SCA/HPCAsia 2026
Jan. 26 - 29, 2026, Osaka, Japan

Preparing for the Post-Fugaku Era: A Performance Evaluation of High-Performance Molecular Dynamics and Electronic Structure Codes on GPU Architectures

Kanako Yoshizawa, Daisuke Matsuoka, Gilles Gouaillardet, Akira Azami, Asako Terasawa, Masato Matsui,
Satoru Shingu, Tatsuya Nakano, Takaaki Noguchi, Takayuki Otani, Naoki Sueyasu, Yoshinori Kusama
Research Organization for Information Science and Technology (RIST)

RIST User Support

- Provides advanced consulting and tuning support for HPCI (High Performance Computing Infrastructure) users and prospective users.
- Installs application software on HPCI systems, including Supercomputer Fugaku.
- Provides information on the installed application software to make them ready-to-use in the HPCI systems.
- Application software includes:
  - ➢ Commonly used Open Source Software (OSS).
  - ➢ Software developed by Japanese Groups and recommended by application development projects and the computational science community.

Installed Application Software

- RIST installs application software packages listed below on Fugaku and other HPCI systems.

  (Note: Each HPCI Resource Provider also installs its own application software on HPCI systems.)

■ Widely-used OSS installed on Fugaku

- ◆ LAMMPS
- ◆ OpenFOAM
- ◆ GROMACS
- ◆ Quantum ESPRESSO
- ◆ FDS

■ Software developed by Japanese National Projects

- ◆ ABINIT-MP (14)
- ◆ AkaiKKR (13)
- ◆ ALAMODE * (12)
- ◆ mVMC * (12)
- ◆ HΦ * (14)
- ◆ NTChem (12)
- ◆ OpenMX (12)
- ◆ GENESIS (13)
- ◆ PHASE/0 (14)
- ◆ MODYLAS (12)
- ◆ Phonopy (12)
- ◆ SMASH (12)
- ◆ SALMON * (12)
- ◆ FrontFlow/blue (13)
- ◆ FrontISTR (14)
- ◆ FFVHC-ACE (12)
- ◆ FFX (12)

* HPCI Software Award 2024
* HPCI Software Award 2023

() : The number of HPCI systems where the application software is installed.

Toward Post-Fugaku

- It has been decided that the Post-Fugaku will be equipped with NVIDIA GPUs. Toward the user support on the Post-Fugaku, it is extremely important to gain experience in running application software on existing GPU systems.
- As the Post-Fugaku aims to advance “AI for Science,” enabling AI-driven simulation is becoming increasingly important.
- We measured the performance of the OSS on the latest GPU system, Miyabi-G, as well as AI-integrated software, which ensures they will run efficiently on future GPU platforms.

Miyabi-G and Fugaku: Hardware Specifications

- We performed performance evaluations of these application software on Miyabi-G.

Feature	Miyabi-G (Grace-Hopper Superchip)	Fugaku (A64FX CPU-only)
Deployment	JCAHPC	RIKEN R-CCS
CPU Cores (per node)	72 cores	48 cores
GPU Count (per node)	1 × NVIDIA Hopper H100	None (CPU-only)
Memory (per node)	120 GB (CPU) + 96 GB (GPU)	32 GB

AI for Science: Molecular Dynamics with AI

- In the materials science community, Molecular Dynamics with AI has recently gained significant attention as an important form of AI-driven simulation.
- MACE enables efficient creation of machine learning potentials for Molecular Dynamics simulations.
- ◆ The result of performance evaluation of LAMMPS with MACE (classical molecular dynamics + machine learning potential) on Miyabi-G is shown in the table below.
  - ➢ **Test data** : methane molecule (5 atoms, 100steps)
  - ➢ **LAMMPS version 10Sep2025**
    - ・**Language version :** nvidia/24.9, nv-hpcx/24.9
    - ・**Environment variables (C++ compile options) :** -fast -std=c++17 -mp
  - ➢ **MACE version 0.3.3**
    - ・**Language version:** Python 3.11
    - ・**Package used:** Pytorch 2.5.1 (“with GPU” / “without GPU”)

“with GPU” / “without GPU”	Loop time (sec)
1CPU (1Core) (without GPU)	7.86
1CPU (1Core) + 1GPU (with GPU)	1.77

Performance Evaluation of OSS on Miyabi-G

- Execution times, etc. measured under certain conditions using typical sample data are provided in comparison with those on Fugaku.

