NAMD

Summary Version

1.0

Purpose of Benchmark (few sentences)

NAMD is the most scalable and commonly used classical MD code. It is provided here as a throughput benchmark, and intended to be run with multiple copies simultaneously in ensemble.

Characteristics of Benchmark

NAMD is a classical molecular dynamics code that simulates molecular interactions using Newtonian laws of motion. It is developed in an object-oriented style using the asynchronous data-driven language, Charm++. NAMD operates by objects, and objects interact with each other independently. The Charm++ runtime is aware of the object distribution and performs dynamic load balancing within the parallel environment of the underlying platform. The rectangular domain of the NAMD simulation is split into each of the three dimensions according to a given cutoff radius. The interactions between the atoms within the cutoff radius are calculated by directly taking into account both bonded and non-bonded interactions. The non-bonded electrostatic interactions beyond the cutoff radius are approximated by the particle-mesh Ewald method. NAMD and Charm++ support a hybrid MPI+threading model, with the number of worker threads per MPI process configurable for optimal performance on a given platform.

Mechanics of Building Benchmark

The source is in the subdirectory "src". Input data files are in "data". To build namd:

```
cd src
 make
```

Mechanics of Running Benchmark

In general, namd is run as follows:

```
mpirun -np <nodes> namd.exe <input_file> +ppn <threads_per_rank>
```

The input file will generally reference other input files (e.g. `.coor`, `.inter`, `.inter.bin`, `.vel`, `.xsc`) which are expected to be in the same directory.

Example command line parameters or inputs for:

1. Small problem:

   Please see the subdirectory `src/acceptance/namd_test_3M`
   The command "make perf512" runs this case.
   This uses the input file `ribomem-eq03G.namd` (and others) from the file called `namd_test_3M.tar.gz` at the benchmark distribution web site.

2. Medium problem:

   Please see the subdirectory
   `src/acceptance/namd_test_20M_replica`
   The command "make perf1024" runs this case.
   This uses input file `stmv_20M.namd` (and others) from the file called `namd_test_20M` (available upon request at coral-apps@lists.llnl.gov)

3. Large Mira problem:
   Please see the subdirectory `data/namd_test_100M`
   The input file is `100stmv.namd`.
   Example run: `qsub -n <nodes> namd.exe 100stmv.amd +ppn 32`

4. CORAL class problem:

Verification of Results
What FOM should be reported?

The FOM is simulation nanoseconds per wallclock day. This is the inverse of the "days/ns" field from the last occurrence of the output line "Info: Benchmark time:".

How are the benchmark results verified for correct answers?

1) The run must go to completion and output "End of program".
2) In the output, the lines beginning with "ENERGY:" show various values; the 12th field on this line is total energy. The total energy on the last ENERGY line must be less than the total energy on the first ENERGY line.