

Chapter 40

High-Resolution FAMIL

Linjiong Zhou, Qing Bao and Haiyang Yu

Abstract This chapter describes the model speed and model in/out (I/O) efficiency of the high-resolution atmospheric general circulation model the Finite-volume Atmospheric Model of IAP/LASG (FAMIL), investigated at the National Supercomputer Center in Tianjin, China, on its Tianhe-1A supercomputer platform. A series of three-model-day simulations were conducted with the standard Aqua-Planet Experiment (APE) project, designed within FAMIL, to obtain time stamps for the calculation of model speed, simulation cost, and model I/O efficiency. The results of the simulation demonstrate that FAMIL has remarkable scalability below 3,456 and 6,144 cores; the lowest simulation costs were found to be 1,536 and 3,456 cores for 12.5 and 6.25 km resolutions, respectively. Furthermore, FAMIL has excellent I/O scalability and efficiencies of more than 80 and 99 % on 6 and 1,536 I/Os, respectively.

Keywords FAMIL · High resolution · Computational performance · Scalability · Efficiency

40.1 Introduction

Climate modeling is an essential tool that is used widely to improve understanding of the evolution of climate systems and project future climate. Therefore, the field of climate modeling aims to improve both the accuracy and performance of

L. Zhou · Q. Bao (✉) · H. Yu
State Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics (LASG), Institute of Atmospheric Physics (IAP), Chinese Academy of Sciences, Beijing 100029, China
e-mail: baoqing@mail.iap.ac.cn

L. Zhou · H. Yu
University of Chinese Academy of Science, Beijing 100049, China

modeling, helping provide confident simulations across a range of spatial scales, from the local to the global.

In the early days of numerical weather prediction (NWP) and climate simulation, the models used for these two separate applications were very different. NWP emphasized accurate prediction of fluid flow by applying the highest resolution possible, whereas climate simulation emphasized parameterized forcing, with conservation considered essential for very long runs (Williamson 2007). Furthermore, investigating the impact of climate change is a computationally expensive process that requires significant computational resources (Worley et al. 2011). Previously, the use of relatively low-resolution models was considered acceptable for climate simulation. In recent years, however, the advent of petascale computing has enabled the execution of a limited number of very high-resolution simulations (Dennis et al. 2012). From the continental scale to the cloud-resolving scale, the resolution of climate models has improved from several hundred kilometers to several kilometers, blurring the distinctions between climate models and weather forecasting models.

High-performance computers or supercomputers with tens of thousands of cores present opportunities for the development of high-resolution models (Satoh et al. 2008; Gall et al. 2011; Dennis et al. 2012), placing new requirements on high-performance computers: large numbers of cores, enormous storage, and rapid point-to-point communication. As this article is being written, the most powerful supercomputer in the world is Japan's K Computer (<http://top500.org/lists/2011/11>), which can achieve an impressive $10.51 \text{ Pflops}^{-1}$ using 705,024 processing cores. The top supercomputer in China is Tianhe-1A, which can achieve performance of 2.57 Pflops^{-1} with over 186,368 processing cores, 229.4 TB of memory, and a 1 PB In/Out (I/O) storage system; this supercomputer is ranked in second place on the TOP 500 list. High-performance computers allowed the development of high-resolution models in China (http://www.iap.cas.cn/xwzx/zhxw/201109/t20110922_3353034.html).

This chapter provides an overview of the computational performance of the new-generation high-resolution atmospheric model of the Institute of Atmospheric Physics/State Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics (IAP/LASG), run on the Tianhe-1A supercomputer for two representative simulations: model speed experiments, and model I/O efficiency experiments. In [Sect. 40.2](#), we describe the methodology used to perform the experiments, including model descriptions and experimental design. The experimental results are discussed in [Sect. 40.3](#). Finally, we present the conclusions of this study and a discussion of further research in [Sect. 40.4](#).

40.2 Methodology

40.2.1 Model Description

The Finite-volume Atmospheric Model of IAP/LASG (FAMIL) is the third generation of atmospheric models developed at IAP/LASG and was originally developed from the Spectral Atmospheric Model of IAP/LASG (SAMIL; Wu et al. 1996; Wang et al. 2004; Bao et al. 2010).

FAMIL adopts the finite-volume algorithm in the dynamic core; this algorithm is calculated on a cubed-sphere grid system (Figs. 40.1 and 40.2), thus avoiding the pole issue inherent in longitude–latitude grid systems (Lin and Rood 1996, 1997; Lin 1997, 1998, 2004; Putman and Lin 2007). A flux-form semi-Lagrangian transport scheme is used to calculate the advection terms in FAMIL, making it both stable and conservative (Lin and Rood 1996; Wang et al. 2013).

