jh180080-MDJ

Gyrokinetic simulation of divertor heat-load in magnetic fusion devices

Toseo Moritaka（National Institute for Fusion Science）

Abstract
X-point Gyrokinetic Code (XGC) is a gyrokinetic particle-in-cell code developed for whole-volume modeling of axisymmetric magnetic fusion devices, i.e., Tokamaks, and useful for precise estimation of diverter heat load. (a) Significant data output produced by production runs of XGC could be a performance bottleneck. This problem has been addressed with the ADIOS parallel I/O middleware with various techniques developed for high-performance I/O. We have successfully demonstrated XGC build on Tsubame3 and conducted a set of performance study by utilizing multiple GPUs and two I/O systems with the heterogeneous filesystems. (b) We have developed an extended version of XGC for non-axisymmetric magnetic fusion device, i.e., Stellarators. The developed version is validated by basic transport simulations in the core region and particle tracing calculations in the entire region. The kernel part relevant to the particle push calculation has been ported and partially optimized on Tsubame 3.

1. Basic Information

(1) Collaborating JHPCN Centers

Global Scientific Information and Computing Center, Tokyo Institute of Technology

(2) Research Areas

- ■ Very large-scale numerical computation
- ■ Very large-scale data processing
- □ Very large capacity network technology
- □ Very large-scale information systems

(3) Roles of Project Members

Jong Choi, ORNL, USA – porting XGC enabled with GPU and Adios and investigating XGC performance around CPU/GPU and I/O subsystem.

Scott Klasky, ORNL, USA – supervise XGC performance study on Tsubame3.

Toseo Moritaka, NIFS, Japan – development of extended version of XGC for non-axisymmetric fusion devices.

Takayuki Aoki, Tokyo Tech Univ., Japan – supervise optimization and large-scale computing on Tsubame3.

Choong-Seock Chang, PPPL, USA– administration of XGC project.

Yasuhiro Idomura, JAEA, Japan – supervise plasma transport simulation in magnetic fusion devices.

Seiji Ishiguro, NIFS, Japan – supervise edge plasma simulation in magnetic fusion devices.

2. Purpose and Significance of the Research

X-point Gyrokinetic Code (XGC) is a gyrokinetic particle-in-cell code developed for whole-volume modeling of magnetic fusion devices. The whole-volume modeling includes entire region of the device from magnetic axis to material wall and is useful to explore dynamic interaction between plasma transport phenomena in core and edge regions. The whole-volume modeling is also useful for precise estimation of divertor heat load under the influence of plasma kinetic effects. XGC employs finite element method using unstructured mesh for field calculations to take realistic device and field line structures in into account. Mesh resolution is defined by microscopic particle motion, e.g., ion gyration motion. In addition, a large number of computational particles are needed to describe non-thermal plasma dynamics near the material wall. Therefore code optimization to large-scale computing systems is quite important to realize the whole-volume gyrokinetic simulation of magnetic fusion devices.

(1: US-side) XGC has been tested and running on many different USA DOE machines and proven its performance with GPU. Production runs of XGC also can produce significant data output, and this could be a performance bottleneck. Our goal is to conduct a set of performance study for running XGC with GPUs and the heterogeneous filesystems on Tsubame3.

(2: Japan-side) So far, XGC has been applied to Tokamaks with axisymmetric geometries such as ITER (International Thermonuclear Experimental Reactor). Stellarators using

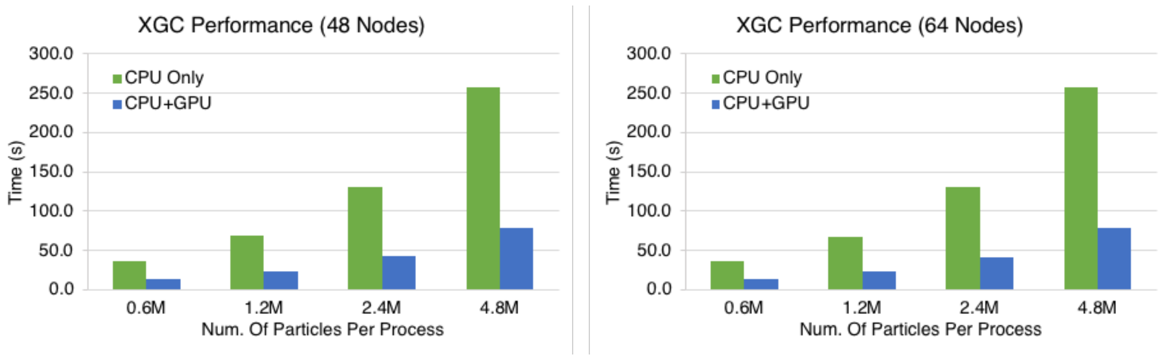

Figure 1. XGC with CPU/GPU performance for D3D case on Tsubame3.

