An AI-Enhanced 1km-Resolution Seamless Global Weather and Climate Model to Achieve Year-Scale Simulation Speed using 34 Million Cores

Xiaohui Duan[a,*], Yi Zhang[b,k,l,*,†], Kai Xu[c,*], Haohuan Fu[d,f,g], Bin Yang[e], Yiming Wang[k], Yilun Han[d], Siyuan Chen[k], Zhuangzhuang Zhou[f], Chenyu Wang[f], Dongqiang Huang[f], Huihai An[a], Xiting Ju[d], Haopeng Huang[d], Zhuang Liu[d], Wei Xue[d,j,f,†] ,Weiguo Liu[a,†], Bowen Yan[d], Jianye Hou[h], Maoxue Yu[c], Wenguang Chen[d], Jian Li[i], Zhao Jing[c], Hailong Liu[c,†], Lixin Wu[c]

[a]School of Software, Shandong University, Jinan, China
[b]Chinese Academy of Meteorological Sciences, Beijing, China
[c]Laoshan Laboratory, Qingdao, China
[d]Tsinghua University, Beijing, China
[e]College of Intelligence and Computing, Tianjin University, Tianjin, China
[f] National Supercomputing Center in Wuxi, Wuxi, China
[g] National Supercomputing Center in Shenzhen, Shenzhen, China
[h] The Chinese University of Hong Kong, Shenzhen, China
[i] Chinese Academy of Meteorological Sciences, Beijing, China
[j] Qinghai University, Qinghai, China
[k]Piesat Information Technology Co.,Ltd., Beijing, China
[l]Beijing Research Institute, Nanjing University of Information Science and Technology, Beijing, China

Abstract

Global Storm Resolving Models (GSRMs) is crucial for understanding extreme weather events under the climate change background. In this study, we optimize Global-Regional Integrated Forecast System (GRIST), which is a unified weather-climate modeling system designed for research and operation, for the next-generation Sunway supercomputer, incorporating AI-enhanced physics suite, OpenMP-based parallelization, and mixed-precision optimizations to enhance both efficiency and performance portability, as well as the unified modeling capability. Our experiments successfully capture significant events during the "23.7" extreme rainfall over northern China influenced by super Typhoon Doksuri, at 1km resolution. Notably, our work scales to 34 million cores, enabling simulation speeds at 491 SDPD (3km) and 181 SDPD (1km).

CCS Concepts: **• Applied computing → Earth and atmospheric sciences**; **• Computing methodologies → Parallel computing methodologies**; **Artificial intelligence**.

*Authors contributed equally to the paper.
† Corresponding authors are Yi Zhang (yizhang@cma.gov.cn), Wei Xue (xuewei@tsinghua.edu.cn), Weiguo Liu (weiguo.liu@sdu.edu.cn), and Hailong Liu (hlliu2@qnlm.ac).

PPoPP '25, March 1–5, 2025, Las Vegas, NV, USA
2025. ACM ISBN 979-8-4007-1443-6/25/03
https://doi.org/10.1145/3710848.3710893

Keywords: Global Storm Resolving Model, the next-generation Sunway supercomputer, performance portability, mixed-precision optimization, AI-enhanced physics suite

1 Introduction

The Earth, which is the home of our human beings, is one of the most complicated subject of our research and exploration. Fig. 1(a), known as the "Blue Marble" photographed by one of the astronauts on Apollo 17 in 1972, is generally recognized as one of the first image we acquire for the entire planet. This unique image, which helped to raise people's concepts about protecting Mother Gaia [19] or Mother Earth, also accidentally captured a severe cyclonic storm crossing Tamil Nandu [3] that led to heavy rain, flood, and severe losses of human lives and properties. Since then, satellite (such as Landsat [37] and MODIS [26]) and communication technologies have been developed to help us achieve constant earth observation from the space, producing an updated 1km resolution of the Earth (Fig. 1(b), merged from MODIS images that cover several months), demonstrating even more details of the global atmospheric system.

Besides Earth observation, Earth system modeling has been another major approach for scientists to understand the climate change mechanisms, and more importantly, to provide predictions that support hazard mitigation of more frequent and severe extreme weather/climate events, and sensible policy making processes that ensure a better future for all humans as one united community. Fig. 1(c) shows the modeling results of the early efforts by S. Manabe et al. [29] to

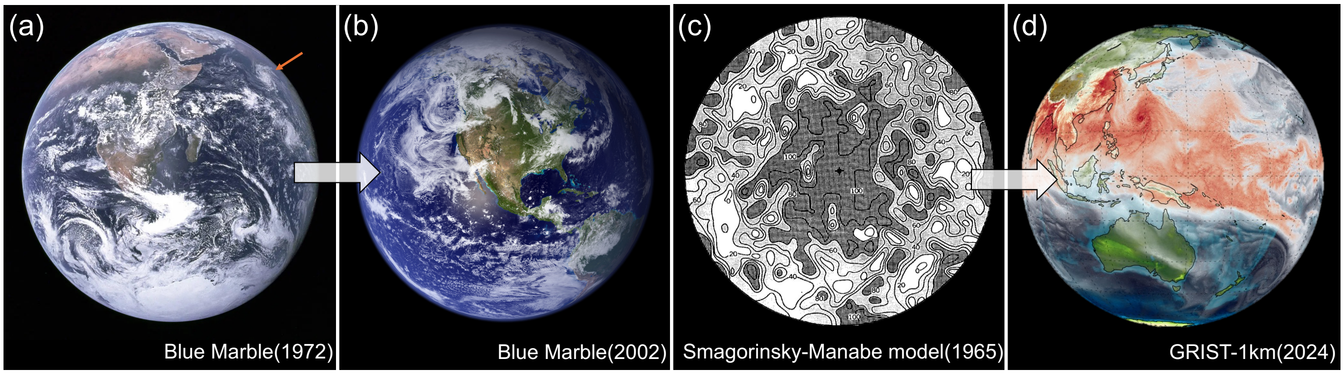

Figure 1. History and progress on observing and modeling the Earth system. (a) the very first Blue Marble photo taken by Apollo 17 astronauts in 1972 (with the red arrow pointing to the Tamil Nandu storm captured in the photo);(b) the 1km Blue Marble merged from MODIS satellite images in 2002; (c) the Smagorinsky-Manabe model [29] (500km resolution) result (relative humidity) in 1965; (d) example of 1km global simulation result (water vapor mixing ratio) of our GRIST model in 2024.

develop a global circulation models (GCM) with a resolution of around 500 km. Over the years, with the improvement of supercomputers' capacities, we see a gradual increase of the model resolution from the scale of 100 km to the scale of 10 km, with an expectation to finally achieve 1km global simulation with the emergence of exascale systems [23].

The benefits of 1km global simulations are multi-fold. Firstly, with a resolution finer than 5 km, we achieve a Global Storm-Resolving Model (GSRM) that can explicitly simulate non-hydrostatic fine-scale atmospheric fluid dynamics, enabling a seamless modeling approach where a single model effectively captures both transient and statistical features of the climate system. This provides the missing parts for hazard prediction and mitigation as in the Tamil Nadu case mentioned above. Secondly, while a 10km resolution already captures many vital atmospheric phenomena, such as mesoscale convective systems and tropical cyclones, an 1km resolution is crucial for accurately resolving smaller convective updrafts and the interactions between multiscale flow, thus removing uncertainties in empirical parameterizations of subgrid cloud and convection processes [30]. Thirdly, 1km modeling resolution (an example result of our work, as shown in Fig. 1(d)), would finally bridge the resolution gap between modeling and the satellite observation (such as the stitched MOIDS result in Fig. 1(b)). Such a seamless bridge, would eventually promote better approaches to fuse the numerical models and the data-driven learning methods [2, 13, 33], and gives better depiction and interaction between the short-term weather processes and the long-term climate processes [20, 25, 31].

While 1km resolution brings all the benefits mentioned above, the involved computational challenges are tremendous. A jump from 10 km to 1 km global simulation would already bring around three orders of magnitude increase in computation amount, eating up all the increased computing capacity from peta-scale machines to exa-scale systems. Moreover, the architectural changes from homogenous CPU clusters to heterogeneous machines with many-core AI-oriented accelerators make it tough to achieve efficient transition of scientific codes to a new system, especially for complicated community codes like climate models (with codes spanning through decades, and a distributed workload among hundreds of modules).

As a pioneering project to address the above challenges and to explore the potential of 1km global simulation, this work builds upon the Global-Regional Integrated Forecast System (GRIST), a unified weather-climate modeling system designed for research and operational applications [40–42]. Such a model enables us to successfully conduct long-term and multiple-resolution climate simulations and gain a reasonable performance in both physical accuracy and computational efficiency. Achieving year-scale simulation speed is of paramount importance for meeting the grand weather and climate prediction challenges. This unified hybrid modeling approach not only improves the weather-climate simulation and forecasting accuracy but also enhances our understanding of weather-climate processes, paving a way to more effective weather hazard mitigation and climate change adaptation strategies. Our contributions can be summed up as follows:

- We achieve an optimal mapping of the 1km GRIST model to the next-generation Sunway supercomputer. Our approach employs a scalable solver design and an OpenMP-based programming model, which facilitates automated porting, thereby maximizing computational efficiency and scalability.
- We integrate an AI-enhanced physics suite and utilizes mixed-precision techniques. These innovations close the performance gap and ensure that the model operates efficiently at such a large scale.

