Quantum Chemistry on GPU: Towards Real World Applications

Ivan Ufimtsev and Todd Martinez

Stanford University, Stanford, CA

Introduction

Two years ago, we started developing a program for quantum chemistry calculations on GPU called *Tera*Chem (standing for "teraflop chemistry"). Since then, the program evolved into powerful general-purpose software that can be used in various fields such as Chemistry, Materials Science, and Drug Discovery. The main goal we pursue is to treat as large molecular systems as possible using main workhorses of quantum chemistry – Hartree-Fock and Density functional theory methods. The code has been parallelized to run on a multi-GPU machine (POSIX threads) as well as on a multiple node GPU cluster (MPI). Herein, we describe some implementation details, demonstrate routine "GPU vs CPU" benchmarks, and present several projects employing *Tera*Chem, which are currently underway in our lab.

Restricted Hartree-Fock method

$$\mathbf{F}(\mathbf{C})\mathbf{C} = \mathbf{S}\mathbf{C}\varepsilon$$

$$\mathbf{F}(\mathbf{C}) = \mathbf{H}_{core} + \mathbf{J}(\mathbf{C}) - \frac{1}{2}\mathbf{K}(\mathbf{C}) \qquad (\mu\nu\,|\,\lambda\sigma) = \iint \chi_\mu(r_1)\chi_\nu(r_1)\frac{1}{|r_1 - r_2|}\chi_\lambda(r_2)\chi_\sigma(r_2)dr_1^3\,dr_2^3$$

$$J_{\mu\nu} = \sum_{\lambda\sigma} P_{\lambda\sigma}(\mu\nu\,|\,\lambda\sigma) \qquad \chi(r,R) = N(x-X)^{nx}(y-Y)^{ny}(z-Z)^{nz}\exp\left(-\alpha|r-R|^2\right)$$

$$K_{\mu\nu} = \sum_{\lambda\sigma} P_{\lambda\sigma}(\mu\lambda\,|\,\nu\sigma) \qquad P_{\lambda\sigma} = 2\sum_i C_{\lambda i}C_{\sigma i}$$

Two-electron integrals $(\mu\nu\,|\,\lambda\sigma)$

The integrals are organized into a $\mu\nu \times \lambda\sigma$ 2D grid
Schwartz inequality arranges them according to the magnitude in an SIMD-friendly way:

$$(\mu\nu\,|\,\lambda\sigma) \le (\mu\nu\,|\,\mu\nu)^{1/2}(\lambda\sigma\,|\,\lambda\sigma)^{1/2}$$

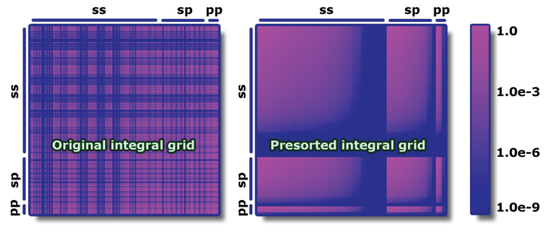

J-matrix: $J_{\mu\nu} = \sum_{\lambda\sigma} P_{\lambda\sigma}(\mu\nu\,|\,\lambda\sigma)$
Column-wise sum leads to one $J_{\mu\nu}$

K-matrix: $K_{\mu\nu} = \sum_{\lambda\sigma} P_{\lambda\sigma}(\mu\lambda\,|\,\nu\sigma)$
Sum over $(\mu\ldots|\nu\ldots)$ leads to one $K_{\mu\nu}$

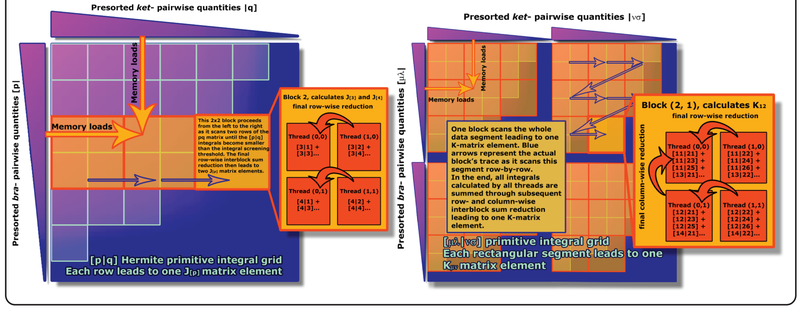

GPU performance and accuracy

Left: *Tera*Chem runs 44-650X faster on a GeForce 280GTX card in comparison with a mature third party quantum chemistry package GAMESS running on a Pentium D 3GHz core.

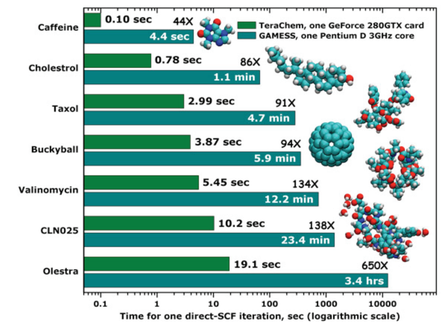

Right: Conservation of the total energy during molecular dynamics is a common metric for assessing accuracy. In our case, the energy is conserved very well (green line).

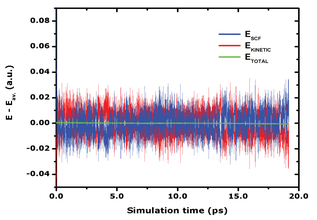

Stress sensitive materials

The main goal of this project is to develop a new generation of mechanosensitive polymers, which would change physical properties (color) in response to mechanical stress.