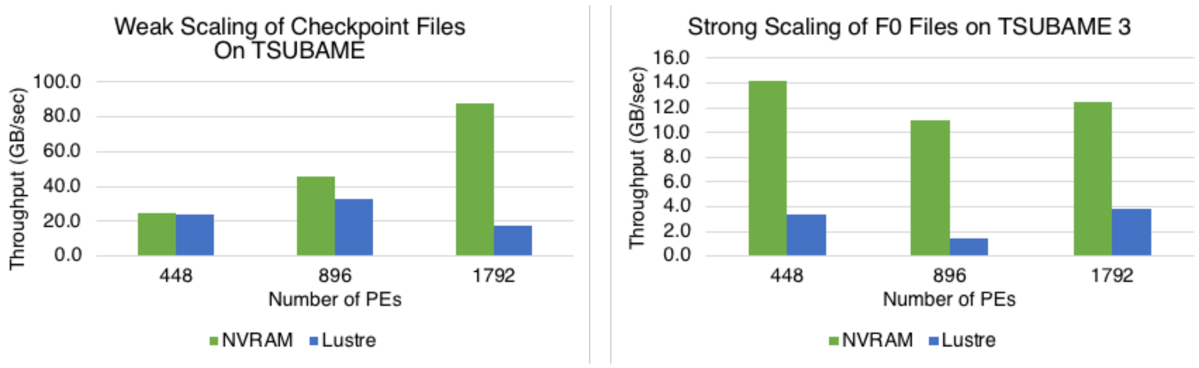

Figure 2. XGC I/O performance on NVRAM and Lustre parallel filesystem.

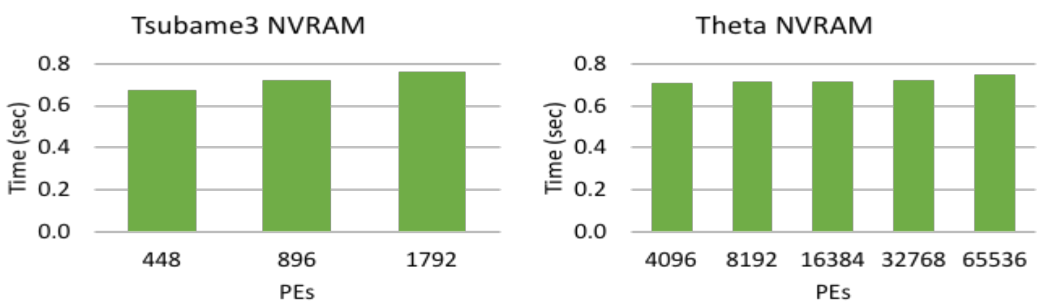

Figure 3. NVRAM-based I/O performance compared with Theta, ALCF, USA.

modular or helical coils are another candidate of magnetic fusion devices. For whole-volume modeling of Stellarators, we are extending XGC to non-axisymmetric geometries with stochastic magnetic field line structures in the Stellarator edge region.

3. Significance as a JHPCN Joint Research Project

(1) Accessing the capability and the performances of XGC on Tsubame3, JHPCN, is an essential task before conducting large scale physics studies. Not only CPU/GPU performance of XGC but I/O performance is also a key in the XGC performance study since production runs of XGC can produce significant data output which can be a performance bottleneck. This problem has been addressed, adequately, with the ADIOS parallel I/O middleware. With ADIOS, the large checkpoint files from, for example, an ITER simulation are handled efficiently with various techniques developed for high-performance I/O, including aggregating I/O operations, delaying meta-data collection, exploiting staging area, and bonding multiple filesystems. Currently, ADIOS can obtain over 180GB/s to the two parallel file systems with XGC on Titan and 210 GB/s on Cori's single Lustre filesystem. Comparing the performances with other DOE machines will be a meaningful research task.

(2) Because of complicated magnetic field structure and particle motion, higher spatial resolution and more computational particles are required for whole-volume modeling of Stellarators. The extended version has been developed on CPU machines such as Plasma Simulator (FX100) in NIFS. For production runs, it is useful to combine the developed part with the latest version of XGC and ADIOS optimized for GPU machines in USA. Tsubame 3 is the largest GPU machine in Japan and could be a good platform for further US-Japan collaboration.

4. Outline of the Research Achievements up to FY2017

5. Details of FY2018 Research Achievements

(1) We have conducted a set of performance experiments to investigate how XGC and ADIOS can work on Tsubame3 and compare performances with other DOE HPC machines in the USA.

First of all, we built XGC on Tsubame3 and measured performance by using multiple parallel nodes. In each node, we ran 4 MPI processes to match 4 GPUs on Tsubame3 node and utilized 7 OPEN MP threads.

Figure 1 shows the performance results of D3D case with XGC on Tsubame3 by using 48 and 64 nodes. We observed 3.25x performance increase with CPU+GPU version of XGC compared with CPU only version.

