Performance Optimization and Evaluation of a Global Climate Application using a 440m Horizontal Mesh on the K computer

RIKEN K computer

Masaaki Terai[1], Hisashi Yashiro[2], Kiyotaka Sakamoto[3], Shin-ichi Iga[2], Hirofumi Tomita[2], Masaki Satoh[4] and Kazuo Minami[1]

1) Software Development Team, Operations and Computer Technologies Division, Advanced Institute for Computational Science, RIKEN
2) Computational Climate Science Research Team, Advanced Institute for Computational Science, RIKEN
3) Fujitsu Systems East Limited
4) Atmosphere and Ocean Research Institute, The University of Tokyo

Introduction

- High-resolution atmospheric simulation provides more precise cloud behavior and an insight into extreme weather events, such as typhoons and heavy rains associated with climate change.
- The K computer provides the opportunity to conduct high-resolution simulations.
- To perform more effective computations with high-resolution in a production run of Nonhydrostatic Icosahedral Atmospheric Model (NICAM)[1][2], we optimized single node performance, and evaluated the performance by weak scaling using almost all the compute nodes on the K computer.
- Furthermore, we analyzed the application bottleneck derived from scheme.

NICAM

- A weather/climate application for high-resolution global atmospheric simulation studies
- Nonhydrostatic and fully compressive equation system, discretization on the icosahedral grid arrangement (Figure 1)
- NICAM uses a finite volume method, requiring only point-to-point communication instead of global communication in spectral transform method.

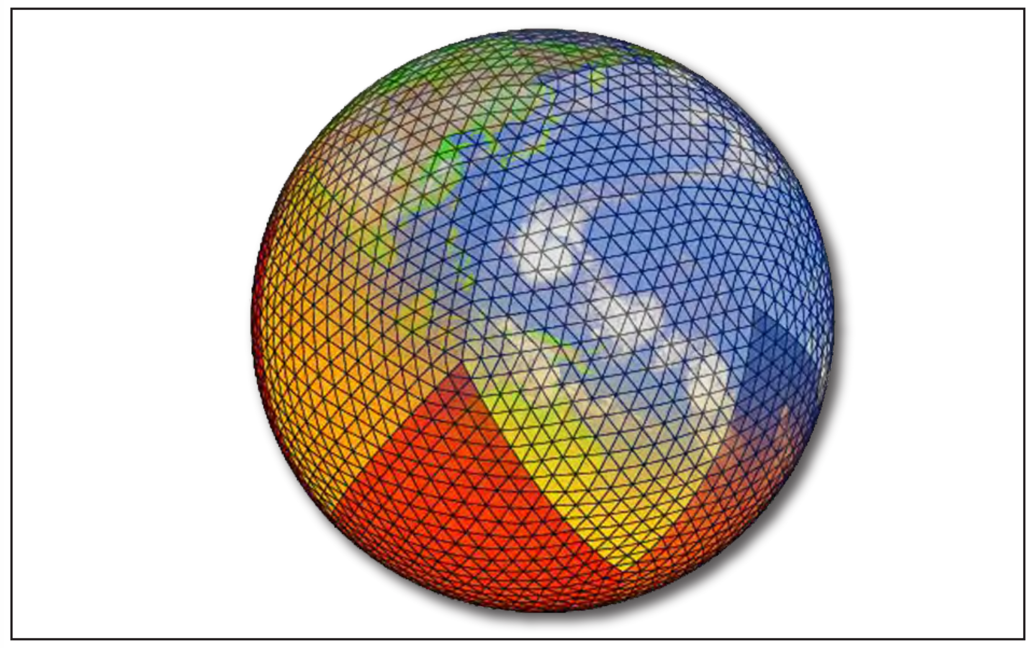

Figure 1 Domain decomposition using the icosahedral grid

K computer

- The K computer[3] is first 10 peta-flops massively parallel computer with 82,944 nodes. (Each node has a SPARC64VIIIfx octa-core chip and 16GB memory.)
- Each node is connected by Tofu interconnect, 6D mesh/torus.

