jh240029

Innovative Computational Science by Integration of Simulation/Data/Learning on Heterogeneous Supercomputers

Kengo Nakajima（The University of Tokyo, Japan）

This project advances the frontier of computational science through the seamless integration of Simulation, Data, and Learning (S+D+L), leveraging heterogeneous supercomputing environments such as Wisteria/BDEC-01 (University of Tokyo), Flow (Nagoya University), and mdx. During fiscal years 2021 and 2022, our primary focus was on earthquake simulations incorporating real-time data assimilation. Beginning in FY.2023, we broadened the scope of the S+D+L paradigm, expanding its application to diverse scientific domains and initiating collaborative research and software development with esteemed international partners. Since FY.2019, the project's principal investigator and colleagues have spearheaded the development of the pioneering software platform h3-Open-BDEC, designed to facilitate the integration of S+D+L across heterogeneous computing systems such as Wisteria/BDEC-01. Currently, Wisteria/BDEC-01 ("Odyssey" for Simulations with A64FX, "Aquarius" for Data/ML/AI with NVIDIA A100) operates as a full-scale platform for "S+D+L" integration with h3-Open-BDEC, gaining global attention. In FY.2023, we conducted research in atmospheric simulations, ported international partner applications to Wisteria/BDEC-01, and investigated "S+D+L" integration using h3-Open-BDEC. For FY.2024, we switched to the international project and focus on the fields of earth science, life science, and library/software/tool for integration of "S+D+L".

1. Basic Information

(1) Collaborating JHPCN Centers

Hokkaido University
Tohoku University
The University of Tokyo
Nagoya University
Kyushu University
mdx

(2) Theme Area

Large scale computational science area

(3) Project Members and Their Roles

① : Earth Science, ② Life Science, ③ Library/Software/Tool

U.Tokyo: Nakajima (Leading PI) ① ② ③, Hanawa ③, Sumimoto ① ② ③, Shimokawabe ③ ,Yamazaki ① ③, Furumura (Co-PI) ①, Tsuruoka ①, Nagao ①, Itoh ①, Ichimura ①, Fujita①, Suzumura③, Hanai③, Kuno③
Hokkaido U.: Fukaya③
Nagoya U.: Hoshino③, Kawai①③, Ueno③
Kyushu U.: Ohshima③
U.Hyogo: Shiba②③
RIKEN: Hascoet①, Sugita②, Yagi②, Ito②, Imamura③, Nakao③, Dawson②
NIES: Yashiro①③
CliMTECH: Arakawa①②③
CEA (France): Boillod-Cerneux (Co-PI)②, Badri ①, Foerster①, Lomet①, Genovese②
JSC (Germany): Di Napoli (Co-PI)③, Caviedes-Voullieme①③, Koh②③, Suarez①②③, Wu③, Conrads③
FAU (Germany): Wellein③, Hager③, Afzal③, Ujeniya③, Gruber③
BUW (Germany): Yoda①
TUM (Germany): Chen①
U.Cologne (Germany): Kubicki③
RBI (Croatia): Davidovic ③, Mijic ③, Badrinarayanan③
Fujitsu: Sakaguchi③, Kasai①③, Obinata③
Hitachi: Matsuba③

2. Purpose and Significance of the Research

This project advances computational science by

integrating "Simulation/Data/Learning (S+D+L)" using heterogeneous supercomputers like Wisteria/BDEC-01 and Flow. In FY.2021 and FY.2022, we focused on earthquake simulation with real-time data assimilation. Since FY.2023, we expanded S+D+L to other fields, initiating joint research with international partners. Using the software platform h3-Open-BDEC, Wisteria/BDEC-01 has gained global attention. Recent research includes atmospheric simulations and integrating international applications. For FY.2024, we switch to the international project and focus on the fields of earth science, life science, and library/software/tool for integration of "S+D+L". JSC (Jülich Supercomputing Centre, Germany), one of the international partners of this project, also conducts research on heterogeneous computing based on the idea of Modular Supercomputing Architecture. Some of the target applications in this project by JSC, such as TSMP and ChASE, are already feasible in heterogeneous environments.

