Kilometer-Scale AI-Powered and Performance-Portable Earth System Model (AP3ESM) to Achieve Year-Scale Simulation Speed on Heterogeneous Supercomputers

Kai Xu
Laoshan Laboratory
Qingdao, China
xukai16@foxmail.com

Maoxue Yu
Laoshan Laboratory
Qingdao, China
mxyu@qnlm.ac

Yuhu Chen
Laoshan Laboratory
Qingdao, China
yhchen@qnlm.ac

Jie Gao
State Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics, Institute of Atmospheric Physics
Chinese Academy of Sciences
Beijing, China
PIESAT Information Technology Co., Ltd.
Beijing, China
jiegao.ehime@gmail.com

Shuang Wang
Gosci Technology Group
Qingdao, China
blue2007@163.com

Jiaying Song
Gosci Technology Group
Qingdao, China
jennys_rainbowz@163.com

Xiaohui Duan
Shandong University
Jinan, China
sunrise.duan@sdu.edu.cn

Junwei Wei
Computer Network Information Center, Chinese Academy of Sciences
Beijing, China
Pengcheng Laboratory
Shenzhen, China
University of Chinese Academy of Sciences
Beijing, China
weijunlin@cnic.cn

Jiangfeng Yu
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
University of Chinese Academy of Sciences
Beijing, China
yujiangfeng21@mails.ucas.ac.cn

Hailong Liu
Laoshan Laboratory
Qingdao, China
hlliu2@qnlm.ac

Jinrong Jiang
Computer Network Information Center, Chinese Academy of Sciences
Beijing, China
jjr@sccas.cn

Yi Zhang
State Key Laboratory of Climate System Prediction and Risk Management (CPRM), and School of Atmospheric Sciences
Nanjing University of Information Science and Technology
Beijing, China
yizhang@cma.gov.cn

Pengfei Lin
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
linpf@mail.iap.ac.cn

Tianyi Wang
School of Information Science and Engineering
Shandong Normal University
Jinan, China
wangty@sdnu.edu.cn

Pengfei Wang
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
wpf@mail.iap.ac.cn

Weipeng Zheng
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
zhengwp@mail.iap.ac.cn

Jingwei Xie
Laoshan Laboratory
Qingdao, China
xiejw23@mail3.sysu.edu.cn

Jiakang Zhang
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
University of Chinese Academy of Sciences
Beijing, China
zhangjiakang23@mails.ucas.ac.cn

Zilu Liu
Laoshan Laboratory
Qingdao, China
lzldyx2024@163.com

Xiaoyu Jin
Laoshan Laboratory
Qingdao, China
xiaoyujin@zju.edu.cn

Jilin Wei
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
weijilin@mail.iap.ac.cn

Qixin Chang
Shandong University
Jinan, China
cqx@mail.sdu.edu.cn

Qingxia Lin
Gosci Technology Group
Qingdao, China
lqx13589356860@163.com

Yanzhi Zhou
Institute of Atmospheric Physics, Chinese Academy of Sciences
Beijing, China
University of Chinese Academy of Sciences
Beijing, China
zhouyanzhi23@mails.ucas.ac.cn

Weiguo Liu
Shandong University
Jinan, China
weiguo.liu@sdu.edu.cn

Wei Xue
Tsinghua University
Beijing, China
National Supercomputing Center in Wuxi
Beijing, China
xuewei@tsinghua.edu.cn

Yiwen Li
School of Ocean Sciences
China University of Geosciences
Beijing, China
lyw@cugb.edu.cn

Haohuan Fu
Tsinghua University
Beijing, China
National Supercomputing Center in Wuxi
Beijing, China
haohuan@tsinghua.edu.cn

Yue Yu
Pengcheng Laboratory
Shenzhen, China
yuy@pcl.ac.cn

Xuebin Chi
Computer Network Information Center, Chinese Academy of Sciences
Beijing, China
chi@sccas.cn

Lixin Wu
Laoshan Laboratory
Qingdao, China
lxwu@ouc.edu.cn

*Corresponding authors are Hailong Liu (hlliu2@qnlm.ac), Jinrong Jiang (jjr@sccas.cn), Yi Zhang (yizhang@nuist.edu.cn), and Pengfei Lin (linpf@mail.iap.ac.cn).

SC '25, St Louis, MO, USA

Abstract

Kilometer-scale Earth system models (ESMs) necessitate exascale supercomputers to facilitate realistic simulations of weather phenomena and climate variability over a time span ranging from days to decades. We present AP^{3}ESM, an ultra-high-resolution,

ACM ISBN 979-8-4007-1466-5/25/11
https://doi.org/10.1145/3712285.3771788

AI-Powered, **P**erformance-**P**ortable **ESM** coupling atmosphere, land surface, ocean, and sea ice components. By leveraging the performance portability features of Kokkos and OpenMP, the AP^{3}ESM operates efficiently on two heterogeneous systems while incurring minimal development overhead. Advanced optimization techniques, such as adaptive parallel algorithms, AI-enhanced physical parameterizations, and mixed-precision computations, have been implemented to further boost the computational efficiency. Breaking the 1-km resolution barrier, AP^{3}ESM delivers 0.85 and 1.98 simulated-years-per-day (SYPD) for the standalone atmosphere and ocean components on 34.1 million Sunway cores and 16085 GPUs, respectively; the holistic AP^{3}ESM achieves 0.54 SYPD on 37.2 million Sunway cores. Notably, the forecast experiment successfully captures Super Typhoon Doksuri in 2023 and its associated extreme rainfall across China.

CCS Concepts

• **Applied computing** → **Earth and atmospheric sciences**; • **Computing methodologies** → **Massively parallel algorithms**.

Keywords

Earth system model, AI-Powered, Performance-Portable, Kilometer-scale, Heterogeneous architecture

ACM Reference Format:

Kai Xu, Maoxue Yu, Yuhu Chen, Jie Gao, Shuang Wang, Jiaying Song, Xiaohui Duan, Junwei Wei, Jiangfeng Yu, Hailong Liu, Jinrong Jiang, Yi Zhang, Pengfei Lin, Tianyi Wang, Pengfei Wang, Weipeng Zheng, Jingwei Xie, Jiakang Zhang, Zilu Liu, Xiaoyu Jin, Jilin Wei, Qixin Chang, Qingxia Lin, Yanzhi Zhou, Weiguo Liu, Wei Xue, Yiwen Li, Haohuan Fu, Yue Yu, Xuebin Chi, and Lixin Wu. 2025. Kilometer-Scale AI-Powered and Performance-Portable Earth System Model (AP3ESM) to Achieve Year-Scale Simulation Speed on Heterogeneous Supercomputers. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '25), November 16–21, 2025, St Louis, MO, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3712285.3771788

1 JUSTIFICATION FOR ACM GORDON BELL PRIZE FOR CLIMATE MODELLING

We present a kilometer-scale **AI-P**owered and **P**erformance-**P**ortable **E**arth **S**ystem **M**odel (AP^{3}ESM): standalone 1-km atmosphere and ocean components achieve 0.85 and 1.98 SYPD on 34.1 million cores and 16085 GPUs; the fully coupled 1-km AP^{3}ESM achieves a record 0.54 SYPD on 37.2 million cores.

2 Performance Attributes

Attribute title	Attribute value
Category achievement	Time-to-solution, Scalability
Type of method used	Semi-implicit
Results reported on the basis of	Whole application excluding I/O and initialization
Precision reported	FP64,FP32 mixed precision
System scale	Results measured on full-scale system
Measurement mechanism	Timers, Simulated-Years-Per-Day

3 Overview of the Problem

Climate change is intensifying the magnitude and frequency of extreme weather events, posing severe threats to ecosystems and food security, and profoundly impacting socioeconomic stability and development [19, 50]. Earth system models (ESMs) that integrate the processes of the atmosphere, ocean, land, and sea ice are indispensable tools for unveiling the complexities of climate change [27, 48]. Compared with a single atmospheric or oceanic module, the coupled ESMs can better simulate the complicated multisphere system and their interactions, particularly the feedback mechanisms between the atmosphere and the ocean [6, 23, 34]. This advantage enables a more realistic portrayal of physical processes including heat exchange, momentum transfer, and moisture fluxes, which are critical in climate variability phenomena such as the El Niño event [3, 18].

The horizontal resolution of ESMs has a critical impact on their ability to resolve fine-scale processes [16]. Although coarse-resolution ESMs can capture global climate trends, they are limited in predicting regional extremes such as catastrophic tropical cyclones, and parameterizations for phenomena like deep convective clouds and oceanic eddies inevitably introduce further uncertainties [12, 33, 35, 43]. To address these limitations, kilometer-scale (km-scale) ESMs that operate at resolutions finer than 10 km are currently actively advanced by the scientific community, aiming to bridge the gap between long-term climate simulations and short-term weather forecasts, as well as between global-scale and local-process modeling. Km-scale atmosphere models can resolve the interaction between cloud microphysics and deep convection dynamics based on the first principle [29, 49]. This improves the representation of radiation feedback and multiscale flow interaction, thus boosting the fidelity of weather-climate forecasting. Km-scale ocean models allow for the accurate simulation of oceanic mesoscale and submesoscale motions [15, 47]. These motions contain the majority of the oceanic kinetic energy and have a significant impact on the intrinsic ocean dynamics and the climate variability [14, 36, 45, 46]. Km-scale ESMs can combine these advantages and further capture the subtle fine-scale interactions among different spheres, providing a new means to advance understanding of the Earth system.

The emergence of exascale computing offers tremendous potential to significantly advance km-scale ESMs simulation capabilities. Nevertheless, effectively adapting ESMs to leverage these advanced computational capabilities remains a significant challenge. With the gradual decline of Moore's Law and Dennard scaling, the increases in computing power come primarily from a larger number of compute cores and innovations in heterogeneous architecture [7, 13, 22]. Modern exascale supercomputers contain from roughly one to several million compute cores, and presently, 9 of the TOP 10 supercomputers (November 2024 TOP 500 List) are equipped with AMD, Nvidia, and Intel GPUs. Therefore, the enormous number of compute cores, as well as the heterogeneous architecture of exascale supercomputers, make the adaptability of ESMs a critical challenge [32].

