

Coarse-grained hybrid parallelization

Silvia Mocavero

HPC Working Group Meeting

28/07/2017

Hybrid parallelization – coarse grained approach

GOAL: reduce the synchronization overhead of the fine-grained hybrid implementation

METHODOLOGY: select one/two representative kernels to implement and test the coarse-grained approach

ACCURACY TEST: restartability/reproducibility tests

PERFORMANCE TEST: comparison of the hybrid version with the pure MPI one. Test to be performed on different architectures (also including KNL)



ZDF package use case

Restructuring of ZDF (vertical physics) module (from Gurvan)

- Introduction of the ZDF manager that calls all the vertical physics routines
- Moving lateral boundary conditions update and restart writing at the end of the manager execution -> limited data exchange among neighboring processors
- Code re-writing to avoid data dependencies -> natural decomposition in horizontal direction

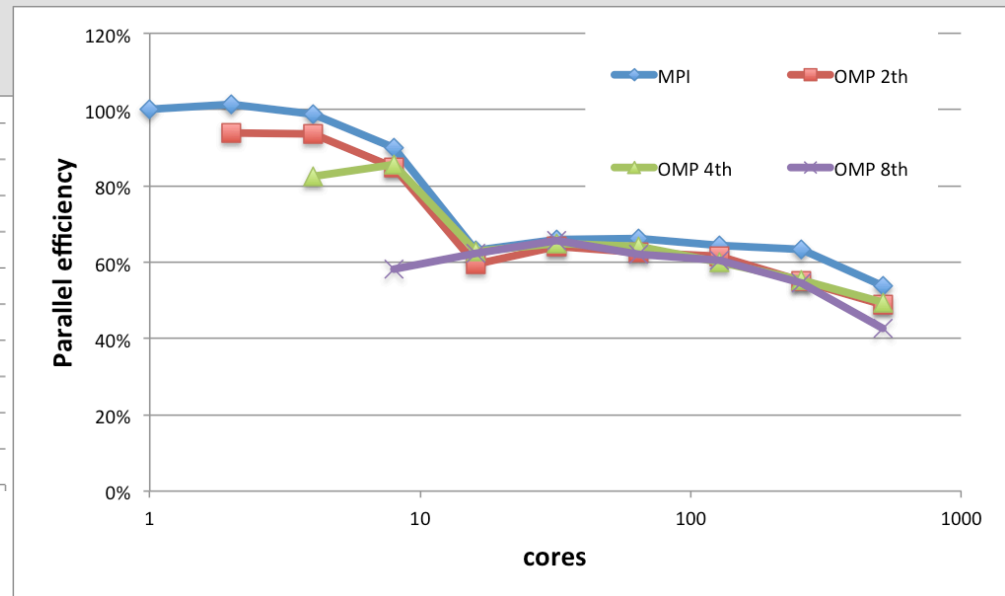
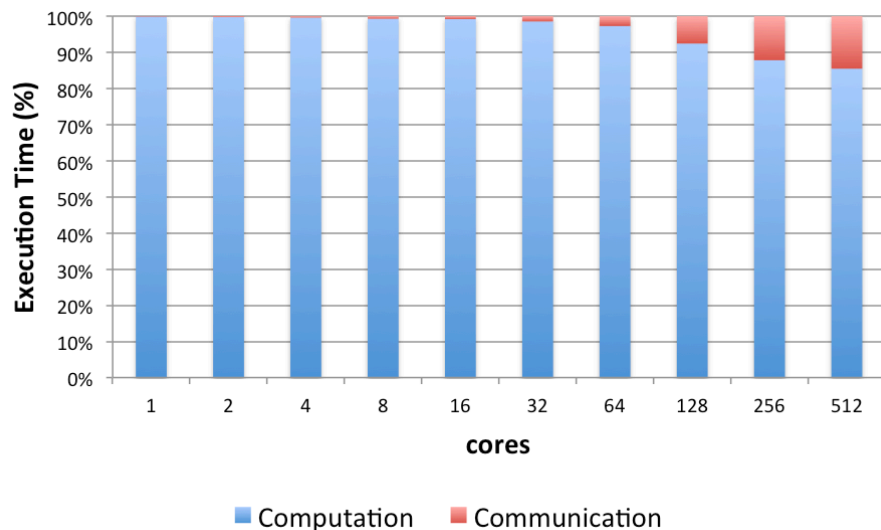
Ideal test case to evaluate coarse-grained performance improvement