3. Significance as JHPCN Joint Research Project

Both of “Wisteria/BDEC-01” and “Flow” are based on heterogeneous architecture, can cover a wide range of workloads, and are optimum as platforms for integration of “S+D+L”. Moreover, the collaboration with mdx promotes the utilization of research outcomes for dissemination and enlightenment. Finally, Osaka U., and Kyushu U. have also introduced heterogeneous systems, which provide an opportunity for the deployment and validation of the developed software and applications in this project.

4. Outline of Research Achievements until FY.2023

In FY.2021 and FY.2022, we focused on earthquake simulation with real-time data assimilation. Since FY.2023, we expanded S+D+L to other fields, initiating joint research with international partners. Using the software platform h3-Open-BDEC, Wisteria/BDEC-01 has gained global attention. Recent research includes atmospheric simulations and integrating international applications.

5. Details of FY.2024 Research Achievements

(a) Earth Science

(a-1) Ensemble Coupling [13]

Ensemble calculation is a technique that reduces computational uncertainty. and enables quantitative evaluation. However, it is difficult to run a large number of ensembles simultaneously due to the constraints of computational resources. Therefore, we proposed a new computational method “ensemble coupling“, that combines low-resolution ensemble calculations with high-resolution single calculations. To realize this idea, we developed a coupling library h3-Open-UTIL-MP(UTIL/MP) and applied it to the non-hydrostatic atmospheric model NICAM. In this case, we couple low-resolution (220 km resolution) ensemble with a high-resolution (14 km resolution) standalone calculation. As a result, we achieved a 112-fold improvement in performance in node time compared to the high-resolution 64 ensemble. Next, we increased the number of ensembles to 1024 and used more realistic input values based on the ERA5 reanalysis data to verify the computational performance and model reproducibility. Table 1 shows the overall execution time for nine hours integration when the number of ensembles is increased to 64, 128, 256, 512, and 1024. It was found that the execution time was almost constant as the number of ensembles increased.

Figure 1 left shows the difference of middle layer air temperature after three days integration between ERA5 and the standalone high-resolution simulation, and Figure 1 right is the difference between ERA5 and the ensemble coupling. It was cleared that the ensemble coupling yields results more closely aligned with the reanalysis values. In this experiment, we conducted a three-day integration period. However, to substantiate the advantages of ensemble coupling, more detailed meteorological analysis is required. This necessitates extending the integration period to several months.

Tab.1 Execution time for each ensemble size [13]

ensemble size	64	128	256	512	1024
time(sec)	346.3	344.9	357.5	346.5	352.5

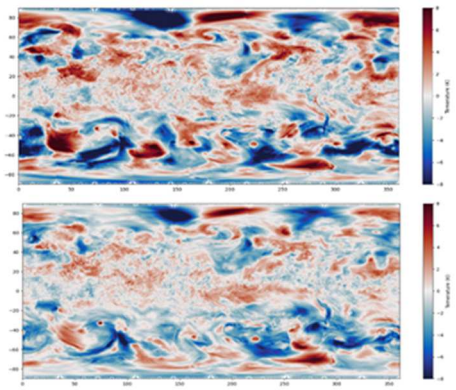

Fig.1 Difference from ERA5 data (upper: standalone, lower: ensemble coupling) [13]

(a-2) Earthquake Simulation by PSD with Machine Learning through Causality [7,8]

This part investigates the integration of high-performance computing (HPC) and artificial intelligence (AI) techniques to enhance large-scale earthquake simulations based on Finite Element Methods (FEM). These simulations reproduce the spatial and temporal evolution of seismic events to generate synthetic time series at multiple spatial locations. The robustness of these simulations can be enhanced through data assimilation, which includes the external measurement data within the simulations. Traditional assimilation techniques often rely on correlation-based methods, which may misrepresent the directional structure of interactions between variables.

