nu_{iz} f_n - \nu_{cx} n_n \left(\frac{f_n}{n_n} - \frac{f_i}{n_i} \right) + \nu_{rec} f_i \quad (7)$$

- **Method of characteristics** to obtain the formal solution of f_n
 - **Two assumptions**, $\tau_{neutral} \text{ losses} < \tau_{turbulence}$ and $\lambda_{mf}, \text{ neutrals} \ll L_{\parallel, \text{plasma}}$, leading to a 2D steady state system for each x-y plane
 - **Linear integral equation** for neutral density obtained by integrating f_n over \mathbf{v}
 - **Spatial discretization** leading to a linear system of equations
- $$\begin{bmatrix} n_n \\ \Gamma_{out} \end{bmatrix} = \begin{bmatrix} K_{p \rightarrow p} & K_{b \rightarrow p} \\ K_{p \rightarrow b} & K_{b \rightarrow b} \end{bmatrix} \cdot \begin{bmatrix} n_n \\ \Gamma_{out} \end{bmatrix} + \begin{bmatrix} n_{n,rec} \\ \Gamma_{out,rec} + \Gamma_{out,i} \end{bmatrix} \quad (8)$$
- This system is solved for neutral density, n_n , and neutral particle flux at the boundaries, Γ_{out} , with the threaded LAPACK solver.


Scalability of MPI+OpenMP GBS

- Hybrid MPI+OpenMP with `MPI_THREAD_FUNNELED` (MPI calls only by thread 0)
- Basic OpenMP directives: `parallel, do, single, master, barrier, simd`
- Simple clauses: `schedule(static), collapse`
- Scalings performed on the Helios Supercomputer system at IFERC-CSC, two 8-core Sandy-bridge processors and 64 GB memory on each node, InfiniBand network

Strong scaling (grid size 256 x 2048 x 128)



Weak scaling (grid size 256 x 2048 x nodes*8)

