"HPC Asia 2018 :
Internatinal Conference on High Perfomance Computing in Asia-Pacific Region"

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

Graduate School of System Informatics Kobe University, Kohei Shimamura

Introduction

The lean divide and conquer density functional theory (LDC-DFT) algorithm [1,2] 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 on a multi-core processor architecture on Blue Gene/Q Mira (16 cores/node)) has been achieved [1,2]. However, the parallelization efficiency in the many-core processor architectures, which will be the mainstream of the next generation, is still unknown. Thus, we investigate the optimization work on the Oakforest-PACS (OFP) system (64 cores/node).

While the parallel performance of the LDC-DFT method on Mira realized a high efficiency of 98.4% (weak scaling) and 80.3% (strong scaling) up to 786,432 cores, the strong scaling on OFP is 60.0% up to 256 nodes (17,408 cores). Furthermore, when the number of nodes exceeded 512 nodes (34,816 cores) on OFP, the parallelization efficiency sharply decreased (35.0%). Although we are investigating the reason for this, probably the thread parallel performance of the FFTW library [3] has not yet been exploited sufficiently.

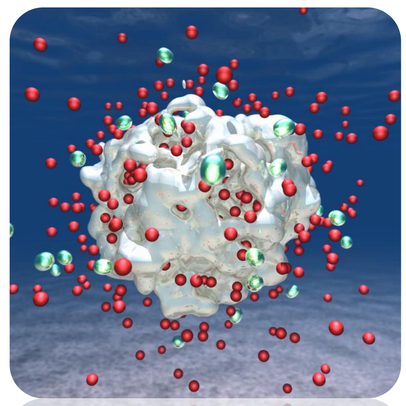

LDC-DFT based MD simulation [1,2,4] (LiAl alloy + H_2O system (16,661 atoms))

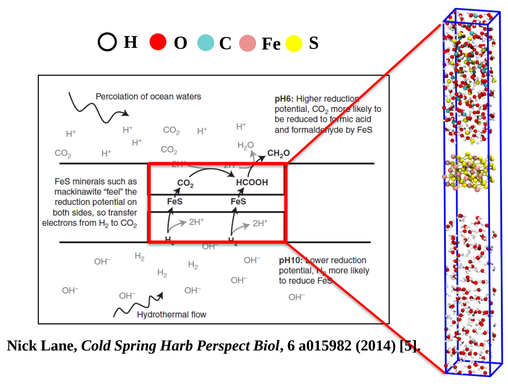

Nick Lane, *Cold Spring Harb Perspect Biol*, 6 a015982 (2014) [5].

LDC-DFT based MD simulation (Alkaline deep-sea hydrothermal vents (996 atoms))

Lean Divide and Conquer Density Functional Theory (LDC-DFT) Method

Original DC-DFT $O(N^3) \rightarrow O(N)$ [6]

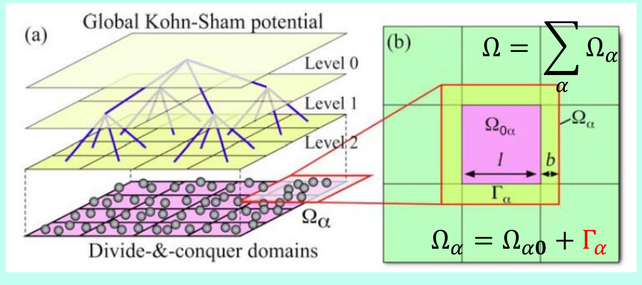

【Global charge density $\rho(\boldsymbol{r})$】

$$\rho(\boldsymbol{r}) = \sum_\alpha \rho_\alpha(\boldsymbol{r})$$

【Local charge density $\rho_\alpha(\boldsymbol{r})$】

$$\rho_\alpha(\boldsymbol{r}) = p^\alpha(\boldsymbol{r})\rho(\boldsymbol{r})$$

【Partition Function p^α】

$$p^\alpha(\boldsymbol{r}) = 0\ (\boldsymbol{r} \notin \Omega_\alpha), \quad \sum_\alpha p^\alpha(\boldsymbol{r}) = 1$$

【Local Hamiltonian】

$$\hat{H}_\alpha = \int_{\Omega_\alpha} d\boldsymbol{r} \int_{\Omega_\alpha} d\boldsymbol{r}' |\boldsymbol{r}\rangle\langle\boldsymbol{r}|\hat{H}|\boldsymbol{r}'\rangle\langle\boldsymbol{r}'|$$