To address this limitation, the project focuses on coupling FEM with causal inference methods. The objective is to improve earthquake simulations by identifying directional relationships between sensor locations and using this structure to guide the assimilation of new data. Causal discovery methods are applied and adapted to extract graphs that capture the influence of one signal on another across time and space to follow propagation patterns that reflect the internal structure of the system.

A proof of concept developed in FY.2024 demonstrated the feasibility of modeling the sensor network using causal graphs derived from FEM outputs. Constraint-based algorithms such as PC iteratively test for (conditional) independencies to prune edges from a fully connected graph. According to [Kalisch et al. 2007], the worst-case computational complexity of the PC algorithm is bounded by $O(p^q)$, where p is the number of variables (sensors and lagged sensors) and q is the maximum size of the neighborhoods encountered during the search. Furthermore, to ensure theoretical guarantees of the approach, we explore kernel-based independence tests [Zhang et al. 2012], which make fewer assumptions about the data distribution and allow for the detection of non-linear relationships. The work in FY.2024 focused on developing methods to model the causal link between time series simulated by FEM. These time series represent surface sensors--data assimilation links--used to measure the seismic

response of the system.

Causal discovery aims to identify cause-effect relationships between variables based on observational or simulated data. Unlike correlation-based approaches, it seeks to uncover the underlying structure that governs how changes in one variable influence others. In time series settings, causal discovery methods try to determine whether past values of one signal contribute to predicting future values of another, often relying on statistical dependencies and conditional independence tests. These conditional independence tests for time series aim to assess whether the future of one variable remains independent of another, given a set of conditioning variables—typically, the past of both. In linear settings, the Partial Correlation (ParCorr) test is widely used due to its simplicity and efficiency; it assumes Gaussianity and captures linear dependencies. In contrast, the CMIknn test estimates conditional mutual information using nearest neighbors, allowing the detection of nonlinear relationships without strong distributional assumptions. While ParCorr is appropriate when linearity is a valid assumption, CMIknn is preferred for capturing more complex, potentially nonlinear interactions that often arise in simulated or real-world dynamical systems. This makes CMIknn particularly relevant for modeling causal interactions in seismological time series. However, this improved expressiveness comes at a higher computational cost.

Another factor that contributes to the computational cost is the choice of the causal discovery framework, such as PCMCI, PCMCI+, or LPCMCI. These methods combine conditional independence tests with a two-stage selection procedure to infer causal graphs from time series data. PCMCI (Peter and Clark Momentary Conditional Independence) applies a preliminary variable selection step (PC1) followed by momentary conditional independence tests (MCI) to limit the conditioning set, improving both robustness and scalability. PCMCI+ extends this by performing additional iterations to better resolve ambiguities and detect hidden confounding, at the expense of increased computational time. LPCMCI (Lag-PCMCI) further refines the approach by explicitly controlling for lag-specific effects in the conditioning sets, which enhances interpretability in systems with known temporal structure. While these refinements improve the accuracy and interpretability of the inferred causal graph, they also increase the overall computational burden as illustrated in Fig.2. This figure shows computational times for PCMCI, PCMCI+, and LPCMCI when inferring causal links in simulated earthquake data, using 800 time steps for 8 sensors, represented in Fig.3. The figure highlights how runtime increases with the maximum allowed time lag and especially for LPCMCI.

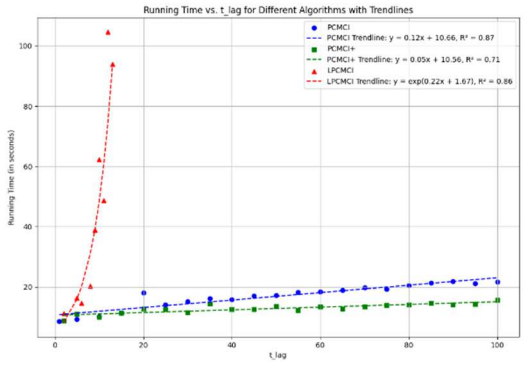

Fig.2 Running time for PCMI, PCMCI+, and LPCMCI following time lag value for 2 simulated earthquakes

