



Evaluation of the computational performance of the finite-volume atmospheric model of the IAP/LASG (FAMIL) on a high-performance computer

Jin-Xiao LI, Qing BAO, Yi-Min LIU & Guo-Xiong WU

To cite this article: Jin-Xiao LI, Qing BAO, Yi-Min LIU & Guo-Xiong WU (2017) Evaluation of the computational performance of the finite-volume atmospheric model of the IAP/LASG (FAMIL) on a high-performance computer, Atmospheric and Oceanic Science Letters, 10:4, 329-336, DOI: [10.1080/16742834.2017.1331111](https://doi.org/10.1080/16742834.2017.1331111)

To link to this article: <http://dx.doi.org/10.1080/16742834.2017.1331111>



© 2017 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 02 Jun 2017.



[Submit your article to this journal](#)



Article views: 298



[View related articles](#)



[View Crossmark data](#)

Evaluation of the computational performance of the finite-volume atmospheric model of the IAP/LASG (FAMIL) on a high-performance computer

LI Jin-Xiao^{a,b}, BAO Qing^a, LIU Yi-Min^a and WU Guo-Xiong^a

^aState Key Laboratory of Numerical Modeling for Atmospheric Sciences and Geophysical Fluid Dynamics (LASG), Institute of Atmospheric Physics, Chinese Academy of Sciences, Beijing, China; ^bCollege of Earth Science, University of Chinese Academy of Science, Beijing, China

ABSTRACT

High computational performance is extremely important for climate system models, especially in ultra-high-resolution model development. In this study, the computational performance of the Finite-volume Atmospheric Model of the IAP/LASG (FAMIL) was comprehensively evaluated on Tianhe-2, which was the world's top-ranked supercomputer from June 2013 to May 2016. The standardized Atmospheric Model Inter-comparison Project (AMIP) type of experiment was carried out that focused on the computational performance of each node as well as the simulation year per day (SYPD), the running cost speedup, and the scalability of the FAMIL. The results indicated that (1) based on five indexes (CPU usage, percentage of CPU kernel mode that occupies CPU time and of message passing waiting time (CPU_SW), code vectorization (VEC), average of Gflops (Gflops_AVE), and peak of Gflops (Gflops_PK)), FAMIL shows excellent computational performance on every Tianhe-2 computing node; (2) considering SYPD and the cost speedup of FAMIL systematically, the optimal Message Passing Interface (MPI) numbers of processors (MNPs) choice appears when FAMIL use 384 and 1536 MNPs for C96 (100 km) and C384 (25 km), respectively; and (3) FAMIL shows positive scalability with increased threads to drive the model. Considering the fast network speed and acceleration card in the MIC architecture on Tianhe-2, there is still significant room to improve the computational performance of FAMIL.

摘要

高性能计算指标对于气候模式的发展，尤其是高分辨率气候模式的发展至关重要。本文系统性的评估了IAP / LASG新一代有限体积大气环流模式FAMIL在Tianhe-2超算平台上的计算性能。通过标准的AMIP试验，本文重点评估FAMIL在每个节点的高性能指标、每天所积分的模式年(SYPD)、计算成本加速比、可扩展性五方面的性能。结果表明：1) FAMIL在CPU使用率、CPU节点间信息传输等待时间、代码向量化、Gflops平均值、Gflops峰值五个方面表现出优异性能。2) 综合考虑SYPD和计算成本加速比，对于C96(100km)和C384(25km)水平分辨率的FAMIL积分，最优进程使用数(MNPs)分别为384和1536。3) FAMIL具有良好的可扩展性，并且随着Tianhe-2网络传输速度的提升及MIC加速的使用，FAMIL的计算性能可以进一步得到提升。

ARTICLE HISTORY

Received 26 December 2016
Revised 31 March 2017
Accepted 19 April 2017

KEYWORDS

FAMIL; scalability;
computational performance;
Tianhe-2

关键词

FAMIL; 可扩展性; 计算性能;
Tianhe-2

1. Introduction

To better represent topographic forcing, reduce uncertainty from physical parameterization and capture small-scale phenomena, such as tropical cyclones, middle-scale vortices, extreme precipitation, atmospheric convection, SST-wind feedback, mid-latitude blocking systems, etc., high resolution modeling is a crucial aspect of climate system modeling development (Haarsma et al. 2016).

With the development of computing science and scientific research, new research directions have appeared, such as computational chemistry (Kendall et al. 2000), computational biology (Goldberg 1989), and computational fluid dynamics (Versteeg and Malalasekera 1995). Using

mathematical and physical models to perform scientific computing has become an important research method. On the other hand, the level of high performance computing available has already become a vital index of comprehensive national power. The TOP500 project (<http://www.top500.org/>), which introduces and ranks the 500 fastest supercomputers in the world, was started in 1993. However, not until the appearance of Tianhe, the supercomputer system developed by the National University of Defense Technology, did the hardware specifications of high performance computing in China rank in the top 10 globally. Particularly, Tianhe-2, a new generation supercomputer located in the Guangzhou National Supercomputer Center (NSCC-GZ), was ranked as the fastest from June 2013 to May 2016.

