SSS27-P03 Room:Convention Hall Time:May 23 17:15-18:30

Performance tuning of the Seism3D, the seismic wave propagation code, for large-scale parallel simulation using K comput

MAEDA, Takuto[1*], FURUMURA, Takashi[1], Shunsuke Inoue[2], Kazuo Minami[2]

[1]CIDIR, III, the University of Tokyo, [2]RIKEN AICS

The "K computer" is the world's fastest super computer at present which has the maximum peak speed of more thatn 10 PFlops (Yonezawa et al., 2011). By using the huge computer resource of K, it is expected that there are massive improvements on numerical simulations such as seismic wave propagation or plate motion simulation in seismology. However, the K computer has more than 80,000 nodes (CPUs) that require us to highly tuned parallel computation. Additionally, the K's CPU has lower memory access speed relative to their flops. To overcome these technical difficulties, we measured the theoretical performance of our seismic wave propagation code "Seism3D" and did performance tunings from the aspects of the parallelization performance and single-CPU performance.

The Seism3D is a numerical simulation code based on the staggered-grid, finite difference method (FDM) explicit solver of the equation of motion of elastic medium (Furumura and Chen, 2005). We have improved the code in several aspects. First, we substituted the constitutive equation between stress and strain from the liner elastic solid to the generalized Zenner body to take broadband intrinsic attenuation into account. Second, the split-PML technique is used for an absorbing boundary condition at the model edges. This choice also enables us to simulate coseismic deformation at the sea floor and/or ground surface in addition to the seismic waves. The code can be extended naturally by incorporating the gravity and equilibrium between gravity and pressure field (Maeda and Furumura, 2011).

Simulation using the Seism3D requires large amount of memory access for referring inhomogeneous earth medium and wave/stress field at each time step. However, recent supercomputers including the K have relatively lower memory access speed compared to their CPU speed. Therefore, the ration between amount of the computation and amount of the memory access determine the maximum performance of the Seism3D as about 16 %.

For massive parallel simulations based on the MPI we also investigated the parallel balance between computational nodes. We found that the horizontal 2D partitioning, not 3D, is best for the computational performance. The single node assigns a rectangular parallelepiped shape having extremely long in the vertical direction. To obtain better performance, we exchanged computational loops among space so that the computation along z-direction is in the innermost loop. By this change, the continuous memory access is assured.

The data read from the memory are placed on L1- or L2- cache temporally. Because the access speed to cache is far faster than that to memory, the effective use on cache is mandatory for effective computation. In this point of view, we also tuned the cache access up (Minami et al., 2012). The K computer share the L2 cache among 8 CPU cores inside the nodes. By using this characteristic, we re-arranged the openMP-based parallelization inside a node so that at least the neighbor core can reuse the data on cache.

As a result, we achieved the 16 % of the maximum performance, which is almost the maximum speed of the Seism3D theoretically, on the K computer by the abovementioned tuning on benchmark codes. We will present the technical detail and some example of large-scale simulation by using the K computer.

Acknowledgements
This study is supported by the SPIRE field 3, MEXT. Part of this result is obtained by early access to the K computer at RIKEN Advanced Institute for Computational Science.

Keywords: Seismic wave propagation, numerical computation, numerical simulation, parallel computation, tuning