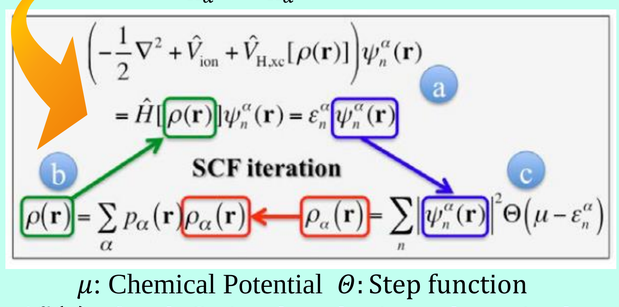

μ: Chemical Potential Θ: Step function

$\psi_n^\alpha(\boldsymbol{r})$ of each Ω_α is selected using μ determined so that the number of electrons $N = \int d\boldsymbol{r}\rho(\boldsymbol{r})$ is preserved. The global information is transmitted to each Ω_α via μ and the calculations at each Ω_α are performed in parallel to achieve $O(N)$ calculation.

LDC-DFT optimized prefactor of $O(N)$ ($(l+2b)^9 \approx b^9$)

Hardwall Potential

$$V_\alpha^{bc}(\boldsymbol{r}) = \begin{cases} 0\ (r \in \Omega\alpha) \\ \infty\ (r \notin \Omega\alpha) \end{cases}$$

Density-Template Potential[7]

$$V_\alpha^{\text{bc}}(\boldsymbol{r}) \approx \frac{\Delta\rho_\alpha(\boldsymbol{r})}{\xi}$$
$$(\Delta\rho_\alpha(\boldsymbol{r}) = \rho_\alpha(\boldsymbol{r}) - \rho(\boldsymbol{r}))$$

Linear response formula for $\Delta\rho_\alpha(\boldsymbol{r})$

$$V_\alpha^{\text{bc}}(\boldsymbol{r}) = \int d\boldsymbol{r}' \frac{\partial V(\boldsymbol{r})}{\partial\rho(\boldsymbol{r}')}\Delta\rho_\alpha(\boldsymbol{r}')$$

Local approximation for response kernel $(\frac{\partial V(\boldsymbol{r})}{\partial\rho(\boldsymbol{r}')})$

$$\frac{\partial V(\boldsymbol{r})}{\partial\rho(\boldsymbol{r}')} \cong \frac{\delta(\boldsymbol{r}-\boldsymbol{r}')}{\xi} \quad (\xi > 0)$$

Regarding the parameter ξ, we adjusted the ξ for SiC, CdSe semiconductors, LiAl and other metals. As a result, the electron density and the forces acting on the atoms did not vary greatly depending on the value, It was found that the same value (0.0333 [7]) could be employed for the materials.

Periodic boundary conditions were imposed on each Ω_α. In addition, instead of the finite difference method of real space used in the oritinal DC-DFT method, calculation within each Ω_α was performed by efficient numerical implementation using plane wave basis set based on Fast Fourier Transform [1,2].

Benchmarks for BlueGene/Q Mira (16 cores/node)

【Calculation Time】

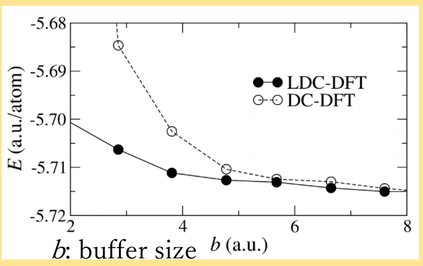

b: buffer size

- CdSeAmorphous (512 atoms)
- L = 24.164 Å
- l = 6.041 Å

When error tolerance is 10^{-3} a.u., b_O = 4.73 a.u. in the DC-DFT method and b_L = 3.57 a.u. in the LDC-DFT method are appropriate values.

Since $T \propto (l+2b)^9$ [1,2], $T_{\text{DC-DFT}}/T_{\text{LDC-DFT}} = 2.89$.

Therefore, it is possible to accelerate by about 3 times.

【Parallelization efficiency】

Using Mira (IBM Blue Gene/Q) of Argonne National Laboratory, SiC crystal 64 atoms were allocated for each core, changing from 16 to 786,432 cores, and weak-scaling was conducted [1, 2]. The parallelization efficiency (weak scaling) reached 98.4%.

On the other hand, strong scaling reached 80.3% using 786,432 cores for LiAl alloy + H_2O system.