- We achieve a highly-scalable and highly-efficient hybrid global weather and climate model, which can employ up to 34 million cores to achieve 0.5 simulated-year-per-day (SYPD) for 1km resolution scenarios.

The remainder of this paper is structured as follows. Section 2 provides related works on porting climate models to supercomputers. Section 3 presents implementation details of our work. Section 4 evaluates the accuracy and efficiency of our implementation. Section 5 concludes the paper.

2 Related Works

Weather and climate models, being one of the first few applications run on the first electronic computer ENIAC [24], have long been the major applications on supercomputers. With the enormous yet complex Earth system as the simulation target, each step of resolution improvement would translate into 3 orders of magnitude in compute increase, and generally require decades' improvement in hardware and software evolution.

As a result, even though km-scale GSRM represents significant advancements in simulation detail and accuracy (reduction of convective parameterization uncertainties[21, 27], and more accurate representations of cloud-to-meso scale dynamics, precipitation patterns, and their interactions with atmospheric circulation [31]), GSRM efforts only become popular in the most recent decade, with the dawning of exascale systems, such as Frontier [32].

Fig. 2 summarizes most of the major GSRM efforts on recent supercomputer systems, such as E3SM on Summit [1] and Frontier [32], IFS on Summit [35], NICAM on K computer and Fugaku [28, 38], CESM on Sunway TaihuLight [6] and its successor [39], and COSMO (a regional model) on Piz Daint [9]. Note that, in our discussion, we focus on series of continuous efforts that try improving the model with the update of leading-edge supercomputers. The DYAMOND project [31] , an international GSRM model intercomparison initiative, provides a more complete list of modeling activities that target horizontal resolutions ranging from 1.5 km to 5 km, such as GEOS, gSAM, ICON, MPAS, NICAM, SHiELD, and GRIST.

Among the efforts, E3SM is the sole project that is specifically targeting an Exascale system, and has gone through years of development iterations [1, 11, 17, 32]. On the system of Summit, E3SM achieves a highly-scalable nonhydrostatic dycore at the resolution of 1 km and 3 km [1]. The dycore itself achieves a simulation speed of 0.97 SYPD at 3 km and 0.049 SYPD at 1 km. Following the dycore work, in 2023, Taylor et al.[32] presented a complete global atmoshpere model, the Simple Cloud Resolving E3SM Atmosphere Model (SCREAM) with a performance of 1.26 SYPD for 3.5km cloud-resolving simulations using 32768 GPUs on 8192 Frontier computing nodes, which won the 2023 Gordon Bell Prize for Climate Modelling.

On the Sunway series of supercomputers, there have also been a long-going project that tries to refactor [7], redesign [6] the CAM model for Sunway architectures, and eventually upgrade it for 5km-atmosphere and 3km-ocean coupled simulations [4], shown as a continuous line of efforts in Fig. 2. In contrast to the E3SM project, tremendous efforts on the tools were dedicated to migrating existing code to a new many-core architecture. The tools and a nearly complete porting of the model prove to be key elements for achieving close to 1 SYPD speed of CAM.

NICAM is another high-resolution effort on both K Computer [28] and Fugaku [38]. Through a careful redesign for the transition from K to Fugaku, NICAM scaled to 512 nodes with a performance of 0.089 and 0.027 SYPD for 14km and 3.5km resolution cases respectively[38]. As the work on Fugaku focused on large-scale ensemble runs, NICAM itself was not running on a large scale. We can further project the performance of 3.5km resolution case over 131,072 nodes of Fugaku to 0.36 SYPD.

ICON-Sapphire achieved a simulation speed of 4 SDPD with a 1.25km resolution case of atmospheric component over 908 nodes of Levant Supercomputer[15], but it only retained parameterizations for the physical processes of radiation, microphysics, and turbulence, impacting its accuracy at finer resolutions[5, 34]. The atmospheric component of ICON has also been tuned and evaluated over GPU architecture[10] and demostrates the performance of 0.58 SYPD for the 5km resolution case with 256 nodes of JUWELS Booster Supercomputer.

The Consortium for Small-Scale Modeling (COSMO) model is the first regional weather model that gets fully migrated to a GPU platform and achieves a performance of 0.043 SYPD for a near-global 1km resolution configuration by using up to 4,888 GPUs on the Piz Daint supercomputer [9]. We show the line of different COSMO configurations in Fig. 2, which demonstrates fairly efficient performance on the 26-Pflops Piz Daint system, due to a reduced complexity in regional forecasting, and a well-tuned GPU implementation.

A recent work demostrates that the CPU-version of ECWMF IFS can be scaled to all the CPUs of the entire Summit supercomputer with a performance of 0.3 SYPD for a global 1.4km resolution configuration with hydrostatic configuration (0.09 SYPD with non-hydrostatic configuration on the PizDaint supercomputer) [35].

As shown in these different projects demonstrated in Fig. 2, moving up the models in resolution and simulation speed, even with the world's fastest computers, is still an extremely tough challenge, which involves a wide range of issues including scaling, code porting, architecture mapping, as well as the development of science itself.

Besides, a recent work ports an ocean model, LICOM, to the Sunway platform with a port of Kokkos performance portable framework [36]. This work translated hotspots from

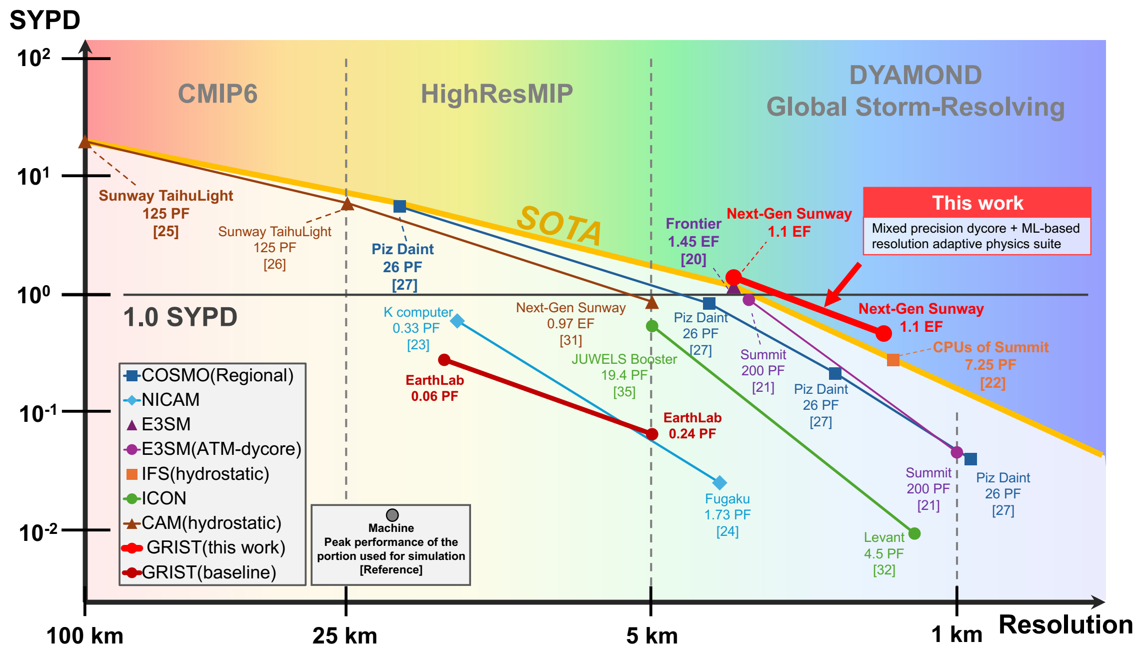

Figure 2. A summary of recent high-resolution weather and climate modeling efforts on supercomputers: a continuous journey towards affordable global storm-resolving modeling.

Fortran to C++, which is the prerequisite of applying Kokkos-based parallelization. In contrast, atmosphere models has significantly larger codebase, using Kokkos requires heavy human efforts to translate the code. So our work employs OpenMP offloading, which can parallelize Fortran code with CPEs by simply adding directives.

In contrast to existing efforts, our AI-enhanced GRIST model is, as far as we know, one of the first efforts that try to combine a highly-scalable mixed-precision dynamical core, with a machine learning(ML)-based resolution-adaptive physics suite, for an 1km global simulation. As shown in Fig. 2, built upon the previous GRIST work on CPUs (0.07 SYPD with a 5km resolution case when using 30,720 cores of Chinese EarthLab supercomputer), the AI-enhanced GRIST has made a concrete step towards year-scale speed for GSRM (1.35 SYPD for 3km global simulation, and 0.5 SYPD for 1km global simulation), and also has demonstrated a clear improvement in the efficiency of converting computing performance into simulation capabilities (in terms of ratio between peak performance utilized, and the simulation speed achieved). Our approach is innovative in three aspects: coarse-grained strategies, multi-variable machine learning models, and the GSRM dataset. Moreover, for the first time, the ML-based resolution-adaptive physics suite improves a GSRM in terms of the unified modeling capability that can approach the fast and reasonable weather and climate simulations across resolutions.

