AI for Science on Modular Supercomputers

Performance Evaluation of Coupled Atmospheric–AI Workflows

Takashi ARAKAWA[1,3*] (arakawa@climtech.jp), Hisashi YASHIRO[2], Shinji SUMIMOTO[1], Kazuya YAMAZAKI[1], Kengo NAKAJIMA[1]
1.The University of Tokyo, Japan, 2.National Institute for Environmental Studies, Japan, 3.CliMTech Inc., Japan

1 Background/Objectives

- **AI surrogates** are increasingly used to accelerate physical processes in Earth system (climate) models
- Atmospheric simulations still rely mainly on **CPU-based models**, requiring efficient coordination with **GPU-based AI**
- **Modular supercomputers** enable heterogeneous CPU–GPU execution, but practical coupling workflows remain underexplored

- Develop a **loosely coupled Atmospheric Model(NICAM)–AI workflow** for AI-for-Science applications
- Demonstrate **asynchronous CPU–GPU coupling** using h3-Open-**UTIL/MP** and h3-Open-SYS/**WaitIO** on modular systems
- Evaluate execution performance and resource balance on Wisteria/BDEC-01 and Miyabi

2 Methods

2.1 WaitIO and UTIL/MP

- **WaitIO**[1] is a **communication library** that enables asynchronous CPU–GPU data transfer across heterogeneous nodes
- **UTIL/MP**[2] is a **coupling library** that manages data exchange and grid remapping between NICAM and the AI module

2.2 NICAM

- **NICAM**[3] is a **global nonhydrostatic** atmospheric model that covers the whole globe with a uniform grid

2.3 NICE

- **NICE**(NICAM Cloud Emulator) is an AI-surrogate model that replaces the cloud microphysics scheme in NICAM
- The AI model consists of a **three-layer multilayer perceptron (MLP)** implemented using **PyTorch**

2.4 Workflow

1. **Data extraction**: Eight microphysics input variables and eight output tendencies are passed from NICAM to UTIL/MP
2. **Communication**: Data are asynchronously transferred to the AI module via WaitIO
3. **Grid remapping**: NICAM's 14 km fields are internally remapped by UTIL/MP to the 224 km grid used for AI training
4. **Layer-wise training**: The AI model is trained to reproduce output tendencies from input variables for each vertical layer

