

# Flat MPI vs. Hybrid: Evaluation of Parallel Programming Models for Preconditioned Iterative Solvers on Multicore Clusters

Kengo NAKAJIMA

Information Technology Center, The University of Tokyo  
nakajima@cc.u-tokyo.ac.jp

In this work, parallel preconditioning methods based on “Hierarchical Interface Decomposition (HID)” and hybrid parallel programming models were applied to finite-element based simulations of linear elasticity problems in media with heterogeneous material properties. Reverse Cuthill-McKee reordering with cyclic multicoloring (CM-RCM) was applied for parallelism through OpenMP [1]. The developed code has been tested on the “T2K Open Supercomputer (Todai Combined Cluster, T2K/Tokyo) [2]” using up to 512 cores. T2K/Tokyo is an AMD Quad-core Opteron-based combined cluster system with 952 nodes. Each node includes four “sockets” of AMD Quad-core Opteron processors (2.3GHz). Because T2K/Tokyo is based on *cache-coherent* NUMA (cc-NUMA) architecture, careful design of software and data structure is required for efficient access to local memory.

Performance of Hybrid 4x4 parallel programming model is competitive with that of flat MPI using appropriate command lines for NUMA control. Furthermore, reordering of the mesh data for contiguous access to memory with first touch data placement provides excellent improvement on performance of Hybrid 8x2 and 16x1, especially if the problem size for each core is relatively small. Finally, improvement of memory copy for communication provides excellent scalability for Hybrid 8x2 and 16x1.

Generally speaking, hybrid parallel programming model provides excellent performance for strong scaling cases where problems are less memory-bound, and could be a reasonable choice for large-scale computing of sparse linear solvers on multi-core/multi-socket architectures, such as “T2K Open Supercomputer”.

Performance of the developed code is also evaluated on Cray XT4 at NERSC [3] with AMD Quad-core Opteron using Flat MPI and Hybrid 4x4. In the poster presentation, the most recent results by Cray XT4 will be demonstrated, and compared with those of T2K/Tokyo.

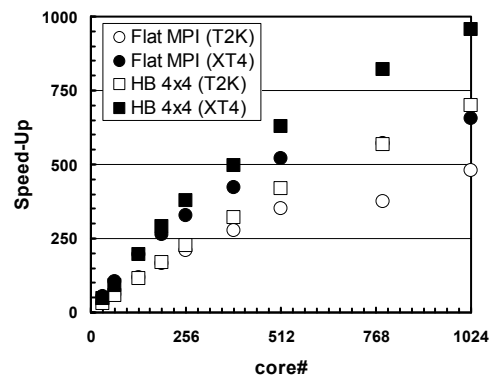


Fig. Strong Scalability of SGS/CG with HID on the T2K/Tokyo and Cray XT4 (NERSC) for the 3D linear elasticity model with 6,440,067 DOF, from 32 cores up to 1,024 cores, normalized by 32-core case for flat MPI on T2K

## References

- [1] K. Nakajima, Flat MPI vs. Hybrid: Evaluation of Parallel Programming Models for Preconditioned Iterative Solvers on “T2K Open Supercomputer”, IEEE Proceedings of ICPP 2009 (The 38th International Conference on Parallel Processing) (2009) (in press)
- [2] Information Technology Center, the University of Tokyo: <http://www.cc.u-tokyo.ac.jp/>
- [3] National Energy Research Scientific Computing Center (NERSC): <http://www.nersc.gov/>