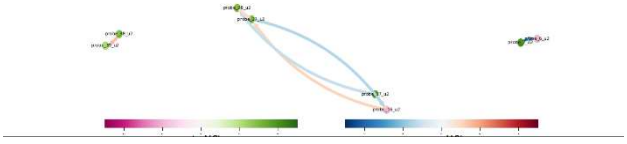

Fig.3 Causal graph obtained with LPCMCI on 2 earthquakes

As mentioned, enhancing the FEM solver with data assimilation remains a key objective of this project. In this context, PSD (a massively parallel (MPI-based) FEM earthquake solver developed at CEA) was first tuned and optimized on the Wisteria/BDEC-01 system. For FY.2024, strong scaling tests of PSD using up to 72 nodes were conducted on the Odyssey nodes, which hosts the FUJITSU PRIMEHPC FX1000. A large-scale earthquake simulation problem with 170 million degrees of freedom was selected. Post tuning, quasi-optimal scaling was achieved for the major FEM phases, including matrix assembly and PETSc operations. This milestone confirms that PSD is now ready to be extended with causality-driven data assimilation techniques.

Given the observed computation times (Fig.2), future work in FY.2025 will aim to reduce the overall runtime of causal discovery methods. Although the conclusions of the tests remain consistent and the differences in p-values are minimal, the high computational demand highlights the need for optimization. One promising direction is to refine the selection of conditioning sets by focusing on informative and computationally efficient subsets. This strategy could reduce the number of required tests without degrading the reliability of causal inference, thereby making large-scale applications more tractable. Hence, future developments will also explore alternative conditional independence tests better suited for specific earthquake simulations, as well as improve code portability and scalability on high-performance computing (HPC) architectures.

(a-3) Construction of Sophisticated 3D Underground Model [1,2,5,10]

This study investigates a method for estimating subsurface structures by data assimilation from a theoretical seismic motion program in one-dimensional (vertical) models and actual observed data. The theoretical seismic motion is calculated with parameters related to the shape of the subsurface structure, seismic wave propagation velocity, and the earthquake source. Suppose parameters can be derived such that the theoretical seismic motion is similar to the observed motion, those parameters can be assumed to represent the shape and characteristics of the subsurface structure. A key advantage of using theoretical seismic motion is that it allows for seismic motion to be calculated with a relatively small number of parameters, making it suitable for estimation via data assimilation.

Up to the early part of this fiscal year, parameter estimation was conducted using the Replica Exchange Monte Carlo (REMC) method. However, due to the strong nonlinearity in the theoretical seismic motion calculations, we could not assimilate parameters effectively. The computation of the theoretical seismic motion is obtained by integrating Green's function in the frequency domain. Then, due to divergences near singularities in the complex number domain, a mechanism for switching calculation methods using conditional branching has been implemented to suppress divergence. As a result, the behavior of all parameters changes nonlinearly in response to modifications in specific parameters. Since REMC is a Bayesian estimation method that assumes Gaussian distributions for output variations in response to inputs, it is presumed that REMC is not effective for such nonlinear problems. For parameter