3 Implementation

3.1 An AI-Enhanced GRIST Model

3.1.1 Architecture of the Hybrid Model. Fig. 3 illustrates the architecture of the GRIST AI-hybrid model, which comprises two primary components: the dynamical core and the physics suite. The dynamical core operates on a globally decomposed horizontal mesh and is numerically resolved. The physics suite includes the conventional parameterization modules, as well as an ML-based physics suite, which leverages multiscale interactions distilled from global storm-resolving simulations to improve the coarse-resolution hydrostatic-scale simulations, reduces computational demands and addresses the load balancing issues. A detailed description about the ML physics suite is provided in section 3.2.

3.1.2 Highly-Scalable Solver for the Dynamical Core. The GRIST nonhydrostatic model is driven by its layer-averaged unstructured hexagonal-C grid dynamical core [41]. The horizontal domain decomposition, facilitated by the METIS library [16], optimizes both load balancing and scaling. A horizontally explicit and vertically implicit approach is used to discretely solve the nonhydrostatic compressible equation set, requiring minimal data exchange procedures across the horizontal computations without the need for global communication. The discretization employs the staggered finite-volume method, approximately second-order, leading to moderate computational load for basic operators (e.g., divergence and gradient). The dynamical core utilizes MPI+OpenMP

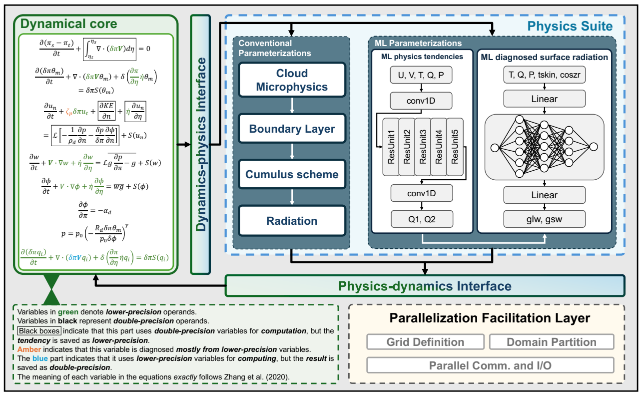

Figure 3. The architecture of our AI-enhanced GRIST model.

programming model, which is further extended to the Sunway architecture (section 3.3). A mixed-precision dynamical core (Fig. 3) has been developed to reduce the computational load, which maintains the stability and accuracy of the original double-precision code based on a careful iterative development procedure and a hierarchy of tests. This mixed-precision dynamical core is detailed in section 3.4.

3.1.3 Parallelization Facilitation Layer.

To provide a consistent support of the solver across different degrees of parallelism, we include a parallelization facilitation layer to handle mapping of grids, data communication, and parallel I/O.

To support a smooth and balanced mapping of the unstructured grid, we perform the mapping through indirect addressing, and optimize the index sequence using the breadth-first-search method to enhance the cache hit rate. To refine the granularity of data exchange and minimize inter-process communications, a linked list is utilized to gather variables for exchange, and a single call to the communication interface efficiently completes the data exchange for all listed variables. In such a way, the model has achieved approximately 83% parallel efficiency scaling from 1920 to 30720 CPU cores. Lastly, a grouped parallel I/O strategy was designed and implemented to ensure efficient data I/O across a large number of MPI processes.

3.2 A Resolution-adaptive ML-based Physics Suite

3.2.1 Generation of Training Data.

Due to the extremely expensive cost of 1km simulation, in our methodology design process, we utilize hourly generated 5km GRIST-GSRM data for training and testing our ML-physics suite. To ensure stability and high precision, experiments are conducted over 80 days, with 20 days to cover each of the four seasons, and varying ENSO and MJO events across different cases. To mitigate overfitting, the testing set consists of three randomly selected time steps per day, while the remaining time steps are allocated for training, maintaining a training/testing ratio of 7:1. The details of data are listed int Table 1.

Table 1. Selected time periods and climate characteristics

Time period	Oceanic Niño Index	Real-time Multivariate MJO index
1-20 January 1998	2.2(El Niño)	0.69 to 1.98
1-20 April 2005	0.4(neutral)	2.72 to 3.71
10-29 July 2015	-0.4(neutral)	0.17 to 1.05
1-20 October 1988	-1.5(La Niña)	0.67 to 2.98

3.2.2 Selection of Physical Tendency Outputs.

Before training the ML-physics suite, it is crucial to carefully select the coarse graining method and output variables for the parameterization. We introduces a novel approach by using residual calculations to derive Q_1 and Q_2 as outputs for our ML-based parameterization physics suite, replacing traditional methods for calculating all temperature and humidity tendencies caused by physical processes. Zhang et al. [43] have indicated that Q_1 and Q_2 calculated from coarse-grained 5km GRIST-GSRM data using the residual method are essentially compatible to theory. This provides significant potential for our resolution-adaptive physics package. Our experiments demonstrate that the ML-based physics suite trained with 30km coarse-grained data perform well not only at G8 (Table 2) resolution but also at G6 (Table 2) resolution.

3.2.3 Configuration of the ML-based Physics Suite.

While Q_1 and Q_2 can replace the sum tendencies of all physical processes, to ensure the stable and efficient operation of the model, it is necessary to additionally predict some of the diagnostic processes within the physical parameterizations. Through several trials, we find that the radiation diagnostic scheme has the most significant impact on both model stability and operational efficiency. Therefore, we separately construct the tendencies of all physical processes (ML physical tendency module) and the radiation diagnostics (ML radiation diagnostic module) to enhance the stability, accuracy, and efficiency of the machine learning physics suite.

The structure of the ML tendency module is shown in Fig. 3, which employs one-dimensional convolutional layers to capture the vertical characteristics of temperature, humidity, and other atmospheric variables during events such as convection and atmospheric instability. With the incorporation of residual connections, this structure is proven to be stable and accurate [12]. To balance computational efficiency with performance, the module incorporates five ResUnits, culminating in an 11-layer deep Convolutional Neural Network (CNN) with a parameter count close to half a million.

For the ML radiation diagnosis module, we additionally train a deep neural network to compute surface downward shortwave radiation (gsw) and longwave radiation (glw), which are provided to the land surface model and surface layer scheme. In order to mimic the radiation process, we add skin temperature (tskin) and cosine of solar zenith angle (coszr) as inputs to provide physical features of the model top insolation and surface state [33]. This can improve the stability of ML radiation diagnostic module coupled with physical diagnosis modules. Because the radiation diagnostic variables are only scalars for the model surface, we introduce a 7-layer Multilayer Perceptron (MLP) with residual connections to process one-dimensional vector computation. It can significantly improve computational efficiency by replacing conventional radiative transfer calculations with continuous matrix multiplication.

3.2.4 Coupling the ML-based Physics Suite with the Dynamical Core.

The online coupling process involves computing the dynamical core and passing input variables (U, V, T, Q, P, tskin, coszr) from the physics-dynamics coupling interface of GRIST model to our trained ML-physics suite, which includes an ML physical tendency module, an ML radiation diagnostic module, and a conventional physics diagnostic module. They together form the new model physics suite, which returns full physical tendencies and diagnostic variables back to the physics-dynamics coupling interface of GRIST for the next-step dynamical core integration. Compared to the conventional physics suite, the ML-physics suite offers significant computational advantages: it features a simplified, unified computational pattern (primarily matrix multiplication), reduced code complexity, ease of optimization, and greater flexibility for adaptation across different architectures.

3.3 A Generalized Programming Model for Sunway Architecture

Our work is based on the next-generation Sunway supercomputer that succeeds TaihuLight [8]. To address the huge amount of efforts required for porting and tuning the various kernels of the 272,000 lines of code in GRIST, we develop a generalized programming model for the upgraded SW26010P processor. One advantage of SW26010P, as compared with SW26010 in TaihuLight, is that we have a 256KB local device memory (LDM), half of which can be configured as a one-level 4-way group associated cache (LDCache), and resolves the issues of writing manual DMA instructions in previous SW26010 cases.

In contrast to SW26010, the computing power of SW26010P is also increased from 4 core groups (CGs) to 6 CGs, each of which still consists of one management processing element (MPE) and 64 computing processing elements (CPEs) organized as an 8×8 array, in total 390 cores per processor. Each CG has a DDR4 shared main memory of 16 GB with a bandwidth of 51.2 GB/s.

Previously, SW26010 and SW26010P faced compatibility issues due to their distinct architectures. However, in this work, we try to achieve a balance between performance and portability. To maintain the generalized MPI+OpenMP programming model in the GRIST model, we have developed SWGOMP, a compatibility layer for OpenMP applications

```
!$omp target !Just add this
!$omp parallel private(ie,v1,v2,ilev)
!$omp do
   do ie = 1, mesh%ne
      v1      = mesh%edt_v(1, ie)
      v2      = mesh%edt_v(2, ie)
      do ilev = 1, nlev
         tend_grad_ke_at_edge_full_level(ilev,ie) = &
         -edt_edpNr_edtTg(ie)*(kinetic_energy(ilev,v2) &
         -kinetic_energy(ilev,v1))/(rearth*edt_leng(ie))
      end do
   end do
!$omp end do nowait
!$omp end parallel
!$omp end target !and this, and enjoy CPEs

!$omp target parallel workshare !or for fortran arrayop
kinetic_energy(:,:) = 0
!$omp end target parallel workshare
```

Figure 4. An example of swgomp usage. In parallel region, loops can be distributed to CPEs via `!$omp do`, and array operations can be distributed via `!$omp workshare`.

running on Sunway architecture. SWGOMP is compatible with the OpenMP 4.5 specification and includes a backport of the unified shared memory feature from OpenMP 5.0, aligning with the shared main memory configuration of SW26010P.