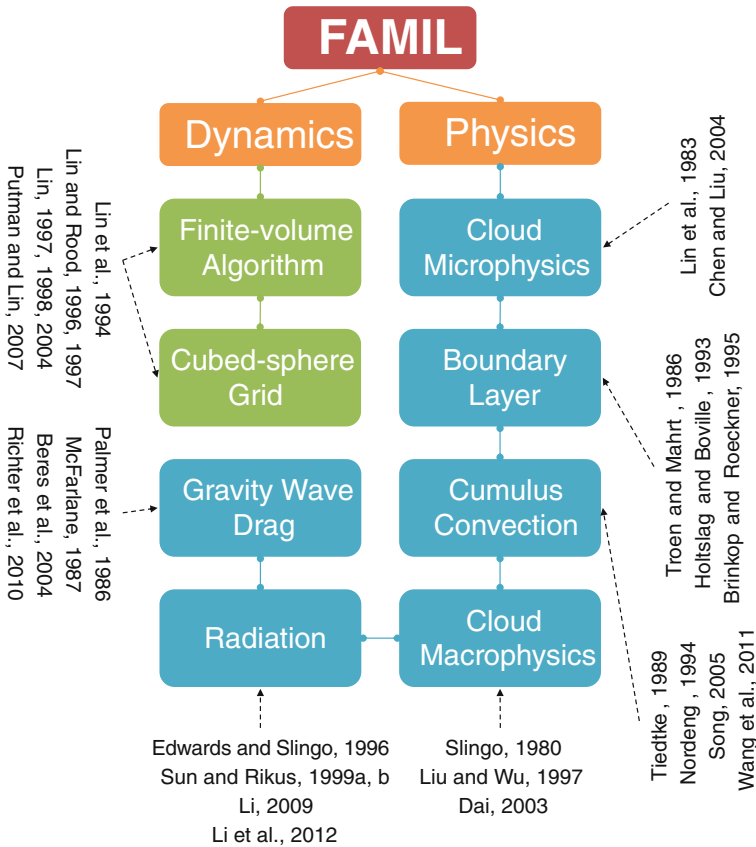


Fig. 40.1 Structure of FAMIL




Fig. 40.2 Cubed-sphere grid in FAMIL

In terms of the physical parameterization, the radiation scheme (Sun and Rikus 1999a, b; Sun 2005, 2008; Li et al. 2009a; Sun 2011; Li et al. 2012a) has been modified from a radiation scheme developed by the UK Meteorological Office (Edwards and Slingo 1996) based on the two-stream equation approach that is now used in HadCM3 (Martin et al. 2006). The mass flux cumulus parameterization of Tiedtke (1989) is employed with the modification by Nordeng (1994), in which the closure for deep convection is based on convective available potential energy instead of large-scale moisture convergence. The planetary boundary layer part of the model is a “non-local” scheme (Holtslag and Boville 1993) that computes the turbulent transfer of momentum, heat, and moisture. The cloud scheme is a diagnostic method based on vertical motion and relative humidity (Slingo 1980), but with a modification by Dai (2003). The gravity wave drag (referred as GWD) scheme is a multi-source GWD scheme from the Whole Atmosphere Community Climate Model (WACCM) (Richter et al. 2010), which considers not only conventional orography-induced GWD (Palmer et al. 1986; McFarlane 1987), but also convection- and frontogenesis-induced GWD (Beres et al. 2004; Richter et al. 2010).

The whole model is built on the Flexible Modeling System (FMS; <http://www.gfdl.noaa.gov/fms>), which is characterized by the flexibility of its resolution adjustment and its excellent parallelized algorithm (Williamson 2007; Donner et al. 2011). FAMIL will be further coupled with the land model to construct a full model able to perform an AMIP run.

In FAMIL, the core number can be set flexibly to 24, 54, 96, 216, 384, 864, 1,536, 3,456, 6,144, or 13,824 when different resolutions are used (e.g., 200, 100, 50, 25, 12.5, and 6.25 km). Meanwhile, the I/O number is also changeable under the rule that it must divide evenly by the core number. Based on these requirements, FAMIL has been tested with the large numbers of dynamical core standard cases proposed by Held and Suarez (1994), Williamson et al. (1992), Jablonowski and Williamson (2006), and APE, as proposed by Neale and Hoskins (2000). The results show that FAMIL exhibits excellent simulation performance.

