

# QMCPACK Bechmark OLCF-6

This is the OLCF-6 benchmark for QMCPACK.

This repository describes the QMCPACK benchmarks and workflow for OLCF-6. The general OLCF-6 benchmark run rules apply except where explicitly modified in this document and should be reviewed before running this benchmark.

## QMCPACK Overview

QMCPACK is an open-source many-body ab initio Diffusion Quantum Monte Carlo code for computing the electronic structure of atoms, molecules, and solids. Ab-initio Quantum Monte-Carlo is one of the leading methods that allow the calculation of many-electron interactions in solids that goes beyond the capabilities of the most widely used density functiona approaches in being able to capture the many body effects that are inaccessible less computationally demanding methods. The QMCPACK code is described and available from <https://www.qmcpack.org> (<https://www.qmcpack.org>).

Electron positions are randomly sampled by a large number of Markov chains (or “walkers”). The benchmark uses a diffusion quantum Monte Carlo method that uses a diffusion-and-branch algorithm and a target population of walkers to evolve the walkers in time. During the simulation, the fluctuating population of walkers is load balanced between nodes at every time step of the simulation. In production simulations, measurements/statistics are collected at each time step of the simulation and for all the walkers.

The benchmarks are particularly sensitive to floating point, memory bandwidth and memory latency performance. To obtain high performance, the compiler’s ability to optimize and vectorize the application is critical. Strategies to place more of the walker data higher in the memory hierarchy are likely to increase performance.

## Code access and compilation

The process of building the benchmark has three basic steps: obtaining the source code, configuring the build system, and compiling the source code.

### Obtaining the QMCPACK source Code

The QMCPACK used for this benchmark is provided here as QMCPACK-OLCF\_6-ref.tar.gz. The source code is commit e99431b from QMCPACK GitHub repository at: <https://github.com/QMCPACK/qmcpack> (<https://github.com/QMCPACK/qmcpack>)

### Configuring the QMCPACK build system

QMCPACK uses Cmake to configure its build. Configuration options are described in the documentation at <https://qmcpack.readthedocs.io/en/develop/installation.html#configuration-options> (<https://qmcpack.readthedocs.io/en/develop/installation.html#configuration-options>).

## QMCPACK configuration used in Frontier benchmarks

### QMCPACK version

This benchmark was constructed using the QMCPACK as distributed in QMCPACK-OLCF\_6-ref.tar.gz.

### Software prerequisites

Building QMCPACK has the following compiler and library requirement:

- C/C++ compiler. **C++ compilers are required to support the C++ 17 standard.**
- An MPI library.
- BLAS/LAPACK, numerical, and linear algebra libraries. Use platform-optimized libraries where available.
- CMake, build utility (<http://www.cmake.org> (<http://www.cmake.org>)).
- Libxml2, XML parser (<http://xmlsoft.org> (<http://xmlsoft.org>)).
- HDF5, portable I/O library (<http://www.hdfgroup.org/HDF5/> (<http://www.hdfgroup.org/HDF5/>)). Good performance at large scale requires parallel version >= 1.10.
- BOOST, peer-reviewed portable C++ source libraries (<http://www.boost.org> (<http://www.boost.org>)). Minimum version is 1.61.0.
- FFTW, FFT library (<http://www.fftw.org/> (<http://www.fftw.org/>)).

### Building QMCPACK

When using the provided QMCPACK distribution QMCPACK-OLCF\_6-ref.tar.gz, follow these steps:

```
tar -xzf QMCPACK-OLCF_6-ref.tar.gz
cd QMCPACK/build
cmake -DCMAKE_C_COMPILER=mpicc \
      -DCMAKE_CXX_COMPILER=mpicxx \
      ..
make -j 8
```

Possible configuration options can be found in the QMCPACK manual at <https://qmcpack.readthedocs.io/en/develop/installation.html#configuration-options> (<https://qmcpack.readthedocs.io/en/develop/installation.html#configuration-options>).

### Testing the build

As the benchmark runs are very short Monte Carlo runs that are not executed to convergence the results from the benchmark runs themselves will not be sufficient to validate the code. The build of QMCPACK should be validated using the procedure described in the QMCPACK documentation

found at <https://qmcpack.readthedocs.io/en/develop/installation.html#testing-and-validation-of-qmcpack> (<https://qmcpack.readthedocs.io/en/develop/installation.html#testing-and-validation-of-qmcpack>)

Ideally these tests should be used:

```
ctest -R unit
ctest -R deterministic -LE unstable
ctest -R short -LE unstable
```

## Running the benchmark

Input files and batch scripts for the various systems are provided in the benchmarks directory. Each problem has its own subdirectory within the benchmarks directory. The input file for the benchmark problems is provided as `NiO.in.xml`. This problem can be run with different number of walkers per node or device and the best performance number obtainable should be reported.

The input files for the benchmark runs are the pseudopotentials (`Ni.opt.xml` and `O.xml`), the orbital file (`NiO-fcc-supertwist111-supershift000-S64.h5`) and the main qmcpack input (`NiO.in.xml`). To run the benchmark, `$walkers` in the input file has to be replaced by the number of walkers per MPI rank and `$orbpath` with the path to the orbital file.

Within the respective benchmark directory run the job through the respective equivalent of

```
srun -n #mpi_ranks /path/to/qmcpack/qmcpack --enable-timers=fine NiO.in.xml >qmc.out
```

(Example script for automatic replacement and job submission on Frontier for multiple walker and node counts are provided as `submit_runs_nodes_N.py`)

The provided script `qmc_throughput.py` extracts the throughput number for a benchmark run from the QMCPACK stdout, i.e. the `qmc.out` from above.

The results reported should be **monte carlo steps per second for the proposed system, monte carlo steps per second for a single node and number of walkers per node** used to achieve these results.

## Results

### Reference performance: OLCF Frontier

The maximum number of walkers on a node is limited by the available device memory. Here we report the reference performance for 1, 256, 4096, 6120 and 8192 nodes on OLCF's Frontier system.

#### Benchmark. Nodes Walkers / rank Steps/second Steps/second / node

NiO_256	1	24	11.396	11.396
NiO_256	256	24	2885.753	11.272
NiO_256	4096	24	46316.130	11.308
NiO_256	6120	24	69359.202	11.333
NiO_256	8192	24	92604.912	11.304

As illustration we provide the steps/second on a single Frontier node up to the memory limit.

NiO 256; one Frontier Node