Single Node Performance

Kernel Programs

- NICAM has two main components in the main loop: dynamics and physics.
- To optimize single node performance, we extracted kernel programs from the components based on profiling results (Table 1).

Table 1 Kernel programs

	Kernel name	Outline	blocks
dynamics	divergence	horizontal divergence operators	1
	diffusion	horizontal diffusion operators	1
	gradient	horizontal gradient operators	2
	divdamp3D	3D divergence damping	4
	advlim	tracer advection with limiters	6
	divdamp	horizontal divergence damping	2
physics	cloud microphysics	1-moment bulk scheme of the cloud microphysics	3
	radiation	broadband longwave/shortwave radiative transfer scheme	3

Tuning and Result

- Some kernel programs are divided into blocks corresponding to loop/loops.
- We took optimization processes of trial and error for each block (Figure 2).

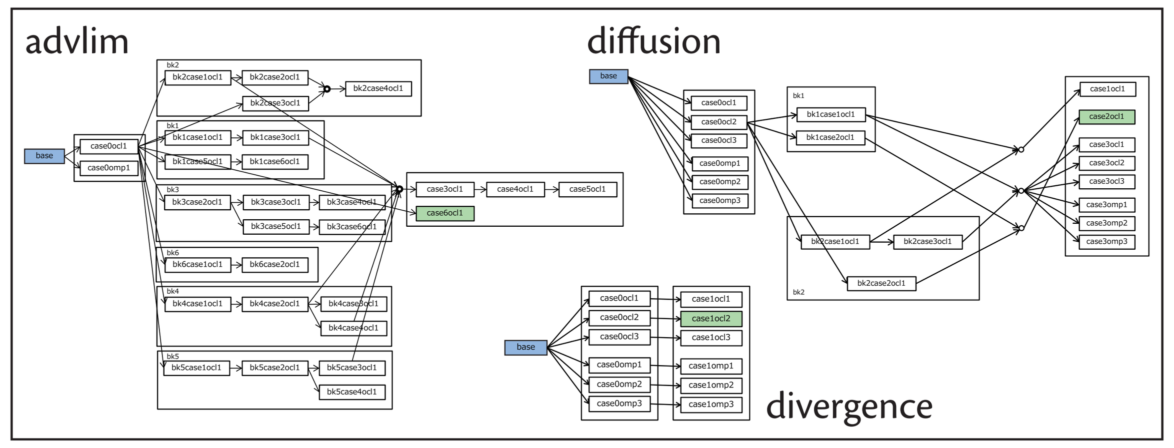

Figure 2 Tuning History of the divergence, diffusion, and advlim

- Optimizing single processor performance improved the peak performance ratio (Figure 3).