3.3.1 Implementation of OpenMP Offload.

SWGOMP is developed as a compiler-plugin-based tool, enabling the adoption of code parallelized in the OpenMP Offloading scheme. Also, non-Offloaded code in OpenMP can be easily ported to CPEs by simply adding `!$omp target` directive and leveraging the unified shared main memory feature to eliminate the need for data copy directives. An example is shown in Fig. 4.

To align with the "teams parallel" scheme of OpenMP Offload, we have developed a flexible thread launching scheme built on top of the Athread library, as shown in Fig. 5. The job server exhibits a high flexibility, allowing new tasks to be assigned to CPE by either the MPE or another CPE. The job server is initialized by MPE using the Athread library. The MPE spawns team-head threads via the job server to execute target portions. These team-head CPEs have the capability to spawn threads on other CPEs within the team to execute parallel code pieces.

3.3.2 Enabling LDM and DMA Usage via `omnicopy`.

The current design of SW26010P memory hierarchy leaves half of the 256KB LDM as cache, the other half as user-programmable buffer to facilitate fine-grained optimizations. While such fine-grained tuning brings significant performance benefits in certain cases, it violates our target to keep a generalized OpenMP model, and reduces the scope of the code that our work could cover. Therefore, to further utilize the rest 128KB LDM, we use the `device` clause to enable functions to allocate their stack and private variables in LDM, and implement a cross-platform `omnicopy` function as a replacement for `memcpy`. This function can determine whether

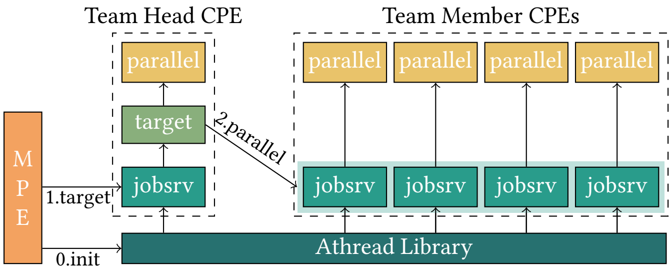

Figure 5. The job spawning hierarchy of SWGOMP. Job servers are initiated on CPEs through the Athread library. MPE can launch target portions on team head CPEs via the job server, while team head CPEs can further distribute parallel tasks to team members via the job server.

data transfer occurs between main memory and LDM, utilizing DMA automatically when feasible. On non-Sunway platforms, `omnicopy` functions identically to `memcpy`, ensuring compatibility of the optimized code with other platforms.

3.3.3 Memory Address Distribution for LDCache.

We have noted a significant performance decline in specific OpenMP parallelized kernels, attributed to elevated LDCache miss ratios in some specific processes. Investigation revealed that many of these kernels access more than four arrays within a single loop, surpassing the number of LDCache ways. Arrays, when well-aligned to a size larger than one cache way and accessed with similar indices, are mapped to the same cache lane, leading to cache thrashing, as depicted in Fig. 6(a). To address this issue, we have implemented a memory-address-distributor enabled pool-based memory allocator to replace the original `malloc` function. This allocator ensures that the starting addresses of arrays are uniformly distributed across cache lanes, thus improving cache performance, as illustrated in Fig. 6(b).

3.3.4 Applying OpenMP Offload in GRIST.

GRIST uses finite volume scheme in dynamics core and column model in physics scheme. Most of loops are conflict-free, so adding directives like Figure 4 is enough for the parallelization on CPEs. Loops with a symmetric update on a vertice and its neighbor are split into a edges-from-vertices loop and an vertices-from-edges loop. For loops identified with cache thrashing, we copy a number of variables onto CPE stack with `omnicopy` function until the cache thrashing is eliminated.

3.4 Optimization with Mixed Precision Computing

With a maximum of 4 times performance benefits on SW26010P processors, mixed-precision computing is another approach to close the gap towards year-scale simulation speed when performing 1km global simulation.

Traditionally, weather and climate models make use of double-precision (64-bit) floating-point arithmetic. Converting them to single-precision (32-bit) arithmetic offers a promising avenue for improving computational efficiency. However,

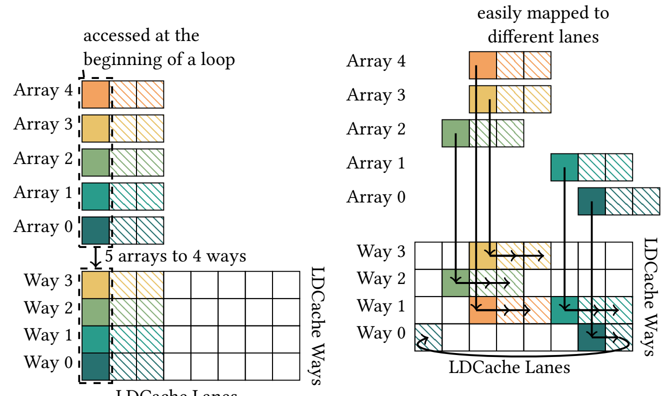

(a) Without distribution, arrays accessed with similar indices may be mapped into the same cache lane. If the number of arrays exceeds the number of cache ways, at least one array will be swapped out. However, it will be accessed again soon, leading to cache thrashing.

(b) With distribution, array starting addresses are distributed into different cache lanes, reducing the risk of swapping out array data in the cache. During loops, the mapped cache addresses move along with the loop index increase, thereby reducing the risk of cache thrashing.

Figure 6. LDCache mapping with and without address distribution.

using single precision for all calculations in our model leads to unacceptable loss of accuracy. Since exploiting a mixed-precision scheme for ML-based parameterizations is straightforward at the operator level due to the model's compact design (refer to Section 3.2 for details), we focus on reducing variable precision within the dynamical core, which accounts for approximately 60% to 80% of the total runtime in high-resolution scenarios.

3.4.1 Overall evaluation metric.

To investigate the delicate balance between precision and performance, we designate surface pressure (*ps*) and relative vorticity (*vor*) as pivotal observation points for tracking deviations within the mass and velocity fields, which serve as indicators of the overall status and the regional dynamics respectively. Subsequently, we gauge error discrepancies resulting from varied precisions using the relative L_2 norm. This metric harmonizes with existing convergence criteria within the model, streamlining the process of comparing precision reductions effectively.

For accuracy evaluation, we use the original double-precision outcomes as the gold standard. Through experiments, we establish a 5% error threshold to ensure the dynamical core's reliability during extended simulations in this investigation.

3.4.2 Identification of Insensitive Variables.

We evaluate all six prognostic equations (Fig. 3) in the GRIST dynamical core and identify viable single-precision or mixed-precision terms. This process reveals that terms associated with pressure gradient and gravity exhibit notable precision sensitivity. Most advective terms utilized in high-order operators demonstrate precision insensitivity, yet they contribute significantly to computational load. Reducing their precision results in substantial speedup for the dynamical core. The passive tracer transport equation (depicted at the bottom of the left panel in Fig. 3) encompasses six prognostic variables. This equation can be computed almost entirely using lower precision. The sole exception is the mass flux $\delta\pi V$, which is accumulated from the dry mass equation (illustrated at the top of the left panel in Fig. 3) and requires double precision information. We have performed a hierarchy of tests ranging from idealized tropical cyclone, supercell, baroclinic waves to real-world long-term climate simulations and GSRM simulations. The stability and accuracy of the mixed-precision code remain robust in all the tests.

3.4.3 Implementation of Mixed-Precision Schemes.

We employ a custom Fortran type, designated as `ns`, to efficiently manage precision switching for insensitive variables. When `ns` is configured to lower precision, the code seamlessly conducts mixed-precision computations; otherwise, it executes the original code unchanged in double precision.

To ensure a smooth transition to mixed precision with minimal alterations to the code structure, only the solver subroutines undergo modification, while the model initialization remains in double-precision. Moreover, if the solver necessitates single-precision operands, double-precision variables are converted to single precision after initialization. Exploiting on-the-fly precision conversion alongside, SWGOMP entails no additional modifications, and further reduces the memory footprint on the Sunway platform.

4 Evaluation

4.1 Hardware Platform Details

Our performance evaluations have been conducted on the next-generation Sunway supercomputer [18] that succeeds TaihuLight. The next-generation Sunway supercomputer has more than 107,520 nodes. Each node has one SW26010P 390-core CPU, resulting in a parallel scale of 41,932,800 cores. In addition, each node in the next-generation 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.

