jh250055

ReplicaEUJP: Research on the replicability of climate simulations across different computing environments

Hisashi YASHIRO (National Institute for Environmental Sciences)

Earth System Models (ESMs) are an indispensable tool to project the next 50 to 100 years of our climate, employing greenhouse gas (GHG) emission and anthropogenic aerosol-based scenarios. Reproducibility is important in this collaborative effort, tracing back simulations to specific configurations, model versions, and compilation flags to reproduce the same simulation in the same environment again, achieving identical results. Equally important is replicability, achieving "identical" results when performing the same experiment using different clusters, computing environments, or compilers. Achieving replicable results is much more challenging, and in practice, bit-to-bit replicability can rarely be achieved. However, results can be replicable in the sense that one simulated climate is statistically indistinguishable compared to another. This becomes important when comparing simulations with the same model conducted in different computing environments. If replicability does not hold, differences between contrasting scenarios performed on different clusters cannot be interpreted exclusively in terms of the changes in external forcing. In ReplicaEUJP, we will conduct multi-member multi-decadal climate simulations on Miyabi-C while conducting the same set of simulations on Fugaku and Marenostrum5 in separate computing projects and investigate the differences in their climates. The study will be based on our recently developed replicability methodology, and the climate model for conducting the simulations will be NICAM coupled to COCO.

1. Basic Information

(1) Collaborating JHPCN Centers

The University of Tokyo

(2) Theme Area

Data science/data usage area

(3) Project Members and Their Roles

Hisashi Yashiro (Representative), National Institute for Environmental Studies (NIES): Project Supervision

Mario Acosta (Deputy Representative), Barcelona Supercomputing Center (BSC): Workflow system development

Kai Keller, Barcelona Supercomputing Center (BSC): Workflow system development

Leo Arriola, Barcelona Supercomputing Center (BSC): Workflow system development

Erick Fernando Doria, Barcelona Supercomputing Center (BSC): Workflow system development

Bruno de Paula, Barcelona Supercomputing Center (BSC): Workflow system development

Marta Alerany, Barcelona Supercomputing Center (BSC): Climate model replicability assessment

Kodama Chihiro, Japan Agency for Marine-Earth Science and Technology (JAMSTEC): Climate model replicability assessment

Mohamed Attia Wahib, RIKEN Center for Computational Science (R-CCS): Workflow system development

Kengo Nakajima, Information Technology Center, The University of Tokyo: Workflow system development

Kazuya Yamazaki, Information Technology Center, The University of Tokyo: Workflow system development

2. Purpose and Significance of the Research

The sixth assessment report (AR6) issued by the Intergovernmental Panel on Climate

Change (IPCC) projects that, in a 1.5-degree warmer climate compared to the pre-industrial period, heat waves become about 8 times more frequent every 50 years, and extreme precipitation events become twice as frequent every 10 years. We need to prepare and adapt to those changes in our global climate. The current best way to do this is to use Earth System Models (ESMs) to project our climate over the next 50 to 100 years, employing greenhouse gas emission and anthropogenic aerosol scenarios. The most prominent initiative dedicated to this aim is the Coupled Model Intercomparison Project (CMIP). Reproducibility is important in this collaborative effort, tracing simulations back to specific configurations, model versions, and compilation flags to reproduce the same simulation in the same environment and achieve identical results. Equally important is replicability: achieving "identical" results when running the same experimental configuration across different clusters, computing environments, or compilers. Achieving replicable results is much more difficult; in practice, bit-to-bit replicability is almost never achieved. However, results can be replicable in the sense that the model's climate in one computing environment is statistically indistinguishable from that in another. The purpose of this research project is twofold. On the one hand, we are contributing to an important research question in computational climate science: Are climate projections replicable across different computing environments, and how can we test whether they are? On the other hand, the project contributes to coordinated collaboration between European and Japanese climate science research. The latter is important, as climate research is often driven by regional interests. For instance, in Japanese climate models, it is important that events in the Pacific region, such as tropical cyclones, are well represented. European models might emphasize other aspects. Bringing together the expertise will improve the overall quality and coverage of our climate models.

3. Significance as JHPCN Joint Research Project

In 2022, the European Commission launched digital partnership initiatives to foster international collaboration. One of the first digital partnerships has been established between Europe and Japan. The Hpc AlliaNce for Applications and supercoMputing Innovation (HANAMI) is the implementation of this digital partnership. European centers receive funding for travels, workshops, and research visits in Japan to strengthen the European-Japanese collaboration. The JHPCN joint research project represents an excellent opportunity to collaborate on a project with shared funding.

