NC STATE UNIVERSITY

Work Supported By

Simultaneous Evolution of Multiple Molecular Dynamics Simulations

Cory Slep and Jacob Eapen
Department of Nuclear Engineering, North Carolina State University, Raleigh, NC 27695
The Consortium for Advanced Simulation of Light Water Reactors

Abstract

The need to generate statistically significant data from time intensive molecular dynamics (MD) simulations drives the search for algorithms that can take advantage of inherent parallelism in computer architectures. CUDA is an ideal platform for performing multiple MD simulations for ensemble averaging. We demonstrate a proof of concept highlighting the potential of CUDA in performing multiple MD simulations with different initial conditions.

CUDA Implementation

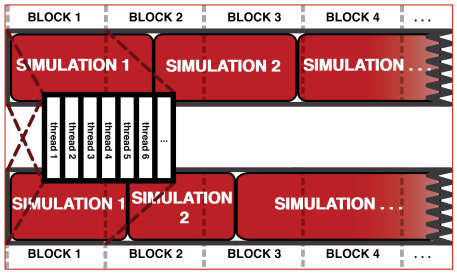

CUDA blocks are not used for memory access optimizations in this proof of concept, thus threads in each block are assigned to each simulation dynamically and do not necessarily need to have similar parameters.

CUDA Implementation Continued

Within each independent MD simulation, a cell (linked) list is implemented to account for short ranged forces that act on the atoms in the system.

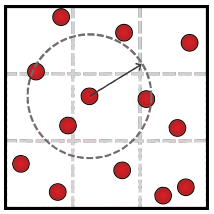

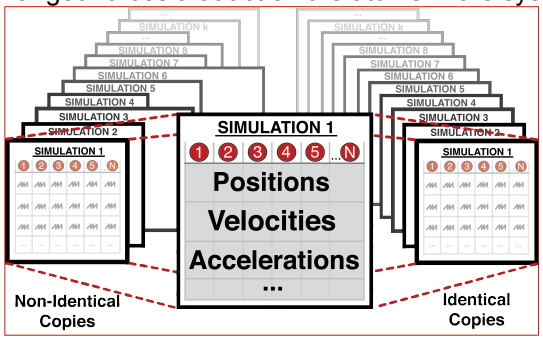

Each simulation has an independent location in the memory address which maintains the simulation data (position, velocity, and acceleration of atoms) as necessary.

Memory requirements for each independent simulation is not necessarily identical.

Results

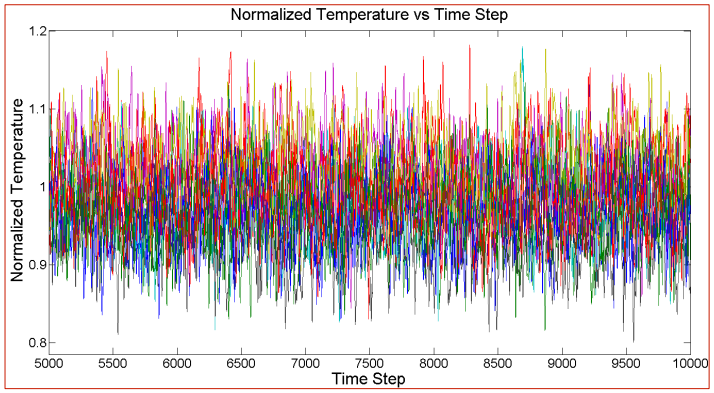

The MD simulations are performed in the NVT ensemble with a thermostat to maintain constant temperature. The above figure shows that the temperature is adequately controlled in all CUDA copies (only results from ten independent simulations are shown).

Results Continued

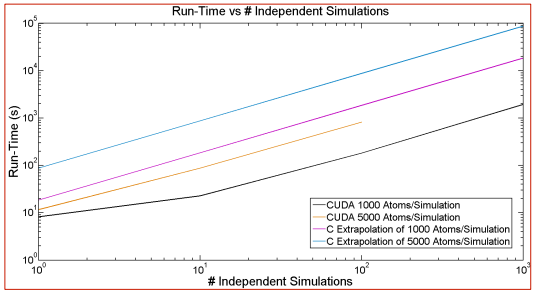

The above figure shows that the CUDA implementation is an order of magnitude faster compared to the traditional (serial) implementation of MD. It is assumed that the traditional MD has at least a linear scaling with system size (number of atoms).

Conclusion

We have performed a set of molecular dynamics simulations that take advantage of the inherent parallelism in CUDA architecture. Compared to the traditional implementation, CUDA is able to deliver the output ten times faster. The advantage of CUDA comes from the faster run times as well as the vastly reduced need to have multiple processors. Work is in progress for improving the performance through memory optimization.

Cory Slep is a III year undergraduate student in Nuclear Engineering.
Images courtesy of Natalie Kerby.