CONTACT BAO Qing  baqing@mail.iap.ac.cn

 The supplemental data for this article is available online at <https://doi.org/10.1080/16742834.2017.1331111>.

© 2017 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The Finite-volume Atmospheric Model of the IAP/LASG (FAMIL) is a new generation of an atmospheric model in the LASG. The aim of FAMIL is to build a high performance Atmospheric General Circulation Model (AGCM) with super high resolutions in both the horizontal and vertical directions. Compared with the previous version, SAMIL (Spectral Atmospheric Model of the IAP/LASG) (Bao et al. 2010, 2013), the most significant changes are the model's dynamical core, advection scheme, and grid system. The dynamic core of FAMIL adopts a finite-volume method and an FFSL advection scheme (Wang et al. 2013), which are calculated on a cubic sphere grid system (Lin 1997, 1998, 2004; Putman and Lin 2007). With a better solution in which grid spacing is narrow in polar regions, it is therefore logical that FAMIL improves horizontal resolution efficiently. Additionally, it is also possible to improve the parallel efficiency of FAMIL. Furthermore, C96 is a standard resolution of FAMIL in CMIP6 DECK (Eyring et al. 2016), and C384 is the greatest resolution of FAMIL for CMIP6 HighresMIP (Haarsma et al. 2016). The notation C, for example, C96, means 96×96 grid points in each tile of the cube sphere.

Previous work has been done to test the computational performance of FAMIL with the Tianhe-1 supercomputer. However, all previous test cases were in the idealized scenario of an Aqua Planet condition (Zhou et al. 2012), which does not consider land-surface processes, real SST, real topography, etc. (Neal and Hoskins 2001a, 2001b). During recent years, FAMIL has been successfully coupled with the NCAR coupler 7 (Craig, Vertenstein, and Jacob 2012), which consists of a driver that controls the top level sequencing, the processor decomposition, and communication between the components and the coupler. Simultaneously, coupler operations, such as mapping and merging, are run under the driver on a subset of processors as if there was a unique coupler model component, which means faster speeds and greater flexibility when using FAMIL. Consequently, the standard AMIP type of experiment (Gates 1992; Gates et al. 1999) can be carried out with multi-resolutions. Many development and numerical FAMIL experiments, for example, using CMIP6 and developing an intra-seasonal forecast system, will be based on AMIP type designs in the future. It is necessary, therefore, to perform a systematic test of FAMIL's computing performance on Tianhe-2.

This paper gives a comprehensive introduction, which includes the computing performance of each node, simulation year per day (SYPD), running cost speedup, and scalability of the atmosphere-land coupled Finite-volume Atmospheric Model of the IAP/LASG (FAMIL) on HPC Tianhe-2. Then, in Section 2, we introduce the methods, including the experimental design and the quantitation

standard. The results are given in Section 3, and lastly, a discussion and concluding remarks are provided in Section 4.

2. Methods

The model used in this study is FAMIL. With respect to the physical parameterizations, vertical diffusion, cumulus convection, cloud microphysics, radiation transfer, and gravity wave drag are included in FAMIL. The vertical layers of FAMIL can be set to 26, 32, 48, or 55, and different Message Passing Interface (MPI) numbers of processors (MNPs), such as 24, 54, 96, 216, 384, 864, 1536, 3456, 6144, or 13,824, can be used to drive FAMIL at different resolutions, such as C48 (~200 km), C96 (~100 km), C192 (~50 km), or C384 (~25 km) (Zhou et al. 2014). The FAMIL model adapted the dynamic core of the finite-volume method and is calculated on a cube-sphere grid system. The global grid system of FAMIL is split into six tiles, and each tile is composed of N^2 grids (where 'N' stands for the horizontal resolution of FAMIL, usually $N = 48$ is approximately 200 km and $N = 384$ is approximately 25 km). For parallelization, 2-D parallel decomposition is used for each of the the six tiles of cube-sphere grid, and the physical processes still use this way of parallel decomposition. Besides, each tile uses the same number of Message Passing Interface (MPI) number of processors (MNPs). If we use 16, 36, 64, or 144 MNPs on each tile, the total MNPs are 96, 216, 384, or 864, respectively. As illustrated in Figure S1, the horizontal resolution is Cube-sphere 48 (C48), which is approximately 200 km. If 96 MNPs are used in this case, then 16 MNPs are used in each tile, and each MNPs handles $48 \times 48 / 16 = 144$ horizontal grids.

