Computational Performance of the High-Resolution Atmospheric Model FAMIL

ZHOU Lin-Jiong[1,2], LIU Yi-Min[1], BAO Qing[1], YU Hai-Yang[1,2], and WU Guo-Xiong[1]

[1] *The 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*
[2] *Graduate University of Chinese Academy of Sciences, Beijing 100049, China*

Received 27 March 2012; revised 12 April 2012; accepted 13 April 2012; published 16 September 2012

Abstract This paper describes the model speed and model In/Out (I/O) efficiency of the high-resolution atmospheric general circulation model FAMIL (Finite-volume Atmospheric Model of IAP/LASG) at the National Supercomputer Center in Tianjin, China, on its Tianhe-1A supercomputer platform. A series of three-model-day simulations were carried out with standard Aqua Planet Experiment (APE) designed within FAMIL to obtain the time stamp 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 3456 and 6144 cores, and the lowest simulation costs are 1536 and 3456 cores for 12.5 km and 6.25 km resolutions, respectively. Furthermore, FAMIL has excellent I/O scalability and an efficiency of more than 80% on 6 I/Os and more than 99% on 1536 I/Os.

Keywords: FAMIL, high-resolution, computational performance, scalability, efficiency

Citation: Zhou, L.-J., Y.-M. Liu, Q. Bao, et al., 2012: Computational performance of the high-resolution atmospheric model FAMIL, *Atmos. Oceanic Sci. Lett.*, **5**, 355–359.

1 Introduction

Climate modeling is an essential tool that is widely used to provide an understanding of the evolution of climate systems and to make climate prediction. Therefore, the field of climate modeling is responsible for improving both its accuracy and performance to 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 of these two applications were very different. NWP emphasized accurate predictions of fluid flow by applying the highest resolution, while climate simulation emphasized parameterized forcing, with conservation considered essential for very long runs (Williamson, 2007). In addition, investigating the impact of climate change is a computationally expensive process that requires significant computational resources (Worley et al., 2011). A relatively low resolution was acceptable in climate simulation. In recent years, the advent of petascale computing 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 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 the opportunity to develop high-resolution models (Satoh et al., 2008; Gall et al., 2011, Dennis et al., 2012), which places new requirements on high-performance computers for a large amount 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), achieving an impressive 10.51 Pflop s^{-1} using 705024 processing cores. In China, the top supercomputer is Tianhe-1A with 2.57 Pflop s^{-1} performance, over 186368 processing cores, 229.4 TB of memory, and a 1 PB of In/Out (I/O) storage system, ranking No. 2 in the TOP500 list. High-performance computers gave birth to the development of high-resolution models in China (http://www.iap.cas.cn/xwzx/zhxw/201109/t20110922_3353034.html).

This paper provides an overview of the computational performance of the new-generation, high-resolution atmospheric model FAMIL on the Tianhe-1A supercomputer for two representative simulations: model speed experiments and model I/O efficiency experiments. In section 2, we describe the methodology used to perform the experiments, including model description and experiment design. The experimental results are discussed in section 3. Finally, we present the conclusions of this study and a discussion of further research in section 4.

2 Methodology

2.1 Model description

FAMIL is the latest atmospheric model in IAP/LASG (Institute of Atmospheric Physics/State Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics) developed based on SAMIL (Spectral Atmospheric Model of IAP/LASG) in recent years. The earliest version of SAMIL was a nine-level rhomboidally truncated spectral model designed by William Bourke in Australia (Bourke, 1974). It was improved by Simmonds (Simmonds, 1985) and introduced to China by Yuanbi Lin in 1987. In 1991, this spectral model was introduced to the Institute of Atmospheric Physics and

Corresponding author: LIU Yi-Min, lym@lasg.iap.ac.cn

underwent considerable developments in both its dynamic framework and physical progresses (Wu et al., 1996; Liu and Wu, 1997; Liu et al., 1998; Shao et al., 1998; Wang et al., 2000). Despite its coarse horizontal and vertical resolution, SAMIL was capable of simulating climate mean states, monsoon onsets, and some inter-annual variability (Wu et al., 1996).

SAMIL has experienced many developments in the past decade, including improvements to its horizontal and vertical resolutions from R15 and 9 levels to R42 (Wu et al., 2003) and 26 levels (Wang et al., 2004) and conversion from a serial model to parallelized one (Wang and Wang, 2006) called SAMIL2. Thus far, SAMIL2 has been widely used in the areas of climate change and dynamic studies, and it has been included in coordinated AGCM inter-comparison projects (e.g., AMIP: Atmospheric Model Intercomparison Project; APE Project: Aqua Planet Experiment Project).