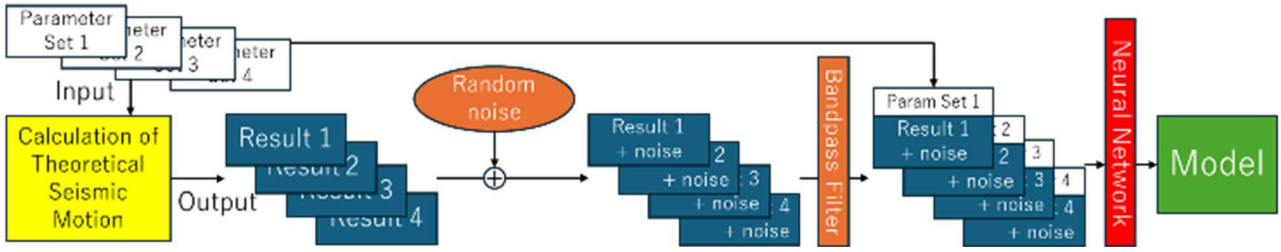

Fig4. Model creation process for underground structure estimation

optimization problems with nonlinear characteristics, a genetic algorithm (GA) approach is considered effective. However, GA typically requires more trial iterations than REMC. Given that the current theoretical seismic motion calculation takes about two minutes per case, completing evaluations for approximately 300 observation points within a practical execution time remains difficult.

Therefore, the current focus is on the use of machine learning models. A sufficient number of output waveforms for arbitrary parameters are prepared using theoretical seismic motion calculations, and a training dataset with added observation noise is created (Fig.4). Then, we have considered two machine learning strategies: one in which the input is a seismic waveform and the output is the theoretical seismic motion parameters, and another in which the input is parameters and the output is the waveform. Both approaches are currently under investigation.

(a-4) TSMP [3,6]

The Terrestrial System Modeling Platform (TSMP) is an integrated soil-vegetation-atmosphere modeling system. TSMP currently comprises three component models: The Consortium for Small-scale Modeling (COSMO) atmospheric model, the Community Land Model (CLM), and the hydrologic model ParFlow. These models are coupled with OASIS3-MCT coupler. Our objective is to replace the OASIS3-MCT coupler with UTIL/MP and evaluate its performance on heterogeneous environments. As a preliminary investigation, we developed a Toy program (ToySMP) that emulates each component of TSMP. These components were coupled using UTIL/MP, and their performance was evaluated. The execution pattern of the components and the data flow are also set to be similar to TSMP. Two components are executed on CPU nodes (cluster module in DEEP, and Odyssey in Wisteria/BDEC-01), the one is executed on GPU nodes (booster module in DEEP, and Aquarius in Wisteria/BDEC-01). In this experiment, the grid size was set to 1024 × 1024, and the number of exchange fields was set to 10. The number of processes assigned to each component varied from 1 to 128. The results on DEEP in Figure 2 left side demonstrate shorter execution times and better scalability compared to those of Wisteria/BDEC-01. UTIL/MP is equipped with the capability to aggregate multiple data transfers into a single transmission. We configured UTIL/MP to aggregate 10 field data sets and exchanges all at once. According to the measurement, performance scaled effectively up to 128 processes on Wisteria/BDEC-01 with Hybrid Mode .

(b) Life Science

(b-1) Brain Aneurysm Simulations [3,6]

In this part, we developed mAIA for MSA (Modular System Architecture) simulation using the WaitIO library on both systems (DEEP and Wisteria/BDEC-01), and compare the

performance. mAIA consists of multiple solver modules and coupler modules, and users can choose the solver and coupler according to their purpose. In FY.2024, we combined Lattice Boltzmann solver and Post-Processing Module. Fig.5 shows the performance on Odyssey and DEEP system. Performance of the applications with WaitIO were competitive with those with native MPI.

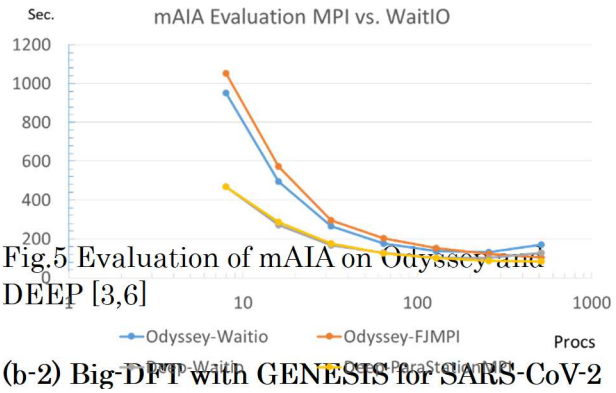

Fig.5 Evaluation of mAIA on Odyssey and DEEP [3,6]