In addition to focus on high horizontal resolution, we also wish to concentrate on high vertical resolution, which necessitates increasing both the vertical levels and model top. In the current version of FAMIL, there are 26 vertical levels with the top level at 2.19 hPa. Alternatively, FAMIL can use 32, 48, and 55 vertical levels with model top uplited to 0.01 hPa (Fig. 40.3).




Fig. 40.3 Four types of vertical level scheme built in FAMIL

40.2.2 Experimental Design

APE applies AGCMs with their complete parameterization packages to an idealization of the planet Earth that has a greatly simplified lower boundary consisting only of an ocean. It has no land or associated orography and no sea ice. The ocean is represented by sea surface temperatures (SSTs), which are specified throughout with simple idealized distributions. Therefore, in the hierarchy of tests available for AGCMs, APE falls between tests with simplified forcings, such as those proposed by Held and Suarez (1994) and Boer and Denis (1997), and the Earth-like simulations of the AMIP. APE aims to provide a benchmark for current model behaviors and to stimulate research toward an understanding of the cause of inter model differences (Williamson et al. 2012).

FAMIL is currently designed as a standard aqua planet model. The basic model configurations are recommended as follows.

- (1) Prescribed idealized SST distribution. No sea ice. Minimum SST is set at 0 °C.
- (2) Equinoctial insolation is fixed to be symmetric about the equator but vary with a diurnal cycle. Eccentricity and obliquity are set to zero. The solar constant is set to $1,365 \text{ Wm}^{-2}$.
- (3) Radiatively active gases (CO_2 , CH_4 , and N_2O) are globally fixed to 348, 1,650, and 306 ppbv, respectively. No radiatively active aerosol.
- (4) A zonally symmetric latitude–height distribution of ozone is specified, symmetrized about the equator, corresponding to the annual mean climatology used in AMIP II.

The experiments can be divided into two groups: model speed experiments, and model I/O efficiency experiments. In the model speed experiments, the model speed¹ (units are model years or months pre-wall-clock day, or MYPD or MMPD) and simulation cost (units are CPU hours pre-model year or month, or HPMY or HPMM) without I/O are calculated and evaluated as a function of the core number. In the model I/O efficiency experiments, the model I/O efficiency (units: %) for a specified resolution and core are calculated and evaluated according to the I/O number. Here, the definitions of model speed, simulation cost, and model I/O efficiency are as follows.

$$\text{Model_Speed} = \frac{\text{Model_Time}}{\text{Wall_Clock_Time}}$$

$$\text{Simulation_Cost} = \frac{\text{Wall_Clock_Time}}{\text{Model_Time}} \times \text{Core_Number}$$

¹ Model speed is different with model speedup. Model speedup is defined as the ratio between the wall-clock time of serial execution and the wall-clock time of parallel execution (Barney 2012).

$$\text{Model_IO_Efficiency} = \frac{\text{Wall_Clock_Time (without_IO)}}{\text{Wall_Clock_Time (with_IO)}}$$

Generally, climate modeling requires a large program that can be divided into three main processes between the initiation of the model run and its end, as follows:

- (1) model initialization, including reading the initial data, setting up common arrays, initializing the parallel calculation environment, identifying the time stamp, reading the name list file, and reading the restart file if it is a restart run;
- (2) model integration, indicating that the model is simulating atmospheric movement by numerically solving the atmospheric primary equations;
- (3) model termination, including writing out the restart file and simulation result, exiting the parallel calculation environment, and releasing the array space.

Of these processes, model integration is the most time consuming and, as the model time increases, it clearly occupies the majority of the wall-clock time. It is also notable that the computational requirement grows exponentially with increases in the resolution. Generally speaking, the increment of the computational requirement is eight to ten times larger when the resolution doubles (Song et al. 2010). In high-resolution experiments, computation times would be too long to perform model integrations over long time periods.

Based on the considerations above, some APE requirements, such as the six-model-month spin-up and the 3.5-model-year simulation, are not considered here to avoid any conflict with our experimental aims. All experiments were conducted at the National Supercomputer Center in Tianjin, China, on the Tianhe-1A supercomputer. Only three model days were performed, with 1,800-second time steps with resolutions of 12.5 and 6.25 km, using 216, 384, 864, 1,536, 3,456, 6,144, and 13,824 cores and I/O numbers of 0, 6, 24, 96, 384, and 1,536. After a three-model-day run, a time stamp was produced with all of the time information necessary to derive the model speed, simulation cost, and model I/O efficiency. In the model I/O efficiency experiments, zonal wind (u), meridional wind (v), specific humidity (q), air temperature (t), and surface pressure (p_s) were output at each time step. Of these, u , v , q , and t are three-dimensional variables with 26 levels on the third dimension, and p_s is a two-dimensional variable. For accuracy, the time costs of the model initialization and model termination were deducted from the calculation.