All experiments in this paper were carried out on a new generation supercomputer, Tianhe-2, which is located at the Guangzhou National Supercomputer Center (NSCC-GZ) and was ranked as the fastest supercomputer from June 2013 to May 2016. Tianhe-2 consists of 16,000 computer nodes, each of which includes two 12-core Intel Ivy Bridge Xeon Phi CPUs (Xeon E5-2692) running at 2.2 GHz, three Xeon Phi coprocessor chips, and 64 GB memory with a maximum memory bandwidth of 14 GB/s \times 8 lanes (TH2 Express-2). Additionally, Tianhe-2's memory reaches 1024000 GB. Tianhe-2's theoretical floating-point peak calculation reaches 54902.4 Tflops, and its practical floating-point peak calculation is 33862.7 Tflops (<http://www.nscg-z.cn/>).

To investigate the computing performances of each node associated with multitasking jobs, we carried out the experiments using MNPs of 216, 384, 1536, 3456, and 6144 at FAMIL standard resolution C96 (~100 km). Then, we used MNPs of 216, 384, 864, 1536, 3456, and 6144 to drive FAMIL with a C384 (~25 km) resolution to compare the SYPD and running cost speedup between the different

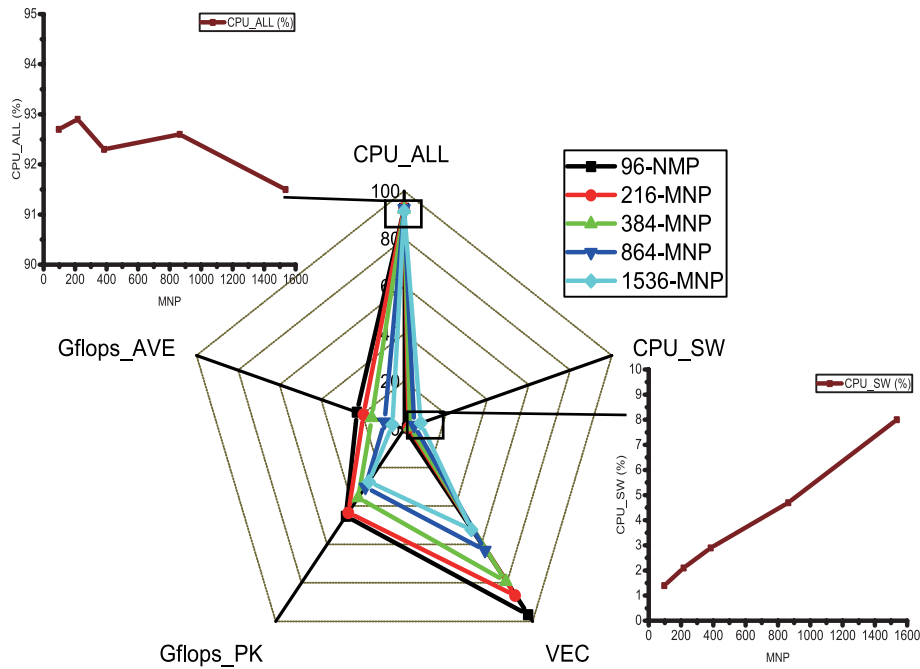


Figure 1. Computing performance for each node for the 100 km resolution FAMIL with different MNPs (36, 216, 384, 864, and 1536). Notes: There are five important indexes regarding computing performance, including CPU_ALL (units: %), CPU_SW (units: %), VEC (units: %), Gflops_AVE (units: times), and Gflops_PK (units: times).

resolutions. Finally, we added one experiment that used 384 MNPs to drive FAMIL with a C192 (~50 km) resolution to test the model's strong scalability and weak scalability among different resolutions on Tianhe-2. The vertical layers of each experiment in this paper were set to 32, and the atmospheric layers extend from the surface to 1 hPa.

The AMIP experiment was designed as follows:

- (1) Prescribe monthly forcing datasets including SST, sea ice, ozone, and aerosols.
- (2) Start from 1 January 1979 and execute seven-day integration, but only choose five-day results in order to avoid occupancy time of I/O.
- (3) Time steps for physical process are set to 600 s.

Simulation year per day (SYPD) was considered to compare the speed of FAMIL with different core numbers and resolutions. To compare the costs of running FAMIL, the cost speedup is defined. Amplification of and varying costs for running the model will be shown directly in the cost speedup equation. Strong scalability (units: %) and weak scalability (units: %) are defined as simple equations that contain the variations of wall clock time and MNPs. The quantitation standards of cost speedup, strong scalability, and weak scalability are defined as follows:

$$\text{Costs Speedup} = \frac{\text{Time(target)}}{\text{Time(standard)}} \times \frac{\text{MNP(target)}}{\text{MNP(standard)}}, \quad (1)$$

$$\text{Strong Scalability} = \frac{\text{Time(standard)}}{\text{Time(target)}} \times \frac{\text{MNP(standard)}}{\text{MNP(target)}} \times 100\%, \quad (2)$$

$$\text{Weak Scalability (\%)} = \frac{\text{Time(standard)}}{\text{Time(target)}} \times 100\%. \quad (3)$$

The notation 'standard' in the equation represents running FAMIL using 96 MNPs at 100 km resolution and using 216 MNPs at 25 km resolution. The notation 'target' in the equations represents changing MNPs from 96 to 1536 at 100 km resolution and changing MNPs from 216 to 6144.

