ParCFD'2020
32^{nd} International Conference on Parallel Computational Fluid Dynamics
May-11-13 2020, Nice FRANCE

PERFORMANCE EVALUATION OF PARALLEL DNS CODES ON THE SUPERCOMPUTER SX-AURORA TSUBASA

Yujiro Takenaka[1], Mitsuo Yokokawa[1], Takashi Ishihara[2], Kazuhiko Komatsu[3], and Hiroaki Kobayashi[4]

[1] Graduate School of System Informatics, Kobe University
Kobe, Japan
e-mail: ytakenaka@stu.kobe-u.ac.jp

[2] Graduate School of Environmental and Life Science, Okayama University
Okayama, Japan

[3] Cybersicence Center, Tohoku University
Sendai, Japan

[4] Graduate School of Information Sciences, Tohoku University
Sendai, Japan

Key words: Turbulence, Direct numerical simulation, Parallel computing, Vector supercomputer

Abstract. Direct numerical simulations (DNSs) of incompressible turbulence have been performed since the 1960s, but simulations that reproduce turbulent flows in the real-world have not been realized. We implemented a parallel Fourier-spectral DNS code by using a two-dimensional domain decomposition (pencil decomposition) for a cutting-edge vector supercomputer in order to carry out larger DNSs than ever before. Vector execution performance of the code was measured on the SX-Aurora TSUBASA vector supercomputer, and was compared to that of a one-dimensinal domain decomposition (slab decomposition) DNS code. The computational time of the code by the pencil decomposition with 512^3 grid points and 64 MPI processes is about 1.3 times faster than that of the code by the slab decomposition.

1 INTRODUCTION

Turbulence exists everywhere around us, and it is very important to elucidate its property. Direct numerical simulations (DNSs) of incompressible turbulence have been performed since the 1960s[1, 2, 3], but simulations that reproduce turbulent flows in the real-world have not been realized. The DNS with 12288^3 grid points on the K computer was performed, in which the Taylor scale Reynolds number is approximately 2300[2]. DNS using GPU with 18432^3 grid points on the Summit at Oak Ridge National lab. was reported in the paper[3].

We developed a parallel Fourier-spectral DNS code by using a two-dimensional domain decomposition (pencil decomposition) and our Fast Fourier Transform(FFT) library. In the code, we used the Fourier spectral method, of which most of the computational time is consumed in calculating the three-dimensional FFTs. Data transfer among processors is necessary to compute FFTs in parallel. The method of data transfer among processors is different between the pencil decomposition and the one-dimensional domain decomposition(slab decomposition).

The vector computers have vector processors that can perform the same calculation for vector data at high speed. Thus, it is suitable for scientific computation such as fluid simulations.

In this paper, the performance of the codes by the slab and the pencil decompositions was compared on the vector supercomputer SX-Aurora TSUBASA.

2 PARALLELIZATION

In the DNS by the Fourier spectral method, the 3D-FFT accounts for more than 90% of the computational time. In a parallel implementation, the 3D-FFT has a global data dependence because of its requirement for global summation over each processor. Thus, data transfer is needed during the computation.

We used our original 2, 4-radix FFT routine without using any open FFT libraries in order to compare the code using FFT library *FFTE-C*[4]. The 4-th order Runge–Kutta–Gill method is used for advancing time.

3 PERFORMANCE EVALUATION

3.1 SX-Aurora TSUBASA

We measured performance of the DNS codes on the vector computer SX-Aurora TSUBASA[5] at Tohuku Univ. The SX-Aurora TSUBASA is a parallel computer system of the distributed memory type, and consists of a vector host(VH) and one or more vector engines(VEs). The VH is a Linux server that provides standard operating system(OS) functions. The VE is built as a Peripheral Component Interconnect Express(PCIe) card, on which a vector processor is mounted. VEOS is an OS for VEs that runs on the VH, and controls VEs. Each VE has a 16MB last-level cache(LLC) and its performance is 268.8Gflop/s for double-precision.

3.2 Performance

We measured computational time and vector execution performance on the SX-Aurora TSUBASA for 50 time steps by changing the number of MPI processes with 512^3 grid points. Figure 1 shows the computational time for two DNS codes with slab and pencil decompositions. The colors represent the time for some parts of the Fourier transforms and inverse Fourier transforms in each direction. The shaded area indicates the time for transposition of the data including both the global communication and the transposition in one core.