40.3 Overall Results

The model speeds of FAMIL with resolutions of 12.5 and 6.25 km are presented in the top panels of Figs. 40.4 and 40.5, respectively. The horizontal axes represent the number of cores, ranging from 0 to 6,500 and from 0 to 15,000, respectively,

Fig. 40.4 Model speed (units MYPD, top panel) and simulation cost (units HPMY, bottom panel) for 12.5 km resolution FAMIL for an aqua planet experiment as a function of Tianhe-1A supercomputer core numbers. From left to right, the circles are located at 216, 384, 864, 1,536, 3,456, and 6,144 cores


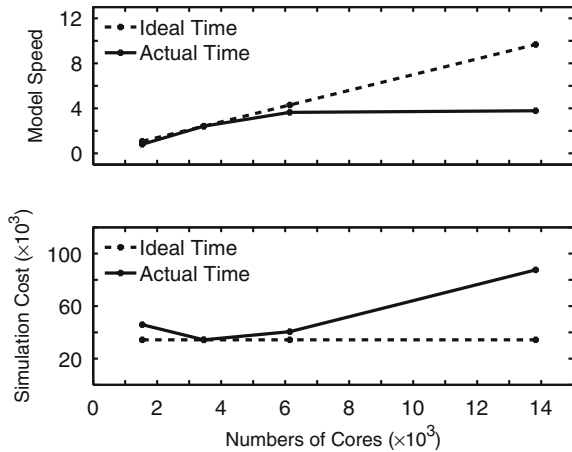


Fig. 40.5 Model speed (units MMPD, top panel) and simulation cost (units HPMM, bottom panel) for 6.25 km resolution FAMIL for an aqua planet experiment as a function of Tianhe-1A supercomputer core numbers. From left to right, the circles are located at 216, 384, 864, 1,536, 3,456, 6,144, and 13,824 cores



whereas the vertical axis represents MYPD and MMPD (ranging from 0.0 to 4.0 and from 0.0 to 12.0 in Figs. 40.4 and 40.5, respectively). The corresponding simulation costs are provided in the bottom panels, where the horizontal axes represent the number of cores (with the same ranges used in the upper panels). The vertical axes represent HPMY and HPMM, with ranges of 15,000–105,000 and 0–120,000 in Figs. 40.4 and 40.5, respectively. The solid lines represent actual model speed and actual simulation cost, indicating the experimental results, whereas the dashed lines indicated the ideal model speed and ideal simulation cost, derived according to the best results in the corresponding group of experiments. The circles on the solid and dashed lines represent actual and ideal values, respectively.

The process of seeking the best results among the experiments, which allows the determination of an ideal model speed and ideal simulation cost, is simple and straightforward. In the top panels of Figs. 40.4 and 40.5, assuming the model speed is equal to zero for zero cores, an ideal point $(x_0, y_0) = (0, 0)$ can be obtained. The straight line joining the actual point (x_i, y_i) , $[I = 1, 2, 3, \dots]$ and the ideal point (x_0, y_0) has a slope of λ_i , $[i = 1, 2, 3, \dots]$. Thus, the actual point (x_n, y_n) with the largest slope λ_n is the best result in this experiment, and the corresponding line indicates the ideal model speed. In the bottom panels of Figs. 40.4 and 40.5, the ideal simulation cost can be derived from the ideal model speed.


For a resolution of 12.5 km, the experimental results (Fig. 40.4, top panel) demonstrate that the model speed increases in almost linearly as a function of core number for less than 1,536 cores, approaching 1.0 MYPD at 1,536 cores. Despite the negligible difference between the actual model speed and the ideal model speed, they match well. However, the model speed slows down gradually as the core number increases above 1,536, resulting in 0.3 and 1.6 MYPD less than the ideal model speeds at 3,456 and 6,144 cores, which are 1.7 and 1.9 MYPD, respectively. Assuming that 0.3 MYPD accounts for only 17.6 % of 1.7 MYPD at 3,456 cores, it can be concluded that FAMIL has remarkable scalability when using fewer than 3,456 cores for a simulation resolution of 12.5 km.