  (Note: The performance will not be the same for different compilation and execution conditions.)

◆ LAMMPS version 22Jul2025 update 1 (classical molecular dynamics)

- ➢ **Test data** : Lennard-Jones (LJ) particles [potential:LJ], Metal (Copper) [potential:EAM], Biomolecule (rhodopsin)) [potential:RHODO], (256,000 atoms)
- ➢ **Miyabi-G**
  - ・**Language version :** gcc/11.4.1, cuda/12.6, ompi-cuda/4.1.6-12.6, kokkos/4.5.01 (Only KOKKOS package)
  - ・**Environment variables (C++ compile options)**
    GPU package : -g -O2 -DNDEBUG -mcpu=neoverse-v2 -msve-vector-bits=128 -std=c++17 -fPIC
    KOKKOS package: -DKokkos_ARCH_ARMV9_GRACE=yes -DKokkos_ARCH_HOPPER90=yes -DKokkos_ENABLE_OPENMP=yes -DKokkos_ENABLE_CUDA=yes
- ➢ **Fugaku**
  - ・**Used spack**: lammps@20250722.1 %fj@4.12.0

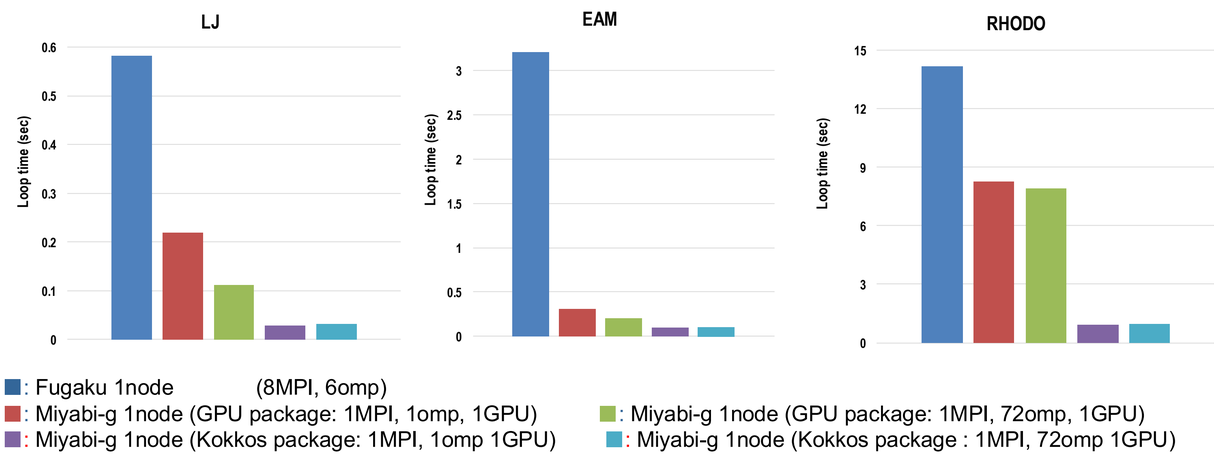

◆ GROMACS version 2025.3 (classical molecular dynamics)

- ➢ **Test data** : lysozyme (water-soluble protein), DPPC (lipid bilayer) (5,000 nsteps)
- ➢ **Miyabi-G**
  - ・**Language version :** nvidia/24.9, nv-hpcx/24.9
  - ・**Environment variables (C++ compile options) :** -msve-vector-bits=128 -fast -Mnoflushz -Mnodaz -Wno-cast-qual SHELL:-mp -fast -O3 -DNDEBUG
- ➢ **Fugaku**
  - ・**Used spack**: gromacs@2025.3 %fj@4.12.0