Adapting a large number of codes to new architectures is another challenge for km-scale ESMs [2]. Each component - atmosphere, ocean, ice, and coupler - contains large legacy codes, which bring productivity and portability challenges in optimizing the code on

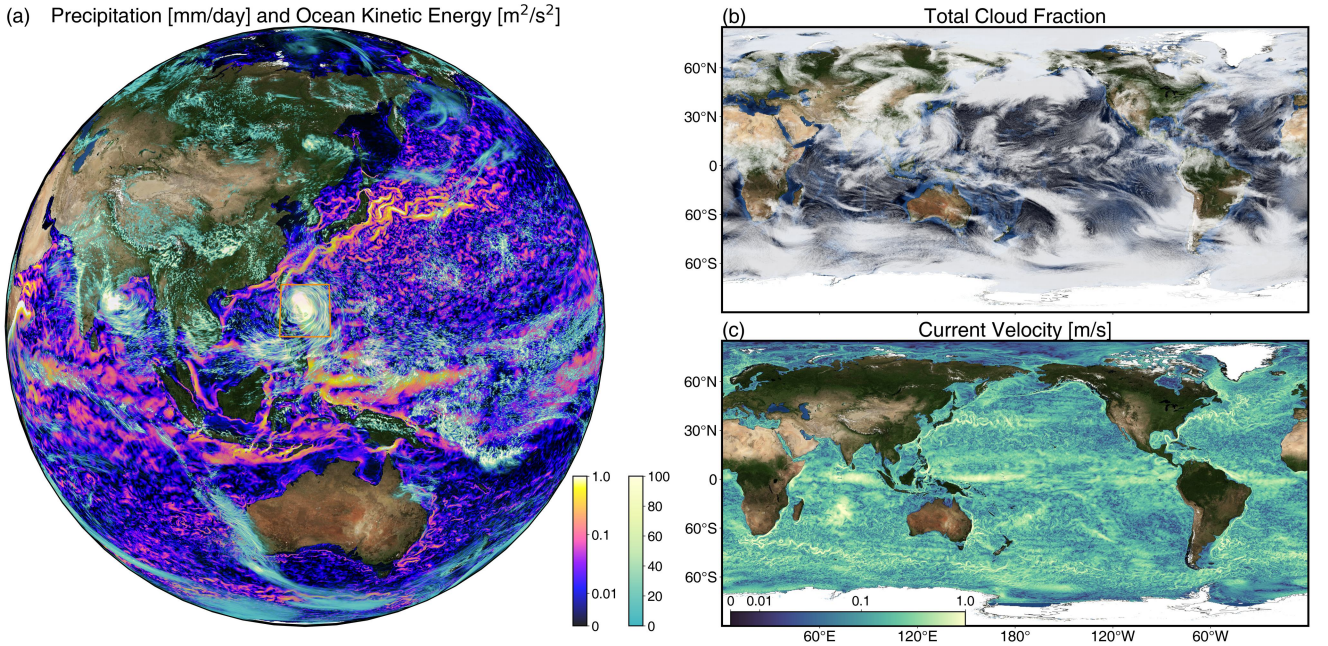

Figure 1: (a) Precipitation and sea surface kinetic energy on 25 July, 2023, from the 3v2 AP^{3}ESM. The orange rectangle highlights the region of typhoon Doksuri. (b) Total cloud fraction snapshot from the 1-km-resolution GRIST atmosphere-only simulation. (c) The magnitude of sea surface velocity snapshot from the 1-km-resolution ($1/100^\circ$ tripole grid) LICOMK++ ocean-only simulation. Kinetic energy in (a) and velocity magnitude in (c) are displayed with logarithmic colorbars.

heterogeneous architectures. Combining intensive communication patterns (both intra and inter models) and extreme process counts creates performance and scalability barriers for km-scale ESMs. Input/output (I/O) bottlenecks and data handling complexities arise from the exponential growth of data size in the km-scale ESMs, overwhelming storage and processing capabilities [1]. The early designs did not account for such high resolutions and large-scale parallelism of modern systems, leading to initialization that consumes much time, which becomes the bottleneck during a porting process, and much memory resources, which leads to the failure of the initialization stage. These challenges collectively pose a significant threat to fully harnessing the power of exascale computing to advance km-scale ESMs, demanding urgent and coordinated efforts to address them effectively.

In 2024, two research teams achieved milestones in km-scale climate modeling: the implementation of an AI-enhanced physics package and the adoption of a performance-portable programming model. These achievements include the development of the Global-Regional Integrated Forecast System (GRIST), a kilometer-scale atmospheric general circulation model enhanced with AI and scalable to 30 million cores [10], and LICOMK++, a performance-portable ocean general circulation model scalable to about 40 million cores with a global mean horizontal resolution below 1 km, the latter shortlisted for the 2024 ACM Gordon Bell Prize for Climate Modeling [40]. Figs. 1b and 1c show the simulated fields of the above two models with the 1-km horizontal resolution. These results demonstrate that the models can provide a solid foundation for developing a fully coupled km-scale Earth System Model, as illustrated in Figure 1a. In addition, another team ported Community Earth System Model (CESM) 2.2 to the Sunway supercomputer with the flux coupler optimized and developed a fully coupled ESM with the atmospheric component of 5 km resolution and the oceanic component of 3 km resolution [9]. **Building on these successes, the three teams are now collaborating to create a novel ultra-high-resolution ESM integrating components of atmosphere, land surface, ocean, and sea-ice at an unprecedented km-scale resolution. This paper describes how the challenges encountered in developing the km-scale ESM are addressed to enhance computational speed and efficiency.**

4 Current State of the Art

The advent of exascale computing and advances in numerical modeling have opened new frontiers in earth system modeling, enabling global simulations at km-scale resolutions. These high-resolution models explicitly resolve critical processes while also revealing potential shifts in large-scale circulation patterns under climate change [21, 30]. However, to date, the km-scale ESM works are still lacking due to their tremendous requirements for computational and storage resources, as well as the challenges of porting and optimizing millions of legacy codes. Fig. 2 summarizes some recent works on high resolution ESM or climate models, providing key developments across major modeling frameworks [5, 8, 9, 17, 26, 31, 38, 44]. Works collected in this paper cover the High-Resolution Model Intercomparison Project (HighResMIP), the Next Generation Earth Modeling Systems (nextGEMS) project with higher resolution, and the Energy Exascale Earth System Model (E3SM) series.

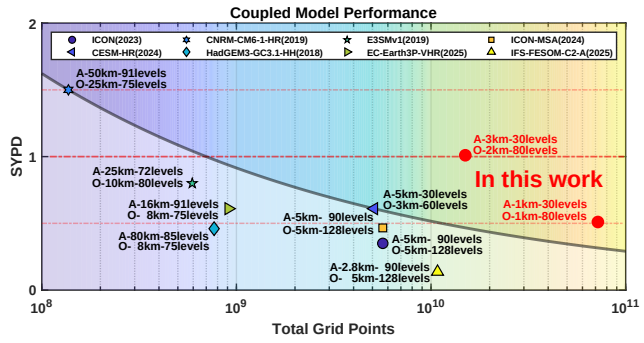

Figure 2: A summary of a recent high-resolution coupled climate model. The case with the highest resolution in each work is selected. The dividing line, representing the state-of-the-art, results from a log-linear fit between the CNRM (2019) and the CESM (2024), identified as the most favorable cases in the 10^8 and 10^9 order-of-magnitude ranges of total grid points.

HighResMIP was established under the Coupled Model Intercomparison Project Phase 6 (CMIP6) framework to provide a benchmark for global high-resolution climate simulations, with typical resolutions of 25–50 km for the atmosphere and 10–25 km for the ocean. For instance, HadGEM3-GC3.1-HH, a high-resolution configuration of the UK Met Office's ESM within the CMIP6 ensemble, features an ultra-high-resolution grid of 0.08° for the ocean and 0.8° for the atmosphere with a lower simulation speed of 0.49 SYPD and a significantly higher resource demand of 588,931 core hours per simulated year [44]. The EC-Earth consortium developed their very high resolution variant model (EC-Earth3P-VHR), with 16 km atmosphere and 8 km ocean grids, achieving ~2.8 SYPD on MareNostrum[26]. With advances in computational resources and modeling techniques, HighResMIP2 aims to further refine model resolutions, targeting 10 km for the atmosphere and 5 km for the ocean, to enhance the representation of critical climate processes while addressing previous limitations[28].

In the nextGEMS project, researchers achieved a breakthrough by performing multi-decadal climate simulations at km-scale resolutions [31]. They attained resolutions as fine as 2.8 km atmosphere and 5 km ocean (i.e., ICON), and 4.4 km atmosphere and 5 km ocean (i.e., IFS-FESOM). For the final 30-year production runs under SSP3-7.0 scenarios, the models were configured at 5 km ocean and 9 km atmosphere resolutions, achieving a throughput of 600 SDPD on the Levante supercomputer. This demonstrates the feasibility of sustained coupled simulations with a global km-scale horizontal resolution, with ICON and IFS-FESOM: the former minimizes parameterizations and the latter leverages operational weather model tuning. The ICON has seen particularly significant advancements in coupling capabilities. Bishnoi et al. (2023) implemented a novel modular supercomputing architecture (MSA) approach on the JUWELS system, where the atmospheric component ran on GPU nodes while the ocean and I/O operated on CPU clusters, achieving 170 SDPD (about 0.47 SYPD) at 5 km global resolution [5]. The ICON-Sapphire team tried global coupling at 2.5 km on monthly timescales and 1.25 km on daily timescales but without public information for computation speed and efficiency [17].

The work of the Energy Exascale Earth System Model (E3SM) in 2019 represents the first simulation in its high-resolution configuration. As a fully coupled ESM (including atmosphere, ocean, sea ice, and land), it coupled advanced earth system components, achieving ~0.8 SYPD using ~640000 cores [8]. The E3SM team also carried out large-scale experiments with the atmospheric component (not fully coupled). They achieved a performance of 0.97 SYPD at a realistic cloud resolution of 3 km on Summit and 1.26 SYPD at a simple cloud-resolving resolution of 3.25 km on Frontier. The latter won the Gordon Bell Prize for Climate Modeling in 2023 [4, 37].

By porting CESM 2.2 to the 40-million-core Sunway supercomputer, Duan et al. (2024) achieved performance metrics: 340 SDPD (0.93 SYPD) for a 5 km atmosphere, 265 SDPD (0.73 SYPD) for a 3 km ocean, and 222 SDPD (0.61 SYPD) for the fully coupled ESM. These advancements were made possible through a non-intrusive optimization approach that minimized manual code modifications while maximizing efficiency [9].

This work introduces a novel **AI**-**P**owered and **P**erformance-**P**ortable **E**arth **S**ystem **M**odel (AP^{3}ESM) comprising atmosphere, ocean, land surface, and sea ice components. It couples the latest versions of GRIST and LASG IAP Climate system Ocean Model (LICOM) within the CESM 2.2. The AI-enhanced GRIST has achieved 1.35 SYPD for a 3-km global simulation and 0.5 SYPD for a global 1-km simulation [10], while LICOMK++ achieves 1.70 SYPD for a global 1-km simulation [40, 47]. Based on these preliminary series of works, breakthroughs have been made in this study:

- AP^{3}ESM is a novel global ultra-high-resolution ESM achieving 1-km resolution, with the highest total number of grid points reported to date in the literature, utilizing AI-powered physical parameterizations and performance-portable technologies.
- AP^{3}ESM at 1-km resolution achieves 0.54 SYPD with a parallel efficiency of 90.7% when scaling to 37.2 million cores on the Sunway OceanLight supercomputer.
- AP^{3}ESM's standalone atmospheric and oceanic components at 1-km horizontal resolution achieve 0.85 SYPD on 34.1 million cores for the atmosphere and 1.98 SYPD on 16085 GPUs for the ocean.
- AP^{3}ESM successfully reproduces the evolution process of super typhoon Doksuri, validating its high capacity to forecast extreme weather events.