By setting monitor toolkits, which include Paramon and Paratune developed by Paratera Ltd. (<http://www.paratera.com/>), for each node in Tianhe-2, we could capture all the information regarding the computing resources that FAMIL occupies, including CPU usage (units: %), percentage of CPU kernel mode that occupies CPU time and of message passing waiting time (CPU_SW (units: %)), code vectorization (VEC) (units: %), average Giga floating-point operations per second (Gflops_AVE) (units: times), and peak Giga floating-point operations per second percentage (Gflops_PK) (units: times), for all running processes (Starting, integration, and results output) (Purohit et al. 1999).

Table 1. Comparison of Gflops_AVE, Gflops_PK and Gflops (%) for FAMIL and top five most frequently used applications in Tianhe-2.

Application name	Gflops_AVE	Gflops_PK	(Gflops_AVE/ Gflops_PK) × 100%
FAMIL	19.6	43.5	45.1
VASP	22.7	263.2	8.6
Gaussian	16.5	189.1	8.7
LAMMPS	5.2	25.7	20.2
NAMD	11.2	33.1	33.8
GROMACS	13.4	113.7	11.8

3. Results

3.1. Computing performance of each node

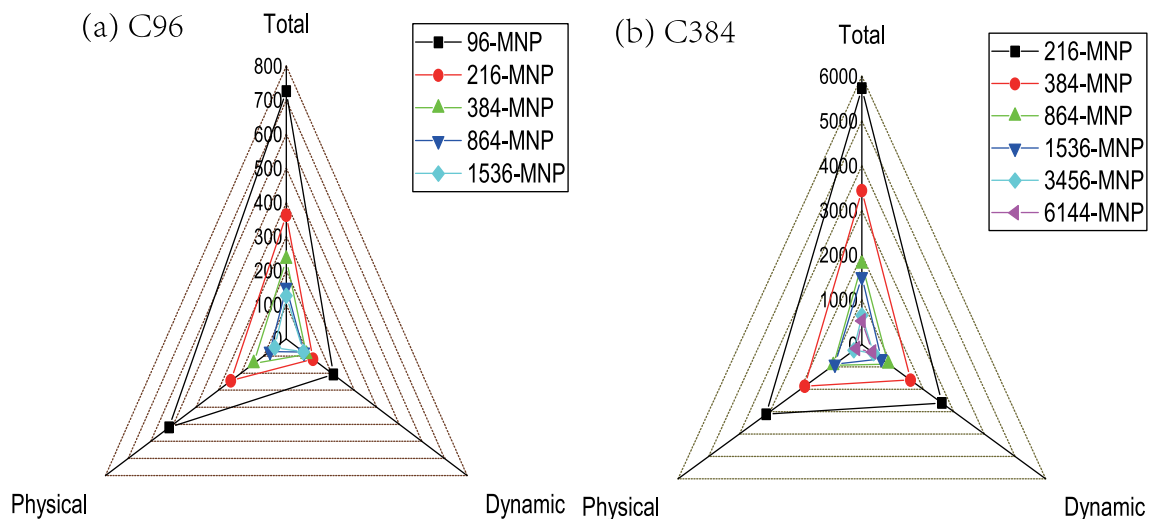
As illustrated in Figure 1, at 100 km resolution, as MNPs increase from 96 to 1536, the CPU usage maintained a very high level (higher than 90%), along with a slightly rising trend of CPU_SW. However, as MNPs increased from 96 to 1536, Gflops_AVE and Gflops_PK both showed a downward trend. When MNPs are increasing, more time is wasted allocating nodes and transforming messages. Optimistically, compared to the top five most frequently used applications, including the Vienna Ab initio Simulation Package (VASP), the computer program for computational chemistry initially released with the name Gaussian, the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS), the Nanoscale Molecular Dynamics Program (NAMD), and the GRoningen Machine for Chemical Simulations (GROMACS) in Tianhe-2, FAMIL shows excellent performance on the Tianhe-2 platform (Table 1). VEC shows a declining trend with increasing MNPs. When 1536 processors are set, the value of VEC decreased to 53%. That is, even though the Intel compiler that FAMIL used can vector code automatically, incremental increase of

MNPs cause FAMIL to be irregular, which affects the computational efficiency of our models. However, compared to other atmospheric models (Dennis et al. 2012), FAMIL shows good computing performance, especially in the Gflops_AVE and Gflops_PK aspects.