Table 2. Configuration of our grids and timesteps

Label	Resolution (km)	Layers	Timestep				Number of		
			Dyn	Trac	Phy	Rad	Cells	Edges	Vertices
G12	1.47~1.92	30	4	30	60	180	167M	503M	336M
G11W	2.94~3.83	30	4	30	60	180	41.9M	126M	83.9M
G11S	2.94~3.83	30	8	60	120	360	41.9M	126M	83.9M
G10	5.87~7.66	30	4	30	60	180	10.5M	31.5M	21.0M
G9	10.6~14.6	30	4	30	60	180	2.62M	7.86M	5.24M
G8	21.7~28.4	30	4	30	60	180	655K	1.97M	1.31M
G6	92.5~113.	30	4	30	60	180	41.0K	123K	81.9K

4.2 Configuration of Modeling Experiments

With successful G12 (30 Layers) scientific simulations using the conventional physics suite, we have evaluated GRIST under different grids and scheme combinations. The grid and timestep configurations of our experiments are shown in TABLE 2. Here G11 has two timestep configurations: G11W is for evaluating weak scalability, it uses the same timestep as G12, and G11S is used to test strong scalability with its largest possible timestep. The configurations mainly vary in different precisions of the dynamical core and physical schemes, which are shown in TABLE 3. For weak scaling tests, we use the same timestep configuration as the G12L30 case to keep the computational cost only related to the number of grids.

Table 3. Configuration of our schemes

Label	Dycore	Physics
DP-PHY	double precision	Conventional
DP-ML	double precision	ML-physics
MIX-PHY	mixed precision	Conventional
MIX-ML	mixed precision	ML-physics

4.3 Major Performance Metrics

In terms of timing, the performance is measured using the average time recorded by running the same case for three times. For the climate modeling scenario, the simulation speed that we can achieve is obviously a more important metrics to consider. For most performance results, we describe the speed of simulation using SDPD (simulated-days-per-day). Note that, these metrics need to be evaluated jointly with the corresponding resolution configurations.

4.4 High-Resolution Simulations Using 1km GRIST

The highest resolution of the GRIST model is configured to icosahedral Grid level 12 (G12, 1.47 to 1.92 km). Prescribed sea surface temperature and sea ice concentration data and an active land surface model [22] have been coupled to the atmosphere model. The atmosphere and land initial field data are taken from ERA5 [14]. The model top is kept the same in all setup (2.25 hPa; ~40 km).

To demonstrate that the model configuration works reasonably, we select a "23.7" extreme rainfall case over North China in 2023. Super Typhoon Doksuri moved northward

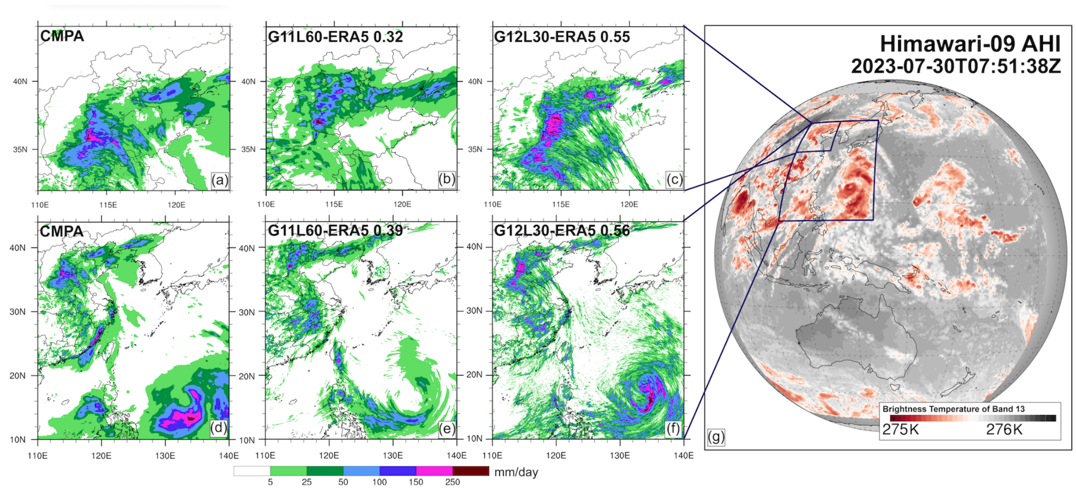

Figure 7. Super Typhoon Doksuri and the "23.7" extreme rainfall event over North China. The left panel shows the mean rainfall rate (mm/day) during UTC00, 29th-UTC00,30th, July, 2023, for CMPA, G11L60 and G12L30, the right panel shows the cloud top temperature (unit: K).

and weakened the influence of low pressure, and an extremely rare heavy rainstorm weather process occured over North China. We ran two GRIST-GSRM simulations, G11L60 (G11, 2.94 to 3.83 km, 60 vertical layers) and G12L30 (G12, 30 vertical layers). Horizontal resolution is vital to this simulation. As shown in Fig. 7, compared with G11L60, G12L30 better simulates the Typhoon rain band, and the extreme rainfall magnitude over North China, closer to that in the CMPA observational data, as quantified by G12L30's higher spatial correlation coefficients.

4.5 Evaluation of the ML-based Physics Suite

For the ML-based parameterization, limited to the computational resource, we firstly derive from the 5km GRIST modeling results (using conventional physics) to produce a coarse-grained 30km modeling data. We obtain the functional relationship between dynamical inputs and Q_1, Q_2. While we plan to feed more training data at different resolutions in the future to further improve the performance of the AI part, the workflow of the AI-enhanced GRIST model remain consistent across all resolutions. Therefore, even though the lack of high-resolution input might impair the modeling accuracy in 3km or 1km cases, the computational performance and efficiency of the ML-based physics suite would be valid in our scaling tests.

Physically, the currently trained ML-physics suite is in principle only justified for modeling at similar horizontal resolutions (e.g., G8; Table 2). Experiments indicate that it also works for lower resolutions (e.g., G6; Table 2), probably because a 30km grid serves as a sub-grid to a 120km grid. Fig. 8 shows the annual mean precipitation rate over United States from one-year climate simulation. The AI-enhanced GRIST also demonstrate stable runs for at least 10 years. Both G6 and G8 with the ML-based physics suite generates the

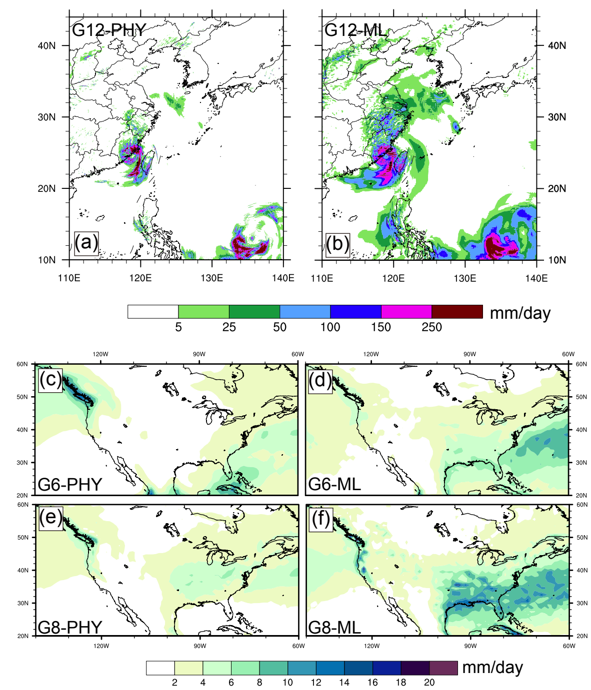

Figure 8. (a) (b) predicted rainfall rate during 3-hour integration simulated by GRIST with the conventional and ML-based parameterization; (c)(d) one-year annual mean rainfall rate (mm/day) over North America simulated by G6 with conventional and ML-based parameterization; (e)(f) same as (c)(d) but for the G8 model resolution.

observed rainfall band well. GRIST with the conventional physics suite can only simulate this rainfall band at G8. While using the present ML suite is theoretically problematic for resolutions higher than 30 km, we attempt to perform G12 (1 km) experiments , and obtain reasonable short-term (3-hour) weather simulations (Fig. 8). We anticipate that further training based on GRIST results of a varied set of resolutions will produce a more physically accurate and resolution-adaptive ML-based parameterization suite.

4.6 Accelerations of Major Kernels Over Multiple CPEs

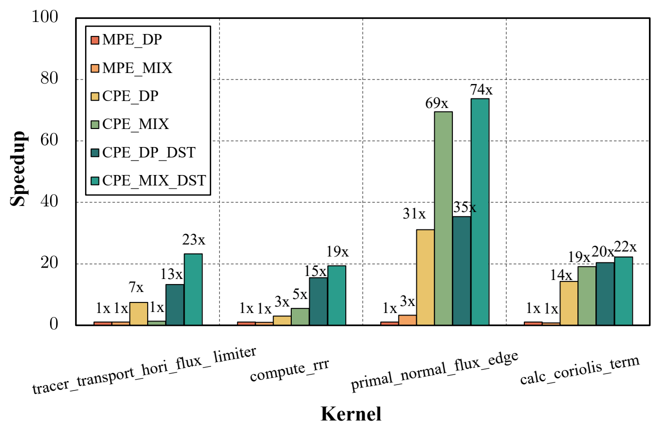

Figure 9. Performance improvements on CPEs for major kernels, DP and MIX represents double precision and mixed precision, DST means the adoption of the memory address distribution strategy.

We have evaluated the acceleration of major kernels over 64 CPEs wthin a CG, under G6 grid with DP-PHY and MIX-PHY configurations in one node, with results shown in Fig. 9. Kernels such as tracer_transport_hori_flux_limiter and compute_rrr feature mixed precision optimization and involve a large number of arrays, showcasing clear speedup with mixed precision and address distribution on CPEs. On the other hand, primal_normal_flux_edge involves numerous division, power, and other computationally expensive calculations, resulting in significant mixed precision speedup. However, calc_coriolis_term, lacking mixed precision optimization and accessing relatively few arrays, derives minimal benefit from mixed-precision and memory address distribution.

