ParCFD'2020
32nd International Conference on Parallel Computational Fluid Dynamics
May 11-13 2020, Nice FRANCE

A DOMAIN DECOMPOSITION METHOD BASED ON A MULTI-PHASE-FIELD MODEL TO REDUCE INTER-NODE COMMUNICATION

SEIYA WATANABE[*], TAKAYUKI AOKI[†], CHANGHONG HU[*], SHINTARO MATSUSHITA[†], AND TOMOHIRO TAKAKI[††]

[*] Kyushu University
Research Institute for Applied Mechanics
6-1 Kasuga-koen Kasuga-city, Fukuoka, Japan
e-mail:swatanabe@riam.kyushu-u.ac.jp, hu@riam.kyushu-u.ac.jp

[†] Tokyo Institute of Technology
Global Scientific Information and Computing Center
2-12-1-i7-3 Ookayama, Meguro-ku, Tokyo, Japan
Email: taoki@gsic.titech.ac.jp, matsushita@sim.gsic.titech.ac.jp

[††] Kyoto Institute of Technology
Faculty of Mechanical Engineering
Matsugasaki, Sakyo-ku, Kyoto, Japan
Email: takaki@kit.ac.jp

Key words: Domain decomposition, Multi-phase-field, Adaptive mesh refinement.

Summary. In order to reduce the inter-node communication cost, we propose a domain decomposition method based on the multi-phase-field (MPF) model. The MPF model for polycrystalline growth minimizes the interfacial energy and forms a convex shape for each crystal grain. In our method, each phase of the MPF model represents a computational sub-domain of each MPI process. We apply the proposed partitioning method to a block-based AMR application for an interface capturing on multiple GPUs. The proposed method can successfully reduce communication costs than the SFC approach.

1 INTRODUCTION

In distributed computations of memory-bound applications, the inter-node communication time often represents a major performance bottleneck. To optimize the communication cost, we propose a domain decomposition method based on the multi-phase-field (MPF) model. The MPF model simulates the process of material solidification, during which, each crystal grain grows to minimize its interfacial energy, thus minimizing its surface area. We can regard the crystal grain as a sub-domain of the computational domain and utilize its structure for the domain decomposition. The communication costs (proportional to the surface areas of each sub-domain) are expected to be minimized. For performance evaluation, we apply the proposed method to a block-based AMR application for an interface capturing and compare the

partitioning performance with the SFC approach.

2 DOMAIN DECOMPOSITION BASED ON MULTI-PHASE-FIELD MODEL

The multi-phase-field (MPF) method is used to study the growth of microstructures in polycrystalline materials. For a system consisting of N different phases (crystal grains), we define the phase-field variables $\phi_i(\boldsymbol{x}, t)$ $(i \in \{1, 2, \cdots, N\})$ representing the existence $(\phi_i = 1)$ and absence $(\phi_i = 0)$ of phase i at position $\boldsymbol{x}$ and time t. At any point of the computational domain, the phase field variables satisfy the following constraint:

$$\sum_{i=1}^{N} \phi_i = 1\,. \tag{1}$$

The MPF model naturally tries to minimize each phase's surface and it is thus well-suited for a partitioning problem with these requirements. However, larger crystal grains grow fast and smaller crystal grains disappear to reduce the total energy of the system. Thus the MPF method cannot be directly applied to partitioning problems. Therefore, we introduce a correction term into the MPF equation in order to balance computational load (volume correction) which penalizes vanishing phases. The volume correction term has been originally proposed by Uehara [1] to balance the volume of each crystal grain in the MPF method. The time evolution equation of the phase-field variable ϕ_i is:

$$\frac{\partial \phi_i}{\partial t} = -\frac{2}{n}\sum_{j=1}^{N} \omega_{ij} \left[\sum_{k=1}^{N} (\omega_{ik} - \omega_{jk}) \left(\phi_k + \frac{4\delta^2}{\pi^2} \nabla^2 \phi_k \right) + k(C_i - C_j) |\nabla \phi_i|^2 \left| \nabla \phi_j \right|^2 \right], \tag{2}$$