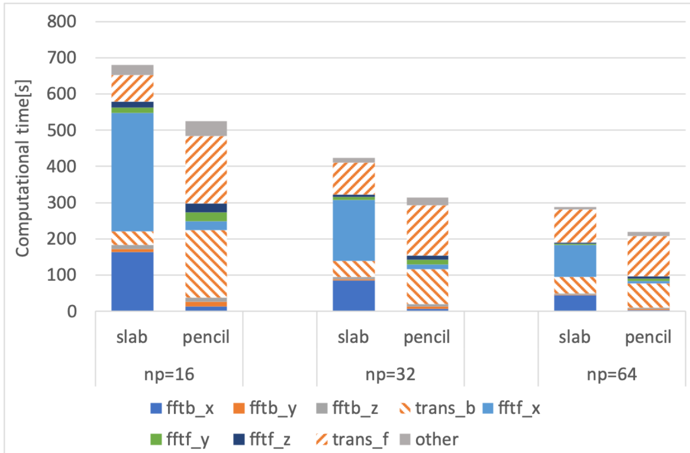

Figure 1: Computational time measured on SX-Aurora TSUBASA with 512^3 grid points by changing the number of MPI processes.

The computational time of the code by the pencil decomposition is about 1.3 times faster than that of the code by the slab decomposition. Since we use the phase-shift method in calculating nonlinear terms, the number of the Fourier transforms is twice as much as that of the inverse Fourier transforms. For fftb_x and fftf_x in the figure, which denotes the computational time of the x-direction Fourier transforms, it is found that the computational time is longer than those of the other directions. Because the second element of the three-dimensional arrays is accessed continuously in the most inner loop for the x-direction FFTs and its access sequences to the memory become non-contiguous. In the pencil decomposition, the time of transposition is longer than that of the slab decomposition since the computational area is divided into small portions, and the vector length becomes shorter.

Table 1 shows average vector length, vector operation ratio and LLC hit ratio for two DNS codes by changing the number of MPI processes. It is found that the average vector length of the slab decomposition is larger than that of the pencil decomposition, since the computational area is devided into smaller areas. Also, the LLC hit ratio is not so different between the two methods.

Table 1: Average vector length(A.V.length), vector operation ratio(V.Op.ratio), and LLC hit ratio on SX-Aurora TSUBASA for two DNS implementations with the different number of MPI processes *np*.

	np	A.V.length	V.Op.ratio	LLC Hit ratio [%]
slab	16	242.2	99.1	23
	32	224.9	98.6	26
	64	212.8	97.9	33.4
pencil	16	146.8	98.1	26.5
	32	137.5	97.7	28
	64	131.4	96.9	32.5

4 CONCLUSIONS

In this paper, we evaluated the performance of the parallel DNS codes by two decompositions, the slab and pencil decompositions on the vector supercomputer SX-Aurora TSUBASA. The computational time of the code by the pencil decomposition is 1.3 times faster than that of the code by the slab decomposition. Reducing bank conflict time by setting the size of the first element of the array appropriately should be considered for executing larger DNSs.

ACKNOWLEDGEMENTS

This work is supported partially by MEXT Next Generation High-Performance Computing Infrastructure and Applications RD Program, entitled RD of a Quantum Annealing-Assisted Next Generation HPC Infrastructure and Its Applications. This work is also supported, in part, by Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures (Project ID: jh190068).

REFERENCES

[1] Orszag, S. A., Numerical methods for the simulation of turbulence, Phys. Fluids, Supp1. Ⅱ, 12, 250-257 (1969).

[2] T.Ishihara, K.Morishita, M.Yokokawa, A.Uno, and Y.Kaneda, *Energy spectrum in high-resolution direct numerical simulations of turbulence*, Phys. Fluids. 1 (2016).

[3] K.Ravikumar, D.Appelahms, and P.K.Yeung, *GPU acceleration of extreme scale pseudo-spectral simulations of turbulence using asynchronism*, SC'19, November 17-22, (2019).

[4] T.Imamura, M.Aoki, and M.Yokokawa, *Batched 3D-distributed FFT kernels towards practical DNS codes*, Proc. of ParCo 2019, Prague, Czech Republic, September 10-13, (2019) (in press).

[5] K. Komatsu et al., *Performance Evaluation of a Vector Supercomputer SX-Aurora TSUBASA*, SC'18, International Conference for High Performance Computing, Networking, Storage and Analysis, Dallas, TX, USA, 2018, pp. 685-696. doi: 10.1109/SC.2018.00057