3.2. SYPD and cost speedup of each node

As illustrated in Figure 2(a), in the case of a 100 km resolution, the physical process occupies more wall clock time than the dynamical core. In contrast, in the case of a 25 km resolution (Figure 2(b)), the proportion of the dynamical core rapidly increases. This occurs because the time step of the dynamic core for C384 is much smaller than the time step of the dynamic core for C96, while the time step of model's physics is the same for both resolutions. Therefore, for the same number of CPU cores, the dynamic core of C384 needs much more computational time, and it utilizes a much higher proportion of the entire execution time. Furthermore, regardless of whether the resolution is 100 km or 25 km, as the MNPs that drive the model increase, the wall clock times of the dynamical core and the physical process both decrease.

SYPD can be compared for a 100 km resolution (Figure 3(a)) and a 25 km resolution (Figure 4(a)). In the 100 km resolution condition, when MNPs increase from 96 to 384, the total SYPD changes from 2.17 to 5.98. However, once the MNPs reach 384, the upward change in SYPD is not obvious. From 864 to 1536 MNPs, the SYPD changes from 7.97 to 8.08. In the 25 km resolution condition, when MNPs increase from 216 to 1536, SYPD rises sharply from 0.29 to 1.15. When the MNPs are more than 1536, SYPD begins to flatten. Additionally, as MNPs rise from 216 to 1536, total

**Figure 2.** Wall clock time (units: sec) of FAMIL for total, dynamical core, and physical processes at (a) the C96 resolution and (b) C384 resolution.

Notes: At (a) the C96 resolution, target MNPs are 96, 216, 384, 864, and 1536. At (b) the C384 resolution, target MNPs are 216, 384, 864, 1536, 3456, and 6144.

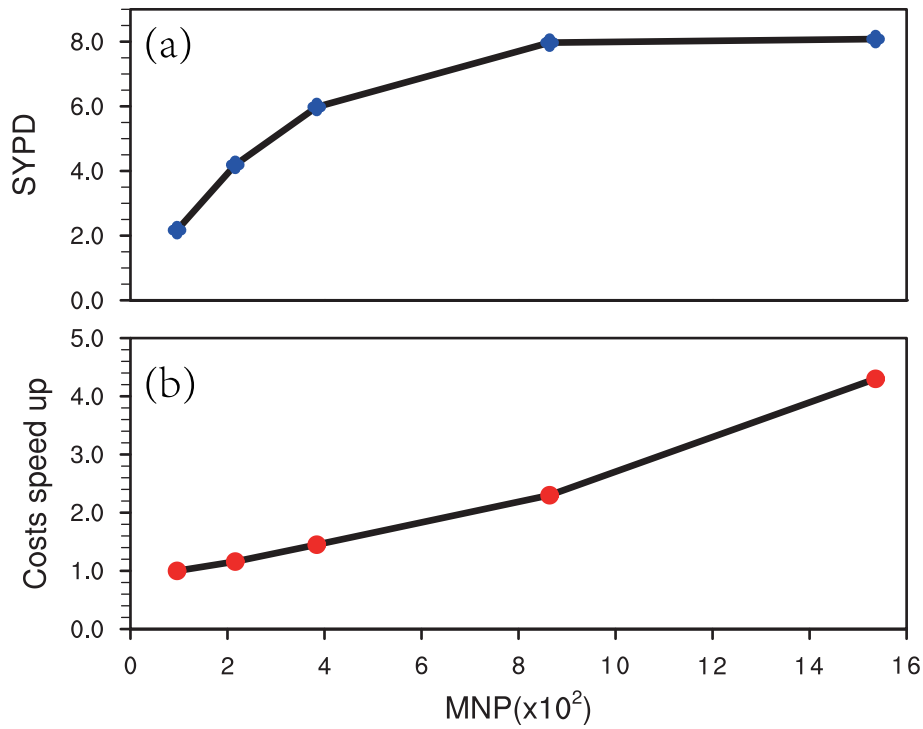


Figure 3. (a) SYPD and (b) cost speed up for FAMIL in C96 resolution with different MNP (96, 216, 384, 864, and 1536).