Code restructuring cannot easily extended to the whole code



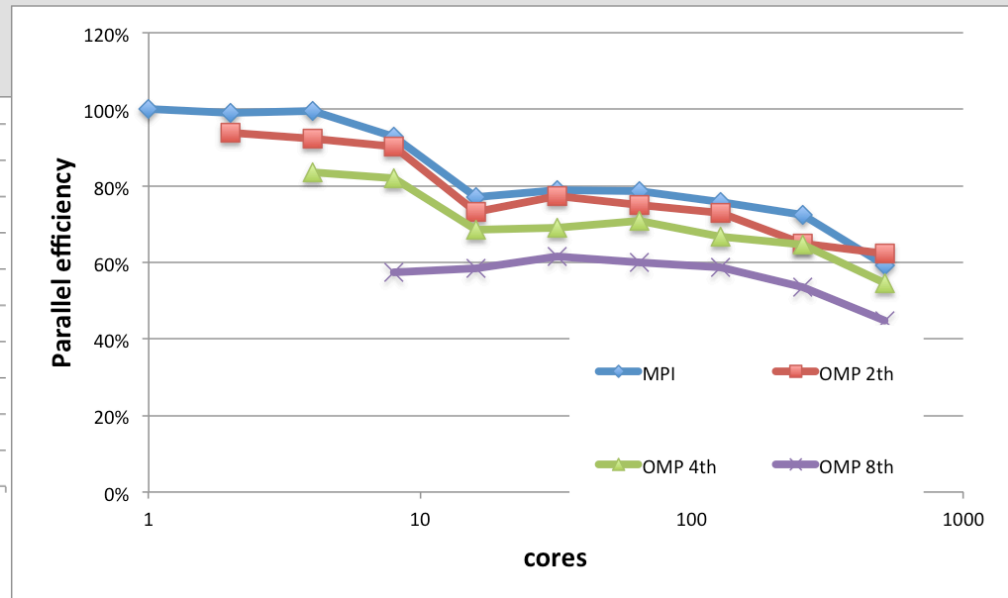
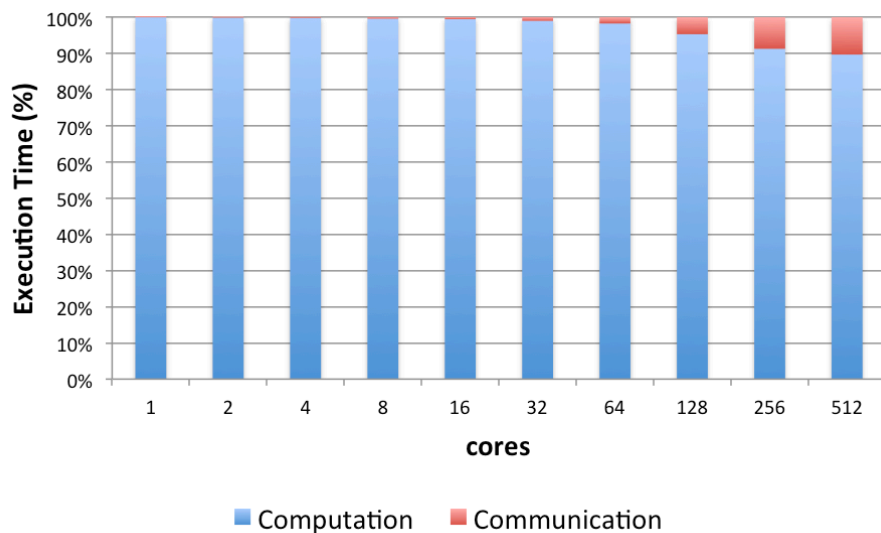
ZDF kernel: coarse-grained implementation

- Tiling of MPI region along the horizontal direction (one tile for each thread)
- Single OpenMP parallel region that includes all the ZDF routines
- Only a synchronization is needed before viscosity and diffusivity coefficients computation
- Restartability and reproducibility tests on GYRE_PISCES have been successfully executed
- Preliminary computational results on CMCC system (**Turbulent Kinetic Energy** parametrisation, $nn_GYRE=24$)



ZDF kernel: coarse-grained implementation

- Tiling of MPI region along the horizontal direction (one tile for each thread)
- Single OpenMP parallel region that includes all the ZDF routines
- Only a synchronization is needed before viscosity and diffusivity coefficients computation
- Restartability and reproducibility tests on GYRE_PISCES have been successfully executed
- Preliminary computational results on CMCC system (**Generic Length Scale** parametrisation, $nn_GYRE=24$)



Summary and future plans

- ZDF package restructuring allows to reduce data dependencies, thus to reduce OMP synchronization overhead
- Communication overhead in ZDF is limited (~10-15% on 512 cores – subdomain size ~20x30)
- Communication time does not increase with the number of MPI procs
- Fixing the number of cores, computation time is almost the same with the two versions

- Memory access overhead to be evaluated for both MPI and OpenMP versions
- Increase tiles number to improve load balance
- Test coarse-grained version on other architectures (MetO plans to test the kernel on KNL)
- Other more significative kernels could be considered