The mechanical sensitivity is usually achieved by incorporating small molecules, so called mechanophores, into the polymer matrix. Expansion of the material caused by mechanical stress selectively breaks or re-forms bonds in mechanophores, changing their properties including the color.

The "activation", however, is a dynamic process depending on many factors such as orientation of the molecule with respect to the external force. Thereby, extensive sampling of the molecule's conformational space is required. In addition, the method must reliably breaking and formation of chemical bonds, which is possible only in quantum chemistry methods.

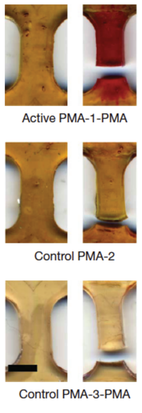

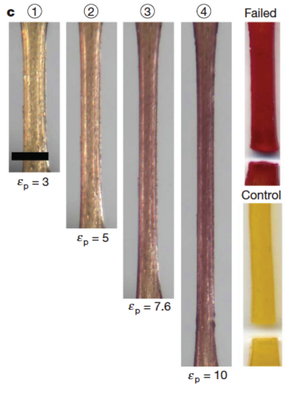

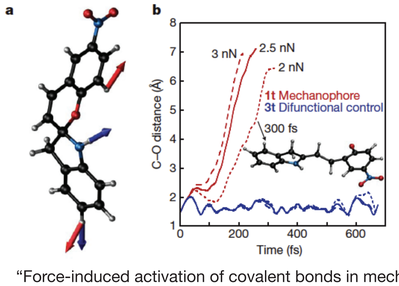

First principles molecular dynamics is a natural way to solve this problem. However, until recently, effectiveness of this method was largely stymied by its computational cost. The use of GPU brought tremendous progress to the field. The number of configurations sampled on a GPU in less than an hour takes a week on a CPU node.

"Force-induced activation of covalent bonds in mechanoresponsive polymeric materials," D.A. Davis et al., *Nature* vol. 459, 2009, p. 68

HIV-1 Protease

The HIV-1 life cycle consists of several steps such as binding of the virus to a cell's CD4 receptor, replicating the viral genome, incorporating it into the cell's DNA etc. One of these steps is responsible for cleaving the newly synthesized viral polyprotein chain at the appropriate places to create mature protein components. Without the cleavage the virus remains uninfectious. Therefore, inhibition of this process is of much pharmaceutical importance.

The cleavage is performed by the HIV Protease enzyme. HIV Protease inhibitors mimic the target site on the polyprotein chain, bind to the active site of the enzyme and finally clog it, aborting the protein maturation. The Protease has two "flaps" that can modulate accessibility of the active site by opening and closing the gate. This project is directed towards studying how different chemical reactions in the active site affect the gate state.

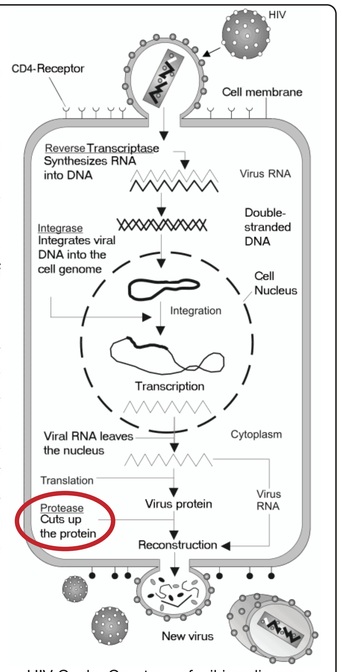

HIV Cycle. Courtesy of wikimedia.com

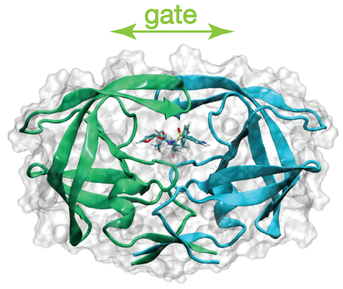

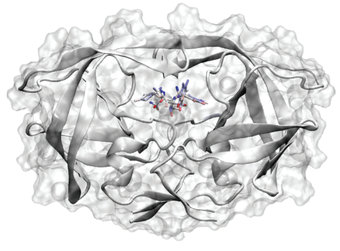

Cartoon representation of the Protease with a bound inhibitor drug molecule. The gate is closed. Upon binding, the system experiences charge transfer (red-blue coloration on the right panel). The charge transfer increases binding strength and serves as a precursor for chemical reactions.

Acknowledgements

This work was supported by the National Science Foundation (CHE-06-26354). *Tera*Chem development was carried out at the University of Illinois at Urbana-Champaign. I.S.U. is an Nvidia fellow.

References

"Quantum Chemistry on Graphical Processing Units. 1. Strategies for Two-Electron Integral Evaluation," Ivan S. Ufimtsev and Todd J. Martinez, Journal of Chemical Theory and Computation vol. 4, no. 2, 2008, p. 222

"Quantum Chemistry on Graphical Processing Units. 2. Direct Self-Consistent-Field (SCF) Implementation," Ivan S. Ufimtsev and Todd J. Martinez, J. Chemical Theory and Computation vol. 5, no. 4, 2008, p. 1004

"Quantum Chemistry on Graphical Processing Units. 3. Analytical Energy Gradients, Geometry Optimization, and First Principles Molecular Dynamics," Ivan S. Ufimtsev and Todd J. Martinez, J. Chemical Theory and Computation DOI: 10.1021/ct9003004