

# Refactoring and Optimizing the Community Atmosphere Model (CAM) on the Sunway TaihuLight Supercomputer

Haohuan Fu<sup>1,2</sup>, Junfeng Liao<sup>1,2,3</sup>, Wei Xue<sup>1,2,3</sup>, Lanning Wang<sup>4</sup>, Dexun Chen<sup>3</sup>, Long Gu<sup>5</sup>, Jinxiu Xu<sup>5</sup>, Nan Ding<sup>1,3</sup>, Xinliang Wang<sup>3</sup>, Conghui He<sup>1,2,3</sup>, Shizhen Xu<sup>1,3</sup>, Yishuang Liang<sup>4</sup>, Jiarui Fang<sup>1,2,3</sup>, Yuanchao Xu<sup>3</sup>, Weijie Zheng<sup>1,2,3</sup>, Jingheng Xu<sup>1,2,3</sup>, Zhen Zheng<sup>3</sup>, Wanqing Wei<sup>1,2</sup>, Xu Ji<sup>1,3</sup>, He Zhang<sup>1,3</sup>, Bingwei Chen<sup>1,2</sup>, Kaiwei Li<sup>3</sup>, Xiaomeng Huang<sup>1,2</sup>, Wenguang Chen<sup>3</sup>, and Guangwen Yang<sup>1,2,3</sup>

<sup>1</sup>Ministry of Education Key Laboratory for Earth System Modeling, and Center for Earth System Science, Tsinghua University, Beijing, China

<sup>2</sup>National Supercomputing Center in Wuxi, Wuxi, China

<sup>3</sup>Department of Computer Science and Technology, Tsinghua University, Beijing, China

<sup>4</sup>College of Global Change and Earth System Science, Beijing Normal University, Beijing, China

<sup>5</sup>National Research Center Of Parallel Computer Engineering and Technology, Beijing, China

## Abstract—

This paper reports our efforts on refactoring and optimizing the Community Atmosphere Model (CAM) on the Sunway TaihuLight supercomputer, which uses a many-core processor that consists of management processing elements (MPEs) and clusters of computing processing elements (CPEs). To map the large code base of CAM to the millions of cores on the Sunway system, we take OpenACC-based refactoring as the major approach, and apply source-to-source translator tools to exploit the most suitable parallelism for the CPE cluster, and to fit the intermediate variable into the limited on-chip fast buffer. For individual kernels, when comparing the original ported version using only MPEs and the refactored version using both the MPE and CPE clusters, we achieve up to 22x speedup for the compute-intensive kernels. For the 25km resolution CAM global model, we manage to scale to 24,000 MPEs, and 1,536,000 CPEs, and achieve a simulation speed of 2.81 model years per day.

**Keywords**—atmospheric modeling, many-core, optimization, tool, OpenACC

## I. INTRODUCTION

Ever since the first generation of supercomputer systems (CDC 6600, Cray-I, etc.), the atmospheric models have been among the major users of computing resources [1], and evolved with the development of supercomputer systems. In the early years, the vector machines, such as the Cray systems (Cray 1, Cray X-MP, Cray Y-MP, etc.), and the famous Japanese earth simulator [2], have been the major computation platform for weather and climate modelers. Then, from the year of 2000 or so, Intel, IBM, and SGI clusters emerged as the replacements of the traditional vector machines, and also made the transition of the atmospheric modeling programs to the style of MPI programs. Along this transition process, with the increase of cores within each processor, we see the introduction of the hybrid parallelization

scheme that combines MPI and OpenMP.

In the recent decade, again, we see the transition of supercomputers from homogeneous systems with only multi-core CPU processors to heterogeneous systems with both CPUs and many-core accelerators [3], [4]. This architectural transition, again, brings significant changes to existing high-performance computing software in various application domains, such as geophysics exploration, sky simulation, and phase-field simulation.

Unlike the above application domains that have made a quick adaptation to the many-core accelerators, the transition of the weather and climate models has been relatively slow. One big reason is the millions lines of legacy code that have been written for multi-core CPUs rather than many-core accelerators. As a result, most existing efforts either focus on standalone physics schemes ([5], [6]), or focus on the dynamic core part ([7], [8]).

In contrast, complete porting projects of entire models onto heterogeneous supercomputers are still few to be seen. Typical examples include the GPU-based acceleration of a next-generation high resolution meso-scale atmospheric model being developed by the Japan Meteorological Agency (JMA) [9], and complete porting of the Princeton Ocean Model (POM) onto GPU devices [10], both of which take a manual rewriting of the code into CUDA. Only in newly-developed climate or weather models, such as NIM [11] and COSMO [12], we see careful considerations for heterogeneous systems, and support for multiple architectures including CPU, GPU, and MIC.

In general, while the weather and climate models are calling for more computing power to support higher resolution and more complex physics [13], there is still a gap between the increasing demand and the increasing supply in the form of many-core accelerators. To fill the gap between the demand and the supply, in our work, we perform an

extensive refactor and optimization of the CAM atmospheric model, for the Sunway TaihuLight Supercomputer, equipped with many-core processors that consists of Management Processing Elements (MPEs) and clusters of Computing Processing Elements (CPEs). We pick CAM [14] as the our target application, as it is one of the most widely used advanced atmospheric model in the world. Note that, to achieve a high scalability over the Sunway system with millions of cores, we select to use the SE dynamic core of CAM [15], and the other dynamic core options are not considered in this work.