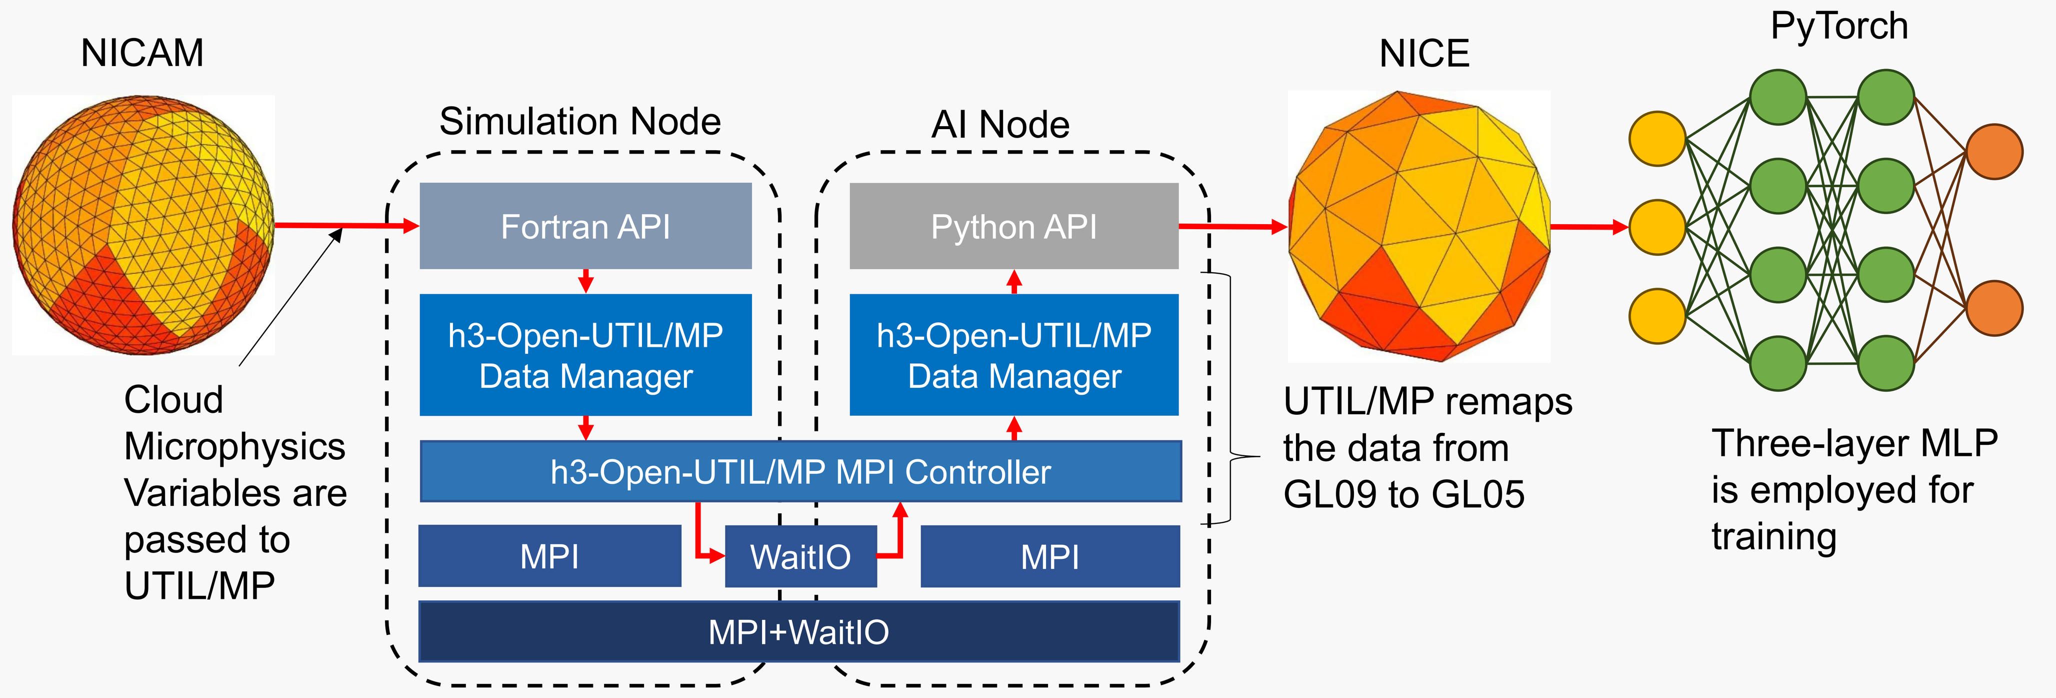

Fig.1 Data flow of NICAM-AI coupling

3 Results

3.1 Reproducibility

- Correlation
  - The AI model was trained using data at **5–6 h** and evaluated by reproducing the fields at **7 h**
  - For three representative variables, the correlation **coefficients exceed 0.92**
  - The regression slopes range from **0.79 to 0.89**, indicating good reproducibility