5 Innovations Realized

The AP^{3}ESM leverages the latest CESM coupler [9] to integrate two state-of-the-art standalone models: the GRIST [10] for atmosphere and the LICOM [40] for ocean, both of which are critical components for climate modelling. Ultimately, this work couples four major components of the ESM at the km-scale: atmosphere, ocean, sea ice, and land, as shown in Fig. 3.

5.1 Engineering Innovation

This research leverages software engineering innovations to realize a hybrid programming paradigm, integrating large, complex legacy Fortran code with incremental C/C++ code for the Sunway system.

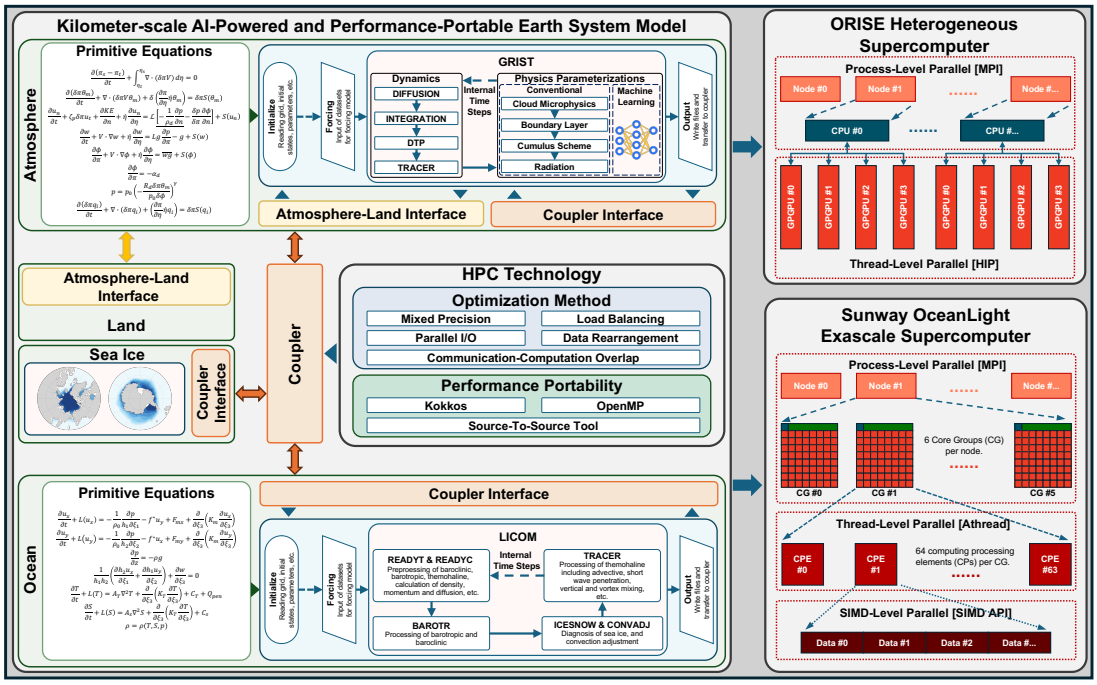

Figure 3: High-level schematics of the AP^{3}ESM model, the ORISE heterogeneous supercomputer, the Sunway OceanLight supercomputer, and their interrelations.

Initially, we refactored the standalone version of each component, such as redesigning the program workflow and restructuring the project tree. Subsequently, we developed a limited run on the management processing element (MPE) prototype that incrementally integrated the above components, driven by the coupler. This process involved addressing two key aspects: 1) resolving numerous engineering challenges, such as naming conflicts, dimension alignment, and redundant memory footprint, and 2) verifying the coupled model's correctness through bit-for-bit and non-bit-for-bit (visualized results) validation at various experimental setups. Finally, to accelerate the prototype model on the Sunway heterogeneous architecture, we employed a performance portability strategy in conjunction with raw kernel optimization. This approach ensures the model's execution on the Sunway OceanLight system and its potential portability to other architectures. Performance portability is addressed by implementing the atmosphere, sea ice, and land components using OpenMP within the Fortran programming language and the ocean component via the Kokkos C++ library.

5.1.1 Components of AP^{3}ESM. The coupler runs on all processors and handles coupler sequencing, model concurrency, and communication between components. The AP^{3}ESM uses the CPL7 coupler, which adopts optimizations such as a quick sort algorithm for rearranging communication [10]. The Model Coupling Toolkit (MCT) is a critical piece of software in the coupler. CPL7 uses MCT-based *init*, *run*, and *finalize* interfaces in each component to control the whole workflow. The three interfaces are implemented for GRIST and LICOM to be integrated into CPL7. Moreover, the *import* and *export* methods are also implemented for GRIST and LICOM to get boundary condition data from other models and provide output boundary condition data for other models in the system. The coupler manages the main clock in the system and maintains a clock that is associated with each component. GRIST and LICOM implement the clock, which is consistent with the coupling clock, and make sure the coupling period is consistent with their internal timestep.

To optimize the atmosphere component GRIST on the Sunway architecture, SWGOMP is developed as a compiler-plugin-based tool. Most of the GRIST loops are conflict-free, so adding !$omp target directives that are compatible with the OpenMP 4.5 specification can achieve good performance [10]. In addition, GRIST integrates a new AI-physics suite. This suite gets the input variables from the dynamical core and returns full physical variables back to the physics-dynamics coupling interface of GRIST for the next-step dynamical core integration. The AI-physics suite features greater flexibility for adaptation across different architectures.

The ocean component LICOM offers two key features enabling a wide range of high-performance computing (HPC) and climate science experiments. As HPC architectures evolve, our team has actively developed architecture-specific versions (CUDA, HIP, and Athread) of LICOM to target diverse platforms [20, 39, 41]. We also implemented a performance-portable version using Kokkos [42], which was enhanced into LICOMK++ [40]. This portfolio of implementations enables AP^{3}ESM to flexibly select the most suitable implementation for each architecture to achieve optimal performance. In this work, we refactor the standalone versions of swLICOM (an Athread-based implementation of LICOM) and LICOMK++ by restructuring their project trees and program architectures to enable

integration with the CESM coupler on the Sunway OceanLight system.

In the present study, GRIST and the land surface model directly exchange data, bypassing the coupler. Consequently, AP^{3}ESM does not currently include a coupler-owned land model component. We couple a sea ice model, CICE4, in our system, which is also the default sea ice model for CESM2.2. The CICE4 has also been optimized on the multi-core system of Sunway. Since these two components are not bottlenecks, the subsequent discussion only focuses on optimizing the atmosphere and ocean components.

5.1.2 Hybrid Task-Data Parallelization Strategy. Executing ultra-high-resolution climate models on the exascale supercomputer poses a significant scientific and engineering challenge. Consequently, designing more efficient parallel algorithms is crucial to adapting to this scenario.

This work implements a scalable hybrid task-data parallelization strategy for the coupled model. For data-level parallelism, model computations are fully parallelized at both process and thread levels. The atmosphere, land, and sea ice components are parallelized with MPI + OpenMP; while the ocean component provides flexible parallelization via MPI with either Kokkos or Athread. SIMD-level parallelism is employed for specific computations. For task-level parallelism, the model provides two strategies. Either all components are executed sequentially within a single domain, or the components are partitioned across multiple domains for concurrent execution; each domain has exclusive access to its own computing resources. Computational resource allocation is adjusted based on the computational profile of each component, thereby optimizing model load balancing for efficient simulation execution.

5.2 Algorithmic Optimization

5.2.1 AI-Powered Resolution-Adaptive Physics Suite For Atmosphere Component. We develop an AI-powered, resolution-adaptive physics suite (Fig. 4)—comprising an AI tendency module, an AI radiation diagnosis module, and a conventional physics diagnostic module—to replace the conventional physical parameterizations suite.

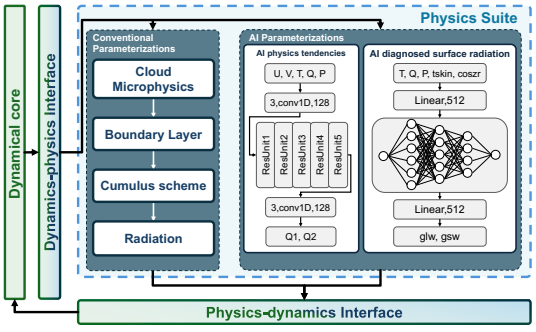

Figure 4: Overview of the AI physics parameterization suite, with two AI parameterization modules and a conventional diagnostic module.

The AI tendency module takes horizontal wind (U and V), temperature (T), specific humidity (Q), and pressure (P) fields as inputs and applies a one-dimensional convolution along the vertical column to capture the characteristics of temperature, humidity, and other atmospheric variables during events such as convection and atmospheric instability. Its architecture comprises five ResUnits within an 11-layer deep Convolutional Neural Network (CNN) totaling approximately 5×10^5 trainable parameters, chosen to balance predictive skill and computational cost.

The AI radiation diagnosis module incorporates skin temperature (tskin) and the cosine of the solar zenith angle (coszr) in addition to the atmospheric inputs to characterize surface state and top-of-atmosphere insolation geometry. A 7-layer multi-layer perceptron (MLP) with residual connections estimates the surface downward shortwave (gsw) and longwave (glw) fluxes, which serve as inputs to the land surface model and surface layer scheme.

The training dataset consists of 5 km GRIST atmospheric fields spanning 80 days (20 from each season). We employ a 7:1 training:test partition, and extract three random time steps per day as a validation subset for hyperparameter tuning for ENSO and MJO events, and reducing overfitting risk. Using high-resolution (5 km) outputs as supervision allows the resulting suite to generalize effectively to lower resolutions, yielding resolution-adaptive performance.

The AI physical parameterization suite achieves computational gains by unifying most operations into highly efficient tensor kernels (principally matrix multiplication)—where modern processors deliver superior floating-point throughput—while also lowering code complexity and simplifying cross-architecture optimization.

5.2.2 Excluding 3D Non-Ocean Grid Points for Optimizing Ocean Component. Oceans cover approximately 71% of the Earth's surface, and 3D realistic scenarios exhibit significant topographical height variations across different regions. A data remapping optimization, incorporating 3D non-ocean point removal, was first implemented in LICOMK++. While previous research utilized this technique for thread-level optimization only in the *canuto* parameterization scheme, we extend its scope to optimize the entire ocean component, as illustrated in Fig. 5. Initially, input data are partitioned, and the total grid points of non-ocean points are removed. Then, an MPI rank mapping ensures correct data access, and a new communication topology optimizes boundary exchange. This results in about 30% computational resource reduction, consistent results, and improved efficiency at the process-level parallelism. Similarly, the optimization method has been applied to the sea-ice model.