We use the Sunway OpenACC compiler (a customized version that expands from the OpenACC 2.0 standard) as the major tool to achieve a suitable mapping of CAM onto the new Sunway heterogeneous many-core processors. Due to large code base developed over the last few decades, and the general demand from climate scientists to maintain a same source, we try to minimize the manual refactoring efforts (to only the dynamic core part), and mainly rely on the source-to-source translation tools to achieve an automated and efficient porting. Besides, compared with GPU and other many-core accelerators, both the on-chip fast buffer and the available memory bandwidth of the Sunway processor are relatively limited (detailed in Section III), which make our porting significantly more challenging. A large part of our tools and optimization strategies would focus on minimizing the memory footprints.

Our major contributions are as follows:

- Through a careful refactor of the SE dynamic core, we manage to combine the distributed loops in the major computational functions into aggregated multi-level loops, and expose a suitable level of both parallelism and variable storage space for the CPE cluster architecture.
- For the physics parts, which includes numerous modules with different code styles by different scientists, we design a loop transformation tool to identify and expose the most suitable level of loop body for the parallelization on the CPE cluster. In addition, we also design a memory footprint analysis and reduction tool, and a number of customized Sunway OpenACC features, to fit the frequently-accessed variables into the local fast buffer of the CPE.
- Comparing the refactored hybrid version using both MPE and the CPE cluster against the ported version using only MPE, we can achieve up to 22x speedup for compute-intensive kernels, and 2x to 7x speedup for kernels involving both computation and memory operations. For the entire CAM model, we achieve a performance improvement of around 2x.
- We manage to scale the refactored CAM model to 24,000 MPEs, and 1,536,000 CPEs, and achieve a simulation speed of 2.81 modeling years per day for the 25km resolution.

While the speedup is not significant, it provides an important base for us to continue tuning the performance of large and complicated scientific programs, such as CAM.

## II. RELATED WORK

When compared with other HPC application domains, the porting of the climate models onto many-core architectures has been relatively slow. Most early-stage efforts focused on the standalone physics modules. The Weather Research and Forecast (WRF) model, which is one of the most widely used numerical weather prediction (NWP) program in the world, was among one of the earliest weather/climate models that integrate GPU-accelerated microphysics schemes [5]. Other typical examples include GPU-based accelerations for the chemical kinetics modules in WRF-Chem [16], and the shortwave radiation parameterization in CAM [6]. As these physics modules are usually compute-intensive and do not involve communications, GPU-based acceleration can generally achieve a speedup of one order of magnitude.

In recent years, we start to see projects that accelerate the dynamic cores (or the major computation parts of the dynamics cores) on GPU devices, such as the efforts on GRAPES [7], CAM-SE [8]. Compared with the physics modules, the dynamic parts generally require communication across different grids, and are more difficult to achieve good parallel performance. The speedup is usually in the range of 3 to 5 times when comparing GPU solutions against parallel CPU solutions.

Complete porting projects that move complete atmospheric or ocean models onto many-core accelerators are still few to see. One example is the GPU-based acceleration of ASUCA, a next-generation high resolution meso-scale atmospheric model being developed by the Japan Meteorological Agency (JMA) [9], with an acceleration of 80-fold when compared against a single CPU core, and an excellent scalability for up to a few thousand nodes. Another example is the complete porting of the Princeton Ocean Model (POM) onto GPU devices [10], which performs a manual porting and optimization of POM onto a hybrid server with 4 GPUs, and achieves an equivalent performance to 408 CPU cores. While the two projects have managed to take advantage of GPU accelerators for complete models, these are relatively less complicated models that involve only tens of thousands of lines of code. Moreover, both projects take the approach of manual rewriting of the program in CUDA, which makes it difficult to keep a same source base and not possible to migrate to other computing architectures.

For newly developed weather or climate models, such as NIM [11] and COSMO [12], we see the emphasis on supporting multiple accelerator architectures. For example, by using OpenMP, OpenACC, and F2C-ACC directives, NIM managed to maintain a single source for application scientists, and the portability over CPU, GPU, and MIC

architectures [11]. Similarly, the COSMO model also maintains both OpenMP and OpenACC directives for the physics schemes in the model, to support both GPU and MIC. For the dynamic core part, the COSMO model uses a C++-based domain specific language to provide both CUDA and OpenMP backends. As newly developed numerical models, NIM and COSMO are relatively ahead in the transition, due to less constraints from the code legacy.

Different from the projects mentioned above, our work focuses on CAM 5.3, which is a complete atmospheric model with 560,000 lines of code, developed over the last few decades. Moreover, our target platform is the new Sunway processor, which is based on a hybrid many-core architecture with a few MPEs and an array of CPEs. Compared with GPU and other many-core accelerators, both the on-chip fast buffer and the available memory bandwidth are more limited, which make our porting significantly more challenging. For such a large code base that have been developed over the decades, we think a manual rewrite would not be a feasible solution. Instead, in our approach, we would minimize the manual refactoring (limited to the dynamic core part), and rely on source-to-source translation tool to remap the code onto the new architecture.

### III. THE SUNWAY TAIHULIGHT SUPERCOMPUTER

#### A. The Hardware System

As one of the two 100PF systems supported by China's National 863 High-Tech Research and Development Program in the 12th five-year plan, the Sunway TaihuLight supercomputer [17] is the successor of the Sunway BlueLight supercomputer hosted in the Jinan Supercomputer Center. Similar to the Sunway BlueLight system, the new Sunway supercomputer is also using China's homegrown processor designs.