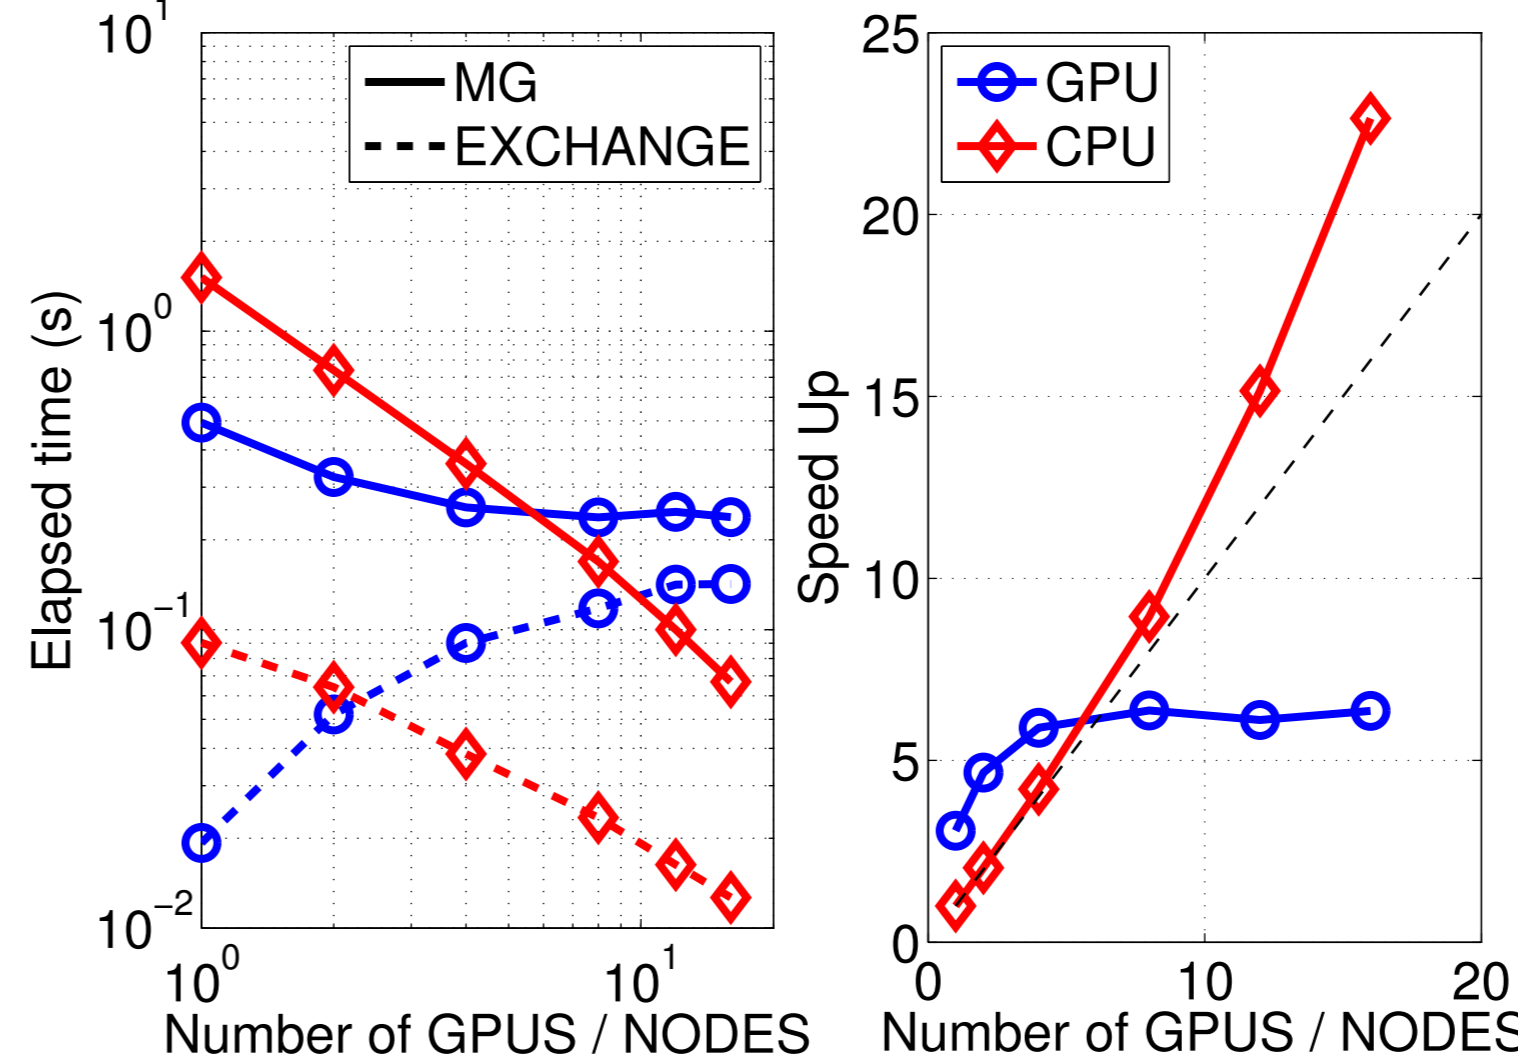


- **Fluid equations** profit from OpenMP directives. Efficiency decreases in the strong scaling due to 3D ghost cell exchange.
- **Parallel multigrid solver** shows very good parallel scalability (2D ghost cell exchange), runs best with 1 or 2 threads, and shows superlinear scaling probably related to cache use.
- Optimal parallelization for **overall timestep** depends on physical case and system architecture