5.2.3 Mixed-Precision Computation Innovations. The double-precision (64-bit) floating-point arithmetic is necessary for ESMs to control computational errors. Low precision can reduce memory, communication, and computational overhead, and is commonly used in AI. To leverage these advantages while maintaining accuracy, we implement a group-wise scaling mixed-precision method (FP64/FP32) for key components of the model, including GRIST (atmospheric) and LICOM (ocean). Since exploiting a mixed-precision scheme for ML-based parameterizations is straightforward at the operator level due to the model's compact design, we focus on reducing variable precision within the dynamical core of GRIST and LICOM. We designed tailored accuracy evaluations for GRIST and

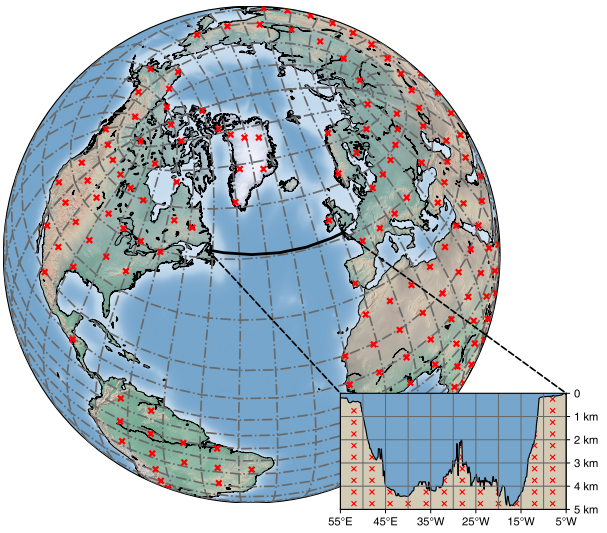

Figure 5: Optimization of the triple-grid ocean component by excluding 3D non-ocean grid points.

LICOM. For GRIST, we measured surface pressure and relative vorticity deviations using the relative L_2 norm against double-precision baselines, with a 5% error threshold for long-term stability. For LICOM, which uses tripolar grids, we incorporated grid area into root mean square deviation (RMSD) calculations. Averaging 30 days of daily data, RMSD values were 0.018°C for temperature, 0.0098 psu for salinity, and 0.0005 m for sea surface height—indicating sufficient accuracy for long-term simulations.

5.2.4 Coupler Optimization. MCT is created to address the challenges of creating a parallel coupler for the Earth system model. Although MCT has optimized its implementation, the memory in a CG of Sunway cannot satisfy the requirements for MCT to construct the GSMap, an MCT datatype used to describe the decomposition in coupling, and the Router table, an $M \times N$ table for transferring a distributed data object from a module running on M processes to another running on N processes. Given two decompositions specified in two GSMaps, the Router table can easily build a mapping between the location of one grid point on a processor and its location on another processor. The process of building two data structures is time- and memory-consuming. To address the problem, the two data structures are generated offline as a preprocessing step. Furthermore, we remove the unnecessary communication variables that are registered in MCT and are not used in GRIST and LICOM. Rearrangement in the coupler generalizes the matrix transpose. The original all-to-all MPI was inefficient; we implemented non-blocking point-to-point MPI, which overlaps communication and computation for improved performance.

5.2.5 Parallel Input/Output (I/O) Strategy. The I/O and storage bottlenecks are challenging for extreme-scale applications running on exascale supercomputers. To address initialization and I/O bottlenecks, a data-partitioning strategy that divides data into smaller subfiles is implemented. We assign groups of MPI ranks to the I/O for a set of subfiles, and leverage a binary format for the I/O data. Compared with the original implementation, this optimized approach enhances the I/O performance and reduces the workload on file and network systems when scaling across tens of thousands of nodes.

5.3 Performance Portability for Sunway OceanLight supercomputer

This work achieves performance portability through two strategies: Kokkos and OpenMP. The former is a C++ template metaprogramming (TMP) library designed for performance portability. For the Sunway architecture, we propose a hash-based function registration and callback mechanism to enable Kokkos execution on TMP-constrained Sunway processors. We enhance processor utilization through a hybrid host-device backend parallelism strategy. Furthermore, Kokkos offers finer-grained tile profiling for multi-dimensional parallel iterations, enhancing algorithmic flexibility. OpenMP, the latter strategy, is a directive-based performance portability solution. For Sunway processors, we utilize a compiler plug-in called SWGOMP to enable OpenMP-driven automatic loop space mapping on Sunway's computing processing elements (CPEs).

6 How Performance Was Measured

6.1 Configuration of Modeling Experiments

Table 1 presents the five primary model configurations employed herein. The atmosphere component GRIST uses five resolutions: 1-km, 3-km, 6-km, 10-km, and 25-km resolution. The dycore, tracer, and model timesteps in the atmosphere component are set to 8 s, 30 s, and 120 s, respectively. For all five configurations, the number of vertical layers is 30.

The ocean component LICOM employs five problem cases with 1-km, 2-km, 3-km, 5-km, and 10-km resolutions, each utilizing 80 vertical levels. Time steps for barotropic, baroclinic, and tracer processes are set to 2 s, 20 s, and 20 s, respectively. The configuration of the sea-ice component is designed to mirror that of the ocean component.

For AP^{3}ESM, the resolution configurations for the atmosphere and ocean components are presented as pairs: 1-km atmospheric resolution with 1-km oceanic resolution (1v1), 3-km with 2-km (3v2), 6-km with 3-km (6v3), 10-km with 5-km (10v5), and 25-km with 10-km (25v10). The first value in each pair denotes the atmospheric resolution, while the second indicates the oceanic resolution. The coupler frequencies are 180, 36, and 180 couplings per day for the atmosphere, the ocean, and the sea ice components, respectively.

6.2 Major Performance Metrics

For our experiments, we describe the speed of simulation using SYPD (simulated-year-per-day). These metrics are calculated by dividing the length of the simulated time interval by the wall-clock time required for execution, ensuring independence from the simulation's time horizon. Wall-clock time measurements are obtained using timers from the GPTL in Coupler 7, with the maximum value across all MPI ranks recorded to account for potential load imbalance. The calculation of simulation speed uses the script tool *getTiming* in Coupler 7. We also differentiate between component performance and full model performance: the former only accounts

Table 1: Major configurations of GRIST, LICOM, and AP^{3}ESM. The labeled "Grids" represents the number of grid points.

GRIST					LICOM				AP^{3}ESM	
Res.(km)	cells	edges	vertices	No. of Grids	Res.(km)	Longitudes Grids	Latitudes Grids	No. of Grids	Label	Total Grids
1	3.4×10^8	5.0×10^8	1.7×10^8	8.6×10^9	1	36000	22018	6.3×10^{10}	1v1	7.2×10^{10}
3	4.2×10^7	1.3×10^8	8.4×10^7	2.1×10^9	2	18000	11511	1.3×10^{10}	3v2	1.5×10^{10}
6	1.1×10^7	3.2×10^7	2.1×10^7	5.4×10^8	3	10800	6907	5.8×10^9	6v3	6.3×10^9
10	2.6×10^6	7.9×10^6	5.2×10^6	1.9×10^8	5	7200	4605	2.1×10^9	10v5	2.3×10^9
25	6.7×10^5	2.0×10^6	1.3×10^6	3.1×10^7	10	3600	2302	5.2×10^8	25v10	5.5×10^8

for time spent inside the AP^{3}ESM atmosphere and ocean component, while the latter includes all time, as well as in the component coupler. Initialization and I/O costs are excluded from our analysis.

6.3 Hardware Platform Details

Our performance evaluations were conducted on the ORISE Supercomputer and Sunway OceanLight supercomputer [25]. The latter is the successor to TaihuLight Supercomputer. The Sunway OceanLight supercomputer has more than 107520 nodes. Each node has one SW26010P 390-core CPU, resulting in a parallel scale of 41932800 cores. In addition, each node in the Sunway OceanLight supercomputer has a dedicated network connection to a leaf switch with 304 ports. Of these, 256 ports are connected to nodes, and 48 are connected to secondary switches. Each 256-processor node group connected to the same leaf switch forms a super node, which enables high-speed communication bandwidth across CPUs. All supernodes are connected through a 16:3 (256:48) oversubscribed multilayer fat tree network. In our evaluation, we assign one process per CG, with the MPE offloading its computation tasks to the CPEs within the same CG.

The ORISE is a heterogeneous supercomputer located at the Computer Network Information Center (CNIC) of the CAS, China. Each computing node comprises one CPU and four single instruction, multiple threads (SIMT) accelerators, which serve as the host and device components. The CPU features a 4-way, 8-core architecture with an x86 instruction set, operates at 2.0 GHz, and has 128 GB of memory. The device side features four GPGPU-like accelerators with performance akin to AMD MI60, utilizing the HIP as the backend programming model, which we will call HIP-based GPU. The CPU and HIP-based GPUs are interconnected through 32-bit PCIe buses featuring Direct Memory Access (DMA) with a bandwidth of 16 GB/s, while a high-speed network with a bandwidth of 25 GB/s is used to connect nodes.

7 Performance Results

7.1 High- and Ultra-high-Resolution Simulations

Typhoons and hurricanes are powerful tropical cyclones accompanied by active air-sea interaction. The fierce winds, torrential rains, and storm surges they bring pose substantial threats to socio-economic situations [11, 24]. We conducted test simulations of super Typhoon Doksuri, which formed in the western Pacific on July 23, 2023, using the configurations of AP^{3}ESM 3v2 and 25v10.

The higher the horizontal resolution, the more details of the typhoon and its related ocean response the ESM can simulate. Fig. 6 shows the spatial distribution of the atmospheric 10 m wind and sea surface Rossby number (Ro, the vertical vorticity normalized by the local Coriolis parameter), as simulated by AP^{3}ESM 3v2 and AP^{3}ESM 25v10 at 00:00 UTC on July 25, 2023. Both configurations effectively capture the fully developed structure of the typhoon, including the low-wind eye and the concentric high-wind eyewall (see Fig. 6a and 6b). Notably, the AP^{3}ESM 3v2 case produces a more compact typhoon eye and resolves significantly finer details in the spatial pattern of the wind field. In addition, the sea surface Ro field clearly depicts the oceanic response to the intense wind field. The AP^{3}ESM 3v2 configuration (Fig. 6c) can resolve a wealth of fine-scale patterns beneath the core and surrounding areas of the typhoon. In contrast, the AP^{3}ESM 25v10 experiment (Fig. 6d) can only depict the localized oceanic response directly beneath the core of the typhoon.