A notable observation is that mixed precision typically does not yield significant speedup on the MPE side but provides notable speedup on CPE-parallelized kernels. Considering that the Sunway architecture generally does not exhibit higher calculation performance in single precision compared to double precision, except for division and elemental functions, we can infer from the results that the MPE code is computation-bound. On CPEs, mixed precision code demonstrates better speedup. One possible reason is that CPE code appears to be constrained by memory bandwidth, and mixed precision reduces data size, conserving memory bandwidth and increasing cache hit ratio.

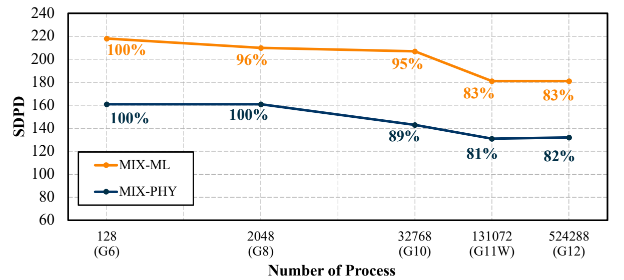

Figure 10. Weak scaling results of the model. Our weak scaling test starts from 128 processes (corresponding to 128 CGs, and 8,320 cores), and scales till 524,288 processes (corresponding to 524,288 CGs, and 34,078,720 cores).

4.7 Weak Scaling

The weak scaling is evaluated under mixed precision, and the results are shown in Fig. 10. All grids use the same timestep settings as G12. Taking 128 core groups running G6 test case as a baseline, all test cases has almost the same number of vertices per core group. Thus, the weak scaling efficiency is calculated by (1):

$$eff_{weak}(N) = \frac{P_N}{P_{128}} \tag{1}$$

Where P_n is the time-to-solution (SDPD) at n processes. The proportion of communication time rises from 19% to 37% as the number of processes increases. This increase is attributed to both the growing number of communicating processes and the computational load unbalance distributed among them. We observe a clear drop of scalability a the scale of 32,768 CGs, possibly due to bandwidth oversubscription in the fat-tree network topology.

Additionally, we observe that the AI-enhanced model (MIX-ML) outperforms the one with conventional parameterizations (MIX-PHY). This does not necessarily imply that the AI-enhanced version reduces computational workload. Instead, we expect learning-based model to exhibit better computation efficiency due to improved memory access patterns and computation vectorization. For instance, ML diagnosed surface radiation requires approximately twice the number of FLOPS operations compared to RRTMG. However, it can achieve peak FLOPS ranging from 74% to 84% during computation, a significant improvement over the 6% in RRTMG, resulting in a substantial improvement of modeling speed.

4.8 Strong Scaling

We have evaluated all configurations of the G12 (1.47~1.92 km) and the MIX-ML configuration of G11S (2.93~3.83 km) for strong scaling, with the number of processes ranging from 32,768 to 524,288. The former value is the minimal number process that can run a G12 test case and the latter value is the largest integral power of 2 less than the total CG number of the next-generation Sunway supercomputer. Since the total workload is fixed, the strong scaling efficiency

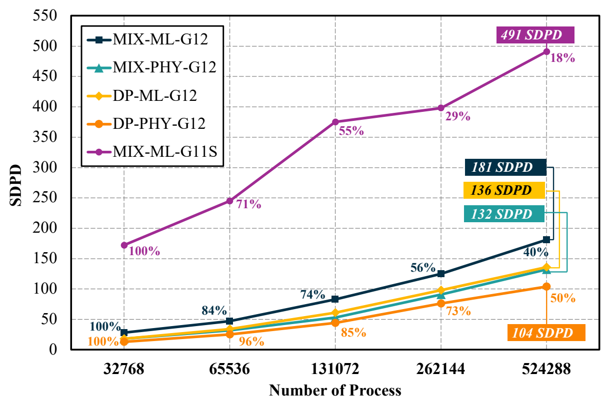

Figure 11. Strong scaling results of the model. Our strong scaling test starts from 32,768 processes (corresponding to 32,768 CGs, and 2,129,920 cores), and scales till 524,288 processes (corresponding to 524,288 CGs, and 34,078,720 cores).

is evaluated as (2):

$$eff_{strong}(N) = \frac{\frac{P_N}{N}}{\frac{P_{32768}}{32768}} \tag{2}$$

Where P_n is the time-to-solution (SDPD) at n processes. The strong scaling results are shown in Fig. 11. In the G12 cases, we have observed a continuous decrease in scalabing efficiency, although the rate of decline decreases over time. This trend can be attributed to the drop of cache hit ratio as the number of processes increases significantly. The G11S grid test demonstrates a marginal increase in computation speed as the number of CGs increases from 131,072 to 262,144. It is worth noting that 131,072 CGs may represent the first plateau in cache hit ratio. However, when the CG count reaches 524,288, the LDCache demonstrates the potential to accommodate several arrays, resulting in another increment in computation efficiency. Ultimately, we achieved a performance of 491 SDPD for the G11S grid and 181 SDPD (around half year per day) for the G12 grid with 524,288 processes.

5 Conclusion

This paper reports our efforts towards 1km global seamless weather and climate model, which is a common goal that the weather and climate community have been exploring throughout the recent decade.

Facing the enormous computing cost of 1km global simulation, we manage to (1) achieve a highly-scalable version of GRIST that can eventually utilize the 34 million cores of the next-generation Sunway Supercomputer; (2) employ AI to achieve more efficient description of the physical parameterizations for different resolutions; (3) derive an OpenMP-compatible programming model, which can map almost all compute kernels to 64 CPEs in a CG, with no extra code needed from the original OpenMP directives; (4) perform effective mixed-precision optimization to further improve the computing speed of the dynamical core part.

With all the above efforts combined, our AI-enhanced GRIST model can finally perform 1km global simulation with an unprecedented speed of 0.5 SYPD, touching the bar of one SYPD and opening the potential chances of performing practical seamless weather and climate studies. Such a model also facilitates us to try performing 1km simulation for extreme weather events, such as the super Typhoon Doksuri. One potential finding from such extreme-scale experiments is that the increase of horizontal resolutions seem to be far more important than the increase of vertical levels (as shown in Fig. 7).

Another important progress that we have made is an AI-enhanced GRIST model that takes a hybrid approach to concretely fuse numerical schemes and data-driven machine learning schemes. Besides the competitive performance in both short-term high-resolution simulation and long-term climate simulation provided by the hybrid model, the resolution-adaptive ML-based parameterization has also demonstrated its scalability and stability across multiple scales, from G6 (92 to 113 km) to G12 (1.47 to 1.92 km). This success not only proves the large-scale applicability of AI in scientific contexts but also opens new pathways for innovative modeling techniques of unified weather-climate modeling in terms of accuracy and speed. As AI continues to evolve, its role in shaping the future of earth system modeling would grow steadily.

The last point we make here is the importance of tools. The prosperity of AI-related research in the recent decade is a typical example, showing that the open-source and flexible deep learning frameworks and data ecosystem play an key role in expanding the scope and population of AI. In contrast, climate modeling, which is both a science and an engineering domain, requires even better tools. In our work, SWGOMP enables a feasible porting of GRIST to a new supercomputer, thus expanding its resolution and speed to an unprecedented level, and incurring interesting research projects to follow. In the near future, we expect that software tools that facilitate extreme-scale simulation, large-scale data-driven learning, and their fusion would be the key to the next generation weather and climate modeling systems.

Acknowledgments

This study was supported by the National Key Research and Development Plan of China (Grant No. 2023YFB3002400), the National Natural Science Foundation of China (Grant No. 92358302, Grant No. U2242210, Grant No. T2125006, Grant No. 62402333) and Jiangsu Key R&D Program (No.BE2022050-1). Computing resources are partially supported by Laoshan Laboratory Project (No.LSKJ202300301).

References

[1] 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.

[2] Christopher S Bretherton, Brian Henn, Anna Kwa, Noah D Brenowitz, Oliver Watt-Meyer, Jeremy McGibbon, W Andre Perkins, Spencer K Clark, and Lucas Harris. 2022. Correcting coarse-grid weather and climate models by machine learning from global storm-resolving simulations. *Journal of Advances in Modeling Earth Systems* 14, 2 (2022), e2021MS002794.

[3] PK Das, CA George, and R Jambunathan. 1973. Cyclones and drpressions of 1972-Bay of Bengal and Arabian Sea. *MAUSAM* 24, 4 (1973), 327–344.

[4] Xiaohui Duan, Yuxuan Li, Zhao Liu, Bin Yang, et al. 2023. Achieving Kilometer-Level Coupled Modeling on a 40-Million-Core Sunway Supercomputer: A Three-Week Porting Experiment. In *HPC China 2023*. CCF, 25–35.

