THE UNIVERSITY OF TOKYO

東京大学情報基盤センター INFORMATION TECHNOLOGY CENTER, THE UNIVERSITY OF TOKYO

JCAHPC

# Long but "Straight" Road to OFP-II


## Information Technology Center, The University of Tokyo

## JCAHPC (Joint Center for Advanced HPC) & OFP-II

筑波大学 University of Tsukuba　東京大学 THE UNIVERSITY OF TOKYO　JCAHPC

In 2013, Center for Computational Sciences, University of Tsukuba (CCS), and ITC agreed to establish the Joint Center for Advanced High-Performance Computing (JCAHPC). JCAHPC consists of more than 20 faculty and staff members of CCS and ITC. Originally, the Primary mission of JCAHPC is designing, installing, and operating the Oakforest-PACS system (OFP) (2016-2022, 8,208 nodes of Intel Xeon/Phi, 25PF, 6th in the 48th TOP500 (November 2016)). In addition, CCS and ITC develop system software, numerical libraries, and large-scale applications for OFP in collaboration. JCAHPC is a new model for collaboration for research and development between multiple supercomputer centers. OFP has contributed significantly to the development of computational science in Japan and around the world. Especially after retirement of K computer in 2019, OFP worked as the National Flagship System in fact until operation of Fugaku started in 2020. OFP retired on March 31, 2022. **JCAHPC is introducing OFP-II with 100+ PFLOPS, the successor of OFP, which starts operation in January 2025.**

## Switching from Many-core CPUs to NVIDIA's GPUs on OFP-II

In order to attain maximum performance under constraint of power consumption, switching to accelerators, such as GPU, is essential in next 10 years. After February 2022, energy cost has been increasing, while that is now stable and reducing. OFP-II consists of 2 types of node groups, Group-A with general-purpose CPUs only, and Group-B with GPUs. Actually, we have already decided on the architecture of GPU in June 2022, more than one year ago. **We need to prepare for porting codes of more than 3,000 users of OFP to GPUs, and we estimated we need 18-30 months for that.** Selection was based on the 7 benchmarks. **Finally, we adopted NVIDIA's GPU, which is H100 or its successor.** Performance itself, and portability of the applications in Fortran are highly evaluated. OpenACC and Standard Parallelism are relatively easy, and suitable for existing applications by users with MPI/OpenMP. We released RFP of OFP-II in late August. We expect peak performance of Group-B is more than 100 PF. CPUs in Group-A and Group-B could be different.

**Group-A CPU only**

**Group-B CPU+GPU**

| Name of the Code | Description | Lang. | Parallelization | GPU |
|---|---|---|---|---|
| P3D | 3-D Poisson's Equation by Finite Volume Method | C | OpenMP | |
| GeoFEM/ICCG | Finite Element Method | Fortran | OpenMP, MPI | |
| H-Matrix | Hierarchical-Matrix calculation | Fortran | OpenMP, MPI | |
| QCD | Quantum-Chromo Dynamics simulation | Fortran | OpenMP, MPI | CUDA |
| N-Body | N-Body simulation using FDPS | C++ | OpenMP, MPI | CUDA |
| GROMACS | Molecular Dynamics simulation | C++ | OpenMP, MPI | CUDA, HIP, SYCL |
| SALMON | Ab-initio quantum-mechanical simulator for optics and nanoscience | Fortran | OpenMP, MPI | (OpenACC) |

List of Seven Benchmarks for Selection of GPUs in 2022

17 Groups for "Supported Porting"

| Category | Name (Organizations) | Target, Method etc. | Language |
|---|---|---|---|
| Engineering (3) | FrontISTR (U.Tokyo) | Solid Mechanics, FEM | Fortran |
| | FrontFlow/blue (U.Tokyo) | CFD, FEM | Fortran |
| | FrontFlow/red (Advanced Soft) | CFD, FVM | Fortran |
| Biophysics (3) | ABINIT-MP (Rikkyo U.) | Drug Discovery etc., FMO | Fortran |
| | UT-Heart (UT Heart, U.Tokyo) | Heart Simulation, FEM etc. | Fortran, C |
| | Lynx (Simula, U.Tokyo) | Cardiac Electrophysiology, FVM | C |
| Physics (3) | MUTSU/iHallMHD3D (NIFS) | Turbulent MHD, FFT | Fortran |
| | Nucl_TDDFT (Tokyo Tech) | Nuclear Physics, Time Dependent DFT | Fortran |
| | Athena++ (Tohoku U. etc.) | Astrophysics/MHD, FVM/AMR | C++ |
| Climate/ Weather/ Ocean (4) | SCALE (RIKEN) | Climate/Weather, FVM | Fortran |
| | NICAM (U.Tokyo, RIKEN, NIES) | Global Climate, FVM | Fortran |
| | MIROC-GCM (AORI/U.Tokyo) | Atmospheric Science, FFT etc. | Fortran77 |
| | Kinaco (AORI/U.Tokyo) | Ocean Science, FDM | Fortran |
| Earthquake (4) | OpenSWPC (ERI/U.Tokyo) | Earthquake Wave Propagation, FDM | Fortran |
| | SPECFEM3D (Kyoto U.) | Earthquake Simulations, Spectral FEM | Fortran |
| | hbi_hacapk (JAMSTEC, U.Tokyo) | Earthquake Simulations, H-Matrix | Fortran |
| | sse_3d (NIED) | Earthquake Science, BEM (CUDA Fortran) | Fortran |