Fig. 7 presents the trajectory and intensity of Typhoon Doksuri from the China Meteorological Administration (CMA) best-track, the ERA5 reanalysis, and the AP^{3}ESM 3v2 simulation. During the initial stage, the simulated track shows close agreement with the CMA best-track. Although the simulated track becomes qualitatively consistent with the CMA best track at the later stages, the AP^{3}ESM 3v2 simulation can reproduce a more intense typhoon compared to the ERA5 reanalysis.

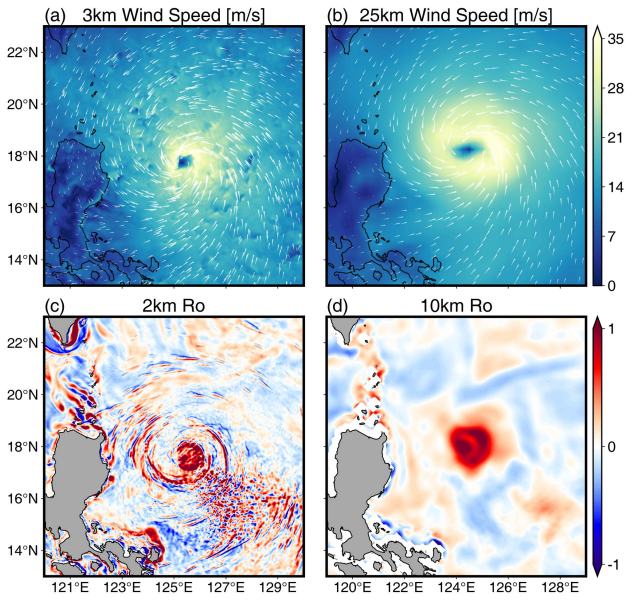

Figure 6: (a) The 10 m wind field and (c) the ocean surface Rossby number from 3v2 AP^{3}ESM at 00:00 UTC, 25 July, 2023. (b) and (d) are the same, but from the 25v10 experiment. The white streamlines depict the 10 m wind vector.

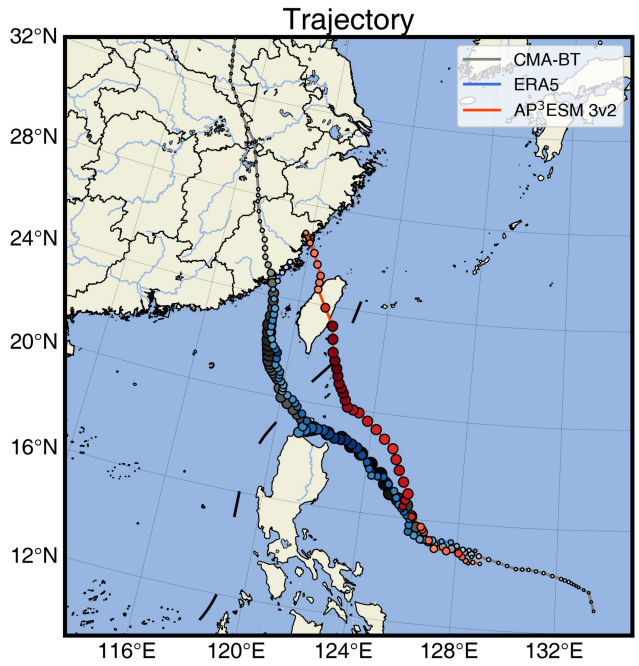

Figure 7: Trajectories of Typhoon Doksuri from the CMA best track, ERA5 reanalysis, and the AP^{3}ESM 3v2 simulation. The size and the color (darker means stronger) indicate the typhoon level and the maximum wind speed, respectively.

7.2 Strong Scalability

The strong scaling results are presented in Fig. 8a, with specific experimental data provided in Table 2. First, we present the results for the atmospheric component (ATM). This component was run on the Sunway OceanLight system at 3-km and 1-km horizontal resolutions. In Fig. 8a, the baseline MPE-only version is represented by the orange dashed line (3km ATM MPE), while the orange curve indicates the CPE-accelerated and optimized version (3 km ATM CPE+OPT). Scaling from 5462 to 43691 nodes, the performance of the MPE code (from 32768 to 262144 cores) improves from 0.0032 to 0.0063 SYPD, achieving a parallel efficiency of 24.6%. In comparison, the CPE+OPT code improves from 0.36 to 1.16 SYPD when scaling from 2129920 to 17039360 cores, with a parallel efficiency of 40.3%. Compared to the MPE code, the CPE+OPT code demonstrates a performance acceleration ranging from 112 to 184 times for the atmospheric component. Because the 1-km ATM configuration runs too slowly on the MPE, we evaluated only the optimized ATM code on the CPEs (1 km ATM CPE+OPT): scaling from 4259840 to 34078270 cores increased performance from 0.20 to 0.85 SYPD, corresponding to a parallel efficiency of 51.5%.

Second, for the oceanic component (OCN), our experiments on the Sunway OceanLight supercomputer mirror those conducted for the atmospheric component, including a comparison between the MPE-only implementation (2 km OCN MPE, green dashed line in Fig. 8a) and a version accelerated and optimized for the CPE (2 km OCN CPE+OPT, green curves in Fig. 8a). The oceanic component was tested at a 2-km global resolution. Both implementations demonstrate scalability from over 3000 to over 50000 nodes. The performance of MPE implementation increases from 0.0014 to 0.019

Table 2: Comparison of the performances of AP^{3}ESM and its major components in different systems and at different resolutions.

ORISE supercomputer						
1 km	OCN model (Original)	Nodes	1000	2000	3000	4000
		GPUs	4000	8000	12000	16000
		SYPD	0.77	1.25	1.49	1.70
		Efficiency	100%	81.6%	64.8%	55.6%
	OCN model (OPT)	Nodes	1015	2015	2982	4022
		GPUs	4060	8060	11927	16085
		SYPD	0.92	1.45	1.76	1.98
		Efficiency	100%	79.4%	65.1%	54.3%
Sunway OceanLight supercomputer						
2 km	OCN model (MPE)	Nodes	3265	6425	12671	50035
		Cores	19608	38550	76026	300210
		SYPD	0.0014	0.0033	0.0060	0.019
		Efficiency	100%	118%	107%	88.6%
	OCN model (CPE+OPT)	Nodes	3265	6425	12671	50035
		Cores	1273415	2505880	4941755	19513780
		SYPD	0.21	0.42	0.72	1.59
		Efficiency	100%	100%	85.7%	49.4%
3 km	ATM model (MPE)	Nodes	5462	10923	21846	43691
		Cores	32768	65536	131072	262144
		SYPD	0.0032	0.0040	0.0052	0.0063
		Efficiency	100%	62.5%	40.6%	24.6%
	ATM model (CPE+OPT)	Nodes	5462	10923	21845	43691
		Cores	2129920	4259840	8519680	17039360
		SYPD	0.36	0.70	0.92	1.16
		Efficiency	100%	97.2%	63.9%	40.3%
1 km	ATM model (CPE+OPT)	Nodes	10923	21846	43691	87382
		Cores	4259840	8519680	17039360	34078720
		SYPD	0.20	0.40	0.71	0.85
		Efficiency	100%	100%	86.5%	51.5%
3v2	AP^{3}ESM (CPE+OPT)	Nodes	8726	17351	34523	93726
		Cores	3403335	6767020	13464165	36553140
		SYPD	0.18	0.34	0.57	1.01
		Efficiency	100%	94.4%	79.2%	52.2%
1v1	AP^{3}ESM (CPE+OPT)	Nodes		22424	44511	95316
		Cores		8745360	17359160	37172980
		SYPD		0.14	0.23	0.54
		Efficiency		100%	82.8%	90.7%

SYPD, utilizing from approximately 20000 to over 300000 cores. Similarly, the CPE+OPT implementation's performance improves from 0.21 to 1.59 SYPD, scaling its core usage from 1273415 to 1951373780. At the largest scale, the parallel efficiency for the MPE and CPE+OPT implementations is 88.6% and 49.4%, respectively. CPE+OPT code obtains a speedup of about 84× to 150× compared to the MPE code. In particular, we designed experiments on the ORISE supercomputer, using the performance record of the 2024 Gordon Bell Prize for Climate Modeling finalist as a benchmark [40], indicated by the yellow dashed line in Fig. 8a (1 km OCN Original). The ocean model scales from 4060 to 16085 GPUs at a global 1-km resolution, with performance increasing from 0.92 to 1.98 SYPD, corresponding to a parallel efficiency of 54.3% (1 km OCN OPT, yellow curves in Fig. 8a). Furthermore, at the largest scale, this work attains a speedup of 1.2× compared to the best record [40]. The parallel algorithm of the ocean model underwent a systematic redesign, specifically removing 3D non-oceanic points, leading to a remapping of computational tasks and resources. As a result of these changes, the GPU count could not be directly matched to the benchmark.

Finally, for AP^{3}ESM, we evaluated the result of the fully coupled ESM, including the atmosphere, ocean, sea ice, and land components. AP^{3}ESM is partitioned into two parallel task domains. The first domain encompasses the coupler, atmospheric, sea ice, and land components. The second domain consists solely of the ocean component. This resource allocation strategy is motivated by the following: 1) The atmosphere component exhibits the highest computational cost, and placing the coupler within the same domain

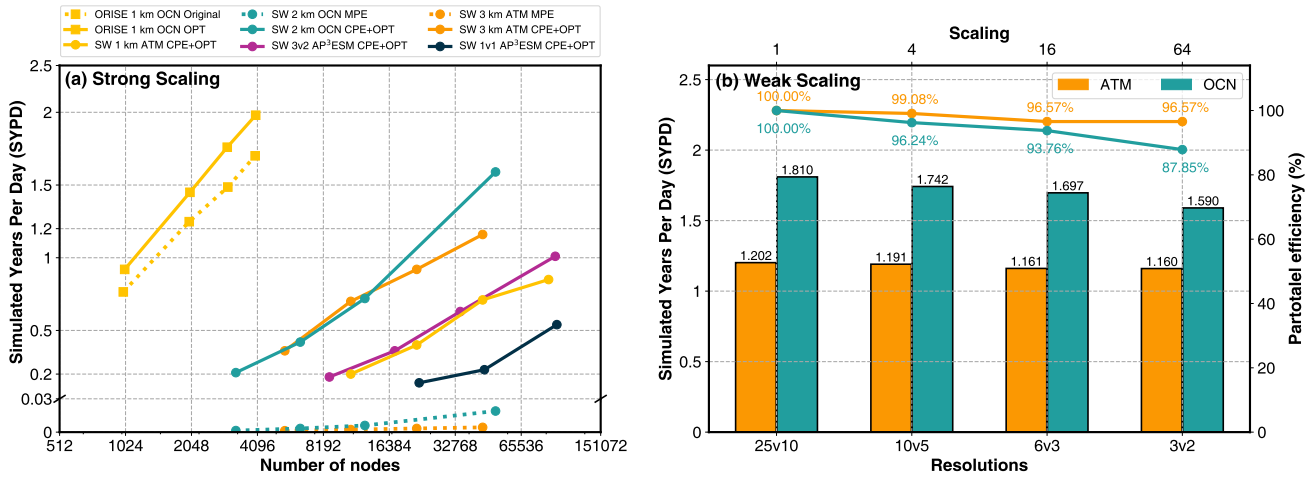

Figure 8: (a) Strong scaling results of the ocean model, the atmosphere model, and the AP^{3}ESM. Throughput measured in simulate-years-per-day, plotted as a function of compute nodes, on the ORISE and Sunway OceanLight Supercomputers. The data are also listed in Table 2. (b) Weak scaling results for the atmosphere and the ocean models of four different resolutions are listed in the Table 1 on the Sunway OceanLight system.