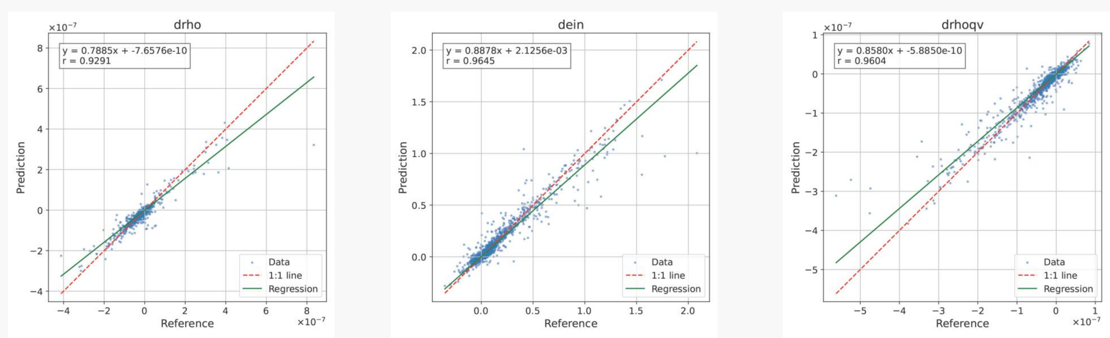

Fig.2 Correlation of simulation vs prediction

- Horizontal fields
  - The model trained on **two days of data** was applied to the **5th day** of the simulation
  - Results are shown for **total air density (drho)** at **5,500 m in the mid-troposphere**
  - The AI prediction shows **excellent agreement** with the simulated field

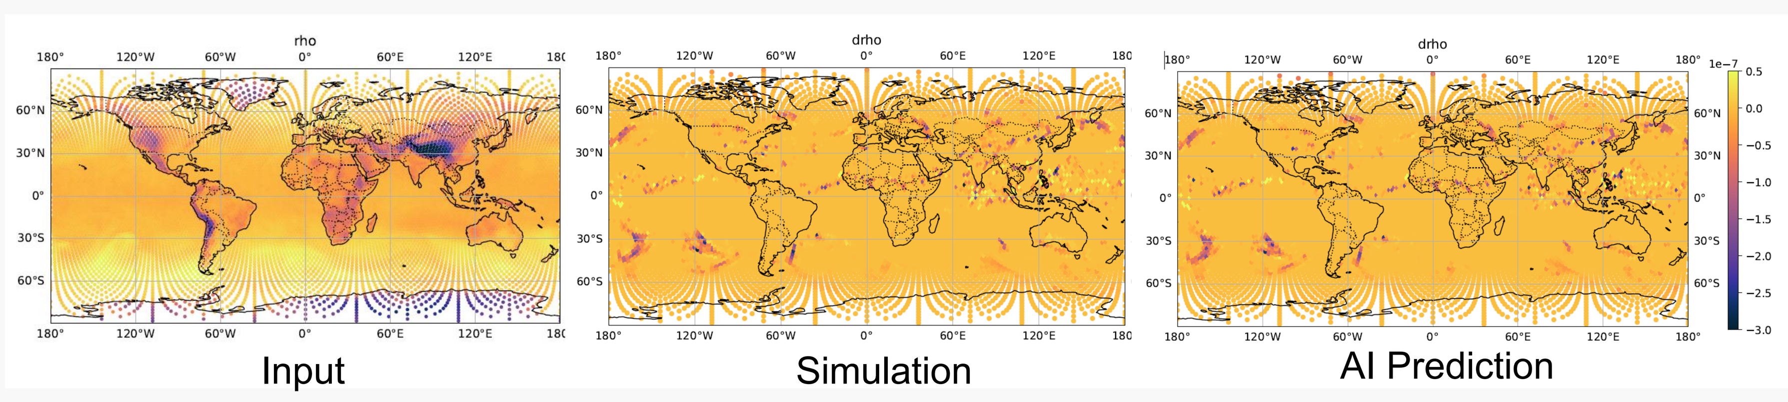

Fig.3 Horizontal field of simulation and AI prediction

3.2 Computational Performance

Table 1 Experimental setting

Model	System	Processes	Nodes	Threads
NICAM	Odyssey	640	160	12
	Miyabi-C	640	32	5
NICE	Aquarius	20	3	None
	Miyabi-G	20	20	None

- Experimental setting
  - NICAM: **640**, AI: **20** processes
  - **Wisteria/BDEC-01:** NICAM **160** and AI **3 nodes**
  - **Miyabi:** NICAM **32** and AI **20 nodes.**
- Overall Execution Time
  - On **Wisteria/BDEC-01**, Dynamics and Physics are **4.4×** and **3.1× faster** than on **Miyabi**, due to larger node counts and the **A64FX high memory bandwidth**
  - On Wisteria, **NICAM outpaces AI**, causing **idle waiting**
  - On Miyabi, **slower NICAM dominates runtime**

