QMCPACK

Summary Version
1.2

Purpose of Benchmark
Multi-node performance of a single QMC instance and multiple instances at different twist angles of a graphite system of 256 electrons. Execute fixed number of MC steps and compute average and statistical error.

Characteristics of Benchmark
QMCPACK is an open-source continuum quantum Monte Carlo simulation code, i.e. the particle positions are randomly sampled according to various QMC algorithms. The particular version this benchmark will focus on is diffusion MC that uses diffusion-and-branch algorithm with a target population of walkers.
Parallelization is done over walkers: each MPI has a number of walkers that are then distributed over threads with OpenMP/MPI hybrid. It can also use CUDA instead of OpenMP per MPI task. The weak scaling (fixed number of walkers per parallel processing unit) is close to perfect and the strong scaling is also high, since the communication-to-computation is low. Multi-level parallelisms are exploited by running multiple instances of the system at different k points, which further reduces the statistical error for the same wall-clock time. Compilers' ability to optimize C++ template constructs and vectorized math library are critical for the performance.

Mechanics of Building Benchmark
Instructions to build and run are available at http://qmcpack.cmscc.org. QMCPACK uses cmake and toolchain files on DOE platforms are available with the code distribution. For I/O, it uses HDF5, FFTW and libxml. Developers' version r5820 on May/3/2013 will be used for the benchmark.

Mechanics of Running Benchmark
For the baseline benchmark, use N_w=199680 total walkers, corresponding to using 780 nodes of Titan and 256 walkers per node. Different N_w can be used for other benchmark runs. The figure of merit is (blocks*steps* N_w /$T_{\text{wallclock time}}$) where blocks and steps are defined in the second qmc section in the input file. It measures the efficiency for any N_w and any number of nodes. For weak scaling, fix the number of walkers per node. Fix N_w, for strong scaling.

Verification of Results
The energies and errors are used to determine the correctness of an implementation. Consult http://qmcpack.cmscc.org/how-to-on-hpc/benchmarks-on-hpc. The reference energy and error for the baseline benchmark are included with the input files for validation. Note that the selected blocks and steps of dmc section in the input file are not sufficient to reach the target error of production-

quality DMC calculations but they are sufficient to measure the performance through the figure of merit and to check the correctness.