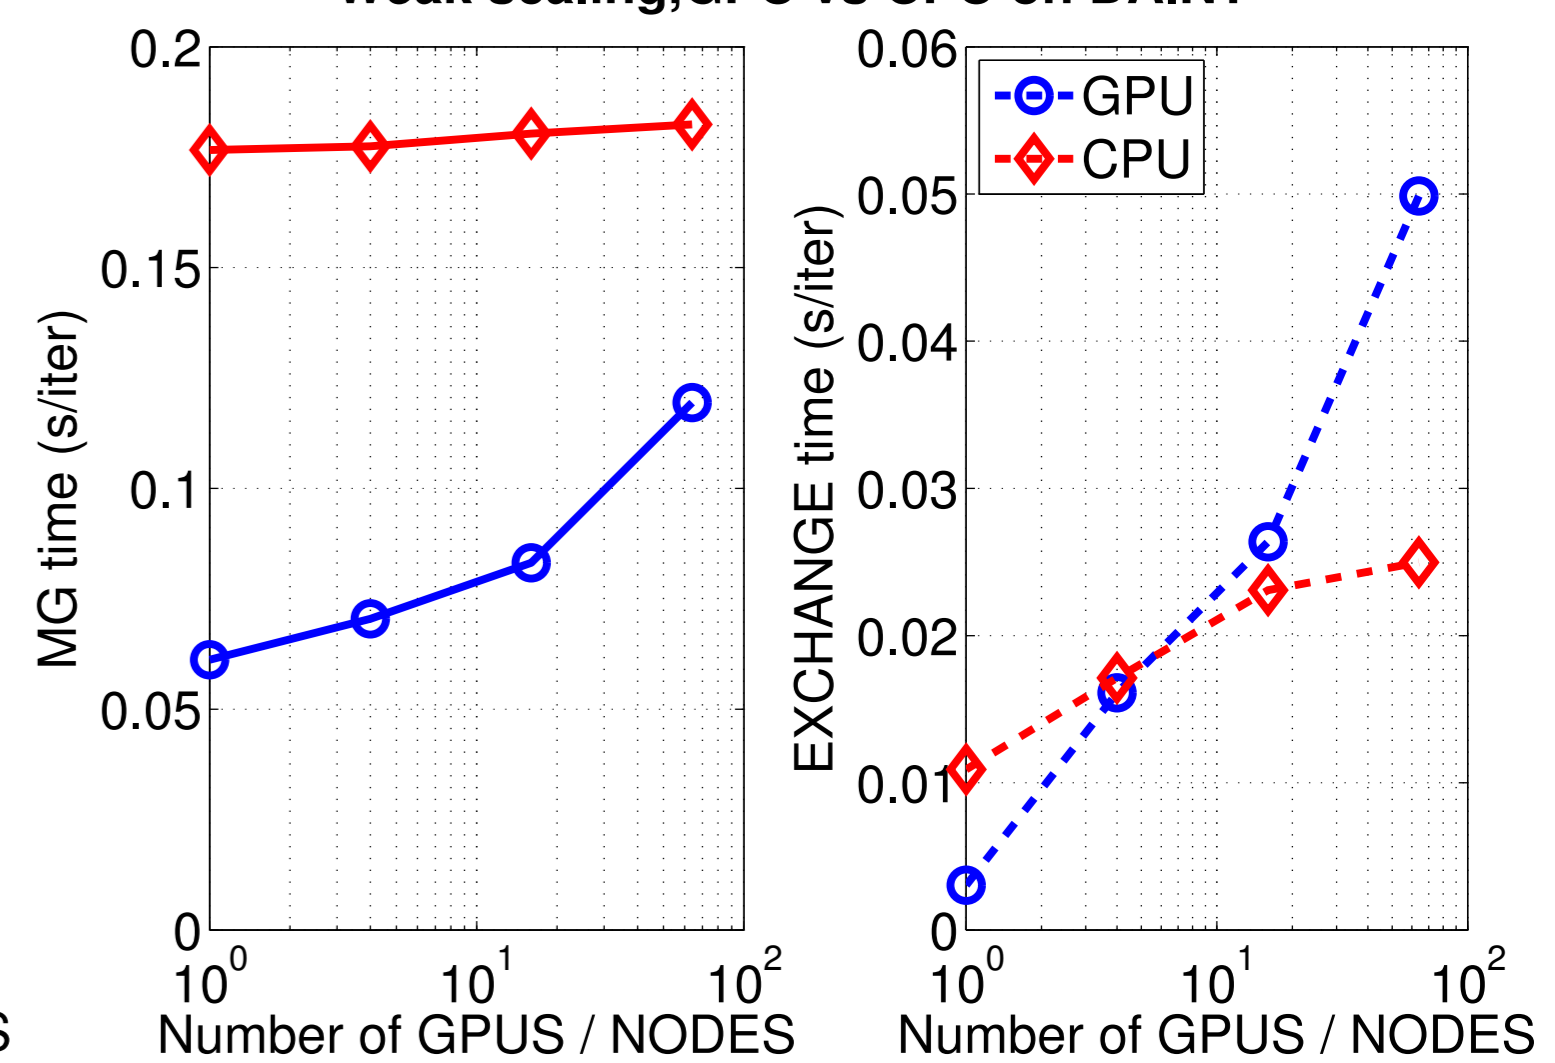
Parallel Multigrid Solver with OpenACC

- Initialize the array `TYPE(grid_2d)::grids(1:levels)` including grids, stencils, and solution arrays on the *host* and offload to the GPU before MG iterations start
 - No support for derived types with `ALLOCATABLE` members in OpenACC-2.0: only *shallow-copy* of such derived types
 - Extension of Cray Fortran: *deep-copy* with the compiler switch `-hacc-model=deep-copy`
- Offloading during MG iterations:
 - Residual norm and discretization error norm (16 Bytes): used in the *stopping criteria*
 - For *multi-node multi-gpu* version, additional offload of 2D domain *boundaries* (1D buffers) for ghost cells *exchange*
- Run on a Cray XC30 (Piz Daint at CSCS) equipped with one 8-core Xeon E5-2670@2.6 GHz and one NVIDIA Tesla K20X per node

Strong scaling, 1024X4096 grid, GPU vs CPU on DAINT



Weak scaling, GPU vs CPU on DAINT



- **3x speed-up** from a single CPU (8 cores) to a single GPU
- Good speed-up in the strong scaling for two GPUs, but saturation above four GPUs
 - In GBS the x-y planes are localized closest in the 3D MPI communicator, so only few GPUs are involved in the 2D parmg solver, and parallel scalability up to 2-4 GPUs is sufficient.
- Super-linear speed-up for the CPU parmg solver, probably due to efficient cache use
- The increase of execution time in the weak scaling is mainly due to the increase in exchange time (offloading and MPI communication)
- Large problem size necessary for efficient use of many GPUs

Summary and Outlook

- The hybrid MPI+OpenMP parallelization implemented recently in GBS leads to performance improvements for the fluid equations in the code.
- The optimal distribution of processors between MPI and OpenMP depends on the chosen problem and platform. The scalability on many-core platforms (Xeon Phi) to be evaluated.
- A hybrid MPI+OpenACC multigrid Poisson solver developed as a first step in porting GBS to mixed CPU+GPU architectures.
- The fluid equations evaluation of GBS still to be ported to MPI+OpenACC. CPUs will be used for MPI ghost cell exchange and diagnostics output, while the main computation is to be carried out on the GPUs.