The general architecture of the new Sunway heterogeneous processor [17] is shown in Figure 1. The processor includes 4 core-groups (CGs). Each CG includes one management processing element (MPE), one computing processing element (CPE) cluster with 8x8 CPEs, and one memory controller (MC). These 4 groups are connected via the network on chip (NoC). Each group has its own memory space, which is connected to the MPE and the CPE cluster through the MC. The processor connects to other outside devices through a system interface (SI).

The MPE is a complete 64-bit RISC core, which can run in both the user mode and the system mode. The MPE supports the complete interrupt functions, memory management, superscalar, and out of order issue / execution. Therefore, the MPE is an ideal core for handling management and communication functions.

In contrast, the CPE is also a 64-bit RISC core, but with limited functions. CPE can only run in the user mode, and does not support interrupt functions. The design goal is to achieve the maximum aggregated computing power, while

minimizing the complexity of the micro-architecture. The CPE cluster is organized as an 8 by 8 mesh, with a mesh network to achieve low-latency register data communication among the 8 by 8 CPEs. The mesh also includes a mesh controller that handles interrupt and synchronization controls.

In terms of the memory hierarchy, each MPE has a 32KB L1 instruction cache and a 32KB L1 data cache, with a 256KB L2 cache for both instruction and data. Each CPE has its own 16KB L1 instruction cache, and a 64KB Scratch Pad Memory (SPM). The SPM can be configured as either a fast buffer that support precise control by the users or a software-emulated cache that achieves automatic data caching in a software manner. However, as the performance of the software-emulated cache is low, in most cases, we need a user-controlled buffering scheme to achieve good performance.

Combining the four CGs of MPE and CPE clusters, each Sunway processor provides a peak performance over 3 Tflops, with a performance-to-power ratio over 10 Gflops/Watt. While the computing performance and power efficiency is among the top when compared with existing GPU and MIC chips, the on-chip buffer size and the memory bandwidth is relatively limited. The four CGs are sharing an aggregated memory bandwidth of around 130 GB/s.

#### B. The Software System

On the software side, the TaihuLight system uses a customized 64-bit Linux as the operating system, with a set of compilation tools to support the development of applications on the new Sunway processor architecture.

The compilation tool set includes the basic compiler components, such as the C/C++, and Fortran compilers. In addition to that, there is also a parallel compilation tool that supports the OpenACC 2.0 syntax and targets the CPE clusters. The customized Sunway OpenACC tool supports management of parallel tasks, extraction of heterogeneous code, and description of data transfers. Moreover, according to the specific features of the Sunway processor architecture, the Sunway OpenACC tool has also made a number of syntax extensions from the original OpenACC 2.0 standard, such as a fine control over buffering of multi-dimensional array, and packing of distributed variables for data transfer (detailed in Section VI-C).

### IV. MAPPING CAM TO THE SUNWAY TAIHULIGHT SUPERCOMPUTER: OUR GENERAL METHODOLOGY

#### A. General Workflow of CAM

CAM, which serves as the atmosphere component of the Community Earth System Model (CESM) [18], is the most computationally expensive component in typical configurations. The computation workflow of CAM can be divided into two phases: the dynamics and physics. The dynamics advances the evolutionary equations for the atmospheric flow, and the physics approximates sub-grid phenomena such

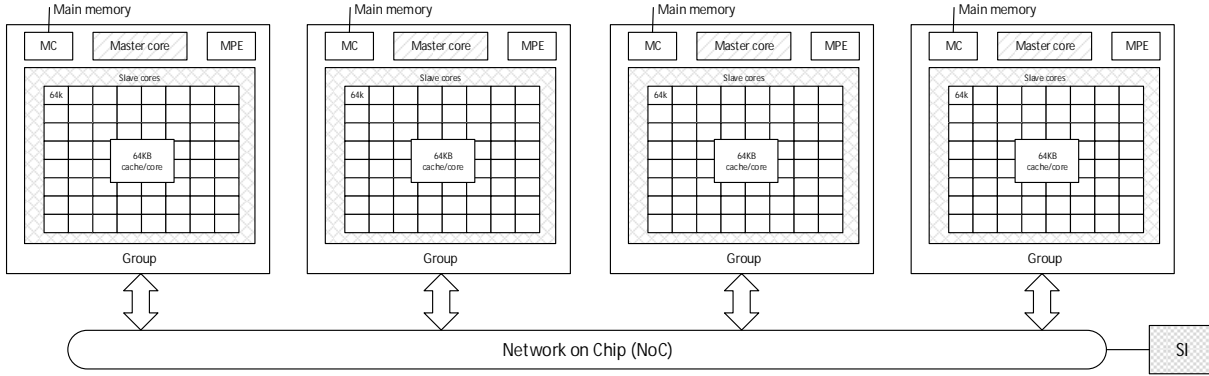


Figure 1. The general architecture of the new Sunway processor.

as clouds, precipitation processes, long/short-wave radiation, and turbulent mixing. As shown in Figure 2, the physics contains two phases. In `phy_run1`, radiation, shallow and deep advection are calculated, and in `phy_run2`, aerosol, and chemistry procedure are computed. State variables, such as temperature and precipitation, are passed through between two physic phases. The physics passes the tracers, such as  $u$ ,  $v$ , to the dynamics. After initialization, the physics and the dynamics are executed in turn during each simulation time-step.




