Benchmarking LAMMPS on Utility Server GPUs

Antoinette Silas

U.S. Army Engineer and Research and Development Center
DoD Supercomputing Resource Center
Vicksburg, MS
Antoinette.R.Silas@usace.army.mil

The High Performance Computing Modernization Program's Technology Insertion (TI) uses a variety of application benchmark codes to assist in determining which HPC vendors are capable of providing the best computing machinery. LAMMPS is a key molecular dynamics code that has been included in the TI benchmark suite for years. With recent technological advances made using graphical processing units (GPUs), we are now exploring the benefit of implementing the aforementioned code within the GPU environment.

By way of the standard LAMMPS benchmarking package given for a typical TI process, we will examine its performance on a utility server that houses GPUs. The resulting times given within a CPU-only environment versus a heterogeneous environment will help determine if the effort for a successful GPU build is worthwhile. One present benefit is that the LAMMPS code already possesses the capability to be run on a GPU with minimal coding effort. Therefore, the primary focus lies within the test cases, computing environments, and results.