minimizes data exchange overhead between them. 2) The land component is inherently coupled with the atmospheric component. 3) The sea ice component contributes minimal computational overhead; its integration with the coupler does not incur additional communication costs when simulating at ultra-large scales. 4) The ocean component represents the second largest computational cost, necessitating its allocation to a separate domain. Under the described experimental setup, AP^{3}ESM (3v2) scales from nearly 10% to approximately 100% of the Sunway OceanLight system capacity, with core counts ranging from over 3.4 million (3403335) to about 36.6 million (36553140), achieving a performance increase from 0.18 to 1.01 SYPD with 52.2% parallel efficiency. Under a 1v1 configuration, scaling the core count from 8745360 to 37172980 improved throughput from 0.14 to 0.54 SYPD, achieving a 90.7% strong-scaling efficiency. This high efficiency stems from optimized component configuration and communication strategies. These results demonstrate that the km-scale AP^{3}ESM achieves year-scale simulation speeds on the Sunway OceanLight supercomputer.

7.3 Weak Scalability

Figure 8b presents the weak scaling results at the horizontal resolution of 25-km, 10-km, 6-km, 3-km for the atmosphere component and 10-km, 5-km, 3-km, 2-km for the ocean component. We evaluated the results on 683, 2731, 10922, and 43691 nodes for the atmosphere component for four resolutions, respectively. The atmosphere component uses 17039360 cores and achieves a parallel efficiency of 87.85%. For the ocean component, we evaluated the results on 2107, 8212, 18225, and 50035 nodes for four resolutions, respectively. The ocean component uses 19513780 cores and achieves a parallel efficiency of 96.57%. The drop in weak scalability at higher resolutions is attributable primarily to increased inter-process communication latency and synchronization overhead at large node counts.

8 Implications

The successful development of the AP^{3}ESM represents a breakthrough in km-scale Earth system modeling. By coupling a global 3-km resolution atmosphere model (GRIST) with a 2-km resolution ocean model (LICOM), along with sea ice and land surface modules, and incorporating innovative technologies such as AI-enhanced physical parameterizations, performance-portability, and mixed-precision computing optimizations, the model achieves year-scale simulation speeds (1.01 SYPD) on the Sunway OceanLight supercomputer. We are advancing AP^{3}ESM toward a fully coupled global 1-km AP^{3}ESM with unprecedented resolution, achieving a performance of 0.54 SYPD. In addition, the standalone atmospheric and ocean components have reached 1-km resolution, achieving 0.85 SYPD on 34.1 million cores and 1.98 SYPD on 16085 GPUs, respectively. This provides a powerful new tool for high-precision climate modeling and extreme weather prediction.

The ultra-high-resolution AP^{3}ESM directly resolves phenomena like deep convective clouds and oceanic mesoscale eddies, traditionally handled through parameterizations, and substantially cuts down the model uncertainties. For example, in simulating Super Typhoon Doksuri, the 3-km atmosphere component captures the eye structure and fine-scale rainband, which is missed in the 25-km simulation. The coupled models also reproduce the sea surface temperature cold trails following typhoon passage. These results highlight the critical importance of km-scale ESMs for studying multi-sphere coupling feedback mechanisms and enhancing the reliability of disaster prevention and mitigation strategies. In addition, AP^{3}ESM shows the potential of establishing a seamless Earth system modeling framework that spans from global to regional scales and from short-term weather forecasting to long-term climate prediction, which will organically combine the scientific exploration with practical application.

Furthermore, the high-performance portability of AP^{3}ESM sets a benchmark for heterogeneous supercomputing applications. Leveraging the Kokkos and OpenMP performance-portable programming models, AP^{3}ESM achieves efficient parallel scaling on both Sunway OceanLight and ORISE heterogeneous supercomputers. This success validates engineering innovations such as hybrid programming paradigms, task-data parallel strategies, and communication topology restructuring, providing a technical blueprint for adapting other ESMs to heterogeneous architectures. Simultaneously, mixed-precision computing and AI-enhanced parameterization schemes dramatically reduce computational costs, demonstrating that deep integration of AI with traditional dynamical models may be the key to overcoming memory and computing power constraints.

Looking forward, the further development of ESMs will necessitate interdisciplinary cooperation and upgrades to computing infrastructure. The km-scale resolution generates vast datasets, imposing higher requirements on I/O efficiency, storage management, and visualization technologies. Moreover, integrating more Earth system components (such as biogeochemical processes) requires overcoming numerical stability challenges in cross-scale coupling. To further scale, we will explore federating geographically distributed HPC clusters through a computing power network, enabling task-level parallel execution of distinct ESM components and thereby improving aggregate performance. Only by continuously optimizing algorithms, expanding AI applications, and engaging in deeper cross-disciplinary collaboration can we fully tap the potential of exascale computing and establish a more solid scientific foundation for combating climate change.

Acknowledgments

This study was supported by the National Natural Science Foundation of China (92358302, U2242214), the National Key R&D Program for Developing Basic Sciences (2022YFC3104802), Strategic Priority Research Program of Chinese Academy of Sciences (XDB0500303, XDB0500101), the Tai Shan Scholar Program (tstp20231237), Laoshan Laboratory Project (LSKJ202 300301), Natural Science Foundation of Shandong Province (ZR2022QF125), Innovation Seed Funding of Computer Network Information Center of Chinese Academy of Sciences (24IF01), and the Startup Foundation for Introducing Talent of Nanjing University of Information Science and Technology. The numerical calculations in this study were carried out on the ORISE and the Sunway OceanLight Supercomputers.

References

[1] Sameh Abdulah, Allison H. Baker, George Bosilca, Qinglei Cao, Stefano Castruccio, Marc G. Genton, David E. Keyes, Zubair Khalid, Hatem Ltaief, Yan Song, Georgiy L. Stenchikov, and Ying Sun. 2024. Boosting Earth System Model Outputs And Saving PetaBytes in Their Storage Using Exascale Climate Emulators . In *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society, Los Alamitos, CA, USA, 1–12. doi:10.1109/SC41406.2024.00008

[2] Peter Bauer, Peter D Dueben, Torsten Hoefler, Tiago Quintino, Thomas C Schulthess, and Nils P Wedi. 2021. The digital revolution of Earth-system science. *Nature Computational Science* 1, 2 (2021), 104–113.

[3] Goratz Beobide-Arsuaga, Tobias Bayr, Annika Reintges, and Mojib Latif. 2021. Uncertainty of ENSO-amplitude projections in CMIP5 and CMIP6 models. *Climate Dynamics* 56 (2021), 3875–3888. doi:10.1007/s00382-021-05673-4

[4] Luca Bertagna, Oksana Guba, Mark A Taylor, James G Foucar, Jeff Larkin, Andrew M Bradley, Sivasankaran Rajamanickam, and Andrew G Salinger. 2020. A performance-portable nonhydrostatic atmospheric dycore for the Energy Exascale Earth System Model running at cloud-resolving resolutions.. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–14.

[5] Abhiraj Bishnoi, Olaf Stein, Catrin I. Meyer, René Redler, Norbert Eicker, Helmuth Haak, Lars Hoffmann, Daniel Klocke, Luis Kornblueh, and Estela Suarez. 2023. Earth system modeling on modular supercomputing architecture: coupled atmosphere–ocean simulations with ICON 2.6.6-rc. *Geoscientific Model Development* 17, 1 (2023), 261–273. doi:10.5194/gmd-17-261-2024

[6] Stuart Bishop, Justin Small, and Frank Bryan. 2020. The global sink of available potential energy by mesoscale air-sea interaction. *Journal of Advances in Modeling Earth Systems* 12 (2020), e2020MS002118. doi:10.1029/2020MS002118

[7] Probir K Bondyopadhyay. 1998. Moore's law governs the silicon revolution. *Proc. IEEE* 86, 1 (1998), 78–81.

[8] Peter M. Caldwell, Azamat Mametjanov, Qi Tang, Luke P. Van Roekel, Jean-Christophe Golaz, Wuyin Lin, David C. Bader, Noel D. Keen, Yan Feng, Robert Jacob, Mathew E. Maltrud, Andrew F. Roberts, Mark A. Taylor, Milena Veneziani, Hailong Wang, Jonathan D. Wolfe, Karthik Balaguru, Philip Cameron-Smith, Lu Dong, Stephen A. Klein, L. Ruby Leung, Hong-Yi Li, Qing Li, Xiaohong Liu, Richard B. Neale, Marielle Pinheiro, Yun Qian, Paul A. Ullrich, Shaocheng Xie, Yang Yang, Yuying Zhang, Kai Zhang, and Tian Zhou. 2019. The DOE E3SM Coupled Model Version 1: Description and Results at High Resolution. *Journal of Advances in Modeling Earth Systems* 11, 12 (2019), 4095–4146. doi:10.1029/2019ms001870

[9] Xiaohui Duan, Yuxuan Li, Zhao Liu, Bin Yang, Juepeng Zheng, Haohuan Fu, Shaoqing Zhang, Shiming Xu, Yang Gao, Wei Xue, Di Wei, Xiaojing Lv, Lifeng Yan, Haopeng Huang, Haitian Lu, Lingfeng Wan, Haoran Lin, Qixin Chang, Chenlin Li, Quanjie He, Zeyu Song, Xuantong Wang, Yangyang Yu, Xilong Fan, Zhaopeng Qu, Yankun Xu, Xiuwen Guo, Yunlong Fei, Zhaoying Wang, Mingkui Li, Yingjing Jiang, Lv Lu, Liang Su, Jiayu Fu, Peinan Yu, Weiguo Liu, Lixin Wu, Lanning Wang, Xin Liu, Dexun Chen, and Guangwen Yang. 2024. Kilometer-Level Coupled Modeling Using 40 Million Cores: An Eight-Year Journey of Model Development. arXiv:2404.10253 [cs.DC] https://arxiv.org/abs/2404.10253

[10] Xiaohui Duan, Yi Zhang, Kai Xu, Haohuan Fu, Bin Yang, Yiming Wang, Yilun Han, Siyuan Chen, Zhuangzhuang Zhou, Chenyu Wang, Dongqiang Huang, Huihai An, Xiting Ju, Haopeng Huang, Zhuang Liu, Wei Xue, Weiguo Liu, Bowen Yan, Jianye Hou, Maoxue Yu, Wenguang Chen, Jian Li, Zhao Jing, Hailong Liu, and Lixin Wu. 2025. An AI-Enhanced 1km-Resolution Seamless Global Weather and Climate Model to Achieve Year-Scale Simulation Speed using 34 Million Cores. In *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Las Vegas, NV, USA) *(PPoPP '25)*. Association for Computing Machinery, New York, NY, USA, 524–538. doi:10.1145/3710848.3710893

[11] James B Elsner, James P Kossin, and Thomas H Jagger. 2008. The increasing intensity of the strongest tropical cyclones. *Nature* 455, 7209 (2008), 92–95.

