SAND2017-8701R

LAMMPS Project Report for the Trinity KNL Open Science Period

Principal Investigator: Stan Moore, stamoor@sandia.gov, (505) 284-0568

Other researchers: Mitch Wood, Aidan Thompson

Abstract

LAMMPS is a classical molecular dynamics code (lammps.sandia.gov) used to model materials science problems at Sandia National Laboratories and around the world. LAMMPS was one of three Sandia codes selected to participate in the Trinity KNL (TR2) Open Science period. During this period, three different problems of interest were investigated using LAMMPS. The first was benchmarking KNL performance using different force field models. The second was simulating void collapse in shocked HNS energetic material using an all-atom model. The third was simulating shock propagation through poly-crystalline RDX energetic material using a coarse-grain model, the results of which were used in an ACM Gordon Bell Prize submission. This report describes the results of these simulations, lessons learned, and some hardware issues found on Trinity KNL as part of this work.

Background, Research Objectives, and Importance

LAMMPS helps gain insight into materials science problems at the atomic, molecular, and mesoscale levels. The objectives of the research performed with LAMMPS during the Trinity KNL Open Science Period were twofold. The first objective was to better understand how to optimize LAMMPS performance on KNL hardware. The second was to better understand how microstructure in energetic materials affects initiation behavior. Realistic energetic materials include defects like voids, and this work helped address safety issues in these materials due to accidental shocks that can lead to degradation of the detonation performance. At an even larger scale, realistic microstructures of energetic materials include multiple grains and grain boundaries, which can also affect initiation behavior under shock loading conditions. The results of this research are described in the following sections.

KNL Performance Benchmarking

LAMMPS is well suited to run on Intel KNL Xeon Phi hardware. LAMMPS has code specifically tuned for Intel hardware, written by an Intel employee, Mike Brown, along with extensive native OpenMP threading support, and growing use of Sandia's Kokkos performance portability abstractions. In order to better understand performance on KNL hardware, three different benchmarks using models of increasing complexity were used. Both strong and weak scaling up to a large number of nodes were investigated. The first benchmark used the Lennard-Jones (LJ) potential, which is computationally inexpensive but still captures the basic physics of a spherical nonpolar molecule (see Figure 1). The second benchmark used the embedded-atom model (EAM), which is commonly used to model metals such as tantalum. The third benchmark used the reactive force field (ReaxFF), which is a very versatile force field than can be used for chemically reactive systems, such as energetic materials. All three of these force fields have threaded versions, potentially allowing them to better take advantage of Trinity's KNL Xeon Phi hardware.

Figure 1 shows weak scaling for the Lennard Jones (LJ) benchmark using different accelerator packages (code specialized for specific hardware) in LAMMPS. Lessons learned from benchmarking LAMMPS on KNL hardware is that vectorization is very important to performance on KNL as seen by the "Intel/KNL/double" results. Also, in general OpenMP threading on KNL is faster than MPI-only oversubscription, except when atomic operations are required for thread safety. The LAMMPS KOKKOS package uses atomic operations in some cases and atomic operations have been found to perform poorly on KNL. The other native threaded packages in LAMMPS instead use data-duplication. We are considering adding a data-duplication option to the Kokkos package as an alternative to atomic operations.

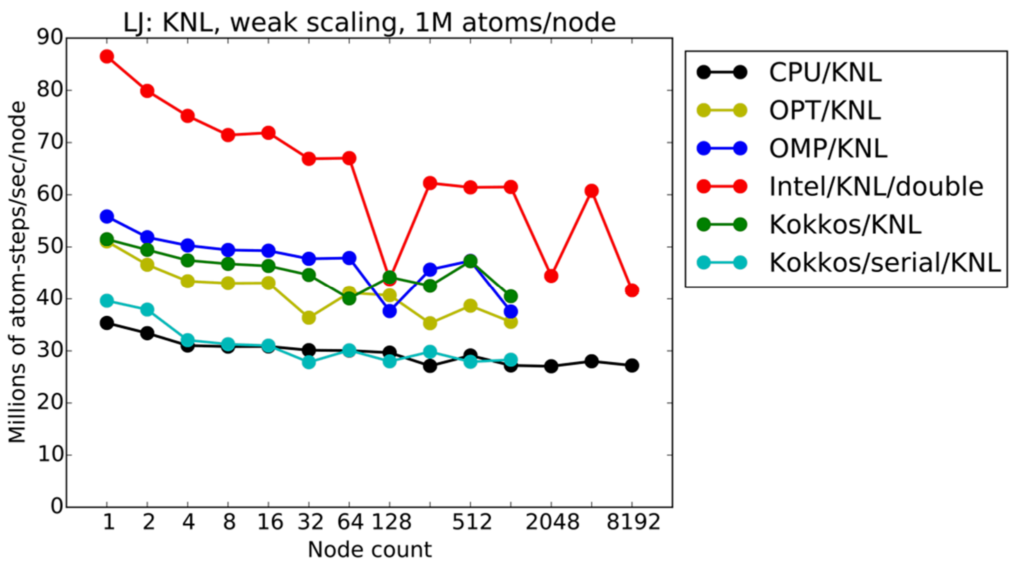

Figure 1 Weak scaling performance on TR2 for the Lennard Jones (LJ) benchmark using different accelerator packages in LAMMPS, higher is better.