where n is the number of coexisting phases at arbitrary position, δ is the interface thickness, C_i is estimates of the computational cost of the i-th sub-domain, and k is a parameter for the correction term. The phase interface is driven by the computation cost difference between neighbor phases. For each sub-domain, we estimate the computation cost as

$$C_i = \frac{N_i^{\mathrm{block}}}{N_{\mathrm{ave}}^{\mathrm{block}}}, \tag{3}$$

where N_i^{block} is the number of blocks in i-th sub-domain and is normalized by the average number of blocks $N_{\mathrm{ave}}^{\mathrm{block}}$.

The time evolution of the MPF equation is an iterative process to reach a steady state solution of the system and is thus independent of the physical time of AMR applications. After initialization, the domain partitioning is computed by iterating on the following steps:

1) Solve the MPF equation for one iteration step.
2) Partition the computational domain based on the phases.
3) Estimate the calculation cost for each sub-domain.

If the load imbalance error is smaller than a certain threshold value or the iteration count reaches the set maximum number of iteration, we stop iterating.

3 PERFORMANCE STUDY

We apply the MPF partitioning to a block-based AMR application of an interface capturing method. We compare the performance of the MPF method with the Morton curve. Our simulations are carried out on multiple GPUs of the TSUBAME3.0 supercomputer at Tokyo institute of technology.

3.1 2D partitioning test

We solve the two-dimensional single vortex problem using the conservative Allen-Cahn equation with the AMR method on 16 GPUs. In this simulation, each block has 16×16 cells, and the effective resolution is 2048×2048 cells. The MPF partitioning with k=800 is computed on a mesh with a maximum resolution equivalent to 128×128 MPF cells. The MPF iteration is stopped if the load imbalance error is 0.05 or less.

Figure 1 shows the result of the domain decomposition. We observe that the complex refined mesh is divided into almost convexly shaped sub-domains. Figure 2 shows the load imbalance error and the boundary block ratio. The MPF method keeps the load imbalance error rather than 0.05 over the entire simulation. The boundary block ratio Br using the MPF method is smaller than that using the Morton curve, it means the MPF method reduces the communication cost that the Morton curve.

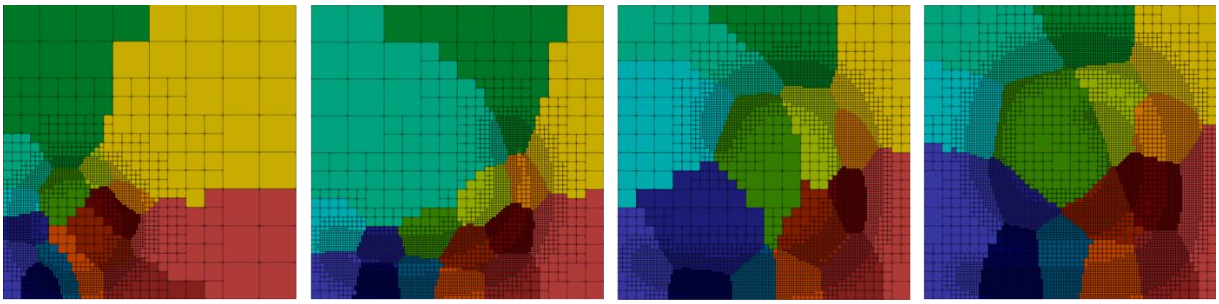

Figure 1: Dynamic domain decomposition based on the MPF model with 16 sub-domains for adaptively refined mesh in a 2D single vortex simulation. The mesh indicates the blocks.