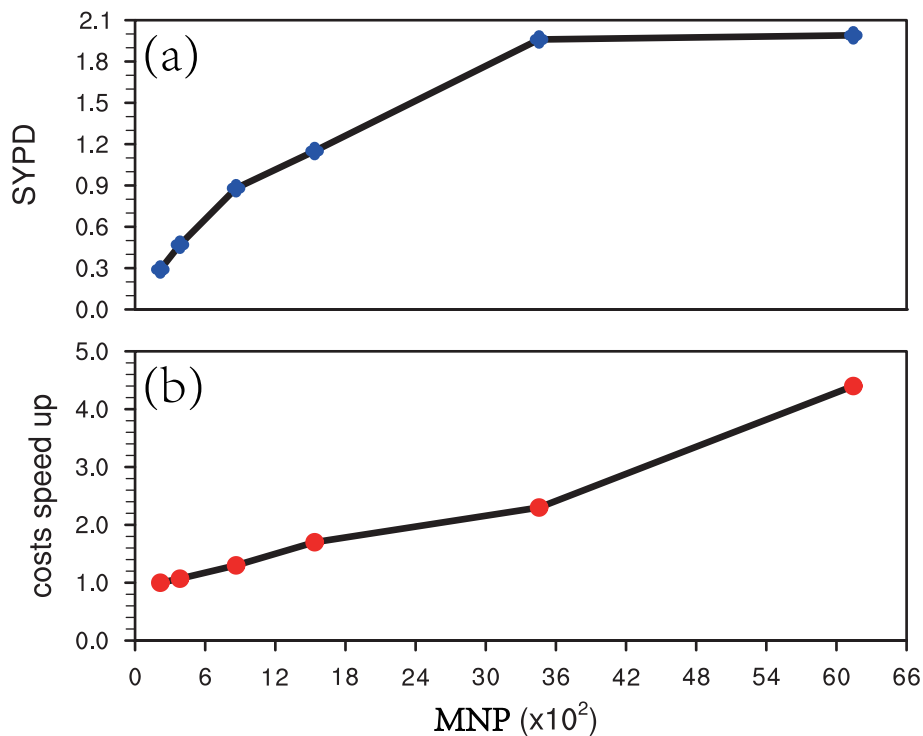


Figure 4. (a) SYPD and (b) cost speed up for FAMIL at the C384 resolution with different MNPs (216, 384, 864, 1536, 3456, and 6144).

SYPD increases from 0.29 to 1.15. The rate of change is close to linear. However, once the MNPs are more than 1536, the rate slows. From 3456 to 6144 MNP, the SYPD increases from 1.96 to just 1.99, and the upward trend of SYPD begins to flatten.

The total cost speedup for a 100 km resolution (Figure 3(b)) and a 25 km resolution (Figure 4(b)) can be compared. In the 100 km condition, a rising but non-linear trend in cost speedup is obvious. When the MNPs are less than 384, the cost speedup is no more than double that seen

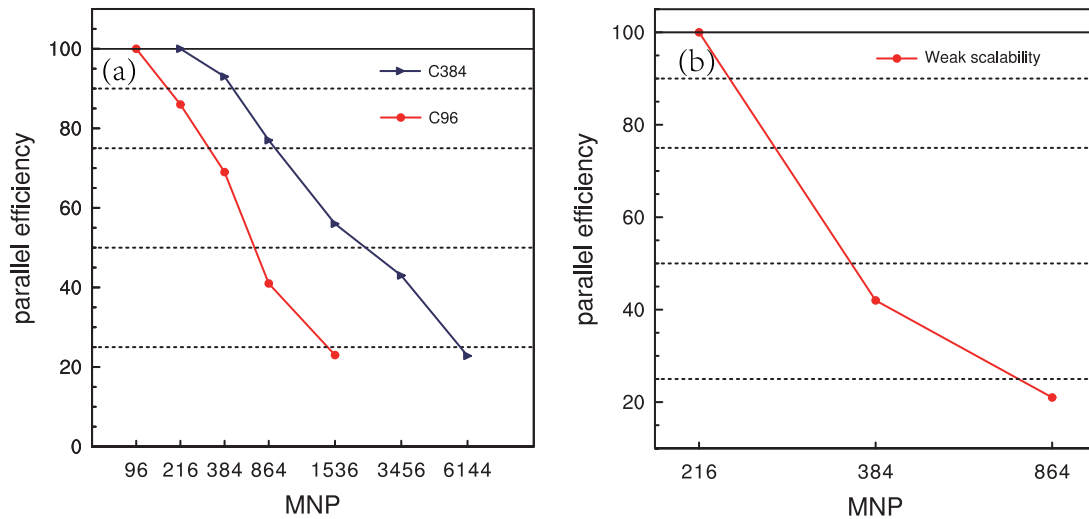


Figure 5. (a) strong scalability of FAMIL at the C96 and C384 resolutions with different MNPs (C96 MNPs are 96, 216, 384, 864, and 1536, and C384 MNPs are 216, 384, 864, 1536, 3456, and 6144.). (b) Weak scalability of FAMIL at the C96, C192, and C384 resolutions with different MNPs (MNPs of C96, C192, and C384 are 216, 384, and 864.).

Table 2. Comparison of Gflops_AVE and Gflops_PK for FAMIL and SAMIL.

Model name	Gflops_AVE	Gflops_PK	(Gflops_AVE/Gflops_PK) × 100%
FAMIL	19.6	43.5	45.1
SAMIL	3.1	5.4	57.4

at 96 MNP. This means that increasing the speed by using greater MNPs, once confined to below 384, appears to be cost effective for FAMIL. However, when MNPs are greater than 384, cost speedup is more than double the cost speedup created by using 96 MNPs. Notably, when using 1536 MNPs, the cost speedup is 4.6 times that at 96 MNPs. Similar results are found in the 25 km resolution condition. When MNPs are less than 1536, cost speedup is less than double the original value, but as the MNPs continue to increase, cost speedup increases dramatically.