4. Outline of Research Achievements until FY2024

This project is not a continuous project.

5. Details of FY2025 Research Achievements

In ReplicaEUJP, we targeted to deploy the atmosphere and ocean fully coupled climate model NICOCO on Miyabi-C. The atmospheric component of NICOCO is the Non-hydrostatic Icosahedral Atmospheric Model (NICAM) [1], and the ocean model is COCO [2]. The two model components are executed with the coupling library Jcup [3], exchanging the necessary data with grid remapping. We aimed to perform multi-member, multi-decadal climate simulations with the model on Miyabi-C using the automatic workflow manager Autosubmit (https://autosubmit.readthedocs.io/).

Aligned with the experiments on Miyabi-C, we perform multi-member, multi-decadal climate simulations on the other two supercomputers, Fugaku in Japan and Marenostrum5 in Spain. The initial conditions, scientific configuration, and boundary conditions are identical across the three simulations on Miyabi-C, Fugaku, and Marenostrum5. Based on our recently developed replicability methodology [4], we will identify differences arising from simulation across the three architectures and software environments.

To enable compilation, a build wrapper was drafted to configure the required environment. This build wrapper relies on the Lmod module system to load NICOCO dependencies: NetCDF (for C and Fortran), HDF5, and the Intel-MPI compiler. It defines compilation settings and specifies the underlying NICAM and COCO Makedef files for the HPC platform. In these system definition files, the compiler commands, flags, and architecture options are set. The NICOCO model was successfully compiled and executed on the Miyabi-C supercomputer, and its output was verified against the provided references to confirm correctness.

The HANAMI workflow (Figure 1) is based on Autosubmit, which enables automated management of climate simulations on HPC systems. Autosubmit combines experiment management with HPC job orchestration and is already fully operational for NICOCO on Marenostrum5 and Fugaku. At the start of this project, the Autosubmit workflow manager did not support Portable Batch System (PBS). Furthermore, it did not support Miyabi's two-factor authentication (2FA) used for SSH connections with one-time passwords (OTP). We carried out additional development to address these issues and completed the porting to Miyabi. An automated workflow was developed to manage the full NICOCO simulation lifecycle, from compilation and configuration through job submission, monitoring, and output handling, improving reproducibility and reducing manual effort. This workflow was then extended to include ensemble run capabilities, enabling the parallel management of multiple simulations with different initial conditions for the atmosphere and ocean. Figure 2 is an example of the GUI monitor of ensemble job management.

To conduct appropriate ensemble climate simulations, we need to prepare initial states with a balanced energy distribution between the atmosphere and the ocean and sufficient variability. This is especially important for the

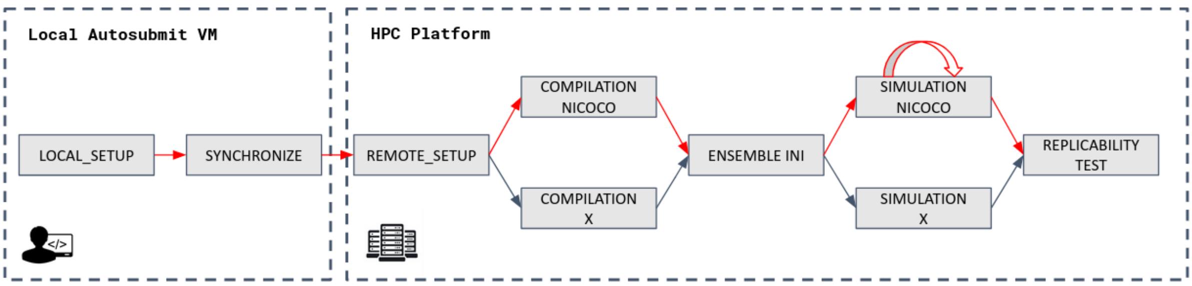

Figure 1. Structure of the jobs that compose the HANAMI workflow and their dependencies, including the platforms where these jobs run.

