


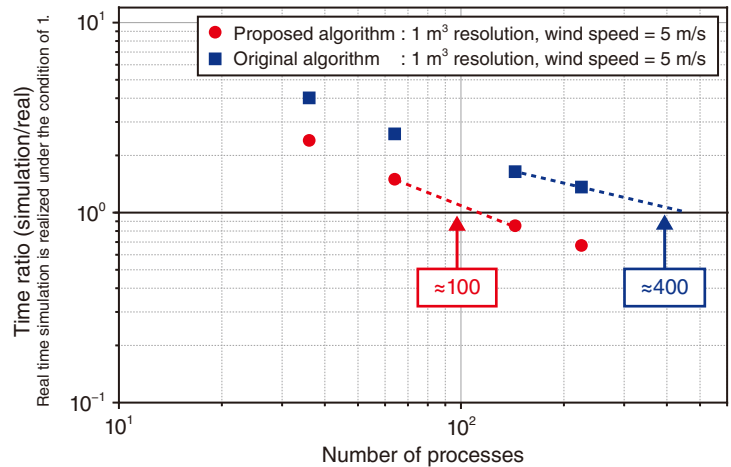
# 9-1 Acceleration of Plume Dispersion Model

## — Reduced-Communication Algorithm for GPU-Based Supercomputers —



**Fig.9-2 Timeline of the reduced-communication algorithm**

Each grid performs time evolution (y-axis) using neighbor grid points (x-axis). The proposed algorithm reduces the required communication by replacing redundant sequential communication with one communication in halo area.



**Fig.9-3 Performance of the reduced-communication algorithm on the supercomputer TSUBAME 3.0**

The reduced-communication algorithm (●) operated significantly faster than the conventional algorithm (■), and real-time simulation was realized using 100 GPUs.

Nuclear security depends on the accurate prediction of environmental dynamics of radioactive substances based on detailed airflow models of plume dispersion. Such wind simulations can be applied to a wide range of industrial needs, including smart city design and nuclear security. Since turbulent airflows in large cities can reach Reynolds numbers of several million, multi-scale computational fluid dynamics (CFD) simulations are needed. Predicting the diffusion of pollutants requires several times faster calculations than real time. Existing wind analysis codes cannot realize real-time simulation with high resolution; new models suitable for next-generation supercomputers must be developed.

As such, an adaptive mesh refinement method was employed that changes the grid resolution according to the flow scale and general-purpose GPU (GPGPU) programming was implemented that enables the use of thousands of GPU cores. Although GPU achieves about 10 times higher computational performance than CPU, the communication performance between processors is comparable; communication time is a severe bottleneck on GPU platforms. Furthermore, GPU computing requires complicated pre- and post-processing steps related to halo data

communications and synchronous processing of data between parallel computers. A reduced-communication algorithm based on a temporal blocking method was thus developed that optimizes data structures to reduce the number of communications (Fig.9-2). In the conventional algorithm, each GPU performs a time evolution only on the calculation area and updates the halo area by communicating with the neighbor GPUs at every step. The proposed algorithm replaces sequential communication processes in the halo area with redundant time evolution calculation (Fig.9-2, →) and MPI communications are executed after the computational procedures (Fig.9-2, →).

The wind airflow model was then evaluated on TSUBAME at the Tokyo Institute of Technology. By implementing the proposed algorithm, the communication cost was reduced by 64% and strong scaling was improved by 200 GPUs (Fig.9-3). The results indicate that real-time airflow simulations for a two square kilometer area with a wind speed of 5 m/s is feasible using 1 m resolution.

This study was partly supported by the Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures (JHPCN) (jh180041-NAH) in Japan.

### Reference

Onodera, N. et al., Communication Reduced Multi-Time-Step Algorithm for Real-Time Wind Simulation on GPU-Based Supercomputers, Proceedings of 9th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA 2018), Dallas, USA., 2018, p.9–16.