3.3. Strong scalability and weak scalability

To analyze the parallel efficiency of FAMIL at different resolutions, both strong scalability and weak scalability were tested (Dennis et al. 2012). From the strong scalability results (Figure 5(a)), the parallel efficiency of FAMIL, at both 100 km and 25 km resolutions, decreases with increasing MNPs. When MNPs reach 864 for the 100 km resolution and 3456 for the 25 km resolution, the parallel efficiency of FAMIL drops to less than 50%. On the other hand, as seen from the results for different resolutions (100 km with 216 MNPs, 50 km with 384 MNPs, and 25 km with 864 MNPs), we can achieve a comprehensive line of weak scalability for FAMIL (Figure 5(b)). A descending trend is obvious for weak scalability. However, other atmospheric models (Dennis et

al. 2012) also show negative trends for strong and weak scalability, which means that there is still huge room for models to improve their parallel efficiency when MNPs are increased.

4. Conclusions and discussion

In the aspect of the computing performance of each node, with the increase of the MPI number of processors (MNPs), CPU usage shows stable performance. Although the Gflops_AVE and Gflops_PK of FAMIL show downward trends with increased MNPs at both 100 km and 25 km resolutions, FAMIL shows excellent computing performance compared to SAMIL, the last generation of AGCM of the IAP/LASG (Table 2). However, the VEC results shows insufficient performance of the Intel Compiler in handling code vectorization automatically, and rather, we should optimize code artificially; for example, we should optimize the structure of code or add a vectoring statement in the appropriate places. According to the increasing trend of CPU_SW with increasing MNPs, we can conclude that the communication overhead is a significant bottleneck for the acceleration of FAMIL. All tests were based on MPI technology; however, if OpenMP technology is adopted, there should be some room for improvement in terms of the computing performance of each node.

In the aspect of SYPD and cost speedup for FAMIL, when MNPs are increased, the SYPD are increased, and with increasing MNPs, usage cost speedup for FAMIL also increases. Optimal MNP choices should take both SYPD and cost speedup into consideration. The recommended MNPs for a 100 km resolution are 384, and the recommended MNPs for a 25 km resolution are 1536 on Tianhe-2.

Both strong scalability and weak scalability of FAMIL decrease with increase of MNPs at all resolutions. Previous research (Zhou et al. 2012) attributes this phenomenon to a platform bottleneck. However, because of a lack of a monitoring method, this is just a hypothesis. In our study, a benefit from arranging monitor toolkits for each node on Tianhe-2 is that we can conclude that the reason for the decline of parallel efficiency can be ascribed not only to the Tianhe-2 platform but also to the model design itself. We need to optimize the parallel structure for FAMIL (Mueller and Scheichl 2013; Wang 2014), and we should make full use of the MIC Cards on Tianhe-2. Considering the acceleration card in the MIC architecture Tianhe-2, there is still great potential to improve the computational performance of FAMIL (Liu et al. 2014; Xue et al. 2015) which might benefit from using the MIC architecture on Tianhe-2.

Acknowledgements

Thanks to Dr. Shian-Jiann LIN and Dr. Bao-Qiang XIANG's encouragement and support during the development of FAMIL.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the National Natural Science Foundation of China [grant number 41675100], [grant number 91337110]; the Third Tibetan Plateau Scientific Experiment: Observations for Boundary Layer and Troposphere [GYHY201406001]; the Key Research Program of Frontier Sciences, Chinese Academy of Science (CAS) (QYZDY-SSW-DQC018); the Special Program for Applied Research on Super Computation of the NSFC-Guangdong Joint Fund (the 2nd phase).