Figure 2. The general workflow of CAM.

### B. Porting and Refactoring

As CAM has not been running on the Sunway architecture before, the first step of our porting is to verify the modeling results. As the current version of CAM has to be run in a coupled mode with the Common Land Model (CLM), using the F compset of CESM 1.2.1, we port both CAM and CLM onto the Sunway system, using only MPE to perform the computation.

After running the coupled CAM and CLM models on the Sunway system for the duration of three years, we compare results of major variables, as well as the conservation of mass and energy, to verify the correctness of our ported version. Compared with the results on the Intel clusters using the same modeling parameter configuration, we see almost identical distribution of key variables and an average relative error in the range of  $10^{-4}$ . Figure 3 shows the variation of

the total mass over one year. We also observe an identical variation cycle when compared with the results on the Intel cluster.




Figure 3. The variation of the total mass of a coupled CAM and CLM run on the new Sunway system.

Using the ported MPE-only version as the starting point, we then perform refactoring and optimization of both the dynamic core and the physics schemes to make utilization of the CPE clusters. During the process of expanding each kernel from MPE-only to MPE-CPE hybrid mode, we take the numerical result of the MPE-only mode as the true value, and ensure that the difference of the results in the MPE-CPE hybrid mode are within a similar range of the floating-point rounding errors.

Considering both the code volume and complexity, for the SE dynamic core, we take a manual approach to refactor and optimize the code. For the physics part, we rely on our source-to-source translation tools to perform loop transformation, and to reduce the memory footprint for the on-chip fast buffer (the SPM of each CPE).

## V. REFACTORING AND OPTIMIZATION OF THE SE DYNAMIC CORE

### A. Major Challenges and Our Solutions

The current CAM code is taking a hybrid parallelization scheme that combines MPI and OpenMP to parallelize the computation. As MPI is used for the inter-node parallelization and OpenMP is used for the intra-node parallelization, we generally apply a configuration with a large number of MPI processes (hundreds to thousands), and small number of OpenMP threads (8 to 16). In the Sunway system, we are dealing with groups with one MPE and 64 CPEs. Therefore, our first challenge is that a direct map from OpenMP to OpenACC would not provide a suitable level of parallelism for the CPE cluster.

The other challenge relates to the buffering of data. As mentioned in Section III-A, 64KB SPM of each CPE needs to be fully controlled by the user, either through a customized buffering scheme, or a software-emulated cache scheme. Our experiments show that the emulated software cache scheme is not efficient enough to provide performance benefits. Therefore, in most cases, we need to design a buffering scheme that loads the proper data into the SPM.

For the first challenge of achieving a suitable level of parallelism for CPE clusters, we do adjustments of both the computational sequence and the loop structures, so as to aggregate enough computations, and to enable the right number of parallel threads. For the second challenge of data buffering, on one hand, we refine the code to minimize the usage of intermediate variables; on the other hand, we design customized buffering schemes to determine when and where to load or unload data arrays.

In the following subsection, we take the *Euler\_step* function (the most time consumption part of the SE dynamic core) as the example to demonstrate our refactorization and optimization schemes in details.

### B. Refactorization of the *Euler\_step* Function

The general structure of the *Euler\_step* function is shown in block ① of Figure 4.

The *Euler\_step* function is the basic forward Euler component used to construct the Strong Stability-Preserving (SSP) Runge-Kutta (RK) methods, consisting of the following three major parts:

- 1) computation of the biharmonic mixing term and its min/max values (later used as limiters);
- 2) the 2D advection step that updates each tracer's Qdp