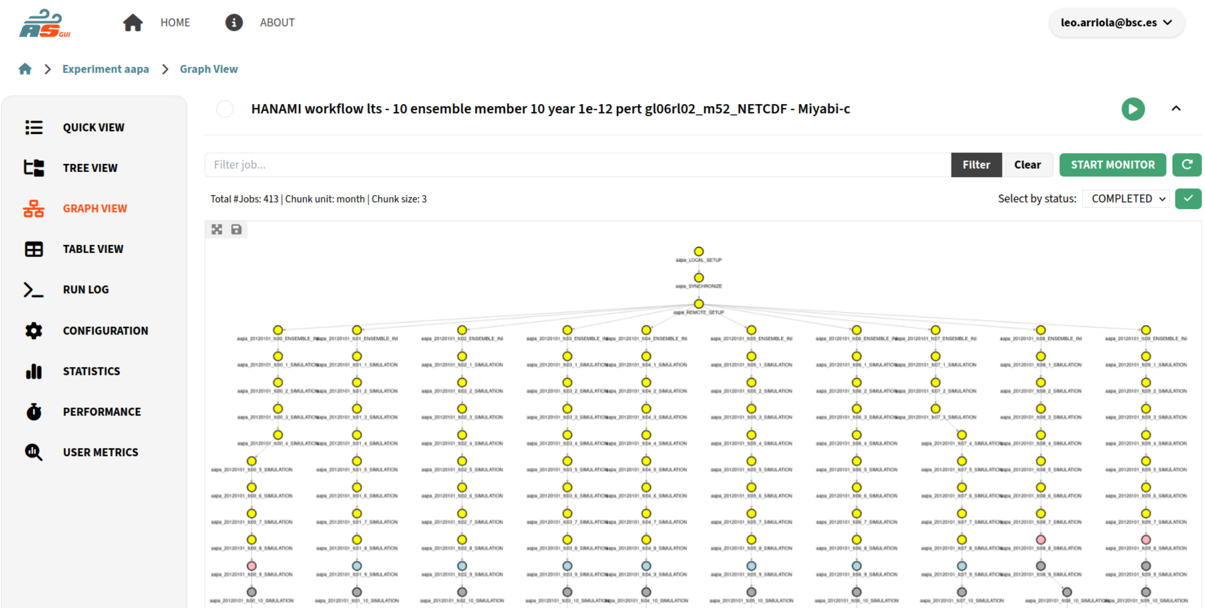

Figure 2. GUI window of Autsubmit workflow manager for the graph view of NICOCO ensemble simulation.

ocean's initial values. The atmosphere is a "fast process" and self-regulates within a few days, even when disturbances are introduced as random noise. On the other hand, the ocean requires several years or more of spin-up to reach a balanced state following the introduction of noise. To solve this problem, we created appropriate initial values from long-term, single-member coupled atmosphere-ocean simulations. These initial values were used consistently in experiments across all supercomputers. We conducted a 150-year NICOCO simulation. The initial state ensemble was generated by sampling 10 ocean states from the long control simulations, spaced 10 years apart, to capture different ocean states (macro perturbation). Additionally, we generated 100 perturbed atmospheric states, adding Gaussian noise to the atmospheric temperature (micro perturbation).

Three two-year test integrations were carried out using the standard 1-degree horizontal-resolution NICOCO configuration, initialized from the same ocean state and three different atmospheric states. Specifically, we performed an unperturbed control simulation, one applying a Gaussian perturbation of 0.001 K, and a third applying a Gaussian perturbation of 1e-12 K to the atmospheric temperature. These runs assessed the climate model's sensitivity to perturbations in the atmospheric initial state. The simulations further confirmed the numerical stability over the integration period and validated the ensemble workflow under realistic conditions.

Finally, a 10-member ensemble NICOCO simulation, configured with a 112-km-mesh atmosphere and a 1-degree-grid ocean, was run, with approximately 2 years per member. Each member was initialized from a different combination of oceanic and atmospheric states, sampling initial-condition uncertainty across both components simultaneously. All members were submitted and managed via the ensemble-aware Autosubmit workflow on Miyabi under the PBS scheduler. Output was stored in NetCDF format and archived for downstream analysis. Figure 3 shows the statistics window