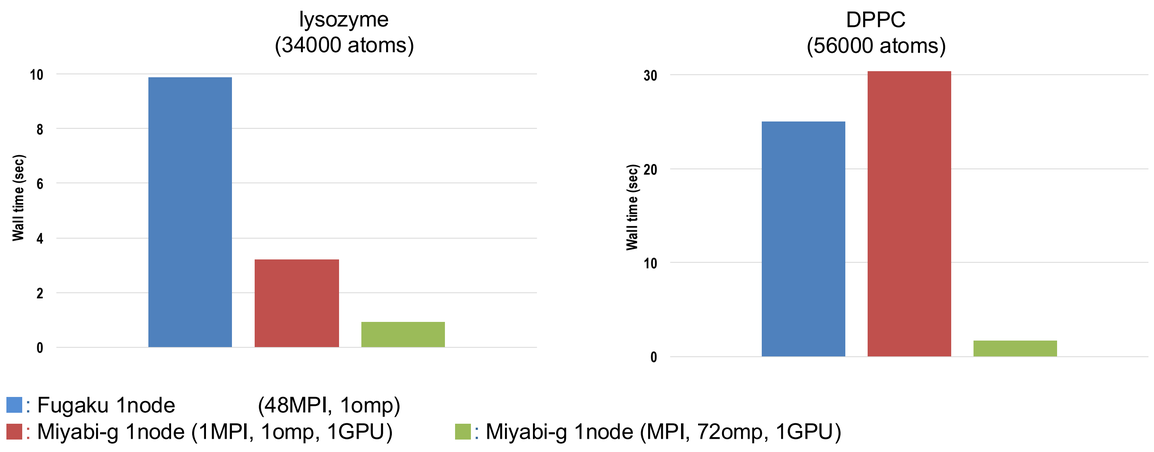

◆ Quantum ESPRESSO version 7.3.1 (electronic-structure calculations)

- ➢ **Test data** : test_3 (95 atoms), AUSURF112 (112 atoms), PSIWAT (586 atoms)
- ➢ **Miyabi-G**
  - ・**Language version :** nvidia/24.9, nv-hpcx/24.9
  - ・**Environment variables (Fortran compile options) :** -fast -mp -Mcache_align -Mpreprocess -Mlarge_arrays
- ➢ **Fugaku**
  - ・**Used spack**: quantum-espresso@7.3.1 %fj@4.12.0

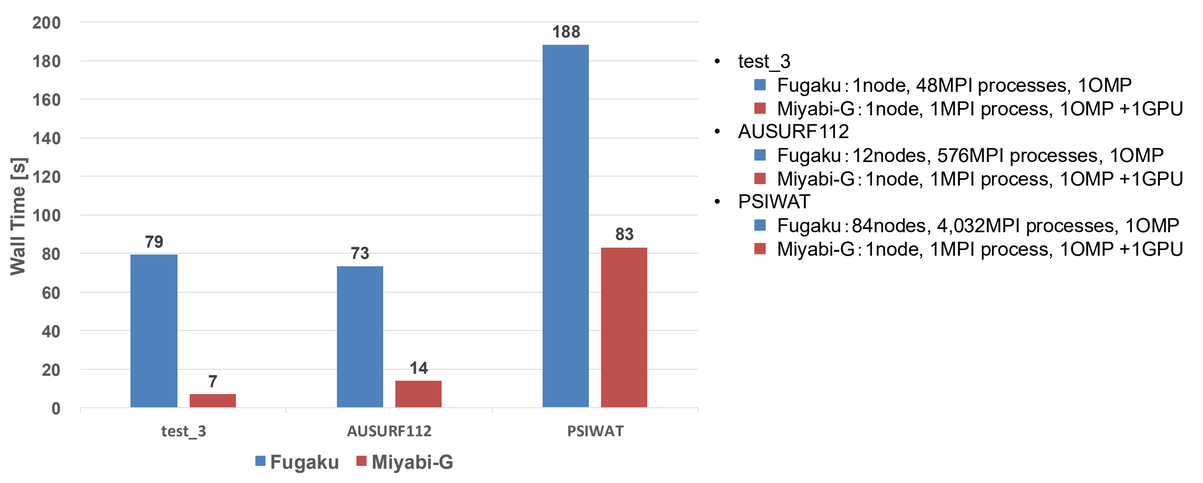

Summary

- In this study, we evaluated the performance of LAMMPS, GROMACS and Quantum ESPRESSO on both Miyabi-G and Fugaku.
- Our evaluation confirms that AI-driven LAMMPS simulation using MACE shows much better performance by using GPUs.
- Our results show that all these applications can achieve high performance on GPUs.