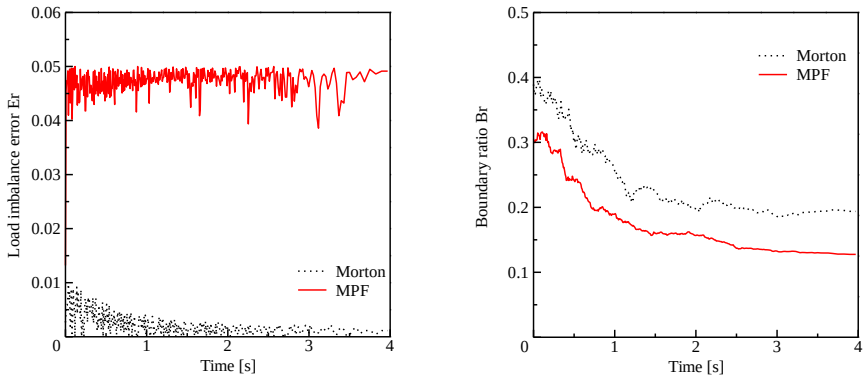

Figure 2: Time history of the load imbalance error (left) and the boundary block ratio (right).

3.2 Strong scalability

We discuss the strong scalability of the 3D AMR application on multiple GPUs using the MPF method. In strong scaling measurement, we perform simulations for two different problem sizes using 16 - 256 GPUs. In the lower- and higher- resolution case, the finest resolutions are corresponding to 2048×2048×2048 cells and 4096×4096×4096 cells, respectively.

Figure 3 shows the strong scaling result. For both lower- and higher-resolution runs, we observed that the MPF method works more efficiently than the Morton approach. In the result of the higher-resolution case on 128 GPUs, the performance of the MPF case has increased to 1.16 times that of the Morton curve.

The breakdowns of both the MPF and the Morton partitioning on 128 GPUs are shown in Figure 4. In both the MPF and the Morton partitioning, the computation time of each MPI process (each GPU) is almost the same, which means that both partitioning methods achieve good load balance. We confirmed that the MPF method reduces the halo communication time of each MPI process than that of the Morton curve. The communication time of each process is not uniform because the MPF model minimizes the total surface area corresponding to total communication volume but does not equalize the surface area of each sub-domain.

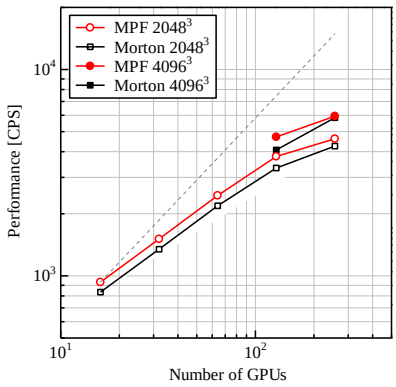

Figure 3: Strong scaling results.

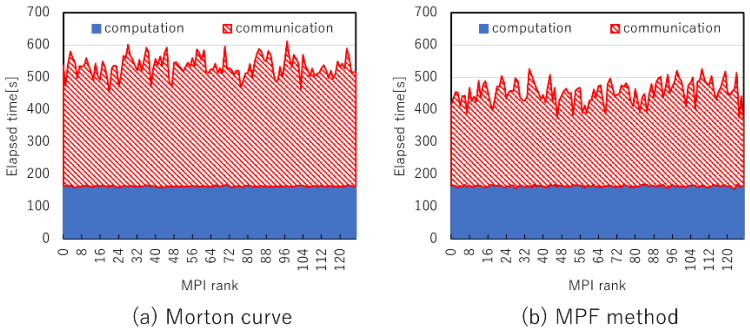

Figure 4: Breakdown in higher-resolution run on 128 GPUs.

4 CONCLUSIONS

We presented a domain decomposition method based on the MPF model. The additional term for balancing computation cost is introduced in the MPF equation, and the communication cost is implicitly optimized by minimizing the interface energy of the MPF model. We applied the proposed partitioning method to a block-based AMR application on multiple GPUs for an interface capturing method. We showed that the proposed method can successfully reduce communication costs than the SFC approach.

REFERENCES

[1] T. Uehara, “Numerical simulation of foam structure formation and destruction process using phase-field model,” *Adv. Mater. Res.*, **1042**, 65–69 (2014).