Multigrid method on Intel Xeon Phi (MIC)

Kab Seok Kang

NMPP,
Max-Planck Institute for Plasma Physics,
Boltzmannstrasse 2,
D-85748 Garching, Germany
kskang@ipp.mpg.de

The multigrid method is a well-known, fast and efficient algorithm to solve many classes of problems. In general, the ratio of the communication costs to computation costs increases on the coarser level, i.e., the communication costs are high on the coarser levels in comparison to the computation costs. So, reducing the costs of the communication is a major issue in implementation of the parallel multigrid method. Using a hybrid programming model which uses OpenMP for parallelization inside node and MPI for message passing between nodes and can reduce the number of MPI tasks, we have better scaling properties on a massively parallel computer.

Modern computer architectures have highly hierarchical system design, i.e., multi-socket multi-core shared-memory computer nodes which are connected via high-speed interconnects, and now accelerators such as Intel Xeon Phi (MIC) coprocessors or GPUs (Graphics Processing Units) are emerged. The OpenMP 4 standard supports to use these accelerators. In this talk, we present the performance results of the parallel multigrid method with hibridization implementation on a MIC partition of the Helios machine which is dedicated machine for Europe and Japan Fusion community.