As a new atmospheric model, FAMIL concentrates on high-resolution and progressive parameterization. In contrast to SAMIL, FAMIL contains the advanced finite volume dynamical core designed by the Geophysical Fluid Dynamics Laboratory (GFDL) (Lin and Rood, 1996, 1997; Lin, 1997, 1998, 2004; Putman and Lin, 2007), which, shares the same set of physical parameterizations as SAMIL temporarily. The crucial characteristics of the finite volume dynamical core are its stable long time steps due to a flux-form semi-Lagrangian transport scheme, its automatic conservation due to monotonicity constraints, its freedom from the pole problem when adapting the cubed sphere grid, 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 compose a full model able to perform an AMIP run.

In FAMIL, the core number can be flexibly set to 24, 54, 96, 216, 384, 864, 1536, 3456, 6144, or 13824 when different resolutions are used, such as 200 km, 100 km, 50 km, 25 km, 12.5 km, and 6.25 km. Meanwhile, the I/O number is also changeable under the rule that the I/O number must be divided 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 (2001). The results show that FAMIL exhibits an excellent simulation performance.

2.2 Experiment 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 the associated orography and no sea ice. The ocean is represented by Sea Surface Temperatures (SST), 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 et al. (1994) and Boer et al. (1997), and the Earth-like simulations of the AMIP. The goals of APE are to provide a benchmark of the current model behaviors and to stimulate research toward an understanding of the cause of inter-model differences (Williamson et al., 2011).

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. The minimum SST is set at 0°C.

2) Fix equinoctial insolation to be symmetric about the equator but with the diurnal cycle. Eccentricity and obliquity are set to zero. The solar constant is set to 1365 W m^{-2}.

3) Radiatively active gases, CO_2, CH_4, and N_2O are globally fixed to 348 ppmv, 1650 ppbv, 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 speed[①] (units: Model Years or Months Pre-wall clock Day, MYPD or MMPD) and simulation cost (units: CPU Hours Pre-Model Year or Month, 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: %) with 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:

$$\text{Model\_Speed} = \frac{\text{Model\_Time}}{\text{Wall\_Clock\_Time}}$$

$$\text{Simulation\_Cost} = \frac{\text{Wall\_Clock\_Time}}{\text{Model\_Time}} \times \text{Core\_Number}$$

$$\text{Model\_IO\_Efficiency} = \frac{\text{Wall\_Clock\_Time(without\_IO)}}{\text{Wall\_Clock\_Time(with\_IO)}}$$

Generally, climate modeling requires a large program that can be divided into three main processes between the initiation of running 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 namelist 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

[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).

occupies the majority of the wall-clock time. It is also notable that the computational requirement grows exponentially along 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, it would cost too much time to perform long time model integrations.

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 purpose. All experiments are carried out at National Supercomputer Center in Tianjin, China on their Tianhe-1A supercomputer, performing only three model days with 1800-second time steps with resolutions of 12.5 km and 6.25 km, using 216, 384, 864, 1536, 3456, 6144, and 13824 cores and I/O numbers of 0, 6, 24, 96, 384, and 1536. After a three-model-day run, a time stamp is produced with all 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) are output at each time step. Among them, 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 in the calculation.

3 Overall results

The model speeds of FAMIL with 12.5 km and 6.25 km resolutions are provided in the top panels of Fig. 1 and Fig. 2. The horizontal axis is the number of cores ranging from 0 to 6500 and from 0 to 15000, while the vertical axis shows MYPD and MMPD ranging from 0.0 to 4.0 and from 0.0 to 12.0. Correspondingly, their simulation costs are provided in the bottom panels. The horizontal axis is the number of cores with the same ranges used in the upper panels. The vertical axis presents HPMY and HPMM ranging from 15000 to 105000 and from 0 to 120000. The solid lines are the actual model speed and actual simulation cost, which indicate the result of the experiments, while the dashed lines are the ideal model speed and ideal simulation cost, derived according to the best results in the corresponding group of experiments. The circles on solid lines and dashed lines represent the actual values and ideal values, respectively.

The process of looking for the best results among the experiments, which leads to the development of an ideal model speed and ideal simulation cost, is simple and straightforward. In the top panels of Fig. 1 and Fig. 2, supposing the model speed is equal to zero at zero cores, we obtain an ideal point $(x_0, y_0)=(0, 0)$. If we draw a straight line joining the actual point (x_i, y_i), (i=1, 2, 3,…) and the ideal point (x_0, y_0), we obtain a slope of λ_i, (i=1, 2, 3,…). 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 Fig. 1 and Fig. 2, the ideal simulation cost is derived from the ideal model speed.

