An Evaluation of Discontinuous Galerkin Method based Global Nonhydrostatic Atmospheric Dynamical Core on A64FX Platform

Xuanzhengbo Ren[1], Yuta Kawai[2], Hirofumi Tomita[2], Takahiro Katagiri[1], Seiya Nishizawa[2]
Tetsuya Hoshino[1], Masatoshi Kawai[1], Toru Nagai[1]
1. Nagoya University
2. RIKEN Center for Computational Science
Japan
ren@hpc.itc.nagoya-u.ac.jp

1 INTRODUCTION

For future high-resolution atmospheric simulations, a dynamical core using discontinuous Galerkin Method (DGM), called SCALE-DG [3], is being developed as an option of high-order fluid schemes in SCALE library [4]. Since the spatial discretization is done locally, we expect the computational performance is highly desirable in modern computer architectures. In this study, we evaluated the scalability and single process performance of SCALE-DG.

2 METHODS AND RESULTS

SCALE-DG is parallelized by both MPI and OpenMP. For MPI parallelization, the computation area is partitioned horizontally into multiple local meshes. In each mesh, the computations are running in multi-threads.

A gravity wave test served as a benchmark in this evaluation. The problem size is determined by the elements in computation area (horizontally and vertically), and the nodes (within the element). The number of nodes depends on the polynomial degree ***P***. As the temporal discretization, there are two options: one is horizontally explicit vertically explicit (HEVE) scheme, the other is horizontally explicit vertically implicit (HEVI) scheme. The performance are investigated on the Fujistu A64FX based supercomputers: Fugaku [2] and Flow [1]. The processor is made up of 4 Core Memory Groups (CMG) where each CMG includes 12 computing cores. In this context, the MPI processes are set to bind with CMGs and the thread number is fixed at 12.

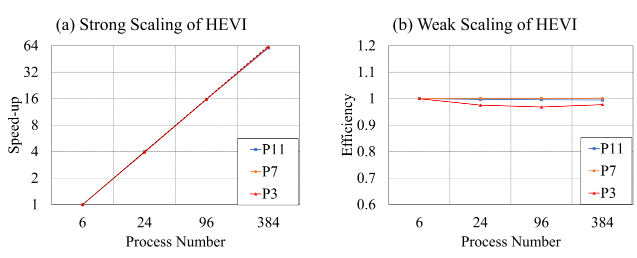

Figure 1: Scalability of SCALE-DG with HEVI scheme

Fig. 1 (a) shows the strong scaling of HEVI scheme in *P*=11, 7 and 3, respectively. The problem size is fixed at 3,538,944 data points in all test cases. Regardless of how *P* changes, it shows optimal speed-ups up to 384 processes. In weak scaling tests the problem size of each process is fixed at 36,864 data points. The result is as shown in Fig. 1 (b). Comparing to *P*=7 and 11, the efficiency slightly decreased with the process increasing when *P*=3, but still performed very well.

Fig. 2 shows a typical issue in single process performance of one of the routines with high computational costs in HEVI scheme for *P*=7, called *hevi_cal_vin_lin* in which band matrix equations are solved. It can be observed that threads 4 to 11 take much longer barrier synchronization time (the purple blocks) than threads 0 to 3. It means there is a server load imbalance problem in this routine. This is because the thread parallelization is applied to a relatively inner loop which treats a $(P+1)^2$ set of one-dimensional vertical columns. Thus, the imbalance problem is evident for small *P*.

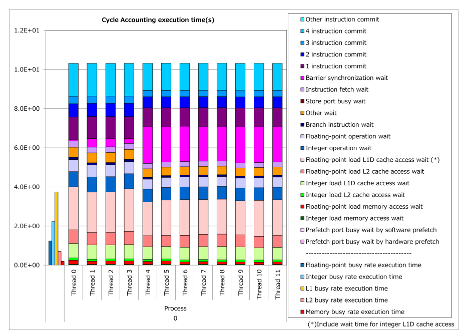

Figure 2: Single process performance of hevi_cal_vi_lin

3 CONCLUSION

SCALE-DG performed excellently in strong and weak scaling, while we observe the load imbalance in the part of solving band matrix equations in HEVI scheme, which needs to apply the thread parallelization to outer loop in the corresponding routine. More details about other performance issues will also be discussed in this poster.

ACKNOWLEDGMENTS

This research was supported by 2023 RIKEN R-CCS HPC Computational Science Internship Program and the supercomputer Fugaku.

REFERENCES

[1] FLOW 2020. Supercomputer, Nagoya University. Retrieved October 26, 2023 from https://icts.nagoya-u.ac.jp/en/sc/
[2] FUGAKU 2020. Fugaku, RIKEN R-CCS. Retrieved October 26, 2023 from https://www.r-ccs.riken.jp/en/fugaku/
[3] Yuta Kawai and Hirofumi Tomita. 2019. FE-Project. Team SCALE, RIKEN R-CCS, Kobe, Japan. https://ywkawai.github.io/FE-Project_web/
[4] Team SCALE. 2012. SCALE. Team SCALE, RIKEN R-CCS, Kobe, Japan. https://scale.riken.jp/