Veerapong Kaewtes 2014: Source-Level Optimization for Molecular Dynamic Simulator. Master of Engineering (Computer Engineering), Major Field: Computer Engineering, Department of Computer Engineering. Thesis Advisor: Mr. Paruj Ratanaworabhan, Ph.D. 45 pages.

LAMMPS is a classical molecular dynamics written in C++. It is used heavily for simulation of solid-state materials and soft matter. It can also be used to model atoms or, more generically, as a parallel particle simulator at the atomic scale. LAMMPS can run on a single processor or on multiprocessors. For the latter, it uses the message-passing techniques for parallel computation. LAMMPS is distributed as an open source code under the terms of the GPL by Sandia National Laboratories, a US Department of Energy laboratory.

This study investigates the runtime behavior of LAMMPS, aiming to further improve its performance. We focus on floating-point operations because they are responsible for a large percentage of LAMMPS execution time. Floating-point operations in LAMMPS have many data (true) dependencies so their execution must necessarily be serialized. Even the superscalar out-of-order processor that extracts instruction-level parallelism (ILP) cannot help much in such a situation. A more sophisticated source code transformation like software pipelining can alleviate the situation. In this paper, however, we focus on a much simpler transformation; instruction reordering that tries to set apart the producers away from the consumers as much as possible. The others optimization methods used in this paper are common techniques; these are loop unrolling, eliminating control flow to reduce branching (an, hence, the penalty from misprediction), and strength reduction.

_____ _____ ___ / ___ / ___

Student's signature Thesis Advisor's signature