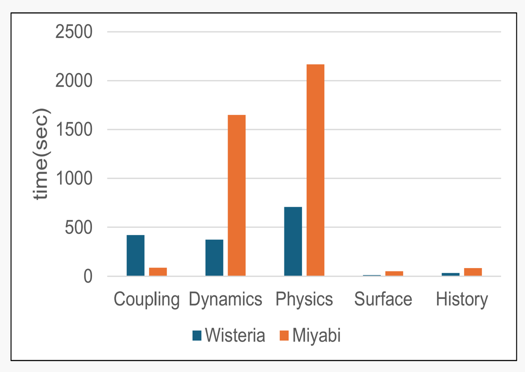

Fig.4 Execution time

- Training Time
  - Training performance is **~2×faster on Miyabi (H100) than on Wisteria (A100)** and is **batch-size–dependent**, indicating a **memory- and data-transfer–bound** regime rather than compute-bound execution

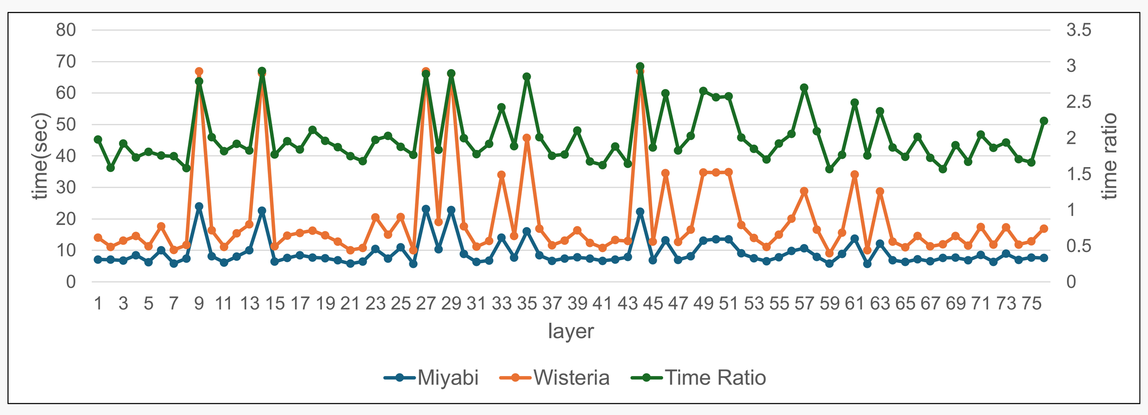

Fig.5 Training time and time ratio

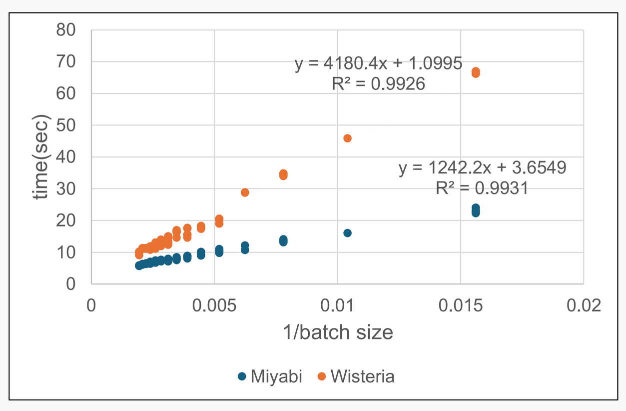

Fig.6 Time vs batch size

4 Conclusion

- We demonstrated a **practical NICAM–AI workflow** using asynchronous heterogeneous coupling on modular supercomputers
- Performance is strongly affected by **resource balance**: faster NICAM causes idle waiting, while faster AI shifts the bottleneck to NICAM execution
- Since training time depends on **batch size**, selecting **appropriate training parameters with training cost in mind** is essential

Reference

1. Sumimoto et al., A System-Wide Communication to Couple Multiple MPI Programs for Heterogeneous Computing, PDCAT 2022, 2022, 10.1007/978-3-031-29927-8_25
2. Arakawa et al., Development of a coupler h3-Open-UTIL/MP, HPC Asia 2022, 2022, 10.1145/3492805.3492809
3. Satoh et al., The Non-hydrostatic Icosahedral Atmospheric Model: description and development, PEPS, 2014, 10.1186/s40645-014-0018-1