[12] Baylor Fox-Kemper, Alistair Adcroft, Claus Böning, Eric Chassignet, Enrique Curchitser, Gokhan Danabasoglu, Carsten Eden, Matthew England, Rüdiger Gerdes, Richard Greatbatch, Stephen Griffies, Robert Hallberg, Emmanuel Hanert, Patrick Heimbach, Helene Hewitt, Christopher Hill, Yoshiki Komuro, Sonya Legg, Julien Le Sommer, Simona Masina, Simon Marsland, Stephen Penny, Fangli Qiao, Todd Ringler, Anne Marie Treguier, Hiroyuki Tsujino, Petteri Uotila, and Stephen Yeager. 2019. Challenges and prospects in ocean circulation models. *Frontiers in Marine Science* 6 (2019), 65. doi:10.3389/fmars.2019.00065

[13] David J Frank, Robert H Dennard, Edward Nowak, Paul M Solomon, Yuan Taur, and Hon-Sum Philip Wong. 2001. Device scaling limits of Si MOSFETs and their application dependencies. *Proc. IEEE* 89, 3 (2001), 259–288.

[14] Jonathan Gula, John Taylor, Andrey Shcherbina, and Amala Mahadevan. 2022. Chapter 8: Submesoscale processes and mixing. In *Ocean Mixing*, Michael Meredith and Alberto Naveira Garabato (Eds.). Elsevier, 181–214. doi:10.1016/B978-0-12-821512-8.00015-3

[15] Helene Hewitt, Baylor Fox-Kemper, Brodie Pearson, Malcolm Roberts, and Daniel Klocke. 2022. The small scales of the ocean may hold the key to surprises. *Nature Climate Change* 12 (2022), 496–499. doi:10.1038/s41558-022-01386-6

[16] Helene T Hewitt, Malcolm Roberts, Pierre Mathiot, Arne Biastoch, Ed Blockley, Eric P Chassignet, Baylor Fox-Kemper, Pat Hyder, David P Marshall, Ekaterina Popova, et al. 2020. Resolving and parameterising the ocean mesoscale in earth system models. *Current Climate Change Reports* 6 (2020), 137–152. doi:10.1007/s40641-020-00164-w

[17] Cathy Hohenegger, Peter Korn, Leonidas Linardakis, René Redler, Reiner Schnur, Panagiotis Adamidis, Jiawei Bao, Swantje Bastin, Milad Behravesh, Martin Bergemann, Joachim Biercamp, Hendryk Bockelmann, Renate Brokopf, Nils Brüggemann, Lucas Casaroli, Fatemeh Chegini, George Datseris, Monika Esch, Geet George, Marco Giorgetta, Oliver Gutjahr, Helmuth Haak, Moritz Hanke, Tatiana Ilyina, Thomas Jahns, Johann Jungclaus, Marcel Kern, Daniel Klocke, Lukas Kluft, Tobias Kölling, Luis Kornblueh, Sergey Kosukhin, Clarissa Kroll, Junhong Lee, Thorsten Mauritsen, Carolin Mehlmann, Theresa Mieslinger, Ann Kristin Naumann, Laura Paccini, Angel Peinado, Divya Sri Praturi, Dian Putrasahan, Sebastian Rast, Thomas Riddick, Niklas Roeber, Hauke Schmidt, Uwe Schulzweida, Florian Schütte, Hans Segura, Radomyra Shevchenko, Vikram Singh, Mia Specht,

Claudia Christine Stephan, Jin-Song von Storch, Raphaela Vogel, Christian Wengel, Marius Winkler, Florian Ziemen, Jochem Marotzke, and Bjorn Stevens. 2022. ICON-Sapphire: simulating the components of the Earth system and their interactions at kilometer and subkilometer scales. *Geoscientific Model Development* 16, 2 (2022), 779–811. doi:10.5194/gmd-16-779-2023

[18] Meiyi Hou and Youmin Tang. 2022. Recent progress in simulating two types of ENSO–from CMIP5 to CMIP6. *Frontiers in Marine Science* 9 (2022), 986780. doi:10.3389/fmars.2022.986780

[19] IPCC. 2023. Climate Change 2023: Synthesis Report. (2023), 35–115. doi:10.59327/IPCC/AR6-9789291691647

[20] Jinrong Jiang, Pengfei Lin, Joey Wang, Hailong Liu, Xuebin Chi, Huiqun Hao, Yuzhu Wang, Wu Wang, and Linghan Zhang. 2019. Porting LASG/ IAP Climate System Ocean Model to Gpus Using OpenAcc. *IEEE Access* 7 (2019), 154490–154501. doi:10.1109/ACCESS.2019.2932443

[21] Falko JUDT, Daniel KLOCKE, Rosimar RIOS-BERRIOS, Benoit VANNIERE, Florian ZIEMEN, Ludovic AUGER, Joachim BIERCAMP, Christopher BRETHERTON, Xi CHEN, Peter DÜBEN, Cathy HOHENEGGER, Marat KHAIROUTDINOV, Chihiro KODAMA, Luis KORNBLUEH, Shian-Jiann LIN, Masuo NAKANO, Philipp NEUMANN, William PUTMAN, Niklas RÖBER, Malcolm ROBERTS, Masaki SATOH, Ryosuke SHIBUYA, Bjorn STEVENS, Pier Luigi VIDALE, Nils WEDI, and Linjiong ZHOU. 2021. Tropical Cyclones in Global Storm-Resolving Models. *Journal of the Meteorological Society of Japan. Ser. II* 99, 3 (2021), 579–602. doi:10.2151/jmsj.2021-029

[22] Hassan N Khan, David A Hounshell, and Erica RH Fuchs. 2018. Science and research policy at the end of Moore's law. *Nature Electronics* 1, 1 (2018), 14–21.

[23] Siqi Li and Changsheng Chen. 2022. Air-sea interaction processes during hurricane Sandy: Coupled WRF-FVCOM model simulations. *Progress in Oceanography* 206 (2022), 102855. doi:10.1016/j.pocean.2022.102855

[24] Yi Li, Youmin Tang, Shuai Wang, Ralf Toumi, Xiangzhou Song, and Qiang Wang. 2023. Recent increases in tropical cyclone rapid intensification events in global offshore regions. *Nature Communications* 14, 1 (2023), 5167.

[25] Rongfen Lin, Xinhui Yuan, Wei Xue, Wanwang Yin, Jienan Yao, Junda Shi, Qiang Sun, Chaobo Song, and Fei Wang. 2023. 5 ExaFlop/s HPL-MxP Benchmark with Linear Scalability on the 40-Million-Core Sunway Supercomputer. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–13.

[26] Eduardo Moreno-Chamarro, Thomas Arsouze, Mario Acosta, Pierre-Antoine Bretonnière, Miguel Castrillo, Eric Ferrer, Amanda Frigola, Daria Kuznetsova, Eneko Martin-Martinez, Pablo Ortega, and Sergi Palomas. 2024. The very-high-resolution configuration of the EC-Earth global model for HighResMIP. *Geoscientific Model Development* 18, 2 (2024), 461–482. doi:10.5194/gmd-18-461-2025

[27] Brian O'Neill, Claudia Tebaldi, Detlef van Vuuren, Veronika Eyring, Pierre Friedlingstein, George Hurtt, Reto Knutti, Elmar Kriegler, Jean-Francois Lamarque, Jason Lowe, Gerald Meehl, Richard Moss, Keywan Riahi, and Benjamin Sanderson. 2016. The Scenario Model Intercomparison Project (ScenarioMIP) for CMIP6. *Geoscientific Model Development* 9 (2016), 3461–3482. doi:10.5194/gmd-9-3461-2016

[28] Malcolm J. Roberts, Kevin A. Reed, Qing Bao, Joseph J. Barsugli, Suzana J. Camargo, Louis-Philippe Caron, Ping Chang, Cheng-Ta Chen, Hannah M. Christensen, Gokhan Danabasoglu, Ivy Frenger, Neven S. Fučkar, Shabeh ul Hasson, Helene T. Hewitt, Huanping Huang, Daehyun Kim, Chihiro Kodama, Michael Lai, Lai-Yung Ruby Leung, Ryo Mizuta, Paulo Nobre, Pablo Ortega, Dominique Paquin, Christopher D. Roberts, Enrico Scoccimarro, Jon Seddon, Anne Marie Treguier, Chia-Ying Tu, Paul A. Ullrich, Pier Luigi Vidale, Michael F. Wehner, Colin M. Zarzycki, Bosong Zhang, Wei Zhang, and Ming Zhao. 2025. High-Resolution Model Intercomparison Project phase 2 (HighResMIP2) towards CMIP7. *Geoscientific Model Development* 18, 4 (2025), 1307–1332. doi:10.5194/gmd-18-1307-2025

[29] Masaki Satoh, Bjorn Stevens, Falko Judt, Marat Khairoutdinov, Shian-Jiann Lin, William M. Putman, and Peter Düben. 2019. Global Cloud-Resolving Models. *Current Climate Change Reports* 5 (2019), 172–184. doi:10.1007/s40641-019-00131-0

[30] Christoph Schär, Oliver Fuhrer, Andrea Arteaga, Nikolina Ban, Christophe Charpilloz, Salvatore Di Girolamo, Laureline Hentgen, Torsten Hoefler, Xavier Lapillonne, David Leutwyler, Katherine Osterried, Davide Panosetti, Stefan Rüdisühli, Linda Schlemmer, Thomas Schulthess, Michael Sprenger, Stefano Ubbiali, and Heini Wernli. 2019. Kilometer-scale climate models: Prospects and challenges Kilometer-scale climate models: Prospects and challenges. *Bulletin of the American Meteorological Society* 101, 5 (2019), E567–E587. doi:10.1175/bams-d-18-0167.1