One issue that made benchmarking difficult was a large startup time when running at large scale. This startup time could potentially be reduced by using the broadcast feature (*sbcast*) available using SLURM (not available during the Open Science period). In addition, some of the benchmark tests should have used more timesteps to reduce variability in the results.

ReaxFF HNS Energetic Material

Realistic energetic materials include defects like voids, and this work helps address safety issues in these materials due to accidental shocks that can lead to degradation of the detonation performance. Reactive molecular dynamics simulations of void collapse in shocked HNS energetic material were performed at unprecedented physical scales as shown in Figure 2. Total CPU time for this work on Trinity KNL was approximately 20 million CPU-hours. This work leveraged the recently-developed Kokkos-enabled implementation of the ReaxFF all-atom model in LAMMPS, which was developed by the LAMMPS ASC P&EM project. The same Kokkos-enabled LAMMPS code base is also being used to run similar simulations on Trinity Haswell nodes and Sandia's Serrano and Ghost (Intel Broadwell) CTS-1 machines. Small-scale testing has indicated that the code-base will also perform well on the LLNL Sierra ATS-2 machine (NVIDIA Volta GPU processors). These types of atomistic simulations are used to train continuum models in the CTH shock physics code developed at Sandia, i.e. using a multi-scale method. From these extremely large atomistic simulations we have gained a more mechanistic and detailed understanding of how defects in energetic materials act as initiation sites which build up into a detonation wave. These mechanisms are now present and accurately captured in Sandia's CTH code.

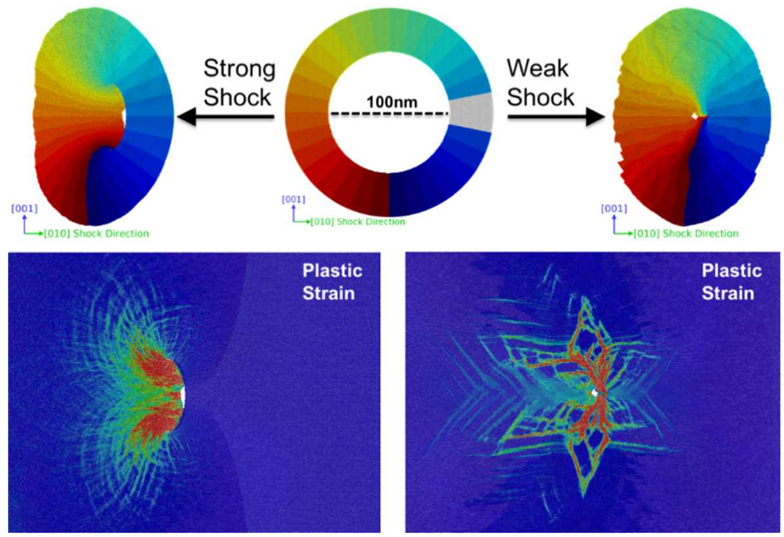

Figure 2 Demonstration of the large scale reactive MD simulations run on Trinity KNL, these simulations provide valuable information about the nature of the shock response of heterogeneous explosives.

Coarse-Grain Model of Poly-Crystalline RDX Energetic Material

At an even larger scale, realistic microstructures of energetic materials include multiple grains and grain boundaries, which can affect initiation behavior under shock loading conditions, see Figure 3. Coarse-grain models for RDX have been developed to increase length and time scales over all-atom models. A 1.1 billion particle dissipative particle dynamics (DPD) simulation of a shock front propagating though poly-crystalline RDX energetic material was run on Trinity KNL using up to 8820 nodes, or nearly the full Trinity KNL machine, see Figure 4. In the coarse-grain model, each particle represents a single RDX molecule that is composed of 21 atoms. A novel stochastic integration scheme is used to allow time steps that are much larger than those used in traditional molecular dynamics integration schemes, leading to an *O*(1000x) speedup over current state of the art reactive molecular dynamics (*e.g.*, ReaxFF). This coarse-grain model was ported to use Sandia's Kokkos performance portability abstraction layer which enables OpenMP threading. An ACM Gordon Bell Prize nomination for special achievements in time-to-solution and scalability was submitted, but unfortunately our submission was not selected as one of the finalists for the prize. This work was in collaboration with the U.S. Army Research Laboratory (ARL) and the DoD High Performance Computing Modernization Program (HPCMP).