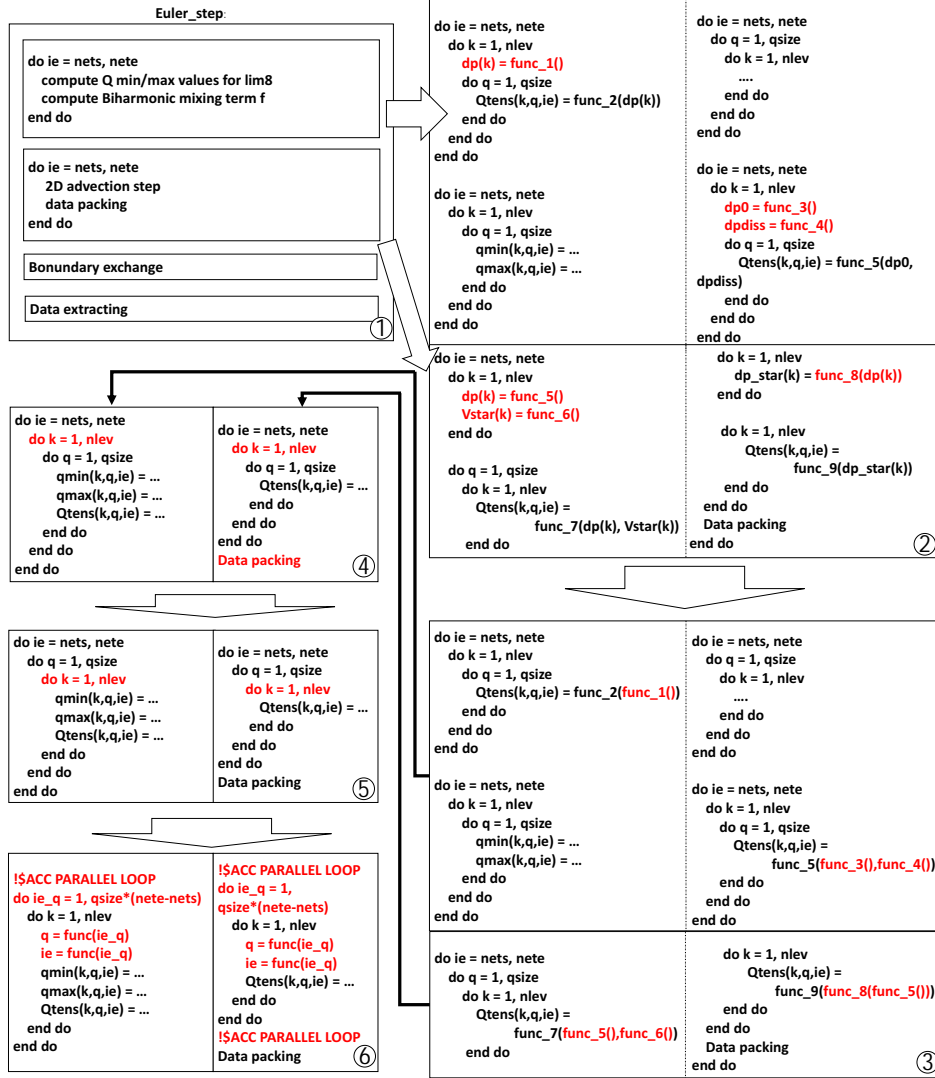


Figure 4. The general structure of the `Euler_step` function, and our corresponding refactor of the biharmonic computation part and the advection part.

the loops into a nested  $(ie, q, k)$  three-level loop. Similarly, we can then collapse the  $ie$ -loop and the  $q$ -loop to expose enough parallelism (as shown in the right part of block ④, ⑤, and ⑥). The required memory space of the intermediate variables of the resulting loop body is also reduced to fit into the SPM.

By performing the above transform of loops and the corresponding optimization for the buffering of local variables, we can achieve 22x speedup for the advection part, 10x speedup for the biharmonic part, and 7x speedup for the advection part, with a 2x to 4x speedup for the entire `Euler_step` function.

We apply similar refactorization and optimization schemes to the `compute_and_apply_rhs` function

that computes a leapfrog timestep, and the `advance_hypervis_dp` function that updates the temperature. In addition, we also refactorize and parallelize the parts that perform packing and unpacking of data elements for the halo communication stage, so as to improve the overall performance of the entire SE dynamic core.

## VI. REFACTORING AND OPTIMIZATION OF THE PHYSICS SCHEMES

### A. The General Parallelization Scheme

As shown in Figure 5, in most physics schemes, the computation is performed on the columns. Each column consists of the different vertical levels of the same surface

location. The computation of each column is completely independent, which provides the parallelism in the physics schemes. In order to achieve dynamic load balancing, CAM generally organizes columns in different locations into one chunk, and assigns each chunk to a different OpenMP thread.



```
do begin_chunk to end_chunk
tphysbc()
{
  convect_deep_tend(6.47%)
  convect_shallow_tend(15.57%)
  macrop_driver_tend(8.38%)
  microp_aero_run(4.29%)
  microp_driver_tend(7.13%)
  aerosol_wet_intr(4.29%)
  convect_deep_tend_2(0.51%)
  radiation_tend(54.07%)
}
enddo
```

```
tphysbc()
{
  do begin_chunk to end_chunk
  convect_deep_tend(6.47%)
  convect_shallow_tend(15.57%)
  macrop_driver_tend(8.38%)
  microp_aero_run(4.29%)
  microp_driver_tend(7.13%)
  aerosol_wet_intr(4.29%)
  convect_deep_tend_2(0.51%)
  radiation_tend(54.07%)
  enddo
}
}
```

```
tphysbc()
{
  do begin_chunk to end_chunk
  convect_deep_tend(6.47%)
  enddo
  .....
  do begin_chunk to end_chunk
  microp_driver_tend(7.13%)
  enddo
  .....
  do begin_chunk to end_chunk
  radiation_tend(54.07%)
  enddo
}
}
```

```
do begin_chunk to end_chunk
convect_deep_tend(6.47%)
{
  zm_conv_tend(6.47%)
  {
    zm_convr(2.03%)
    zm_conv_evap()
    montran()
    convtranc(0.06%)
  }
}
enddo
```



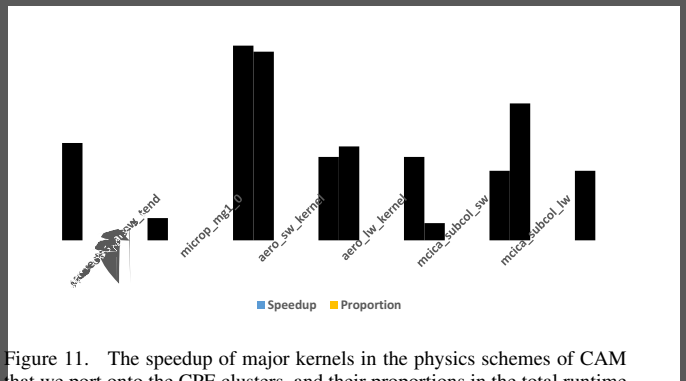


Figure 11. The speedup of major kernels in the physics schemes of CAM that we port onto the CPE clusters, and their proportions in the total runtime of the physics part. The speedup is comparing the performance of the kernel running on 1 MPE and 64 CPEs against the performance of the kernel running on only 1 MPE.

