Large scale *ab initio* calculation using LDC-DFT algorithm on many-core processor architectures

KOHEI SHIMAMURA, Graduate School of System Informatics, Kobe University

The lean divide and conquer density functional theory (LDC-DFT) algorithm realizes $O(N)$ quantum mechanical calculation for an N-electron system, in which the three-dimensional space is represented as a union of spatially localized domains. In spite of the simple methodology and easy implementation, a high parallelization efficiency (98% by weak scaling) on a multi-core processor architecture (786,432 cores on Blue Gene/Q Mira (16 cores/node)) has been achieved so far. However, the parallelization efficiency in the many-core processor architectures, which will be the mainstream of the next generation, is still unknown. Here, we report the optimization work on the Oakforest-PACS system.

CCS Concepts: • **Computing methodologies → Massively parallel and high-performance simulations**

1 INTRODUCTION

The linear scaling calculation methods for the large-scale electronic state have been actively developed due to its wide range of application [1]. In particular, it has become essential for the understanding of life phenomena. The biomolecules such as proteins, nucleic acids of current life are commonly composed of more than 10,000 atoms and the atomistic interactions with the surrounding water molecules should be considered the involving reactions [2]. In addition, many enzymatic reactions involve metals such as iron. In order to accurately reproduce the physical properties of biomolecules containing metal and surrounded by water, the linear scaling method for the large-scale electronic state is indispensable. We recently developed an algorithm based on lean divide and conquer density functional theory (LDC-DFT), which is a method that satisfies the above needs and can also investigate dynamic physical properties by combining with molecular dynamics [3]. In this method, the physical system equally divided into spatial localized domains based on "the principle of electronic nearsightedness [4]", and then quantum computation is executed in parallel in each domain. The global physical properties are reproduced by linearly combination of the spatial domains. While the computational cost of the conventional density functional method is $O(N^3)$ with respect to the number of electrons N, thus, $O(N)$ calculation is achieved in the present method. Parallelization performance realized a high efficiency of 98% by weak scaling using up to 786,432 cores on Blue Gene/Q Mira (16 cores/node) on Argonne Laboratories (meanwhile, strong scaling is 83%) [3]. Due to the simple methodology and easy implementation while high accuracy and low computational cost, the LDC-DFT method could be used for a wider range of applications.

However, the parallel performance for the many-core processors which will be the mainstream of the next generation computer is still unknown. Here, we report the result of benchmark of LDC-DFT for liquid water using the Oakforest-PACS (OFP) system located in the Information Technology Center on the University of Tokyo, which is a massively parallel many-core supercomputer (68 cores/node).

2 COMPUTATIONAL DETAILS

We describes briefly the lean divide-and-conquer DFT method [3]. The algorithm represents the three-dimensional space Ω as a union of overlapping spatial domains, $\Omega = \cup_\alpha \Omega_\alpha$(see Fig. 1). Global charge density ρ is calculated by a real-space multigrid method [5] as linear combinations of domain local densities ρ_α. On the other hand, a plane-wave basis is used to represent local charge density ρ_α within each domain Ω_α, which takes advantage of a highly efficient numerical implementation based on fast Fourier transform. Each domain Ω_α is further decomposed into core domain $\Omega_{0\alpha}$ and the surrounding buffer layer Γ_α. $\Omega_{0\alpha}$ is a non-overlapping with

other core domains (*i.e.* $\Omega_{0\alpha} \cap \Omega_{0\beta}=0$ $(\alpha \neq \beta)$). By introducing a buffer layer of an appropriate size, it is possible not only to obtain accurate ρ_α but also to reduce the calculation cost. Our code based on LDC-DFT algorithm was written in Fortran 90 with MPI for massage passing. The FFTW library [6] was also used. In addition, the code was compiled by the command with auto-parallelization.

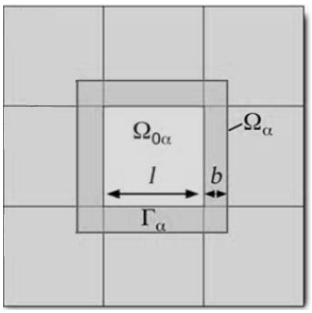

Fig. 1. Two-dimensional schematic of the LDC-DFT algorithm. The physical space Ω is a union of overlapping domains, $\Omega = \cup_\alpha \Omega_\alpha$. Ω_α is further decomposed into a non-overlapping core $\Omega_{0\alpha}$ and a buffer layer Γ_α. The thickness of the buffer layer is b.