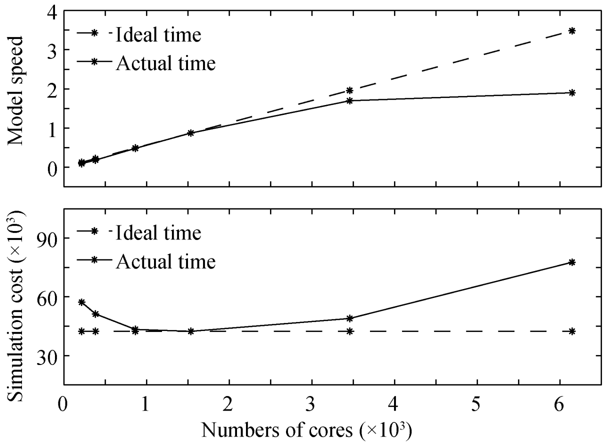

Figure 1 Model speed (units: MYPD, top panel) and simulation cost (units: HPMY, bottom panel) for 12.5 km resolution FAMIL on aqua planet experiment as a function of Tianhe-1A supercomputer core numbers. From left to right, the circles are located at 216, 384, 864, 1536, 3456, and 6144 cores.

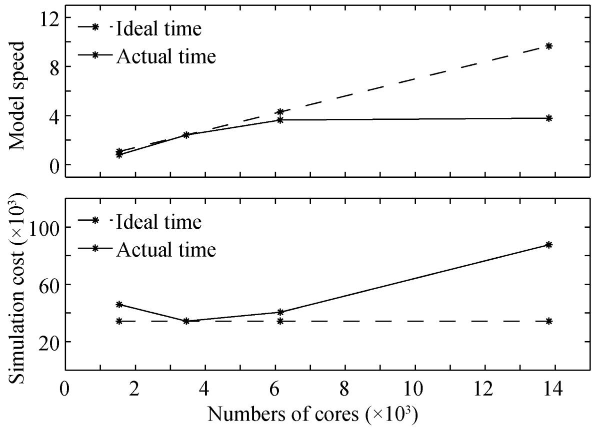

Figure 2 Model speed (units: MMPD, top panel) and simulation cost (units: HPMM, bottom panel) for 6.25 km resolution FAMIL on aqua planet experiment as a function of Tianhe-1A supercomputer core numbers. From left to right, the circles are located at 216, 384, 864, 1536, 3456, 6144, and 13824 cores.

For the 12.5 km resolution, the result of experiments (Fig. 1, top panel) demonstrates that the model speed increases in a nearly linear fashion as a function of a core number with less than 1536 cores, approaching 1.0 MYPD at 1536 cores. Despite this negligible difference, the actual model speed matches the ideal model speed well. However, as the core number increases from 1536, the model speed slows down gradually, resulting in 0.3 MYPD and 1.6 MYPD less than the ideal model speeds at 3456 and 6144 cores, which are 1.7 MYPD and 1.9 MYPD, respectively. Considering that 0.3 MYPD accounts for only 17.6% of 1.7 MYPD at 3456 cores, we can conclude that FAMIL has remarkable scalability when using fewer than 3456 cores for a 12.5 km resolution simulation.

The results for the 6.25 km resolution (Fig. 2, top panel) are approximately the same as those of the 12.5 km resolution, with the exception that the model speed begins to slow when more than 3456 cores are used: 3.6 MMPD at 6144 cores and 3.8 MMPD at 13824 cores. These measurements are 0.7 MMPD and 5.9 MMPD below the ideal model speeds, respectively. Therefore, considering that 0.7 MMPD accounts for only 19.4% of 3.6 MMPD at 6144 cores, we can conclude that FAMIL has remarkable scalability when using fewer than 6144 cores for a 6.25 km resolution.

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 would be independent of the core number (Dennis et al., 2011). However, as the bottom panels of Fig. 1 and Fig. 2 demonstrate, the simulation costs reach their minimum points at 1536 and 3456 cores for the 12.5 km and 6.25 km resolutions and increase toward both sides. On the right side, the simulations using more cores have higher simulation costs, which is consistent with the drop in the model speed in the top panel, while 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. As a consequence, FAMIL performs best at 1536 and 3456 cores for the 12.5 km and 6.25 km resolutions because the lowest simulation costs occur under those conditions.