Second, we focused on measuring I/O performance for XGC checkpoint writing on Tsubame3 file systems, Lustre parallel filesystem and NVME-based BeeGFS. Figure 2 shows the performance comparison between Lustre and NVME for both weak-scale and strong scale I/O scenarios in XGC.

Figure 3 shows the comparison of Tsubame3's NVME and node-local NVME on Theta, ANL, USA. We observed comparative performance results.

(2) We have developed the following new features to apply XGC to non-axisymmetric

magnetic fusion devices, i.e., Stellarators. (a) An interface to three-dimensional magnetic field equilibrium data is developed. The equilibrium data are distributed onto each torus cross section. Three-dimensional spline interpolation is implemented for particle-mesh interpolation using the distributed equilibrium data. (b) Field-following unstructured meshes are generated based on flux coordinates in the core region and numerical field line tracing in the edge region.

Developed numerical schemes and unstructured meshes are validated by several test calculations on collisionless damping of GAM oscillation and neoclassical transport phenomena in the core region of Large Helical Device (LHD). We have also tested high-energy particle orbits in the entire region of LHD. Obtained results are consistent with results obtained from other established simulation codes. In addition, finite element solver is applied to the unstructured mesh and convergence to a smooth solution is confirmed for given charge density profiles independently of a boundary between the core and the edge regions.

We have started to port the extended version of XGC to Tsubame 3. In production runs, particle orbit calculation (particle push) in electron sub-cycling scheme is the most computationally intensive part. Three-dimensional spline interpolation to estimate magnetic field at the particle position is one of the main differences from original XGC for Tokamaks. By using openACC directives, elapse time in the interpolation process with 1 GPU decreases to become approximately 1/23 of that obtained with 1 CPU core. In the present case, we use unified memory and data transfer between CPU and GPU are not specified explicitly. While some other loops relevant to time integration of the gyrokinetic equations of motion are not accelerated effectively at present, elapse time of the particle push part decreases by 1/8 with 1 GPU.

6. Progress of FY2018 and Future Prospects

(1) We have successfully demonstrated XGC build on Tsubame3 and conducted a set of performance study by utilizing multiple GPUs and two I/O systems on Tsubame3.

In the next, we will continue to run XGC with ADIOS to study the in-memory in situ data science which can improve data processing during the simulations. We will investigate the coupling of multi-scale fusion physics codes with XGC to enable long-timescale fusion simulations. The US team will investigate the use of in-node SSDs (Burst Buffer) for this XGC coupling workflow.

(2) We have developed the extended version of XGC for Stellarators. The developed code is validated by basic plasma phenomena and now able to employ for plasma transport simulations in the core region of Stellarators. For whole-volume simulation including the edge region, we needed to improve the field solver to evaluate self-consistent electrostatic field and charge density profiles.

A kernel code of the extended version relevant to the particle push calculation has been ported to Tsubame 3. We confirm that the kernel works correctly on Tsubame 3. One of the main part of the kernel code, i.e., three-dimensional spline interpolation for particle-mesh interpolation, is accelerated with openACC directives. Data transfer between CPU and GPU should be specified explicitly to optimize the other do loops.

We will continue the optimization of the kernel on Tsubame 3 and combine with the original XGC for Tokamaks already optimized for GPU computers by US-team. Local plasma processes such as multi-species binary collision and other ionization processes are modeled in the original XGC and can be used also for Stellarators. Using the binary collision module, we will investigate neoclassical transport phenomena in the core region of LHD including potential variation on flux surfaces relevant to global torus structures and plasma micro-turbulences.

7. List of Publications and Presentations

(1) Journal Papers

[1] Toseo Moritaka, Robert Hager, Michael Cole, Samuel Lazerson , C.S. Chang, S. Ku, Seikichi Matsuoka, Shinsuke Satake and Seiji Ishiguro, Plasma, accepeted (2019)

(2) Conference Papers

[1] T. Moritaka, R. Hager, M.D.J Cole, S. Laserzon, S. Satake, C.S. Chang, S. Ku, S. Matsuoka, S. Ishiguro, Proceedings of the 27th IAEA Fusion Energy Conference, 2018.

(3) Oral Presentations

[1] Scott Klasky and Jong Choi, "XGC1 on TSUBAME 3 with ADIOS 2", US-Japan JIFT Exascale Computing Workshop, Princeton Plasma Physics Laboratory, Princeton, NJ USA, July 30-31, 2018

[2] T. Moritaka, S. Matsuoka, S. Satake, S. Ishiguro, R. Hager, M.Cole, S. Lazerson, C-S. Chang, S. Ku, US-Japan JIFT Exascale Computing Workshop, 2018. 07. 30 - 2018. 07. 31, Princeton USA.

[3] T. Moritaka, ,US-Japan Joint Institute of Fusion Theory Workshop on "Multi-scale Simulations in Plasma Physics", 2019. 01. 11 - 2019. 01. 12, Inuyama, Japan.

(4) Others