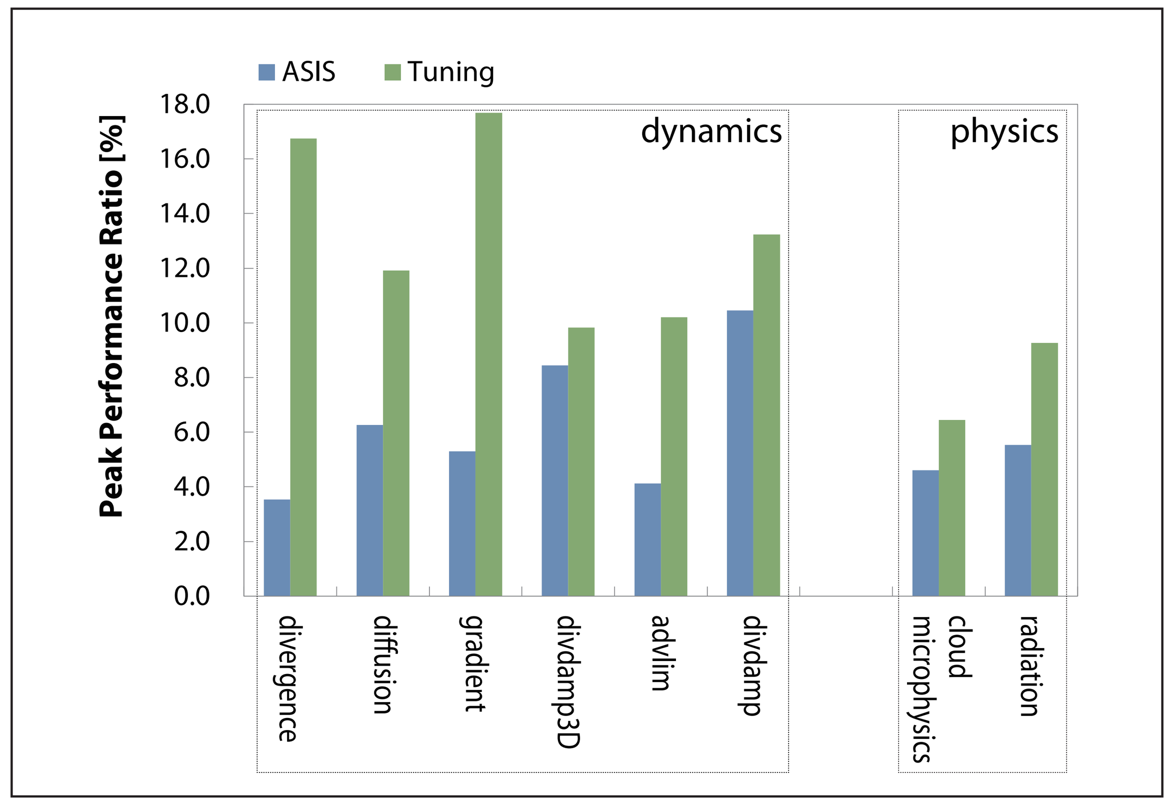

Figure 3 Result of single node performance optimization

Weak Scaling and Performance

Experimental Conditions

- We evaluated the performance by weak scaling the compute nodes from 5 to 81,920 nodes.
- In the K computer, NICAM uses MPI (inter-node) and Thread (intra-node) hybrid parallelization.
- We focused on the main loop of NCAM.
- Setup and I/O time were ignored.
- To maintain same problem size per process, the experiment was carried out under the conditions listed in Table 2.

Table 2 Experimental conditions for weak scaling

glevel	# of horiz. grids	# of procs.	Δx [km]
7	163840	5	56
8	655360	20	28
9	2621440	80	14
10	10485760	320	7.0
11	41943040	1280	3.5
12	167772160	5120	1.7
13	671088640	20480	0.87
14	2684354560	81920	0.44

Result

- The dynamical step is good scalability (Figure 4).
- In the physical step, the peak performance ratio decreased, and elapsed time increased with the number of processes despite there being no communication.