(b-2) Big-DFT with GENESIS for SARS-CoV-2 Main Protease [34,35]

One of the challenges of developing medicines for viruses is drug-resistance. Though a proposed drug may effectively inhibit key viral proteins, such inhibition does not last as evolutionary pressure leads to new variants. Thus, essential medicines like Darunavir, which can be used to treat HIV/AIDS, are at risk of becoming ineffective. It is thus critical to develop methods that can predict future variants of concern, and design rules for new medicines to respond to those variants. However, as even mutations far from the active site – so called "distal mutations" – can have an impact on binding, sophisticated methods that reveal information about the entire protease-drug complex in its many conformations is necessary .

Recently, we have been implementing linear scaling algorithms into to the density functional theory (DFT) based BigDFT code that enable it to easily and robustly treat systems of the size of the HIV-1 main protease (3,000-4,000 atoms). Simultaneously, the Genesis classical molecular dynamics code has been optimized for GPU architectures. By combining these codes, we might hope to realize the in-depth electronic structure picture provided by DFT calculations, while suitably sampling the conformer space, using a combined quantum mechanics + molecular mechanics workflow.

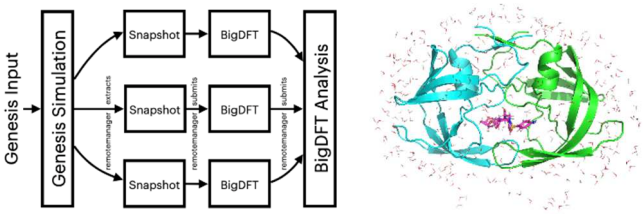

Fig.6 Combined BigDFT and Genesis workflow (left) used for the studying of Darunavir bound to HIV-1 main protease (right). Both codes simulate systems including explicit water molecules, with BigDFT only considering water molecules within 3.0 Å of any atom in the protein / ligand complex.

We installed the BigDFT and Genesis codes on to the Wisteria system, with BigDFT targeting the CPU and Genesis the GPU nodes (Figure 1). We then built a workflow using our remote manager python library [6] to deploy the codes together. As the Genesis calculation runs a long molecular dynamics trajectory over several days, a local daemon watches the output for new samples to be processed by BigDFT. As soon as a new sample is available, a corresponding BigDFT calculation is generated and automatically submitted. The daemon also watches for completed BigDFT runs, which may then be postprocessed to extract the quantum observables needed for studying the binding network of the complex. This daemon approach ensured a low overall time to solution, as BigDFT calculations can begin without having

to wait for the sampling of Genesis to be completed. In practice, we found that Genesis calculations should be performed using one full node of Aquarius, BigDFT calculations with 64 nodes of Odessey, and post processing with a single node of Odessey. We deployed this workflow simultaneously on three different mutants – the wildtype (6OPS), 2 mutation variant (6OPT), and 11 mutation variant (6OPZ) – to understand the effect of mutations both near and far from the binding site.

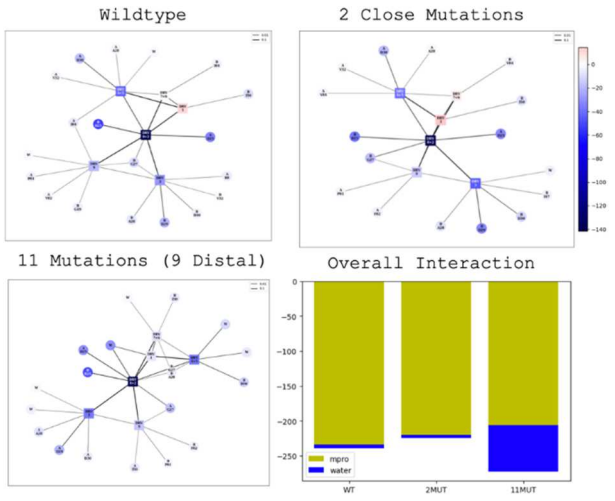

Fig.7 Interaction graphs and overall binding of Darunavir to different variants of the HIV-1 main protease. The interaction network reveals a weakening of the binding as the number of mutations increase. Interaction descriptors are in the unit of kcal / mol.