The results for model with resolution of 6.25 km (Fig. 40.5, top panel) are approximately the same as those for a resolution of 12.5 km, except that the model speed begins to slow when more than 3,456 cores are used, producing values of 3.6 and 3.8 MMPD at 6,144 and 13,824 cores, respectively. These measurements are 0.7 and 5.9 MMPD, respectively, below the ideal model speeds. Therefore, assuming that 0.7 MMPD accounts for only 19.4 % of 3.6 MMPD at 6,144 cores, it can be concluded that FAMIL has remarkable scalability when using fewer than 6,144 cores for a resolution of 6.25 km.

Unlike the model speed, which indicates the average speed of every core it uses, the simulation cost represents the total model speed for all cores. With perfect model scalability and computer performance, the simulation cost should be independent of the core number (Dennis et al. 2012). However, as the bottom panels of Figs. 40.4 and 40.5 demonstrate, the simulation costs reach their minimum points at 1,536 and 3,456 cores for resolutions of 12.5 and 6.25 km, respectively, and increase toward both sides. On the right-hand side, the simulations using more cores have higher simulation costs; this is consistent with the drop in the model speed in the top panel. Conversely, on the left-hand side, those using fewer cores also have higher simulation costs, highlighting the small difference between the ideal model speed and the actual model speed in the top panel. Consequently, FAMIL performs best at 1,536 and 3,456 cores for resolutions of 12.5 and 6.25 km, respectively, because the lowest simulation costs occur under these conditions.

I/O operations are generally regarded as inhibitors of parallelism. Considering the considerable difference I/O that operations can make, model I/O efficiency is also an essential indicator of model performance. We obtained a total of approximately 200 GB data in each case for the model I/O efficiency experiments.

Fig. 40.6 Model I/O efficiency (%) for 12.5 km resolution FAMIL for an aqua planet experiment using 1,536 Tianhe-1A supercomputer cores as a function of the I/O number



The results are presented in Fig. 40.6 and demonstrate that, if there is no I/O, I/O operations are free of cost and efficiency is 100 %. When six I/Os are opened, the model I/O efficiency is maintained at no less than 80 %. The model I/O efficiency also increases as the I/O number continues to increase. If all cores (i.e., 1,536 cores) have I/Os, the model I/O efficiency approaches 100 %, indicating that FAMIL has excellent I/O scalability and efficiency.

40.4 Conclusions and Discussion

This chapter describes the assessment of the computational performance of the high-resolution atmospheric model FAMIL on the Tianhe-1A supercomputer at the National Supercomputer Center in Tianjin, China, based on two groups of experiments: model speed experiments, and model I/O efficiency experiments. Although the scientific results from this three-model-day aqua planet run are still preliminary, they indicate that FAMIL exhibits impressive performance.

Based on the wall-clock time for three model days of runs, we have demonstrated climatologically useful model speeds at approximately 1.0 and 1.7 MYPD for 1,536 and 3,456 cores, respectively, at a resolution of 12.5 km. For resolution of 6.25 km, the model speed was found to drop to approximately 2.5 and 3.6 MMPD for 3,456 and 6,144 cores, respectively. The best scalabilities occur for fewer than 3,456 and 6,144 cores for resolutions of 12.5 and 6.25 km, respectively. The simulation cost results for both resolutions demonstrate that the simulation cost is not independent of the core number but exhibits its minimum value at a specific core number that depends on the resolution: 1,536 and 3,456 cores for resolutions of 12.5 and 6.25 km, respectively.

A number of factors may contribute to the slowing at large core numbers and the dependence of the simulation cost on the core number, all of which can be attributed to issues of scalability. The primary four factors (Barney 2012) are as follows: (1) hardware, particularly memory-CPU bandwidths and network communications; (2) application algorithms; (3) parallel overhead; and (4) characteristics of the specific application and its coding. In our experimental design, the influence of parallel overhead has been determined, but the mechanisms influencing the remaining three factors remain unknown.

Although we have yet not achieved perfect model speed performance, this does not mean that FAMIL is potentially limited. Further efforts will focus on OpenMP parallelized architecture, which is not included in the current version of FAMIL and will provide considerable improvements in model speed.

We performed model I/O efficiency experiments with enormous I/O operations to test the I/O performance. Our experimental results demonstrate that FAMIL has excellent I/O scalability and efficiency, which will allow us to perform many productive simulations for model tuning and other experiments.