H_2 production from water by catalyzed by LiAl alloy [4]

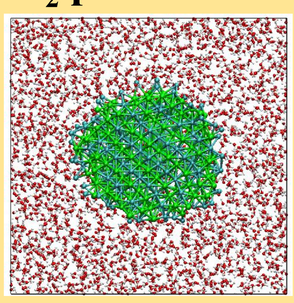

16,611 atoms
($Li_{441}Al_{441}$ + 5,243H_2O system)

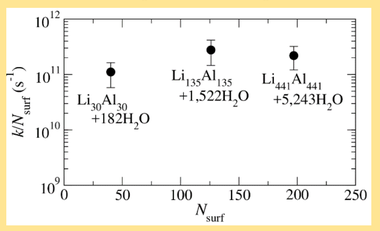

Plotting the normalized H_2 production rate as a function of the number of alloy surface atoms N_{surf} for three systems with k/N_{surf} different numbers of atoms.

The H_2 production rate is constant with respect to N_{surf} The size effect was negligible and it was shown that it reacts regardless of the curvature of the alloy surface.

Therefore, it shows that Li promotes the reaction of Al particles dramatically, and it is expected that it can be scaled up to an industrially appropriate particle size.

Benchmarks for Oakforest-PACS (64 cores/node)

【Benchmark system】

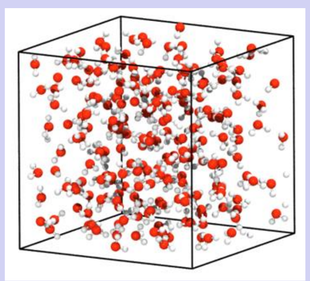

- 216H_2O (648 atoms)
- L = 19.0 Å
- $l + 2b$ = 13 Å

【Parallelization efficiency】

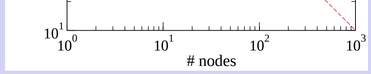

Strong scaling benchmark of LDC-DFT on OFP. The ideal strong scaling line is drawn for eye guide.

The total system Ω was divided into 4×4×4 = 64 $\Omega_{0\alpha}$ core domains.

When error tolerance is 10^{-3} a.u., b_L = 4.125 a.u. in the LDC-DFT method are appropriate values.

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 parallel efficiency up to 256 nodes α≡100×($T(P$=4)/$T(P$=256))/(256/4) is α = 60.0%, while α up to 512 nodes is α = 35.0%.

Organic molecule production in alkaline deep-sea hydrothermal vents

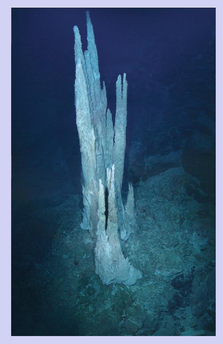

Alkaline hydrothermal vent [2]

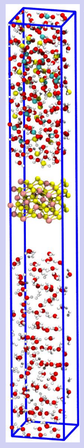

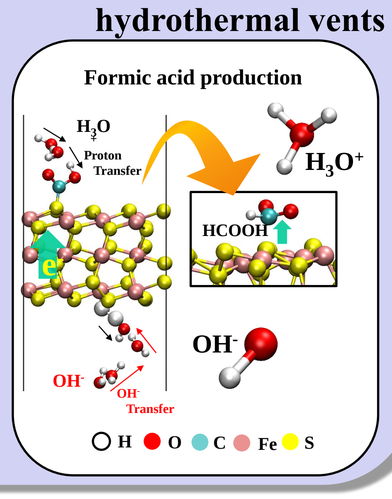

Summary/Achknowledge

In summary, we investigated and showed the parallel performance of the LDC-DFT method [1,2] on many-core supercomputer Oakforest-PACS (OFP) as well as BlueGene/Q Mira. The number of processor (cores) per node are 64 and 16 for OFP and Mira, respectively. Parallelization performance on Mira realized a high efficiency of 98.4% by weak scaling using up to 786,432 cores, while strong scaling is 80.3% [1,2]. On the other hand, parallelization efficiency of strong scaling on OFP is 60.0% up to 256 nodes (17,408 cores) only with Flat MPI and automatic parallelization. However, when the number of nodes exceeded 512 nodes (34,816 cores), the parallelization efficiency sharply decreased (35.0%). Although we are investigating the reason for this, probably the thread parallel performance of the FFTW library [3] has not yet been exploited sufficiently.

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] F. Shimojo, *et al.*, *J. Chem. Phys.*, **140**, 18A529 (2014).
[2] K. Nomura, *et al.*, *Proc. SC**14***, 661 (2014).
[3] M. Frigo *et al.*, *Proc. of IEEE.* **93**, 216 (2005).
[4] N. Lane, *Cold Spring Harb Perspect Biol*, **6**, a015982 (2014).
[5] K. Shimamura, *et al.*, *Nano Lett.* **14**, 4090 (2014).
[6] W. Yang, *Phys. Rev. Lett.*, **66**, 1438 (1991).
[7] N. Ohba, *et al.*, *Comput. Phys. Commun.*, **183**, 1664 (2012).