[31] Hans Segura, Xabier Pedruzo-Bagazgoitia, Philipp Weiss, Sebastian K. Müller, Thomas Rackow, Junhong Lee, Edgar Dolores-Tesillos, Imme Benedict, Matthias Aengenheyster, Razvan Aguridan, Gabriele Arduini, Alexander J. Baker, Jiawei Bao, Swantje Bastin, Eulàlia Baulenas, Tobias Becker, Sebastian Beyer, Hendryk Bockelmann, Nils Brüggemann, Lukas Brunner, Suvarchal K. Cheedela, Sushant Das, Jasper Denissen, Ian Dragaud, Piotr Dziekan, Madeleine Ekblom, Jan Frederik Engels, Monika Esch, Richard Forbes, Claudia Frauen, Lilli Freischem, Diego García-Maroto, Philipp Geier, Paul Gierz, Álvaro González-Cervera, Katherine Grayson, Matthew Griffith, Oliver Gutjahr, Helmuth Haak, Ioan Hadade, Kerstin Haslehner, Shabeh ul Hasson, Jan Hegewald, Lukas Kluft, Aleksei Koldunov, Nikolay Koldunov, Tobias Kölling, Shunya Koseki, Sergey Kosukhin, Josh Kousal, Peter Kuma, Arjun U. Kumar, Rumeng Li, Nicolas Maury, Maximilian Meindl, Sebastian Milinski, Kristian Mogensen, Bimochan Niraula, Jakub Nowak, Divya Sri Praturi, Ulrike Proske, Dian Putrasahan, René Redler, David Santuy, Domokos Sármány, Reiner Schnur, Patrick Scholz, Dmitry Sidorenko, Dorian Spät, Birgit Sützl, Daisuke Takasuka, Adrian Tompkins, Alejandro Uribe, Mirco Valentini, Menno Veerman, Aiko Voigt, Sarah Warnau, Fabian Wachsmann, Marta Wacławczyk, Nils Wedi, Karl-Hermann Wieners, Jonathan Wille, Marius Winkler, Yuting Wu, Florian Ziemen, Janos Zimmermann, Frida A.-M. Bender, Dragana Bojovic, Sandrine Bony, Simona Bordoni, Patrice Brehmer, Marcus Dengler, Emanuel Dutra, Saliou Faye, Erich Fischer, Chiel van Heerwaarden, Cathy Hohenegger, Heikki Järvinen, Markus Jochum, Thomas Jung, Johann H. Jungclaus, Noel S. Keenlyside, Daniel Klocke, Heike Konow, Martina Klose, Szymon Malinowski, Olivia Martius, Thorsten Mauritsen, Juan Pedro Mellado, Theresa Mieslinger, Elsa Mohino, Hanna Pawłowska, Karsten Peters-von Gehlen, Abdoulaye Sarré, Pajam Sobhani, Philip Stier, Lauri Tuppi, Pier Luigi Vidale, Irina Sandu, and Bjorn Stevens. 2025. nextGEMS: entering the era of kilometer-scale Earth system modeling. *EGUsphere* 2025 (2025), 1–39. doi:10.5194/egusphere-2025-509

[32] John Shalf, Sudip Dosanjh, and John Morrison. 2010. Exascale computing technology challenges. In *Proceedings of the 9th International Conference on High Performance Computing for Computational Science* (Berkeley, CA) *(VECPAR'10)*. Springer-Verlag, Berlin, Heidelberg, 1–25.

[33] Steven Sherwood, Sandrine Bony, and Jean-Louis Dufresne. 2014. Spread in model climate sensitivity traced to atmospheric convective mixing. *Nature* 505, 7481 (2014), 37–42. doi:10.1038/nature12829

[34] Justin Small, Frank Bryan, Stuart Bishop, and Robert Tomas. 2019. Air-Sea Turbulent Heat Fluxes in Climate Models and Observational Analyses: What Drives Their Variability? *Journal of Climate* 32 (2019), 2397–2421. doi:10.1175/JCLI-D-18-0576.1

[35] Bjorn Stevens and Sandrine Bony. 2013. What are climate models missing? *Science* 340, 6136 (2013), 1053–1054. doi:10.1126/science.1237554

[36] John Taylor and Andrew Thompson. 2023. Submesoscale dynamics in the upper ocean. *Annual Review of Fluid Mechanics* 55, 1 (2023), 103–127. doi:10.1146/annurev-fluid-031422-095147

[37] Mark Taylor, Peter M Caldwell, Luca Bertagna, Conrad Clevenger, Aaron Donahue, James Foucar, Oksana Guba, Benjamin Hillman, Noel Keen, Jayesh Krishna, et al. 2023. The Simple Cloud-Resolving E3SM Atmosphere Model Running on the Frontier Exascale System. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–11.

[38] A. Voldoire, D. Saint-Martin, S. Sénési, B. Decharme, A. Alias, M. Chevallier, J. Colin, J.-F. Guérémy, M. Michou, M.-P. Moine, P. Nabat, R. Roehrig, D. Salas y Mélia, R. Séférian, S. Valcke, I. Beau, S. Belamari, S. Berthet, C. Cassou, J. Cattiaux, J. Deshayes, H. Douville, C. Ethé, L. Franchistéguy, O. Geoffroy, C. Lévy, G. Madec, Y. Meurdesoif, R. Msadek, A. Ribes, E. Sanchez-Gomez, L. Terray, and R. Waldman. 2019. Evaluation of CMIP6 DECK Experiments With CNRM-CM6-1. *Journal of Advances in Modeling Earth Systems* 11, 7 (2019), 2177–2213. doi:10.1029/2019ms001683

[39] P. Wang, J. Jiang, P. Lin, M. Ding, J. Wei, F. Zhang, L. Zhao, Y. Li, Z. Yu, W. Zheng, Y. Yu, X. Chi, and H. Liu. 2021. The GPU version of LASG/IAP Climate System Ocean Model version 3 (LICOM3) under the heterogeneous-compute interface for portability (HIP) framework and its large-scale application. *Geoscientific Model Development* 14, 5 (2021), 2781–2799. doi:10.5194/gmd-14-2781-2021

[40] Junlin Wei, Xiang Han, Jiangfeng Yu, Jinrong Jiang, Hailong Liu, Pengfei Lin, Maoxue Yu, Kai Xu, Lian Zhao, Pengfei Wang, Weipeng Zheng, Jingwei Xie, Yanzhi Zhou, Tao Zhang, Feng Zhang, Yehong Zhang, Yue Yu, Yuzhu Wang, Yidi Bai, Chen Li, Zipeng Yu, Haoyu Deng, Yaxin Li, and Xuebin Chi. 2024. A Performance-Portable Kilometer-Scale Global Ocean Model on ORISE and New Sunway Heterogeneous Supercomputers. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis* (Atlanta, GA, USA) *(SC '24)*. IEEE Press, Article 3, 12 pages. doi:10.1109/SC41406.2024.00009

[41] Junlin Wei, Jinrong Jiang, Hailong Liu, Feng Zhang, Pengfei Lin, Pengfei Wang, Yongqiang Yu, Xuebin Chi, Lian Zhao, Mengrong Ding, Yiwen Li, Zipeng Yu, Weipeng Zheng, and Yuzhu Wang. 2023. LICOM3-CUDA: a GPU version of LASG/IAP climate system ocean model version 3 based on CUDA. *J. Supercomput.* 79, 9 (Jan. 2023), 9604–9634. doi:10.1007/s11227-022-05020-2

[42] Junlin Wei, Pengfei Lin, Jinrong Jiang, Hailong Liu, Lian Zhao, Yehong Zhang, Xiang Han, Feng Zhang, Jian Huang, Yuzhu Wang, Youyun Li, Yue Yu, and Xuebin Chi. 2024. Accelerating LASG/IAP climate system ocean model version 3 for performance portability using Kokkos. 160 (11 2024), 901–917. doi:10.1016/j.future.2024.06.029

[43] Constanze Wellmann, Andrew Barrett, Jill Johnson, Michael Kunz, Bernhard Vogel, Ken Carslaw, and Corinna Hoose. 2020. Comparing the impact of environmental conditions and microphysics on the forecast uncertainty of deep convective clouds and hail. *Atmospheric Chemistry and Physics* 20 (2020), 2201–2219. doi:10.5194/acp-20-2201-2020

[44] K. D. Williams, D. Copsey, E. W. Blockley, A. Bodas-Salcedo, D. Calvert, R. Comer, P. Davis, T. Graham, H. T. Hewitt, R. Hill, P. Hyder, S. Ineson, T. C. Johns, A. B. Keen, R. W. Lee, A. Megann, S. F. Milton, J. G. L. Rae, M. J. Roberts, A. A. Scaife, R. Schiemann, D. Storkey, L. Thorpe, I. G. Watterson, D. N. Walters, A. West, R. A. Wood, T. Woollings, and P. K. Xavier. 2018. The Met Office Global Coupled Model 3.0 and 3.1 (GC3.0 and GC3.1) Configurations. *Journal of Advances in Modeling Earth Systems* 10, 2 (2018), 357–380. doi:10.1002/2017ms001115

[45] Jingwei Xie, Hailong Liu, and Pengfei Lin. 2023. A multifaceted isoneutral eddy transport diagnostic framework and its application in the Southern Ocean. *Journal of Advances in Modeling Earth Systems* 15 (2023), e2023MS003728. doi:10.1029/2023MS003728

[46] Jingwei Xie, Xi Wang, Hailong Liu, Pengfei Lin, Jiangfeng Yu, Zipeng Yu, Junlin Wei, and Xiang Han. 2024. ISOM 1.0: a fully mesoscale-resolving idealized Southern Ocean model and the diversity of multiscale eddy interactions. *Geoscientific Model Development* 17 (2024), 8469–8493. doi:10.5194/gmd-17-8469-2024

[47] Jingwei Xie, Jiangfeng Yu, Yanzhi Zhou, Hailong Liu, Junlin Wei, Xiang Han, Kai Xu, Maoxue Yu, Zipeng Yu, Pengfei Lin, Jinrong Jiang, Weipeng Zheng, Tao Zhang, Rong Wang, Zhao Jing, and Lixin Wu. 2025. A 1-km resolution global ocean simulation promises to unveil oceanic multi-scale dynamics and climate impacts. *The Innovation* 6, 7 (2025), 100843. doi:10.1016/j.xinn.2025.100843

[48] Mark Zelinka, Timothy Myers, Daniel McCoy, Stephen Po-Chedley, Peter Caldwell, Paulo Ceppi, Stephen Klein, and Karl Taylor. 2020. Causes of higher climate sensitivity in CMIP6 models. *Geophysical Research Letters* 47, 1 (2020), e2019GL085782. doi:10.1029/2019GL085782

[49] Yi Zhang, Xiaohan Li, Zhuang Liu, Xinyao Rong, Jian Li, Yihui Zhou, and Suyang Chen. 2022. Resolution sensitivity of the GRIST nonhydrostatic model from 120 to 5 km (3.75 km) during the DYAMOND winter. *Earth and Space Science* 9 (2022), e2022EA002401. doi:10.1029/2022EA002401

[50] Tianjun Zhou. 2021. New physical science behind climate change: What does IPCC AR6 tell us? *The Innovation* 2, 4 (2021). doi:10.1016/j.xinn.2021.100173