

Short-range $O(N)$ algorithms for First-Principles Molecular Dynamics at extreme scales *

Jean-Luc Fattebert (fattebert1@llnl.gov)

Center for Applied Scientific Computing, Lawrence Livermore National Laboratory

1 Density Functional Theory modeling for chemistry, materials science and molecular biology

Classical Molecular Dynamics (MD) simulations rely on parameterized potentials that directly describe interactions between atoms without calculating the electronic structure. Billions of atoms can be modeled with MD. However, classical MD is not applicable in many physical situations where classical potentials fail or are not even available, *e.g.*, for various conditions of pressure and temperature or in situations when bonds breaking/making occurs. Thus, simulating matter at the atomistic level often requires the much more computationally demanding calculation of the electronic structure — quantum electrons — to build realistic models.

First-Principles Molecular Dynamics (FPMD) is a very general and fundamental predictive tool to study matter at the atomistic scale. It includes calculating the electronic structure and the actual potential describing the “glue” tying atoms together. FPMD typically uses the Born-Oppenheimer approximation (classical ions surrounded by quantum electrons) and requires solving the equations of Density Functional Theory (DFT), the Kohn-Sham (KS) equations to model the quantum electronic wave functions (see *e.g.* [3]). The high computational cost of FPMD currently limits practical simulations of interest to application scientists to $O(500)$ atoms.

2 Computational challenges of quantum models on large scale parallel computers

Unlike classical physics problem where the number of variables (such as temperature, pressure, etc.) is fixed and does not grow with the system size, quantum mechanics models have a number of fields — electronic wave functions — proportional to the system size. This leads to $O(N^2)$ degrees of freedom to represent $O(N)$ electronic wave functions for a problem composed of N atoms and to $O(N^3)$ operations for standard solvers (eigensolvers).

The present practical FPMD computational limit is around 500 atoms, that is a 3D cell of about $8 \times 8 \times 8$ atoms. This is not enough to go much beyond simple atomic geometries involving single atomic species or very small molecules. Many problems need to be simplified to be tractable, and side effects of such simplifications can be quite damaging. With a thousand-fold increase in computer power and the same wall clock time requirements, an $O(N^3)$ algorithm would allow one to handle a system with no more than about 5000 atoms, that is a 3D cell of about $17 \times 17 \times 17$ only. To make the matters worse, many $O(N^3)$ algorithms are poorly scalable and communication hungry, demanding for high data throughput, on parallel computers.

*This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Advanced $O(N^3)$ algorithms have been developed to distribute computational work efficiently on large parallel computers using hybrid distributions — each processor is responsible for a fraction of the coefficients describing a fraction of the electronic wave functions (see [5]). Pushing such a strategy on today largest computers enable very large calculations (100,000 atoms in Ref. [6]), but with a time to solution far too long to be useful for any real application of interest to domain scientists.

3 Requirement for exascale algorithms: $O(N)$ complexity, short-range communications and accuracy

To make an efficient use of tomorrow’s largest exascale computers, algorithms with $O(N)$ complexity and short-range communications are needed, so that one can simulate a number of atoms directly proportional to the number of processors available, for hundreds of thousands of atoms using hundreds of thousands of processors.

A lot of research has been carried out in the last 20 year in the physics and chemistry communities in an effort to develop $O(N)$ algorithms for electronic structure calculations (see [2] for a recent review). Fewer researchers in the applied mathematics community have been working on this topic. A recent exception is the work of Benzi et al. [1], but more involvement by the applied mathematics community in this field could help at various levels, such as error estimations and solvers for instance.

Most $O(N)$ algorithms introduce some approximations or truncations of terms to reduce computational complexity. It thus becomes important to evaluate and control the accuracy of the resulting algorithms [4]. A sufficient level accuracy often means that these $O(N)$ algorithms become competitive only at large scale (more than 500 atoms).

But $O(N)$ complexity is not enough if one hopes to make an efficient use of exascale computers. Optimal algorithms also need to avoid global communications. One category of $O(N)$ algorithms with such a property is the so-called “Divide and Conquer” [7]. In this type of algorithm, the global problem is divided in sub-problems that are each solved independently, with a buffer region around the sub-domain associated with the sub-problem. The solutions of those independent sub-problems are then put together to build the global solution in an iterative manner. Dividing a problem into sub-problems can however be quite tricky and can lead to hard-to-quantify errors and hard-to-solve (unphysical) sub-problems. Other $O(N)$ algorithms are typically difficult to scale beyond a few thousand processors.

Note that in FPMD simulations, the electronic wave functions are just an intermediate quantity needed to calculate accurate forces acting on atoms and propagate these atoms according to Newton’s equation. There is the possibility of developing algorithms robust to soft errors to compute those wave functions, adapting typically used iterative solvers. Atomic positions and velocities on the other hand are likely to be quantities we want to protect from any soft errors. But since since they constitute only a very small fraction of the data, it would be easy to replicate their values.

Finally, existing $O(N)$ algorithms work for systems with a finite band gap (insulators and semiconductors) for which various mathematical properties of the solution can be exploited to reduce complexity. Tackling the very important category of metallic systems is much more challenging. In general, more research is needed to solidify existing $O(N)$ algorithms and turn them into reliable tools to use on future large parallel computers. In addition, new algorithm developments are needed to satisfy all the requirements of scalability and accuracy and to enable materials scientists, chemists and molecular biologist to take advantage of future exascale computers.

References

- [1] M. Benzi, P. Boito, and N. Razouk. Decay properties of spectral projectors with applications to electronic structure. *SIAM Review*, 55:3–64, 2013.
- [2] D R Bowler and T Miyazaki. $O(N)$ methods in electronic structure calculations. *Reports on Progress in Physics*, 75(3):036503, 2012.
- [3] E. Cancès, M. Defranceschi, W. Kutzelnigg, C. Le Bris, and Y. Maday. Computational quantum chemistry: a primer. In Ph. Ciarlet and C. Le Bris, editors, *Handbook of numerical analysis. Volume X: special volume: computational chemistry*.
- [4] J.-L. Fattebert and F. Gygi. Linear scaling first-principles molecular dynamics with controlled accuracy. *Comput. Phys. Commun.*, 162:24–36, 2004.
- [5] Francois Gygi, Erik W. Draeger, Martin Schulz, Bronis R. de Supinski, John A. Gunnels, Vernon Austel, James C. Sexton, Franz Franchetti, Stefan Kral, Christoph W. Ueberhuber, and Juergen Lorenz. Large-scale electronic structure calculations of high-Z metals on the BlueGene/L platform. In *SC '06: Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, page 45, New York, NY, USA, 2006. ACM.
- [6] Yukihiro Hasegawa, Jun-Ichi Iwata, Miwako Tsuji, Daisuke Takahashi, Atsushi Oshiyama, Kazuo Minami, Taisuke Boku, Fumiyoshi Shoji, Atsuya Uno, Motoyoshi Kurokawa, Hikaru Inoue, Ikuo Miyoshi, and Mitsuo Yokokawa. First-principles calculations of electron states of a silicon nanowire with 100,000 atoms on the k computer. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11*, pages 1:1–1:11, New York, NY, USA, 2011. ACM.
- [7] W Yang. Direct calculation of electron-density in density-functional theory. *Phys. Rev. Lett.*, 66(11):1438–1441, 1991.