

11-3 Massively Parallel Fusion Plasma Simulation using 10^5 Processors

— Development of Latency-Hiding Communication Technique —

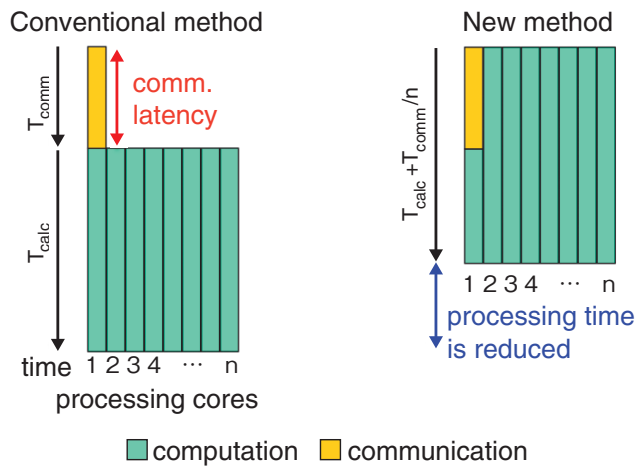


Fig.11-6 Scheduling computation and communication tasks on a processor with n processing cores

In the conventional method, computations are processed after communications, and most of the processing cores are idle during the communication latency. In the new method, these tasks are simultaneously processed on different processing cores, and the computing performance is improved by hiding the communication latency.

Predicting plasma turbulence is a key issue in estimating the performance of fusion core plasmas in ITER, which is determined by turbulent transport phenomena. A calculation model of fusion plasmas is described in five-dimensional (5D) phase space (three-dimensional position \times two-dimensional velocity), which is estimated as requiring $300^3 \times 128 \times 32 \sim 10^{11}$ grids for existing devices such as the JT-60U. The development of modern parallel computers enabled large-scale turbulence simulations of the 5D problem, and computational plasma turbulence research has been significantly advanced. However, ITER is a few times larger than existing devices, and simulating its turbulent transport requires an order of magnitude larger computational resources. Although a next-generation massively parallel computer such as the K computer is a promising solution, the latency due to data communications among $\sim 10^5$ processors has been a bottleneck in improving the computational performance of massively parallel simulations.

To resolve this issue, a novel communication technique that dramatically reduces the communication latency is developed.

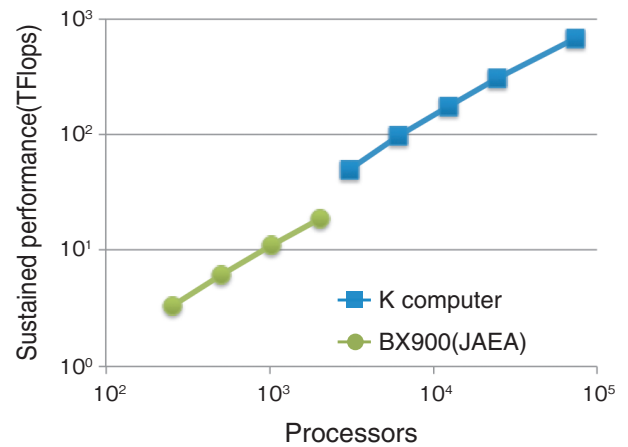


Fig.11-7 Computational performance (TFlops = 10^{12} operations per second) of fusion plasma simulations on the BX900 (JAEA) and the K computer

The maximum sustained performance of the BX900 was ~ 19 TFlops using 4098 processors with four processing cores. When the new latency-hiding communication technique was applied, the K computer achieved ~ 682 TFlops using 73728 processors with eight processing cores.

Turbulence simulations based on grids process calculation models in parallel by assigning decomposed small domains to each processor. This computational process, which requires boundary data in the neighboring domain, is normally performed after the boundary data is communicated. In the new technique, the computational performance is dramatically improved by assigning different tasks, computation and communication, to multiple processing cores on a single processor and processing them simultaneously (Fig.11-6). By applying this technique to 73728 processors on the K computer, massively parallel fusion plasma turbulence simulations were accelerated by ~ 35 times (Fig.11-7). This speedup enables ITER-size simulations, which were estimated to take more than a year on previous machines, to be completed in a week.

This work was supported by a Ministry of Education, Culture, Sports, Science and Technology of Japan (MEXT) grant for the HPCI Strategic Program Field No.4, "Next-Generation Industrial Innovations."

Reference

Idomura, Y. et al., Communication-Overlap Techniques for Improved Strong Scaling of Gyrokinetic Eulerian Code beyond 100 k Cores on the K-Computer, International Journal of High Performance Computing Applications, 2013, doi: 10.1177/1094342013490973.