In Fig. 7, we show interactions networks between Darunavir and the three different HIV-1 main protease variants. When we calculate the total interaction energy between Darunavir and just the protease (i.e. projecting out the interactions with water), our results qualitatively agree with the trend seen in experiment. The interaction network that includes distal mutations show a number of water molecules interacting with the drug. Thus, our preliminary conclusion is that in addition to the binding site mutations reducing affinity for Darunavir, the distal mutations increase solvent accessibility, leading to substantially reduced inhibitory power. Our analysis so far has been limited to these three variants and to a limited set of conformers. Our future work will run the full set of experimentally characterized variants, using the same workflow, to further validate our results. We will also refine our analysis framework for representing the variability found in interaction strengths and network topologies over the course of the Genesis simulation. Based on this extended data, we hope to identify future potential mutants or drug design strategies for targeting the HIV-1 main protease.

(c) Library/Software/Tool for Integration of "S+D+L"

(c-1) ChASE [9]

The Chebyshev Accelerated Subspace Eigensolver (ChASE) presents a scalable alternative by incorporating iterative methods with Chebyshev polynomial filters. However, previous implementations of ChASE faced challenges in multi-GPU environments, such as redundant computations, inefficient communication patterns, and excessive memory consumption. To address these issues, this research introduces significant refinements to ChASE, optimizing its parallelization scheme and computational efficiency. Key enhancements are, (1)Optimized Parallelization Scheme, (2)Communication-Avoiding QR Factorization, and (3) GPU-Native Collective Communication.

The ChASE (https://github.com/ChASE-library/ChASE) codebase has undergone a major restructuring to improve modularity and

separation of concerns. Algorithmic logic is now cleanly separated from hardware-specific backends (CPU, GPU, MPI), enabling easier testing, maintenance, and extensibility. A full suite of unit tests was developed to validate all numerical kernels across supported architectures and data types. ChASE has been successfully integrated into the YAMBO codebase (https://www.yambo-code.eu/) as a high-performance Hermitian eigensolver. This enables improved scalability and performance for solving eigenvalue problems in many-body perturbation theory, with initial validation showing both numerical accuracy and runtime gains compared to existing solvers. ChASE has been extended to support full Bethe–Salpeter Equation (BSE) matrices, which are quasi-Hermitian. This extension enables broader application of ChASE to excited-state calculations.

As part of ongoing performance exploration on GPU architectures, we tested the Ozaki-scheme-based GEMM implementation (https://github.com/enp1s0/ozIMMU) that leverages integer cores on NVIDIA GPUs. The goal was to assess its numerical behavior and potential integration with ChASE GPU builds. These tests are exploratory and aimed at identifying potential future directions for integrating low-level optimized GEMM routines using NVIDIA GPU integer cores in ChASE.

(c-2) h3-Open-BDEC [4,11,12]

h3-Open-BDEC was originally developed for integration of heterogeneous node groups in systems, such as Wisteria/BDEC-01. Currently, we are extending it for the software in QC-HPC Hybrid Environment, where QC's (Quantum Computers) and Supercomputers (HPC's) are utilized simultaneously. WaitIO-Router is a coupling communication software for efficiently implementing and integrating communication and data transfer in QC-HPC Hybrid systems in real time (Fig.8). The ultimate goal is to quickly and safely connect multiple HPC/QC centers connected by SINET-6. There are three challenges to realizing WaitIO-Router:

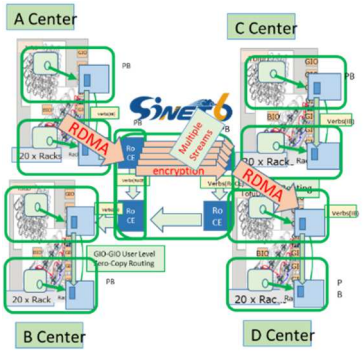

Fig. 8 WaitIO-Router Challenges and Solution

(1) Accelerating conventional WaitIO-Socket (File) communication
(2) Accelerating long-distance communication between multiple centers and strengthening security
(3) Providing seamless communication function between QC and HPC

