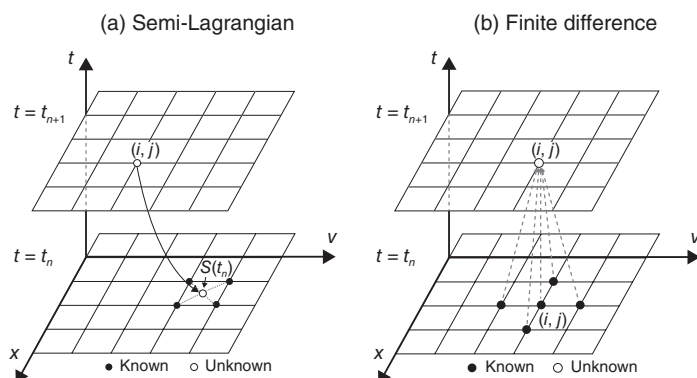


## 9-2 Development of a Low-Power-Consumption Computing Technology Based on Accelerators — Optimization of Nuclear-Fluid-Simulation Kernels on Accelerators —

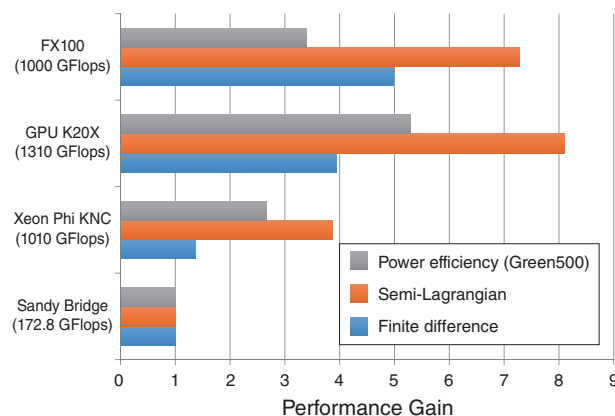


**Fig.9-4 Numerical schemes of fluid-simulation kernels**

Memory-access patterns in the time development of a convection operator using (a) semi-Lagrangian and (b) finite-difference schemes. Here,  $t$  and  $(x, v)$  show time and space, respectively. The semi-Lagrangian scheme traces the development of the solution by data access along streamlines, whereas the finite-difference scheme develops the solution via regular data access.

Higher computing performance is needed to improve nuclear-fluid simulations, such as analyses of severe accidents and the environmental dynamics of radioactive substances, and applying them to real problems. The development of exascale supercomputer is ongoing worldwide, and in Japan, the Post-K computer, which will succeed the K-computer, is being developed. One critical issue for such exascale supercomputers is low-power-consumption computing technology, which enables a performance more than two orders of magnitude higher than current supercomputers with similar power consumption. On the hardware side, accelerators have been developed, and both requirements on low power consumption and improved computing performance have been satisfied by suppressing the processor frequency and accumulating many computing cores within a processor. However, on the software side, it is not clear if conventional numerical schemes can be efficiently processed on such accelerators. In this work, we optimized the computing kernels of nuclear-fluid-simulation codes on accelerators and evaluated their computing performances.

We optimized convection-operator kernels based on the finite-difference and semi-Lagrangian schemes, which are representative schemes in fluid simulations, using graphical processing units (GPUs) and two kinds of many-core processors, namely Xeon Phi and FX100. GPUs have many computing cores designed for graphics processing, while many-core processors are based on conventional CPUs. The finite-difference scheme has regular memory access, and the semi-Lagrangian scheme traces the solution along streamlines,



**Fig.9-5 Performance comparisons of fluid-simulation kernels on accelerators**

Power efficiency and performance ratios of the semi-Lagrangian and finite-difference schemes on accelerators (FX100/GPU/XeonPhi), as normalized by those on a CPU (Sandy Bridge). The power efficiency is based upon computation speed (Flops) per unit power consumption in a matrix-computation benchmark called Green500.

which become random-memory access (Fig.9-4). Although the above accelerators have 3–5 times higher computing performance per unit power consumption than conventional CPUs, they have complicated hierarchical memory structures for supplying data to many computing cores. Therefore, adapting the memory-access patterns to hierarchical memory structures based upon an understanding of both the memory structures and the physical properties of the convection operator is the key to optimization.

It is found that FX100 can attain high computing performance without special optimization, as it is based on a shared in-processor memory, similar to that of conventional CPUs. On the other hand, on Xeon Phi, which is based on a distributed in-processor memory on each core, optimizing data access and computation patterns for the distributed memory improves computing performance by approximately 1.3 times, and on GPUs, an optimized algorithm that fully utilizes a texture memory for graphic processing leads to approximately 1.9 times higher computing performance. Thanks to the above optimization techniques, the finite-difference and semi-Lagrangian kernels are respectively accelerated up to 5 and 8 times compared with conventional CPUs (Fig.9-5), and the use of accelerators to attain low-power-consumption computing capability has been demonstrated.

This research was supported by the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan, Post-K Priority Issue 6, “Accelerated Development of Innovative Clean Energy Systems”.

### Reference

Asahi, Y., Idomura, Y. et al., Optimization of Fusion Kernels on Accelerators with Indirect or Strided Memory Access Patterns, IEEE Transactions on Parallel and Distributed Systems, vol.28, no.7, 2017, p.1974-1988.