3 RESULTS

The scalability of the LDC-DFT algorithm has been tested on OFP. The calculation model is explained below. The total system contains liquid water consisting of 648 atoms in the cubic cell of side length 19.0 Å. We first took time to choose the appropriate size of buffer layer that enable the LDC-DFT potential energy to converge within 10^{-3} a.u. per atom in the same way as our previous study [3]. The total system Ω was divided into 4×4×4 = 64 $\Omega_{0\alpha}$ core domains. The size of the domain containing the buffer layer (*i.e.* Ω_α) was a cube with a side of 13 Å. Each Ω_α domain is solved in parallelism. Fixing the domain size, we investigated the strong scaling property. The above description is the one for 512-node computation, in which the number of the OMP threads is set to be two. The elapsed time was measured for the electronic state calculation for a given atomic structure. The numbers of the used processor nodes are P = 4, 8, 16, 32, 64, 128, 256, and 512. The calculations were carried out with different numbers of OMP threads. The number of MPI processes is set to be (64 × nodes) / OMP threads. It was found that our code shows the fastest result, when the number of OMP threads is two. The elapsed time vs the number of nodes in this case are shown in Fig. 2. The parallel efficiency up to 256 nodes $\alpha \equiv (T(P=4)/T(P=256))/(256/4)$ is α=0.60, while α up to 512 nodes is 0.35.

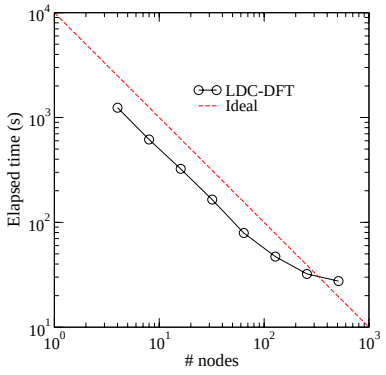

Fig. 2. Strong scaling benchmark of LDC-DFT on OFP. The ideal strong scaling line is drawn for eye guide.

4 CONCLUSIONS

In summary, we investigated the parallel performance of the LDC-DFT method [3] on many-core supercomputer Oakforest-PACS. Parallelization efficiency of strong scaling is $\alpha = 0.60$ up to 256 nodes (17,408 processors) only with Flat MPI and automatic parallelization. However, when the number of nodes exceeded 512 nodes (34,816 processors), the parallelization efficiency sharply decreased ($\alpha = 0.35$). Although we are investigating the reason for this, probably because the thread parallel performance of the FFTW library [6] has not yet been exploited sufficiently.

The result of further optimizing the code will be introduced in the presentation. In addition, we will also report the results of *ab initio* molecular dynamics simulation based on the LDC-DFT method for the biological system containing iron sulfide as an application.

ACKNOWLEDGMENTS

We would like to express our deepest and sincere gratitude to Professor Takeo Hoshi, Tottori University for useful discussions and helpful suggestions. This research used computational resources of HPCI-JHPCN System Research Project (Project ID: jh170058-NAHI, Project Representative: Takeo Hoshi, Tottori University).

REFERENCES

[1] Jean-Luc Fattebert, Daniel Osei-Kuffuor, Erik W. Draeger, Tadashi Ogitsu, and William D. Krauss. 2016. Modeling Dilute Solutions using First-Principles Molecular Dynamics: Computing more than a Million Atoms with over a Million Cores. *Proc. of International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2016*. (Nov. 2016), 12-22.

[2] Bella L. Grigorenko, Alexander V. Rogov, Igor A. Topol, Stanley K. Burt, Hugo M. Martinez, and Alexander V. Nemukhin. 2007. Mechanism of the myosin catalyzed hydrolysis of ATP as rationalized by molecular modeling. *Proc. Natl. Acad. Sci. U.S.A.* 104, (Mar. 2007), 7057–7061. DOI: 10.1073/pnas.0701727104

[3] Fuyuki Shimojo, Shinnosuke Hattori, Rajiv K. Kalia, Manaschai Kunaseth, Weiwei Mou, Aiichiro Nakano, Ken-ichi Nomura, Satoshi Ohmura, Pankaj Rajak, Kohei Shimamura, and Priya Vashishta. 2014. A divide-conquer-recombine algorithmic paradigm for large spatiotemporal quantum molecular dynamics simulations. *J. Chem. Phys.* 140, (Mar. 2014), 18A529, 14 pages. DOI: 10.1063/1.4869342

[4] E. Prodan and W. Kohn. Nearsightedness of electronic matter. *Proc. Natl. Acad. Sci. U.S.A.* 102, (Aug. 2005), 11635–11638. DOI: 10.1073/pnas.0505436102

[5] Fuyuki Shimojo, Rajiv K.Kalia, Aiichiro Nakano, and PriyaVashishta. Embedded divide-and-conquer algorithm on hierarchical real-space grids: parallel molecular dynamics simulation based on linear-scaling density functional theory. *Comp. Phys. Comm.* 167, (May. 2005), 151-164. DOI: 10.1016/j.cpc.2005.01.005

[6] M. Frigo and S.G. Johnson. The Design and Implementation of FFTW3. *Proc. of IEEE.* 93, (Feb. 2005) 216-231. DOI: 10.1109/JPROC.2004.840301