In FY.2024, we have completed the development of WaitIO-Verbs and WaitIO-Tofu, which realize RDMA communication from WaitIO-Socket in (1)), and WaitIO4Py in (3)). The communication performance of WaitIO-Verbs and WaitIO-Tofu achieved 1/10 communication latency, 6 times the communication bandwidth (WaitIO-Verbs), and 29 times the communication bandwidth (WaitIO-Tofu).(Fig. 9).

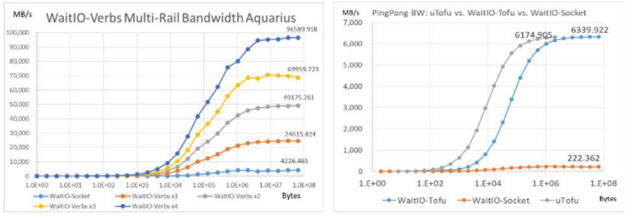

Fig.9 WaitIO-Verbs and WaitIO-Tofu Performance

(c-3) Web-based Framework [57,86,87]

Web-based interactive system for real-time earthquake simulation on Wisteria/BDEC-01 developed in this project has been utilized for education and outreaching, where web-server is on mdx. In FY.2025, we will extend this system for web-based operations of various applications for integration of (S+D+L) on heterogeneous systems, such as Wisteria/BDEC-01 and Miyabi. The first target is QC-HPC Hybrid Environment. Introduction of Open OnDemand (https://openondemand.org/) to such applications will be also investigated continuously.

6. Self-review of Current Progress and Future Prospects

This is an international collaboration with 9 topics by more than 50 researchers from 17 organizations of 4 countries. Since FY.2021, we have been continuously holding monthly online meetings with domestic researchers. We also held online meetings with overseas groups every 3-4 months, and held one or two in-person meetings during international conferences in FY.2024. The most important principle in collaborative research is continuous communication. By faithfully adhering to this principle, we have been able to achieve significant results in FY.2024 that far exceeded our initial expectations. We published international joint papers [3,6], and a project poster on this JHPCN collaboration will be presented in ISC-HPC 2025 in Hamburg, Germany [16]. Generally, our activities have been highly recognized internationally, and we had more than 10 invited talks in FY.2024.

Since November 2023, h3-Open-BDEC has been utilized as software for constructing QC-HPC Hybrid environment in JHPC-quantum project (https://jhpc-quantum.org/), as was mentioned in (c-2) of the previous section [14,15].

In FY.2023 and FY.2024, we have conducted research in earth science, life science, and library/software/tool for integration of "S+D+L". In FY.2025, we add a new field (quantum sciences), and extend the idea of h3-Open-BDEC for QC-HPC Hybrid Environment, while we have already started preliminary works in FY.2024.

We are also passionately promoting the idea of integration of "S+D+L". In addition to the QC-HPC Hybrid tutorial, we regularly conduct h3-Open-BDEC tutorials. We have demonstrated earthquake simulations at the EU-ASEAN/ACM-ASEAN HPC School since 2022, and are invited to provide a 2-Day tutorial at NCTS (National Center for Theoretical Sciences, Taiwan) in February 2025. List of URLs for such activities is as follows:

- QC-HPC Hybrid Tutorial: https://www.cc.u-tokyo.ac.jp/events/lectures/238/
- h3-Open-BDEC Tutorial: https://www.cc.u-tokyo.ac.jp/events/lectures/236/
- 2025 EU-ASEAN School: https://europe.acm.org/seasonal-schools/asean/2025
- NCTS Winter School 2025: https://ncts.ntu.edu.tw/upload/events_3_33624121207012494.pdf

#Section 7, achievements of this FY, should be input on the JHPCN website.