References

- Bao, Q., P. F. Lin, T. J. Zhou, Y. M. Liu, Y. Q. Yu, G. X. Wu, B. He, et al. 2013. "The Flexible Global Ocean-Atmosphere-Land System Model, Spectral Version 2: FGOALS-s2." *Advances in Atmospheric Sciences* 30: 561–576.
- Bao, Q., G. X. Wu, Y. M. Liu, J. Yang, Z. Z. Wang, and T. J. Zhou. 2010. "An Introduction to the Coupled Model FGOALS1.1-s and Its Performance in East Asia." *Advances in Atmospheric Sciences* 27: 1131–1142.
- Craig, A. P., M. Vertenstein, and R. Jacob. 2012. "A New Flexible Coupler for Earth System Modeling Developed for CCSM4 and CESM1." *International Journal of High Performance Computing Applications* 26: 31–42. doi: 10.1177/1094342011428141.
- Dennis, J. M., M. Vertenstein, P. H. Worley, A. A. Mirin, A. P. Craig, R. Jacob, and S. Mickelson. 2012. "Computational Performance of Ultra-High-Resolution Capability in the Community Earth System Model." *International Journal of High Performance Computing Applications* 26: 5–16.
- Eyring, V., S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor. 2016. "Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) Experimental Design and Organization." *Geoscientific Model Development* 9: 1937–1958. doi:10.5194/gmd-9-1937-2016.
- Gates, W. L. 1992. "AMIP: The Atmospheric Model Intercomparison Project." *Bulletin of the American Meteorological Society* 73: 1962–1970.
- Gates, W. L., J. S. Boyle, C. Covey, C. G. Dease, C. M. Doutriaux, R. S. Drach, M. Fiorino, et al. 1999. "An Overview of the Results of the Atmospheric Model Intercomparison Project (AMIP I)." *Bulletin of the American Meteorological Society* 80: 29–55.
- Goldberg, D. 1989. *Genetic Algorithms in Search, Optimization and Learning*, 27–56. Boston, MA: Addison-Wesley.
- Haarsma, R. J., M. Roberts, L. V. Pier, A. S. Catherine, B. Alessio, Q. Bao, P. Chang, et al. 2016. "High Resolution Model Intercomparison Project (HighResMIP V1.0) for CMIP6." *Geoscientific Model Development* 9: 4185. 2016-66. Accepted.
- Kendall, R. A., E. Aprà, D. E. Bernholdt, E. J. Bylaska, M. Dupuis, G. I. Fann, and R. J. Harrison. 2000. "High Performance Computational Chemistry: An Overview of NWChem a Distributed Parallel Application." *Computer Physics Communications* 128: 260–283.
- Lin, S. J. 1997. "A Finite-volume Integration Method for Computing Pressure Gradient Forces in General Vertical Coordinates." *Quarterly Journal of the Royal Meteorological Society* 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.'" *Quarterly Journal of the Royal Meteorological Society* 124: 2531–2533.
- Lin, S. J. 2004. "A 'Vertically Lagrangian' Finite-volume Dynamical Core for Global Models." *Monthly Weather Review* 132: 2293–2307.
- Liu, Y. Q., X. Y. Zhang, C. Yang, F. F. Liu, and Y. T. Lu. 2014. "Accelerating HPCG on Tianhe-2: A Hybrid CPU-MIC Algorithm." *The IEEE International Conference on Parallel and Distributed System* 542–551.
- Mueller, E. H., and R. Scheichl. 2013. "Massively Parallel Solvers for Elliptic PDEs in Numerical Weather and Climate Prediction." *Quarterly Journal of the Royal Meteorological Society* 140: 55–78.
- Neal, R. B., and B. J. Hoskins. 2001a. "A Standard Test for AGCMs including Their Physical Parameterizations. I: The Proposal." *Atmospheric Science Letters* 1: 101–107.
- Neal, R. B., and B. J. Hoskins. 2001b. "A Standard Test for AGCMs including Their Physical Parameterizations. II: Results for the Met Office Model." *Atmospheric Science Letters* 1: 108–114.
- Purohit, S. C., A. Kaginalkar, I. Jindani, and J. V. Ratnam. 1999. "Development of Parallel Climate/Forecast Models on 100 GFLOPS PARAM Computing System." Proceedings of the Eight ECMWF Workshop on the Use of Parallel Processors in Meteorology, WorldScientific, Reading, UK.
- Putman, W. M., and S. J. Lin. 2007. "Finite-volume Transport on Various Cubed-sphere Grids." *Journal of Computational Physics* 227: 55–78.
- Versteeg H. K., and W. Malalasekera. 1995. "An Introduction to Computational Fluid Dynamics – The Finite Volume Method." *Epfl* 20: 400.
- Wang, P. F. 2014. "DIY the Integrated Climate Model and Its Computational Performance." Eprint Arxiv: 1405.7473.
- Wang, X.-C., Y. Liu, S.-J. Lin, Q. Bao, and G.-X. Wu. 2013. "The Application of Flux-form Semi-Lagrangian Transport Scheme in

- aSpectralAtmosphericModel." *Adv. Atm. Sci.* 30(1), 89–100. doi: [10.1007/s00376-012-2039-2](https://doi.org/10.1007/s00376-012-2039-2).
- Xue, W., C. Yang, H. H. Fu, X. L. Wang, Y. T. Xu, J. F. Liao, L. Gan, et al. 2015. "Ultra-scalable CPU-MIC Acceleration of Mesoscale Atmospheric Modeling on Tianhe-2." *IEEE Transactions on Computers* 64: 2382–2393.
- Zhou, L. J., Y. M. Liu, Q. Bao, H. Y. Yu, and G. X. Wu. 2012. "Computational Performance of the High-resolution Atmospheric Model FAMIL." *Atmospheric and Oceanic Science Letters* 5: 355–359.
- Zhou, L. J., Q. Bao, Y. M. Liu, G. X. Wu, W. C. Wang, X. C. Wang, B. He, et al. 2014. "Global Energy and Water Balance: Characteristics from Finite-volume Atmospheric Model of the IAP/LASG (FAMIL1)." *Journal of Advances in Modeling Earth Systems* 7: 1–20.