I/O operations are generally regarded as inhibitors of parallelism. Considering the great difference I/O operations can make, model I/O efficiency is also an essential indicator of the model's performance. We obtained a total of approximately 200 GB data in each case in the model I/O efficiency experiments. The results are provided in Fig. 3, which demonstrate that if there is no I/O, I/O operations are free of cost, and the efficiency is 100%. When six I/Os are opened, the model I/O efficiency is maintained at no less than 80%. As the I/O number continues to increase, the model I/O efficiency also becomes higher. If all the cores have I/Os (1536 cores), the model I/O efficiency approaches 100%, indicating that FAMIL has excellent I/O scalability and efficiency.

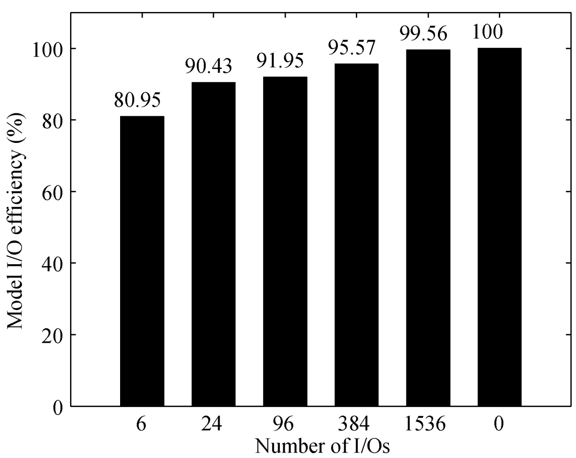

Figure 3 Model I/O efficiency (%) for 12.5 km resolution FAMIL on aqua planet experiment using 1536 Tianhe-1A supercomputer cores as a function of the I/O number.

4 Conclusions and discussion

In this paper, the computational performance of the high-resolution atmospheric model FAMIL is assessed on the Tianhe-1A supercomputer at National Supercomputer Center in Tianjin, China, according to 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 has an 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 MYPD at 1536 cores and 1.7 MYPD at 3456 cores for 12.5 km resolution. At the 6.25 km resolution, the model speed dropped to approximately 2.5 MMPD at 3456 cores and 3.6 MMPD at 6144 cores. The best scalabilities occur at fewer than 3456 cores for a 12.5 km resolution and fewer than 6144 cores for a 6.25 km resolution. The simulation cost results for both the 12.5 km and 6.25 km resolutions demonstrate that the simulation cost is not independent of the core number but has its minimum value at a specific core number, depending on the resolution. They are 1536 and 3456 cores for the 12.5 km and 6.25 km resolutions.

A number of factors may contribute to the slowing at large core numbers and the dependence between the simulation cost and core number, all of which can be attributed to issue 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 of influence of the remaining three factors are still 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 a great improvement to the model speed performance.

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 allows us to perform many productive simulations for model tuning and other experiments.

Acknowledgements. We are indebted to Dr. Shian-Jiann LIN for his great support and encouragement during the development of FAMIL. We would like to thank the National Supercomputer Center in Tianjin, China, for access to their Tianhe-1A supercomputer platform. This study was supported by the Strategic Priority Research Program of the Chinese Academy of Sciences (Grant No. XDA 05110303), the National Basic Research Program of China (973 Program, Grant Nos. 2012CB417203 and 2010CB950404), the National High Technology Research and Development Program of China (863 Program, Grant No. 2010AA012305), and the National Natural Science Foundation of China (Grant No. 41023002).