### B. Performance of the SE Dynamic Core in Different Parallel Scales

Figure 12 demonstrates the execution time of the entire CAM-SE dynamic core that we can achieve with different parallel scales and different simulation configurations. The parallel scale is described by the number of CGs, as we generally run one MPI process for each CG. Each CG includes 1 MPE and 64 CPEs. For the simulation configuration, we change the resolution and the number of tracers. *ne30* and *ne120* refer to the 100-km and 25-km resolution configurations respectively. The number of tracers is denoted by *qsize*, which can be 25 (default configuration), or 108 (heavy chemistry configuration).

While we achieve up to 22x speedup for single kernels within the SE dynamic core, when we consider the entire SE (including computation, memory copy, communication, and the serial parts that simply can not take advantage of the CPEs), the speedup is between 2 times and 4 times. In general, the speedup drops down as the parallel scale increases, which corresponds to the increased portion of communication and reduced portion of computation for each CG. As for the number of tracers, a larger *qsize* would bring a better speedup, as there would be a larger *q* to process for the CPE clusters (as discussed in Section V).




Figure 12. The execution time of the CAM-SE dynamic core that we port onto the CPE clusters with different parallel scales and different simulation configurations. We demonstrate both the time for running the CAM-SE dynamic core on both MPEs and CPE clusters and the time for running on only MPEs. The *x* axis describes the number of CGs (each CG includes 1 MPE and 64 CPEs). *ne30* and *ne120* refer to the 100-km and 25-km resolution configurations respectively. *qsize* denotes the number of tracers used.

### C. Performance of the Physics Part in Different Parallel Scales

In contrast to the dynamic core, as the physics schemes do not involve MPI communications, the performance improvement achieved from the CPE clusters does not change significantly with the parallel scale.

Figure 13 demonstrates the execution time of the entire physics part that we can achieve with different parallel scales

and different resolution configurations. With the number of parallel CGs increasing, we see a constant performance improvement of two times by using both MPE and CPEs.

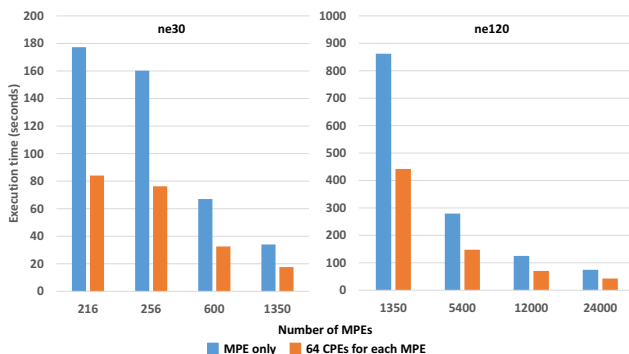


Figure 13. The execution time of the physics part that we port onto the CPE clusters with different parallel scales and different simulation configurations. Similarly, we demonstrate both the time for running on both MPEs and CPE clusters and the time for running on only MPEs. The *x* axis describes the number of CGs (each CG includes 1 MPE and 64 CPEs). *ne30* and *ne120* refer to the 100-km and 25-km resolution configurations respectively.

### D. Speedup of the entire CAM model

Figure 14 shows the simulation speed of the CAM model (measured in Model Years Per Day (MYPD)) on the new Sunway supercomputer, with the number of CGs increasing from 1,024 to 24,000. Similar to previous reported results on other systems, the CAM model demonstrates a good scalability on the new Sunway supercomputer system, with the simulation speed increasing steadily with the number of CGs. For the large scale cases, we demonstrate the performances for using MPE only, and using both MPE and CPE clusters. As shown in the last two points in Figure 14, by using both the MPE and the CPE clusters, we can further improve the simulation speed by another 2x. When scaling the CAM model to 24,000 CGs (24,000 MPEs, and 1,536,000 CPEs), we can achieve a simulation speed of 2.81 MYPD.

### E. Performance Result Analysis

In previous sections, we provide the general results about the performance of the refactored and optimized CAM-SE model on the Sunway TaihuLight system. The MPE and CPE hybrid architecture of the new Sunway system is largely different from previous heterogeneous systems. While MPE is like a CPU core, and the CPE cluster is like the many-core accelerator, both the CPU and the accelerator are now fused into one chip. The speedup of a complete CG (1 MPE and 64 CPEs) over the one MPE provides an important indication about how well our refactoring method and code-transformation tools can port the heavy codes of CAM-SE,

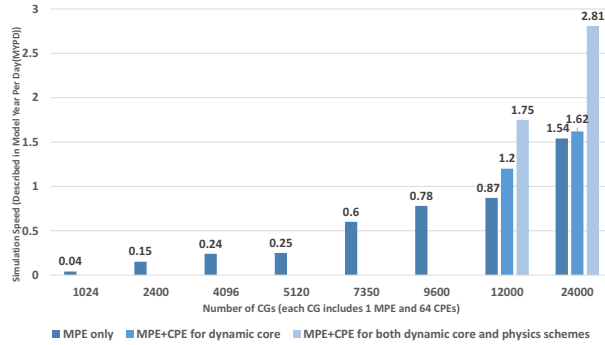


Figure 14. The simulation speed of the CAM model (measured in Model Years Per Day (MYPD)) on the new Sunway supercomputer, with the number of CGs increasing from 1,024 to 24,000. For the large-scale run with 12,000 and 24,000 CGs, we show the performances of the model for three scenarios: (1) using MPE only; (2) using MPE+CPE for the dynamic core; (3) using MPE+CPE for both the dynamic core and the physics schemes.

which were originally written for multi-core CPU, to the new hybrid chip.