[5] Saulo R. Freitas, William M. Putman, Nathan P. Arnold, David K. Adams, and Georg A. Grell. 2020. Cascading Toward a Kilometer-Scale GCM: Impacts of a Scale-Aware Convection Parameterization in the Goddard Earth Observing System GCM. *Geophysical Research Letters* 47, 17 (2020), e2020GL087682. https://doi.org/10.1029/2020GL087682

[6] Haohuan Fu, Junfeng Liao, Nan Ding, Xiaohui Duan, Lin Gan, Yishuang Liang, Xinliang Wang, Jinzhe Yang, Yan Zheng, Weiguo Liu, et al. 2017. Redesigning CAM-SE for peta-scale climate modeling performance and ultra-high resolution on Sunway TaihuLight. In *Proceedings of the international conference for high performance computing, networking, storage and analysis*. 1–12.

[7] Haohuan Fu, Junfeng Liao, Wei Xue, Lanning Wang, Dexun Chen, Long Gu, Jinxiu Xu, Nan Ding, Xinliang Wang, Conghui He, et al. 2016. Refactoring and optimizing the community atmosphere model (CAM) on the sunway taihulight supercomputer. In *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 969–980.

[8] Haohuan Fu, Junfeng Liao, Jinzhe Yang, Lanning Wang, Zhenya Song, Xiaomeng Huang, Chao Yang, Wei Xue, Fangfang Liu, Fangli Qiao, et al. 2016. The Sunway TaihuLight supercomputer: system and applications. *Science China Information Sciences* 59 (2016), 1–16.

[9] Oliver Fuhrer, Tarun Chadha, Torsten Hoefler, Grzegorz Kwasniewski, Xavier Lapillonne, David Leutwyler, Daniel Lüthi, Carlos Osuna, Christoph Schär, Thomas C Schulthess, et al. 2018. Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0. *Geoscientific Model Development* 11, 4 (2018), 1665–1681.

[10] M. A. Giorgetta, W. Sawyer, X. Lapillonne, P. Adamidis, D. Alexeev, V. Clément, R. Dietlicher, J. F. Engels, M. Esch, H. Franke, C. Frauen, W. M. Hannah, B. R. Hillman, L. Kornblueh, P. Marti, M. R. Norman, R. Pincus, S. Rast, D. Reinert, R. Schnur, U. Schulzweida, and B. Stevens. 2022. The ICON-A model for direct QBO simulations on GPUs (version icon-cscs:baf28a514). *Geoscientific Model Development* 15, 18 (2022), 6985–7016. https://doi.org/10.5194/gmd-15-6985-2022

[11] Jean-Christophe Golaz, Luke P Van Roekel, Xue Zheng, Andrew F Roberts, Jonathan D Wolfe, Wuyin Lin, Andrew M Bradley, Qi Tang, Mathew E Maltrud, Ryan M Forsyth, et al. 2022. The DOE E3SM Model Version 2: overview of the physical model and initial model evaluation. *Journal of Advances in Modeling Earth Systems* 14, 12 (2022).

[12] Yilun Han, Guang J Zhang, Xiaomeng Huang, and Yong Wang. 2020. A moist physics parameterization based on deep learning. *Journal of Advances in Modeling Earth Systems* 12, 9 (2020), e2020MS002076.

[13] Yilun Han, Guang J Zhang, and Yong Wang. 2023. An ensemble of neural networks for moist physics processes, its generalizability and stable integration. *Journal of Advances in Modeling Earth Systems* 15, 10 (2023), e2022MS003508.

[14] Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. 2020. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society* 146, 730 (2020), 1999–2049.

[15] Cathy Hohenegger, Peter Korn, Leonidas Linardakis, René Redler, Reiner Schnur, Panagiotis Adamidis, Jiawei Bao, Swantje Bastin, Milad Behravesh, Martin Bergemann, et al. 2023. ICON-sapphire: Simulating the components of the Earth system and their interactions at kilometer and subkilometer scales. *Geoscientific Model Development* 16, 2 (2023), 779–811.

[16] George Karypis and Vipin Kumar. 1998. METIS: A software package for partitioning unstructured graphs. *Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version* 4, 0 (1998), 0–80.

[17] L Ruby Leung, David C Bader, Mark A Taylor, and Renata B McCoy. 2020. An introduction to the E3SM special collection: Goals, science drivers, development, and analysis. *Journal of Advances in Modeling Earth Systems* 12, 11 (2020), e2019MS001821.

[18] 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.

[19] James E Lovelock and Lynn Margulis. 1974. Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis. *Tellus* 26, 1-2 (1974), 2–10.

[20] Hiroaki Miura, Tamaki Suematsu, Yuta Kawai, Yoko Yamagami, Daisuke Takasuka, Yuki Takano, Ching-Shu Hung, Kazuya Yamazaki, Chihiro Kodama, Yoshiyuki Kajikawa, et al. 2023. Asymptotic matching between weather and climate models. *Bulletin of the American Meteorological Society* 104, 12 (2023), E2308–E2315.

[21] Yoshiaki Miyamoto, Yoshiyuki Kajikawa, Ryuji Yoshida, Tsuyoshi Yamaura, Hisashi Yashiro, and Hirofumi Tomita. 2013. Deep moist atmospheric convection in a subkilometer global simulation. *Geophysical Research Letters* 40, 18 (2013), 4922–4926.

[22] Guo-Yue Niu, Zong-Liang Yang, Kenneth E Mitchell, Fei Chen, Michael B Ek, Michael Barlage, Anil Kumar, Kevin Manning, Dev Niyogi, Enrique Rosero, et al. 2011. The community Noah land surface model with multiparameterization options (Noah-MP): 1. Model description and evaluation with local-scale measurements. *Journal of Geophysical Research: Atmospheres* 116, D12 (2011).

[23] TN Palmer. 2016. A personal perspective on modelling the climate system. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 472, 2188 (2016), 20150772.

[24] George W Platzman. 1979. The ENIAC computations of 1950—Gateway to numerical weather prediction. *Bulletin of the American Meteorological Society* 60, 4 (1979), 302–312.

[25] David A Randall and Kerry Emanuel. 2024. The Weather–Climate Schism. *Bulletin of the American Meteorological Society* 105, 1 (2024), E300–E305.

[26] Vincent V Salomonson, WL Barnes, et al. 1989. MODIS: Advanced facility instrument for studies of the Earth as a system. *IEEE Transactions on Geoscience and Remote Sensing* 27, 2 (1989), 145–153.

[27] 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.

[28] Masaki Satoh, Hirofumi Tomita, Hisashi Yashiro, Yoshiyuki Kajikawa, Yoshiaki Miyamoto, Tsuyoshi Yamaura, Tomoki Miyakawa, Masuo Nakano, Chihiro Kodama, Akira T Noda, et al. 2017. Outcomes and

challenges of global high-resolution non-hydrostatic atmospheric simulations using the K computer. *Progress in Earth and Planetary Science* 4, 1 (2017), 1–24.

[29] J Smagorinsky and RF Strickler. 1965. Simulated climatology of a general circulation model with a hydrologic cycle. *Mon. Wea. Rev* 93 (1965), 769–798.

[30] Bjorn Stevens and Sandrine Bony. 2013. What Are Climate Models Missing? *Science* 340, 6136 (2013), 1053–1054.

[31] Bjorn Stevens, Masaki Satoh, Ludovic Auger, Joachim Biercamp, Christopher S Bretherton, Xi Chen, Peter Düben, Falko Judt, Marat Khairoutdinov, Daniel Klocke, et al. 2019. DYAMOND: the DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains. *Progress in Earth and Planetary Science* 6, 1 (2019), 1–17.

[32] 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.

[33] Xin Wang, Yilun Han, Wei Xue, Guangwen Yang, and Guang J Zhang. 2022. Stable climate simulations using a realistic general circulation model with neural network parameterizations for atmospheric moist physics and radiation processes. *Geoscientific Model Development* 15, 9 (2022), 3923–3940.

[34] Nils P Wedi, Inna Polichtchouk, Peter Dueben, Valentine G Anantharaj, Peter Bauer, Souhail Boussetta, Philip Browne, Willem Deconinck, Wayne Gaudin, Ioan Hadade, et al. 2020. A baseline for global weather and climate simulations at 1 km resolution. *Journal of Advances in Modeling Earth Systems* 12, 11 (2020), e2020MS002192.

[35] Nils P. Wedi, Inna Polichtchouk, Peter Dueben, Valentine G. Anantharaj, Peter Bauer, Souhail Boussetta, Philip Browne, Willem Deconinck, Wayne Gaudin, Ioan Hadade, Sam Hatfield, Olivier Iffrig, Philippe Lopez, Pedro Maciel, Andreas Mueller, Sami Saarinen, Irina Sandu, Tiago Quintino, and Frederic Vitart. 2020. A Baseline for Global Weather and Climate Simulations at 1 km Resolution. *Journal of Advances in Modeling Earth Systems* 12, 11 (9 2020). https://doi.org/10.1029/2020ms002192

[36] 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. https://doi.org/10.1109/SC41406.2024.00009

[37] Michael A Wulder, David P Roy, Volker C Radeloff, Thomas R Loveland, Martha C Anderson, David M Johnson, Sean Healey, Zhe Zhu, Theodore A Scambos, Nima Pahlevan, et al. 2022. Fifty years of Landsat science and impacts. *Remote Sensing of Environment* 280 (2022), 113195.

[38] Hisashi Yashiro, Koji Terasaki, Yuta Kawai, Shuhei Kudo, Takemasa Miyoshi, Toshiyuki Imamura, Kazuo Minami, Hikaru Inoue, Tatsuo Nishiki, Takayuki Saji, et al. 2020. A 1024-member ensemble data assimilation with 3.5-km mesh global weather simulations. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–10.