## Long but "Straight" Road to OFP-II

Activities for porting are done under strong support by NVIDIA-Japan. Because we have more than 3,000 users, we have two categories for support. Generally, we recommend users to port their applications to GPU by themselves. We have many options, such as:

- ✓ 1-week Hackathons: every 3 or 4 months, online and in-person
- ✓ Monthly Open Meeting for Consultation via Zoom
- ✓ Excellent portal site with useful information (in Japanese)

[in Japanese]

We have another procedure, "**Supported Porting**". We are supporting 17 groups of community codes and OpenFOAM. While we also prepare special budgets for outsourcing, many members of these groups are also attending Hackathons. **Mostly, our users' codes are parallelized by MPI+OpenMP, therefore OpenACC is recommended.**

## Example: Porting CG Solvers of Lynx code to GPUs [Naruse et al. 2022]

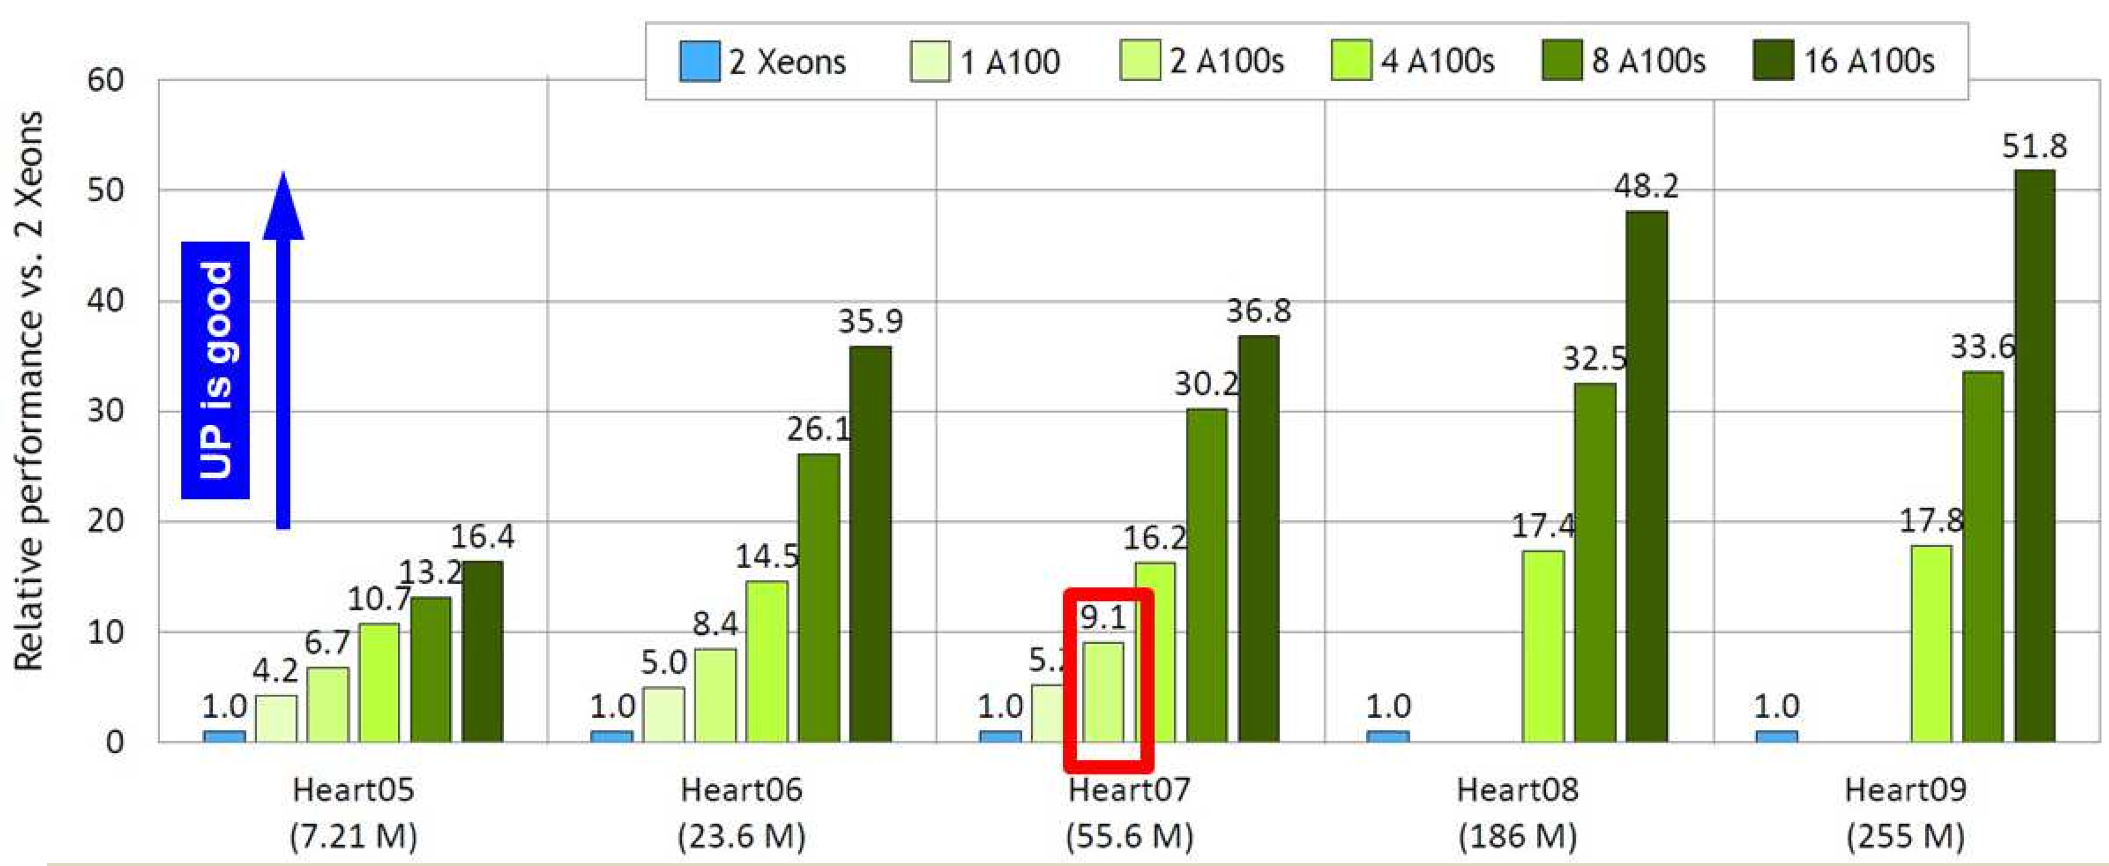


