

LAMMPS

Summary Version

1.0

Purpose of Benchmark

Single node performance with standard optimizations and full-system parallel scaling efficiency.

Characteristics of Benchmark

LAMMPS is a well-known open source code that is capable of simulating a wide variety of “particle-like” systems. For the purposes of the Sequoia procurement, LAMMPS will be used to run a simple classical molecular dynamics problem.

Parallelization is by spatial discretization and the vast bulk of communication is nearest neighbor. This makes LAMMPS scale well, and, more importantly, it makes it possible to predict with confidence the parallel efficiency for large systems.

Mechanics of Building Benchmark

The LAMMPS home page is located at <http://lammps.sandia.gov/>, and the code can be downloaded from <http://asc.llnl.gov/sequoia/benchmarks/>. The version of 22 June 2007 should be used.

The particular problem that will be run does not require that the code be built with its FFT capability.

Mechanics of Running Benchmark

For the Sequoia benchmark, only the classical MD EAM benchmark from the standard LAMMPS benchmark suite (<http://lammps.sandia.gov/bench.html>) will be run. Information about this problem is found at <http://lammps.sandia.gov/bench.html#eam>.

The problem is to be run in the weak scaling mode, i.e., a constant number of particles per processor are used as the processor count is increased.

The figure of merit to be reported should be calculated as $32,000$ (not 32K) \times the number of processors \times 100 (time steps) divided by the wall-clock simulation time.