Both the MPE and the CPE support 256-bit vector instructions. However, as the MPE supports dual issuing of instructions and the CPE only supports single issuing of instructions, the peak performance of one MPE (23.2 GFlops) is two times of the peak performances of one CPE (11.6 GFlops). Therefore, in terms of computing capability, each CG is equivalent to 33 MPEs. In terms of the memory bandwidth, the measured bandwidths for the MPE and the 64 CPEs are 5 GB/s and 20 GB/s respectively. Thus, each CG is only five times better than each MPE. These metrics provide a general guideline about the performance improvement that we can achieve when porting the code from one MPE to the complete CG with additional 64 CPEs.

The speedup numbers we achieve for various dynamic core and physics kernels align well with the ranges defined by the above metrics. For compute-intensive kernels (such as the *advection\_step* function), we see a speedup of 22 times, which is close to the 33 theoretical bound. For memory-intensive kernels (such as *edgeVpack\_euler* and *edgeVpack\_neighbor*), the speedup is only 2 to 3 times, which is more bounded by the bandwidth.

In contrast to our work, NCAR, ORNL, Cray, and NVIDIA’s collaborative effort on accelerating CAM-SE on Titan [8] was focused on the most expensive kernels in the dynamic core part. When moving from the multi-core CPU to the many-core GPU (with the peak computing performance improved by around 10 times and the peak memory bandwidth improved by around 5 times), the GPU accelerator would only bring around  $6\times$  speedup for expensive kernels, and  $2\times$  speedup for the entire model.

We originally expect a higher speedup for the kernels in the physics parts, which are generally more compute-

intensive, and do not involve any communications. However, due to the complexity of the physics codes, we rely on our automated tools to perform the refactoring and the optimization in this part. We think that the current low speedup for certain kernels are partly due to the constraints of our current tools. Therefore, one of our future plans is to further improve our automated transformation tools, and to involve manual optimizations for certain expensive physics schemes, so as to further improve the performance of the physics part.

For the performance of the entire model, as the execution time is distributed in a large number of kernels, which sometimes involve both memory and communication operations, the ported CAM-SE demonstrates a clear memory-bound behavior. In our tested scenarios, porting onto CPEs only improves the performance by roughly 2 times, which we think is mostly constrained by the memory bandwidth and the communication parts that have not yet been fully optimized. The current model provides a simulation speed of 2.81 MYPD for the resolution of 25 km, using 24,000 CGs (24,000 MPEs and 1,536,000 CPEs), which is similar to the speed of around 2 MYPD for the high-resolution CESM runs at the Yellowstone supercomputer of NCAR [18]. While such a speed is still not good enough for scientists to perform experiments of a few hundreds or even thousands of years, we think it is a good starting point for our future efforts that would further improve the speed and efficiency through other algorithmic and architectural redesigns.

## VIII. CONCLUSION

In this paper, we report our efforts on porting the CAM model to the Sunway TaihuLight many-core supercomputer. Due to the differences between the Sunway many-core processor (4 CGs, each of which consists of 1 MPE and 64 CPEs) and the traditional multi-core CPUs, and especially the 64KB SPM that needs to be explicitly controlled by the user, we perform an extensive refactor of CAM to expose the right level of parallelism to the 64 CPEs in each CG, and apply various optimization techniques to fit the involved variables and arrays into the 64KB SPM of each CPE. For single kernels in both the dynamic core and the physics schemes, we achieve 14x to 22x speedup for kernels that provide a suitable fit to both the computational and memory architecture of the CPE cluster. For kernels that are not suitable, we can still achieve around 2x to 7x speedup after applying the loop transformation tool and various variable storage reduction tool. Our refactored CAM model shows a good scalability on the Sunway TaihuLight supercomputer, and can efficiently use up to 24,000 MPEs and 1,536,000 CPEs, and provide a simulation speed of 2.81 MYPD when using the 25-km resolution.

While the speedup of the entire CAM model is not significant, it provides an important base for us to continue optimizing the performance of such a large and complicated

scientific simulation program. In our future work, we will continue working on our source-to-source translation tools to provide an automated workflow. The goal is to refine the refactor and optimization strategies for more suitable mappings between the algorithm and the architecture, and to improve the simulation speed of the model to a level that makes high-resolution simulation an applicable scientific tool on the Sunway system.

#### ACKNOWLEDGMENT

The authors are grateful to the reviewers for valuable comments that have greatly improved the paper. This work is partially supported by National key research and development plan of China under Grant No. 2016YFA0602103, China Special Fund for Meteorological Research in the Public Interest under Grant No. GYHY201306062, Natural Science Foundation of China under Grant Nos. 61303003, 41374113, 91530103, 91530323, 51190101, and 61361120098. The corresponding author is Wei Xue (email: xuewei@tsinghua.edu.cn).

