GTC Workshop Japan 2011 (2011 年 7 月 22 日)

Multi-GPU Computing of Large Scale Phase-Field Simulation for Dendritic Solidification on TSUBAME 2.0

Takashi Shimokawabe[1)], Takayuki Aoki[1)], Tomohiro Takaki[2)], and Akinori Yamanaka[1)]

1) Tokyo Institute of Technology (2-12-1 Ookayama, Meguro-ku, Tokyo, 152-8550)

2) Kyoto Institute of Technology (Gosyokaido-cyo, Matsugasaki, Sakyo-ku, Kyoto, 606-8585)

Key Words: *CUDA, Phase-field method, Multi-GPU computing, Dendritic solidification*

1 Introduction

The mechanical properties and performance of metal materials depend on the intrinsic microstructures in these materials. In materials science and related areas, the phase-field method is widely used as one of the powerful computational methods to simulate the formation of complex microstructures such as dendrites during solidification and phase transformation of metals.

The phase-field model consists of a large number of complex nonlinear terms comparing with other stencil computation. As a result, the phase-field simulations require large-scale computation, such as multi-GPU computing, to perform realistic three-dimensional simulations in the typical scales of the microstructural pattern.

2 Multi-GPU implementation and results

In this simulation, the time integration of the phase field and the solute concentration are carried out by the second-order finite difference scheme for space with the first-order forward Euler-type finite difference method for time on a three-dimensional regular computational grid.

We decompose the whole computational domain in both the y- and z-directions (2D decomposition) and allocate each sub domain to a GPU since 3D decomposition tends to degrade the GPU performance that we are able to achieve. Similar to conventional multi-CPU calculations with MPI, multi-GPU computing requires boundary data exchanges between sub domains. Because a GPU cannot directly access to the global memory of the other GPUs, host CPUs are used to bridge GPUs for the data exchange between the neighbor GPUs. This process is composed of the following three steps: (1) the data transfer from GPU to CPU using the CUDA runtime library, (2) the data exchange between nodes with the MPI library, and (3) the data transfer back from CPU to GPU with the CUDA runtime library.

Figure 1 demonstrates the dendritic growth during the binary alloy solidification using the mesh size of $768 \times 1632 \times 3264$ on 1156 GPUs of TSUBAME 2.0.

To evaluate the performance of our multi-GPU code for the phase-field simulations, we use the TSUBAME 2.0 supercomputer at the Tokyo Institute of Technology. The TSUBAME 2.0 supercomputer in Tokyo Institute of Technology is equipped with 4224 NVIDIA Tesla M2050 GPUs. Each node of TSUBAME 2.0 has three Tesla M2050 attached to the PCI Express bus 2.0 $\times$16 (8 GB/s), two QDR InfiniBand and two sockets of the Intel CPU Xeon X5670 (Westmere-EP) 2.93 GHz 6-core. Figure 2 shows that the performance of the phase-field simulation by the multi-GPU computing using TSUBAME 2.0 in both double- and single-precision. The performance of $1.017\,\mathrm{PFlops}$ is achieved for the mesh size of $4,096 \times 6,500 \times 10,400$ using 4,000 GPUs with 16,000 CPU cores of TSUBAME 2.0 in single precision.

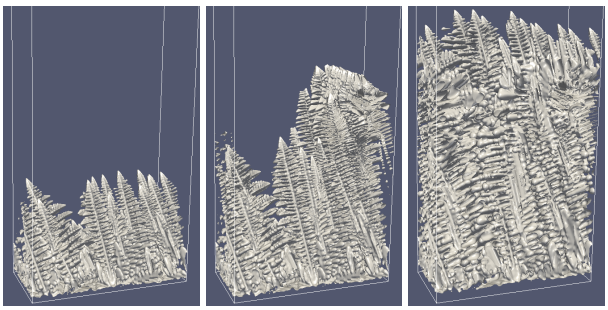

Fig. 1: Dendritic growth in the binary alloy solidification

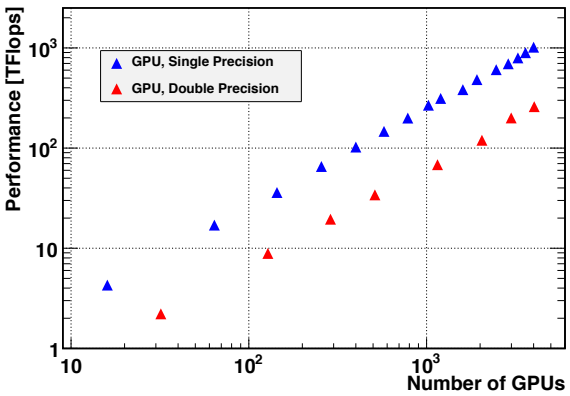

Fig. 2: Performance of multi-GPU computation

REFERENCES

[1] T. Aoki, S. Ogawa, and A. Yamanaka. Multiple-GPU scalability of phase-field simulation for dendritic solidification. *Progress in Nuclear Science and Technology*, in press.