References

Barney, B., 2012: Introduction to parallel computing, in: *Training*

Materials at Lawrence Livermore National Laboratory, available at https://computing.llnl.gov/tutorials/parallel_comp/.
Boer, G. J., and B. Denis, 1997: Numerical convergence of the dynamics of a GCM, *Climate Dyn.*, **13**, 359–374.
Bourke, W., 1974: A multi-level spectral model. I. Formulation and hemispheric integrations, *Mon. Wea. Rev.*, **102**, 687–701.
Dennis, J. M., M. Vertenstein, P. H. Worley, et al., 2012: Computational performance of ultra-high-resolution capability in the community earth system model, *Int. J. High Perf. Comput. Appl.*, **26**, 5–16.
Donner, L. J., B. Wyman, R. S. Hemler, et al., 2011: The dynamical core, physical parameterizations, and basic simulation characteristics of the atmospheric component AM3 of the GFDL global coupled model CM3, *J. Climate*, **24**, 3484–3518.
Gall, J. S., I. Ginis, S.-J. Lin, et al., 2011: Experimental tropical cyclone prediction using the GFDL 25-km-resolution global atmospheric mode, *Wea. Forecasting*, **26**, 1008–1019.
Held, I. M., and M. J. Suarez, 1994: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models, *Bull. Amer. Meteor. Soc.*, **75**, 1825–1830.
Jablonowski, C., and D. L. Williamson, 2006: A baroclinic instability test case for atmospheric model dynamical cores, *Quart. J. Roy. Meteor. Soc.*, **132**, 2943–2975.
Lin, S. J., 1997: A finite-volume integration method for computing pressure gradient forces in general vertical coordinates, *Quart. J. Roy. Meteor. Soc.*, **123**, 1749–1762.
Lin, S. J., 1998: Reply to comments by T. Janjic on "A finite-volume integration method for computing pressure gradient forces in general terrain-following coordinates", *Quart. J. Roy. Meteor. Soc.*, **124**, 2531–2533.
Lin, S. J., 2004: A "vertically Lagrangian" finite-volume dynamical core for global models, *Mon. Wea. Rev.*, **132**, 2293–2307.
Lin, S. J., and R. B. Rood, 1996: Multidimensional flux form semi-Lagrangian transport schemes, *Mon. Wea. Rev.*, **124**, 2046–2070.
Lin, S. J., and R. B. Rood, 1997: An explicit flux-form semi-Lagrangian shallow water model on the sphere, *Quart. J. Roy. Meteor. Soc.*, **123**, 2477–2498.
Liu, H., and G. X. Wu, 1997: Impacts of land surface on climate of July and onset of summer monsoon: A Study with an AGCM plus SSiB, *Adv. Atmos. Sci.*, **14**, 289–308.
Liu, H., X. H. Zhang, and G. X. Wu, 1998: Cloud feedback on SST variability in western equatorial Pacific in a CGCM, *Adv. Atmos. Sci.*, **15**, 410–423.
Neale, R. B., and B. J. Hoskins, 2001: A standard test for AGCMs including their physical parameterizations. I: The proposal, *Atmos. Sci. Lett.*, **1**, doi:10.1006/asle.2000.0020.
Putman, W. M., and S. J. Lin, 2007: Finite-volume transport on various cubed-sphere grid, *J. Comput. Phys.*, **227**, 55–78.
Satoh, M., T. Matsuno, H. Tomita, et al., 2008: Nonhydrostatic icosahedral atmospheric model (NICAM) for global cloud resolving simulations, *J. Comput. Phys.*, **227**, 3486–3514.
Shao, H., Y. F. Qian, and Q. Q. Wang, 1998: The effects of the diurnal variation of solar radiation on climate modeling of R15L9, *Plateau Meteor.* (in Chinese), **17**, 158–169.
Simmonds, I., 1985: Analysis of the "spinning" of a global circulation model, *J. Geophys. Res.*, **90**, 5637–5660.
Song, Z. Y., H. X. Liu, X. Y. Lei, et al., 2010: The application of GPU in ocean general circulation mode POP, *Comput. Appl. Softw.* (in Chinese), **27**, 27–29.
Wang, B., H. Liu, and G. Y. Shi, 2000: Radiation and cloud schemes, in: *IAP Global Ocean-Atmosphere-Land System Model*, X. H. Zhang et al. (Eds.), Science Press, Beijing, 28–49.
Wang, P. F., and Z. Z. Wang, 2006: A Parallelization scheme for atmospheric general circulation model, *Chinese J. Atmos. Sci.* (in Chinese), **30**, 519–525.
Wang, Z. Z., G. X. Wu, T. W. Wu, et al., 2004: Simulation of Asian monsoon seasonal variations with climate model R42L9/LASG, *Adv. Atmos. Sci.*, **21**, 879–889.
Williamson, D. L., J. B. Drake, J. J. Hack, et al., 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry, *J. Comput. Phys.*, **102**, 211–224.
Williamson, D. L., M. Blackburn, B. J. Hoskins, et al., 2011: *The APE Atlas*, NCAR/TN-484+STR, NCAR Technical Note, Boulder, Colorado, 499pp.
Williamson, D. L., 2007: The evolution of dynamical cores for global atmospheric models, *J. Meteor. Soc. Japan*, **85B**, 241–269.
Worley, P. H., A. P. Craig, J. M. Dennis, et al., 2011: Performance of the community earth system model, in: *High Performance Computing, Networking, Storage and Analysis (SC)*, 2011 International Conference, Seattle, WA, 1–11.
Wu, G. X., H. Liu, Y. C. Zhao, et al., 1996: A nine-layer atmospheric general circulation model and its performance, *Adv. Atmos. Sci.*, **13**, 1–18.
Wu, T. W., P. Liu, Z. Z. Wang, et al., 2003: The performance of atmospheric component model R42L9 of GOALS/LASG, *Adv. Atmos. Sci.*, **20**, 726–742.