

9-4 Multi-Phase Fluid Simulations Using the World's Largest GPU Supercomputer — Development of Communication-Avoiding Matrix Solvers on GPU —




Fig.9-7 Communication overlap technique

(a) A computational domain on each GPU is decomposed into core and surface parts; (b) communication and core computation are processed simultaneously by splitting these processes on each GPU and processing them in parallel.

Existing computational fluid dynamics (CFD) simulations must be improved to analyze complex thermal-hydraulic phenomena in nuclear reactors. Whereas conventional CFD codes have been developed on petascale CPU-based computers, exascale GPU-based computers are essential for future simulations with ~ 100 billion grids. However, GPU-based computers have a significantly different performance gap between computation and communication than do conventional CPU-based computers because communication and computation are accelerated by several and several tens of times, respectively; the development of new matrix solvers capable of working at such a performance profile is thus required. In this study, GPU-computing technologies were developed for matrix solvers, which occupy a dominant cost of CFD codes. Additionally, extreme-scale multi-phase CFD simulations were accelerated using 7680 GPUs on Summit, which is the world's largest GPU supercomputer operated at the Oak Ridge National Laboratory in the United States.

CFD codes compute extreme-scale matrices that are obtained by discretizing a fluid model using a finite difference method. Conventionally, an iterative matrix solver based on the conjugate gradient (CG) method has been used. Parallel computation of the CG method requires two types of communication processes per iteration, i.e., halo data communication for finite difference computation on decomposed computational domains, and reduction communication for inner product




Fig.9-8 Scaling of computational performance on the GPU supercomputer Summit

Relationship between the total processing time per time step and number of GPUs in multi-phase fluid simulations (~ 7.5 billion grids) using 7680 GPUs of Summit (Oak Ridge National Laboratory).

operations of distributed vector data. Two approaches were used in this work to reduce these communication processes, including a communication and computation overlap technique and the communication-avoiding CG (CA-CG) method. In the communication and computation overlap technique, shown in Fig.9-7, a computational domain on each GPU is decomposed into core and surface parts. The halo data communication and the finite difference computation for the core, which does not require the halo data, are processed simultaneously to reduce the communication cost. In the CA-CG method, inner product operations for multiple iterations are processed at once by improving the mathematical algorithm of the CG method, thus reducing the number of reduction communications by an order of magnitude.

The performance of the developed method was evaluated in a multi-phase CFD simulation (approximately 7.5 billion grids, corresponding to the scale of four fuel assemblies) on the melt relocation of fuel assemblies on Summit. As shown in Fig.9-8, the total processing time was reduced nine-fold between 480 and 7680 GPUs, thus demonstrating scaling close to the target performance.

This study was conducted on “Exascale CFD Simulations at JAEA (CSC367)”, supported by the Oak Ridge Leadership Computing Facility (OLCF) Director’s Discretion Project.

(Yasuhiro Idomura)

Reference

Ali, Y., Idomura, Y. et al., GPU Acceleration of Communication Avoiding Chebyshev Basis Conjugate Gradient Solver for Multiphase CFD Simulations, Proceedings of 10th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (Scala 2019), Denver, U.S.A., 2019, p.1–8.