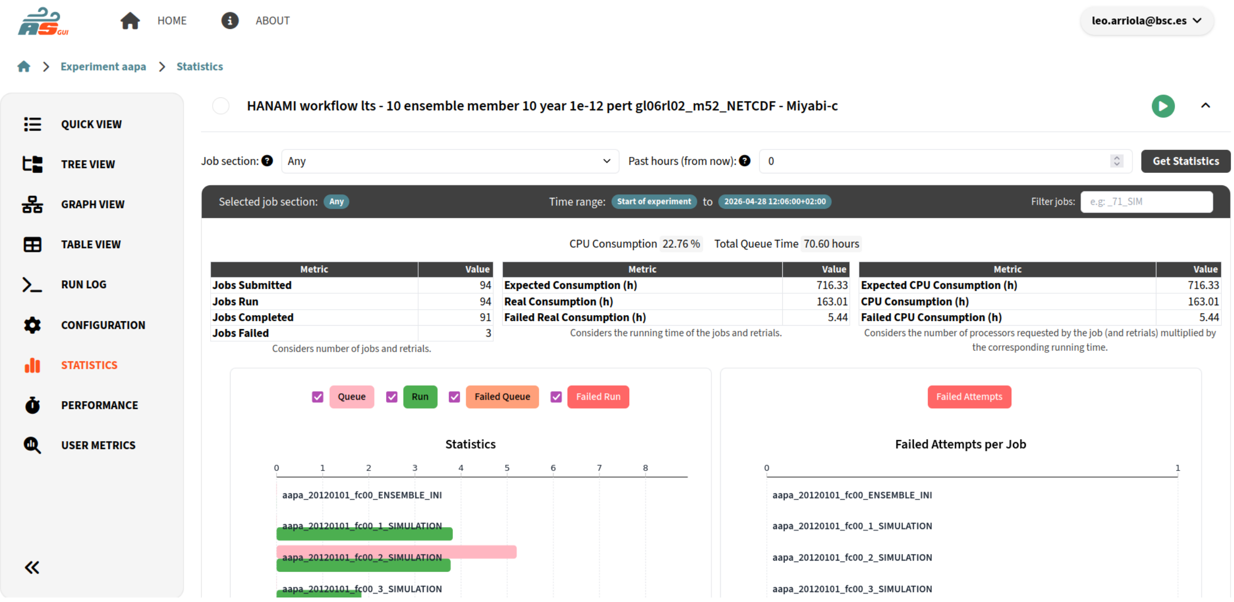

Figure 3. GUI window of Autsubmit workflow manager for the statistics view of NICOCO ensemble simulation.

of the Autosubmit workflow, which manages the NICOCO ensemble experiment.

6. Self-review of Current Progress and Future Prospects

Regarding the original project plan, the model and workflow deployment and initial condition preparation tasks have been completed, and early ensemble experiments are underway. However, several factors added time and effort beyond what was originally planned. Porting Autosubmit to Miyabi required implementing PBS scheduler support and ensuring compatibility with Miyabi's 2FA authentication system, both of which involved nontrivial technical work that delayed the start of production runs.

As a result, while the 150-year spin-up simulation to create the ocean states was successfully completed, the full 30-year historical ensemble runs on Miyabi-C could not yet be completed. The 10-member, 2-year ensemble runs carried out on Miyabi-C serve as the initial validation step, confirming that the workflow is ready to scale up to the full 50-member, 30-year runs. Looking ahead, the priority is to complete those ensemble simulations.

We hope to have the opportunity to complete the task by applying to the next JHPCN call. Meanwhile, we will evaluate the data generated by the simulations on Miyabi-C, together with the output from Fugaku and Marenostrum5. We will be able to use the data to draw first conclusions about the replicability of the climate model on the three clusters, including Miyabi-C.

References

[1] Satoh, M., Tomita, H., Yashiro, H. et al. The Non-hydrostatic Icosahedral Atmospheric Model: description and development. Prog. in Earth and Planet. Sci. 1, 18 (2014). https://doi.org/10.1186/s40645-014-0018-1

[2] Hasumi, H. *CCSR ocean component model (COCO) version 4.0.* Atmosphere and Ocean

Research Institute, The University of Tokyo, 111pp. (2006). [Available at https://ccsr.aori.u-tokyo.ac.jp/~hasumi/COCO/coco4.pdf.]

[3] Arakawa, T., T. Inoue, H. Yashiro, and M. Satoh. Coupling library Jcup3: Its philosophy and application. Prog. Earth Planet. Sci., 7, 6, (2020). doi:10.1186/s40645-019-0320-z

[4] Keller, K. R., Alerany Solé, M., and Acosta, M.: Replicability in Earth System Models, Geosci. Model Dev., 18, 10221–10243, https://doi.org/10.5194/gmd-18-10221-2025, 2025.