#### REFERENCES

- [1] J. Drake, I. Foster, J. Michalakes, B. Toonen, and P. Worley, "Design and Performance of a Scalable Parallel Community Climate Model," *Parallel Computing*, pp. 1571–1591, 1995.
- [2] S. Shingu, H. Takahara, H. Fuchigami, M. Y. Y. Tsuda, M. Yamada, Y. Tsuda, and et. al., "A 26.58 Tflops Global Atmospheric Simulation with the Spectral Transform Method on the Earth Simulator," in *In Proceedings of the ACM / IEEE Supercomputing SC2002 conference*, 2002.
- [3] X.-J. Yang, X.-K. Liao, K. Lu, Q.-F. Hu, J.-Q. Song, and J.-S. Su, "The TianHe-1A Supercomputer: Its Hardware and Software," *Journal of Computer Science and Technology*, vol. 26, no. 3, pp. 344–351, 2011.
- [4] X. Liao, L. Xiao, C. Yang, and Y. Lu, "MilkyWay-2 supercomputer: system and application," *Frontiers of Computer Science*, vol. 8, no. 3, pp. 345–356, 2014.
- [5] J. Michalakes and M. Vachharajani, "GPU acceleration of numerical weather prediction," in *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, April 2008, pp. 1–7.
- [6] R. Kelly, "GPU Computing for Atmospheric Modeling," *Computing in Science and Engineering*, vol. 12, no. 4, pp. 26–33, 2010.
- [7] Z. Wang, X. Xu, N. Xiong, L. Yang, and W. Zhao, "GPU Acceleration for GRAPES Meteorological Model," in *High Performance Computing and Communications (HPCC), 2011 IEEE 13th International Conference on*, Sept 2011, pp. 365–372.
- [8] I. Carpenter, R. Archibald, K. Evans, J. Larkin, P. Micikevicius, M. Norman, and et. al., "Progress towards accelerating HOMME on hybrid multi-core systems," *International Journal of High Performance Computing Applications*, vol. 27, no. 3, pp. 335–347, 2013.
- [9] T. Shimokawabe, T. Aoki, C. Muroi, J. Ishida, K. Kawano, T. Endo, and et. al., "An 80-Fold Speedup, 15.0 TFlops Full GPU Acceleration of Non-Hydrostatic Weather Model ASUCA Production Code," in *High Performance Computing, Networking, Storage and Analysis (SC), 2010 International Conference for*, Nov 2010, pp. 1–11.
- [10] S. Xu, X. Huang, Y. Zhang, H. Fu, L.-Y. Oey, F. Xu, and et. al., "gpuPOM: a GPU-based Princeton Ocean Model," *Geoscientific Model Development Discussions*, vol. 7, no. 6, pp. 7651–7691, 2014.
- [11] M. Govett, J. Middlecoff, and T. Henderson, "Directive-Based Parallelization of the NIM Weather Model for GPUs," in *Accelerator Programming using Directives (WACCPD), 2014 First Workshop on*, Nov 2014, pp. 55–61.
- [12] B. Cumming, C. Osuna, T. Gysi, M. Bianco, X. Lapillonne, O. Fuhrer, and T. C. Schulthess, "A review of the challenges and results of refactoring the community climate code COSMO for hybrid Cray HPC systems," *Proceedings of Cray User Group*, 2013.
- [13] J. M. Dennis, M. Vertenstein, P. H. Worley, A. A. Mirin, A. P. Craig, R. Jacob, and et. al., "Computational performance of ultra-high-resolution capability in the Community Earth System Model," *International Journal of High Performance Computing Applications*, vol. 26, no. 1, pp. 5–16, 2012.
- [14] R. B. Neale and et al., "Description of the NCAR Community Atmosphere Model (CAM 5.0)," Natl.Cent. for Atmos. Res., Boulder, Colo., Tech. Rep. Note NCAR/TN-4861STR.
- [15] J. M. Dennis, J. Edwards, K. J. Evans, O. Guba, P. H. Lauritzen, A. A. Mirin, and et. al., "CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model," *International Journal of High Performance Computing Applications*, vol. 26, no. 1, pp. 74–89, 2012.
- [16] J. Linford, J. Michalakes, M. Vachharajani, and A. Sandu, "Multi-core acceleration of chemical kinetics for simulation and prediction," in *High Performance Computing Networking, Storage and Analysis, Proceedings of the Conference on*, Nov 2009, pp. 1–11.
- [17] H. Fu, J. Liao, J. Yang, L. Wang, Z. Song, X. Huang, C. Yang, W. Xue, F. Liu, F. Qiao, W. Zhao, X. Yin, C. Hou, C. Zhang, W. Ge, J. Zhang, Y. Wang, C. Zhou, and G. Yang, "The sunway taihulight supercomputer: system and applications," *Science China Information Sciences*, pp. 1–16, 2016. [Online]. Available: <http://dx.doi.org/10.1007/s11432-016-5588-7>
- [18] R. J. Small, J. Bacmeister, D. Bailey, A. Baker, S. Bishop, F. Bryan, J. Caron, J. Dennis, P. Gent, H.-m. Hsu, M. Jochum, D. Lawrence, E. Muoz, P. diNezio, T. Scheitlin, R. Tomas, J. Tribbia, Y.-h. Tseng, and M. Vertenstein, "A new synoptic scale resolving global climate simulation using the community earth system model," *Journal of Advances in Modeling Earth Systems*, vol. 6, no. 4, pp. 1065–1094, 2014. [Online]. Available: <http://dx.doi.org/10.1002/2014MS000363>
- [19] D. Quinlan and C. Liao, "The ROSE Source-to-Source Compiler Infrastructure," in *Cetus Users and Compiler Infrastructure Workshop, in conjunction with PACT 2011*, October 2011.