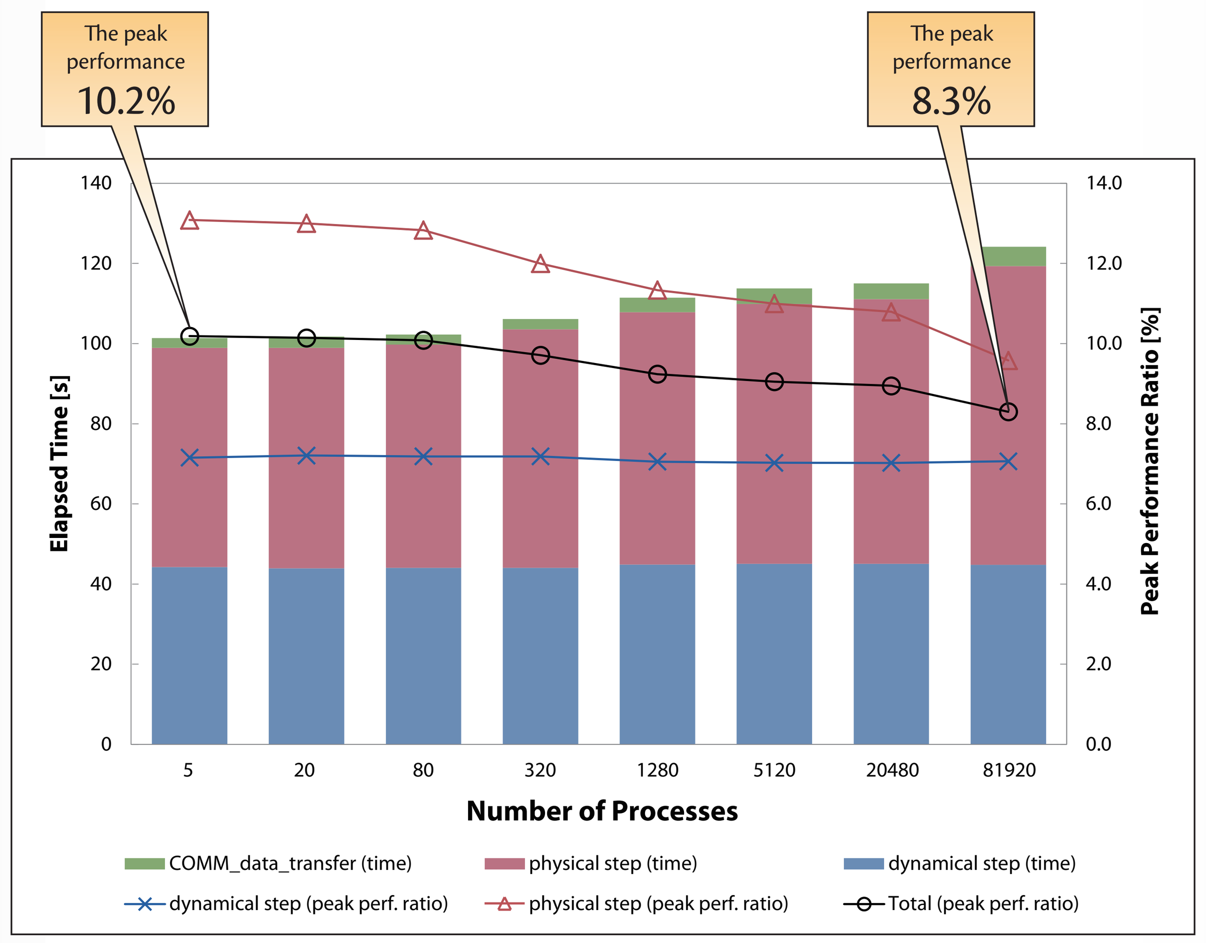

Figure 4 Result of weak scaling and performance

Analysis and Discussion

- We analyzed the performance degradation in the physical step.
- In the physical step, average number of floating-point operations per node is constant with the number of processes (Figure 5).