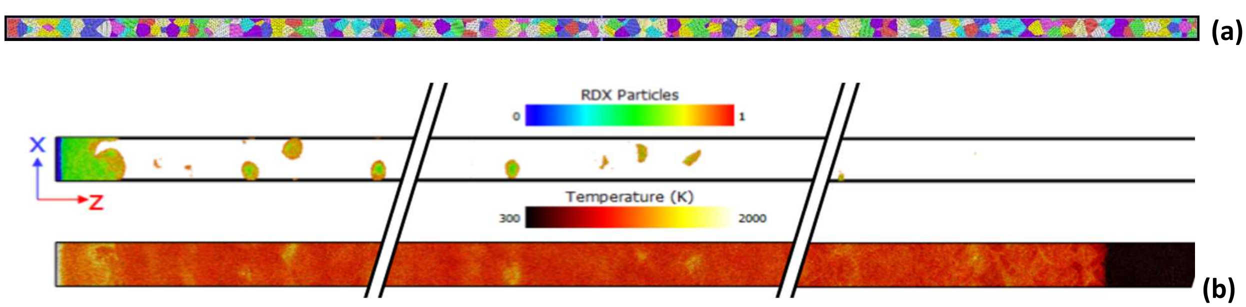

Figure 3: (a) A 3D polycrystalline RDX sample of size 40x40x2500 nm³, where individual crystal grains are delineated by the coloring scheme. (b) The reaction progress and internal temperature of shocked, polycrystalline RDX shown after 0.4 nanoseconds at various positions along the shocked direction. For the reaction progress snapshot, the unreacted material is not shown to depict the surface area of the reaction zones.

Table 1: Summary of three polycrystalline RDX models: Small, Medium, and Large.

Model Name	Sample Size (nm^3)	Average Grain Size (nm)	Number of Grains	Number of Atoms	Number of CG Particles
Small	40 x 40 x 2500	30	243	429,509,220	20,452,820
Medium	100 x 100 x 2500	75	90	2,651,446,707	126,259,367
Large	300 x 300 x 2500	225	31	23,665,453,119	1,126,926,339

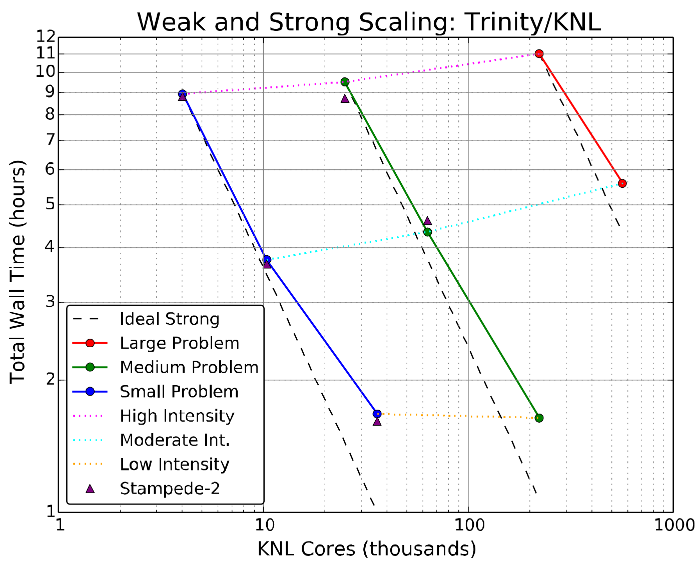

Figure 4: LogLog USER-DPD scaling results for Trinity KNL: strong scaling on three problem sizes as shown in Table 1; weak scaling at three computational intensities

For this work, we had two full-machine dedicated application times (DATs). The first DAT was only partially successful—we were able to run the small and medium problems to completion but the largest 1.1 billion particle simulations failed with MPI issues and out of memory issues. The checkpoint file was found to be corrupt and the causes of the memory issues were identified. We fixed both of these issues, and during the second DAT successfully ran all simulations to completion. However, we were plagued by a node failure issue (described below) and had to restart jobs several times. One lesson learned is that fixing memory issues is important to effectively use KNL's high-bandwidth memory (MCDRAM).

Publications

We submitted a 2017 ACM Gordon Bell Prize Submission in time-to-solution and scalability categories:

Highly-Scalable Discrete-Particle Simulations with Novel Coarse-Graining: Accessing the Microscale, Timothy I. Mattox, James P. Larentzos, Stan G. Moore, Christopher P. Stone, Daniel A. Ibanez, Aidan P. Thompson, Martin Lísal, John K. Brennan, Steven J. Plimpton

However, our submission was not chosen as one of the finalists for the prize. We are planning to submit other publications based on this work in the future.

Hardware Issues

During the Open Science period, we saw 23 node failures in which an unrecoverable hardware error was generated and the node had to be rebooted. The LAMMPS team worked with Mike Davis (Cray Analyst) to try to diagnose the issue, and a reproducer input deck for the problem was given to Cray and Intel, however the root cause is still unknown. Investigation of this issue is continuing. The LAMMPS team also discovered that one of the nodes on TR2 was faulty and was giving incorrect numerics that were crashing simulations. This bad node has been removed, leading to a more stable Trinity KNL platform.

Acknowledgements

Parts of the work were funded by ASC P&EM and DoD Productivity Enhancement, Technology Transfer, and Training (PETTT) and were in collaboration with the U.S. Army Research Laboratory (ARL) and the DoD High Performance Computing Modernization Program (HPCMP).

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.