[39] Shaoqing Zhang, Haohuan Fu, Lixin Wu, Yuxuan Li, Hong Wang, Yunhui Zeng, Xiaohui Duan, Wubing Wan, Li Wang, Yuan Zhuang, et al. 2020. Optimizing high-resolution Community Earth System Model on a heterogeneous many-core supercomputing platform. *Geoscientific Model Development* 13, 10 (2020), 4809–4829.

[40] Yi Zhang, Jian Li, Rucong Yu, Zhuang Liu, Yihui Zhou, Xiaohan Li, and Xiaomeng Huang. 2020. A multiscale dynamical model in a dry-mass coordinate for weather and climate modeling: Moist dynamics and its coupling to physics. *Monthly Weather Review* 148, 7 (2020), 2671–2699.

[41] Yi Zhang, Jian Li, Rucong Yu, Shixun Zhang, Zhuang Liu, Jiahao Huang, and Yihui Zhou. 2019. A layer-averaged nonhydrostatic dynamical framework on an unstructured mesh for global and regional atmospheric modeling: Model description, baseline evaluation, and sensitivity exploration. *Journal of Advances in Modeling Earth Systems* 11, 6 (2019), 1685–1714.

[42] Yi Zhang, Jian Li, He Zhang, Xiaohan Li, Li Dong, Xinyao Rong, Chun Zhao, Xindong Peng, and Yiming Wang. 2023. History and Status of Atmospheric Dynamical Core Model Development in China. *Numerical Weather Prediction: East Asian Perspectives* (2023), 3–36.

[43] 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, 9 (2022), e2022EA002401.

Appendix: Artifact Description/Artifact Evaluation

Artifact Description (AD)

A Overview of Contributions and Artifacts

A.1 Paper's Main Contributions

C_1 An OpenMP based programming model support Sunway many-core heterogeneous systems, which facilitates automated porting, thereby maximizing computational efficiency and scalability.

C_2 We employ mixed-precision computing to address the memory bottleneck and improve the computing speed of the dynamical core part.

C_3 We integrate an AI-enhanced physics suite to achieve more efficient description of the physical parameterizations for different resolutions

A.2 Computational Artifacts

Contributions	Related Supported Paper Elements
C_1	Figure 9 Figure 10 Figure 11
C_2	Figure 9 Figure 10 Figure 11
C_3	Figure 9 Figure 10 Figure 11

B Artifact Identification

B.1 Computational Artifact A_1

Relation To Contributions

In our work, we optimize a global storm resolving model, Global-Regional Integrated Forecast System (GRIST), for the next-generation Sunway supercomputer, incorporating an ML-enhanced physics suite, OpenMP-based parallelization for SW26010P heterogeneous architecture, and mixed-precision code optimizations to improve both efficiency and performance portability. This appendix includes experimental setups and step-by-step instructions to replicate the findings of the paper.

Expected Results

The artifact will provide instructions for running GRIST model on the Sunway System. The experiments with 128 processes(18 nodes) at 100km resolution yield results consistent with Figure 9. The experiments with configuration in Table 1 yield results consistent with Figure 10 and Figure 11.

These experiments demonstrate the feasibility of successfully executing GRIST model on the Sunway system, supporting the claims made in this paper's C1, C2 and C3.

Expected Reproduction Time (in Minutes)

This artifact is compiled using Makefile, which encompass compilation information for various configuration. The compilation time does not exceed 5 minutes.

During the Artifact Execution phase, the execution time for running GRIST depends on the simulation time.

For experiments at high resolution, the increased data volume results in longer processing times. We recommend simulating approximately 3 days, which may take between 3 to 20 minutes depending on the resolution.

GRIST outputs simulation times during execution. You can intuitively observe the computation speed in the output files. This step is expected to take approximately 2 minutes.

Artifact Setup (incl. Inputs)

Hardware. Our experiments are mainly tested on the next-generation Sunway supercomputer. Computation resources for next-generation Sunway supercomputer are not currently publically available, but it is expected to be available after release. The information about the Sunway supercomputer is as follows:

Software. There are two kinds of software in our experiments:

1. GRIST: which is a original version on Sunway. It can run directly on SW processors.
   URL:
   `https://github.com/Xu-Kai/SW-GRIST`
2. GRIST-OPT: which is an optimized version with mixed-precision and AI-enhanced physics suite:
   URL:
   `https://github.com/Xu-Kai/SW-GRIST`

The artifact's source code includes the naive version and optimized version on Sunway. Below, we list the software environments on Sunway as a reference.

New Sunway System

- **Compiler:**
  - GNU Fortran (GCC) 7.1.0 20170502
  - swgcc (GCC) 7.1.0 20170502
- **MPI:** MPICH version 3.2b2
- **NetCDF:** netCDF 4.9.2
- **OpenMP:** SWGOMP 0.2.0

Datasets / Inputs. The 100km (G6) data and the weight of AI-enhanced physics suite along with its corresponding parameter files, are available for download on
`https://github.com/Xu-Kai/SW-GRIST`

Please contact xukai16@foxmail.com to access other data.

Installation and Deployment. The required libraries for compiling or running on Sunway have been listed in the software section above. The project already contains configuration schemes for multiple versions. To compile and run the program on Sunway platform, navigate to the file which path is *SW-GRIST/GRIST/bld* and *SW-GRIST/GRIST-OPT/bld*.

The steps are listed as below:

1. Pull the source code from the repository `https://github.com/Xu-Kai/SW-GRIST`
2. Before building, you can use:
   `make clean`
   to clean up the previous compilation and recompile the source code of software.
3. Build MPE double-precision(MPE-DP) version:
   ```
   cd GRIST
   sh build-mpe-dp.sh
   ```
4. Build MPE mixed-precision(MPE-MIXED) version:
   ```
   cd GRIST
   sh build-mpe-mixed.sh
   ```
5. Build CPE double-precision(CPE-DP) version:
   ```
   cd GRIST-OPT
   sh build-cpe-dp.sh
   ```
6. Build CPE mixed-precision(CPE-MIXED) version:
   ```
   cd GRIST-OPT
   sh build-cpe-mixed.sh
   ```

Artifact Execution

After executing the build command, the executable file of the program *ParGRIST-GCM_AMIPW** will be generated. Then you can cd *demo-g6-aqua* to run the program. There are *run-*.sh* for different versions of GRIST. For example, you can execute *"sh run-cpe-mixed.sh"* to run CPE mixed-precision verison with AI-enhanced physics suite. For more detailed execution instructions, please refer to

`https://github.com/Xu-Kai/SW-GRIST`

Artifact Analysis (incl. Outputs)

The performance will be written by GRIST in the console and log file named "*grist-*.log*". The default output files will be placed in *demo-g6-aqua*. Through these output files, you can obtain the runtime of this task and many kernels. In our experiments, we chose Simulation Year Per Day(SYPD) as the performance measure.

Artifact Evaluation (AE)

The overall process of this artifact can refer to the steps provided in the AD. For AE, we recommend conducting it on New Sunway System to compare with the data provided in the paper. Next, we will provide a detailed introduction on how to run this artifact on these two supercomputers.

B.1 Computational Artifact A_1

Artifact Setup (incl. Inputs)

You can use this command to obtain the artifact's source code:

git clone https://github.com/Xu-Kai/SW-GRIST.git

In the following text, we will use "GRIST" to represent the root directory of the artifact.

The software environment and compilation options for the system have already been prepared in the *GRIST/bld/build*.sh* and *GRIST-OPT/bld/build*.sh*. You only need to execute the shell scripts to build the application.

Artifact Execution

After the build command complete the compilation of the application, AE reviewers can submit the job to New Sunway System. In the directory *GRIST/demo-g6-aqua*, there are *run-*.sh* to run different versions of GRIST. Upon execution of the command, the job will be dispatched to the System with the application's output in current directory. If reviewers want to run high resolution data, please contact xukai16@foxmail.com to prepare the data. The duration of the application's execution, ranging from 3 to 20 minutes, is contingent upon the resolution and the number of processes on the platform.

Artifact Analysis (incl. Outputs)

The default output files will be stored in the directory *GRIST/demo-g6-aqua*. Through the output files, you can obtain the runtime of this task and the corresponding SYPD (Simulation Year Per Day).

The 100km(G6) experiment results from the runs on New Sunway platform should be close to the Figure 9 in the paper, indicating that the work presented in this article can fully utilize Sunway heterogeneous computing resources and has achieved an acceleration ratio of about 20-70x compared to MPE double-precesion version for major kernels.

For weak scaling, we take the number of the process 128 for 100km(G6) resolution. When the resolution is increased by factor of 2, we scale the number of processes by a factor of 4. The performance result of weak scaling in Figure 10 can be measured with the configuration in Table 1.

For strong scaling, the resolutions of 1km(G12) and 3km(G11) are used to evaluate the results. We start from 32768 processes and scale up to 524288 processes which run on near full Next Generation Sunway Supercomputer.The result for strong scaling is shown in Figure 11.

The results of weak scaling and strong scaling indicates that the work presented in this article can fully utilize the heterogeneous computing resources on the Sunway platform, and has demonstrated a high degree of scalability.