Relative "Best" Performance over 2-Sockets of Xeon Icelake, Performance Ratio of 2-A100 over 2-Xeon is up to 9.1 (Heart07), Reasonable based on memory bandwidth (2-Xeon: 0.41 TB/s, 2-A100: 3.1 TB/s)

Lynx is a simulator for Cardiac Electrophysiology by Simula Research Laboratory in Norway. It solves 3D diffusion/reaction equations by FVM with unstructured tetrahedral meshes. The most expensive procedure is conjugate gradient solver for large-scale linear equations. This part was ported to GPU by OpenACC and Standard Parallelism on Wisteria/BDEC-01 (Aquarius) with NVIDIA A100 GPUs. Reasonable speed-up compared to Intel Xeon Icelake was obtained. Performance improvement by GPU is significant for larger problems.

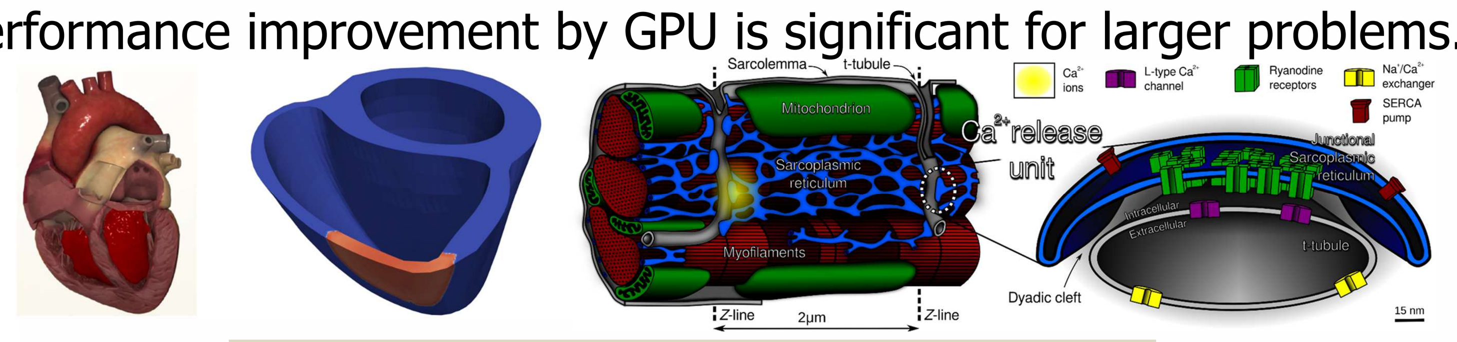


Models in Lynx for Cardiac Electrophysiology

simula　NVIDIA　東京大学 THE UNIVERSITY OF TOKYO


Contact: Kengo Nakajima (nakajima@cc.u-tokyo.ac.jp)
Information Technology Center, The University of Tokyo


Full Slides for this Material