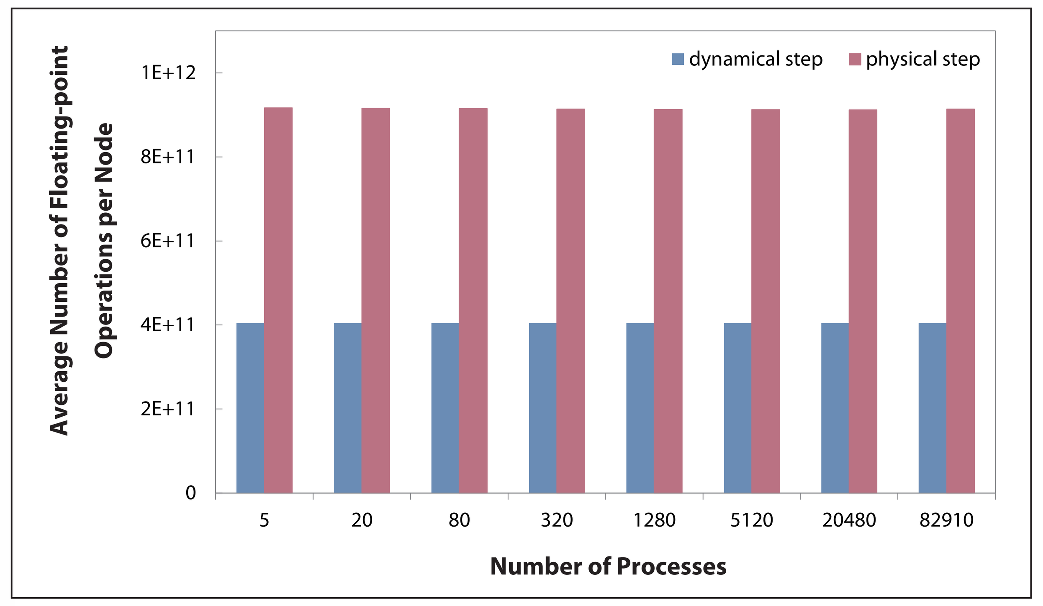

Figure 5 Average number of floating-point operations per node

- Furthermore, we analyzed the number of floating-point operations for subroutines of the physical step, and found inter-node load imbalances in the cloud microphysics.
- Regions with a large number of floating-point operations correspond to the distribution of the ice water (Figure 6).
- The intra-node load imbalances derived from the scheme becomes more remarkable as the total number of processes increases.
- For future work in high-resolution calculation, our results suggest that the physical step has an intrinsic bottleneck that cannot be ignored.

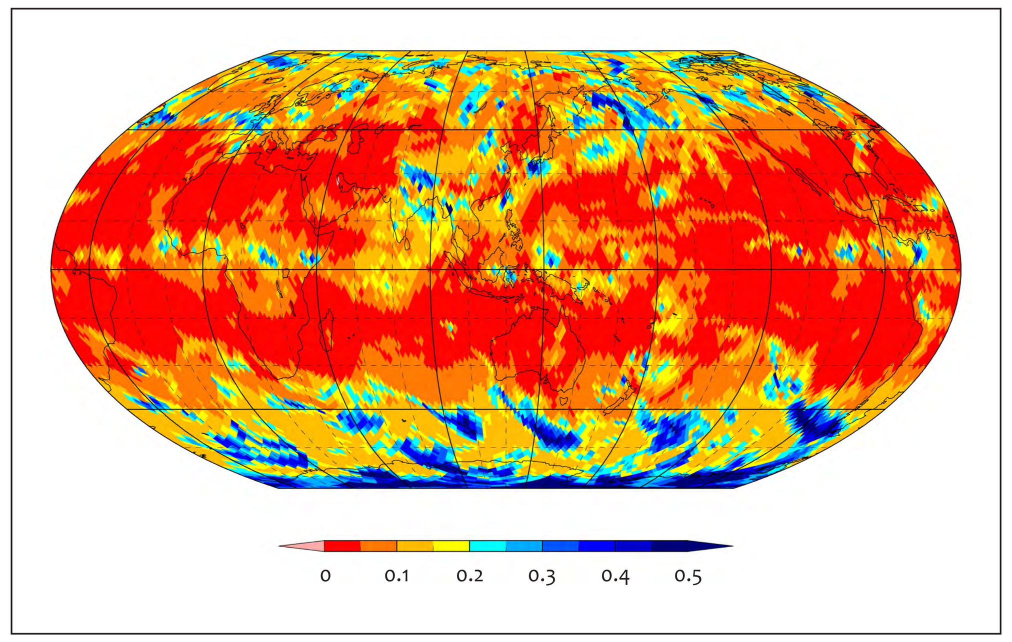

Figure 6 Contour of normalized numbers of floating point operations on a global map in the cloud microphysics (glevel-14)

Acknowledgement

Part of the results is obtained by using the K computer at the RIKEN Advanced Institute for Computationla Science.

References

[1] M. Satoh et al, Nonhydrostatic icosahedral atmospheric model (NICAM) for global cloud resolving simulations, J. of Comp. Phys., 227:3486-3514, 2008.
[2] Y. Miyamoto et al, Deep moist atmospheric convection in a subkilometer global simulation, Geophys. Res. Lett., 40 (18):4922-4926, 2013.
[3] H. Miyazaki et al, K computer: 8.162 petaflops massively parallel scalar supercomputer built with over 548k cores, In ISSCC 2012, pp. 192-194, 2012.