1 OAK RIDGE National Laboratory
2 Argonne NATIONAL LABORATORY
3 Sandia National Laboratories
4 Lawrence Livermore National Laboratory
5 intel
6 NVIDIA
7 STONE RIDGE TECHNOLOGY

QMCPACK

Predictive and systematically improvable quantum-mechanics based simulations of materials

P. Kent (PI)[1], J. Krogel[1], G. Lopez[1], A. Tillack[1], Y. W. Li[1],
L. Shulenburger[2], A. Baczewski[2], C. Trott[2], A. Cangi[2], R. Clay[2], M. Morales-Silva[3], A. Correa[3],
A. Benali[4], N. Romero[4], Y. Luo[4], M. Dewing[4], J. Kim[5], J. Larkin[6], K. Esler[7]

QMC for accurate electronic structure

We seek to accurately predict the quantum-mechanically determined properties of materials using only the atomic structure and composition.

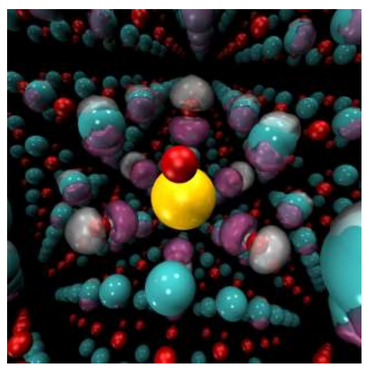

Exascale will enable realistic materials with defects & dopants from across the periodic table, and calculations on simpler materials with enough throughput for upscaling.

(NiO:K charge density)

Quantum Monte Carlo (QMC) techniques use stochastic methods to treat the Schrodinger equation's high dimensionality .

- Broadly applicable from metals, strongly correlated oxides, to isolated molecules.
- Most QMC methods have variational errors, and the few approximations are potentially systematically improvable (only achievable today for molecules).
- Reasonable system size scaling : $O(N^3)$-$O(N^4)$
- Ample opportunities for parallelism.

Goal:

Simulate challenging transition metal oxides of 1000 atoms to 10meV statistical accuracy, in 1 day, at the exascale, with performance portability.

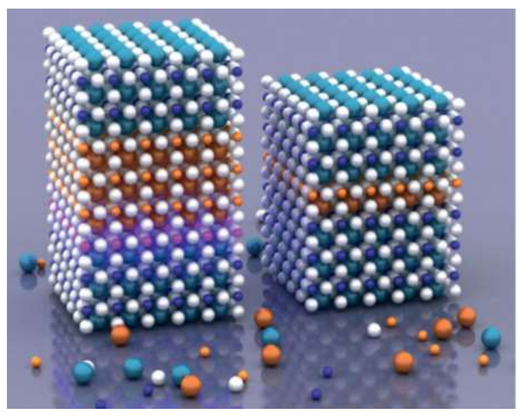

Key Challenges for Exascale

- New levels of on-node concurrency, for speed and memory usage. E.g. Threading >30x current.
- Portable optimization strategies to optimize for bandwidth of the overall memory hierarchy.
- Sufficient algorithmic and data layout flexibility to achieve performance portability across current and near-future architectures.
- New standards-based languages and software tools to express and achieve the above.

ECP Integration Activities

- Kokkos, SOLLVE - Initial performance portability focus
- SPACK – QMCPACK and miniQMC packages
- EZ/ZFP, VeloC – Data compression and checkpointing
- Suggestions for more integrations are welcome!

QMCPACK and QMC Miniapps

QMCPACK is an open source C++ code available at **github.com/QMCPACK**. Two miniapps (miniqmc, miniafqmc) encapsulate the core algorithms for experimentation with new software technologies. Continuous integration speeds development.

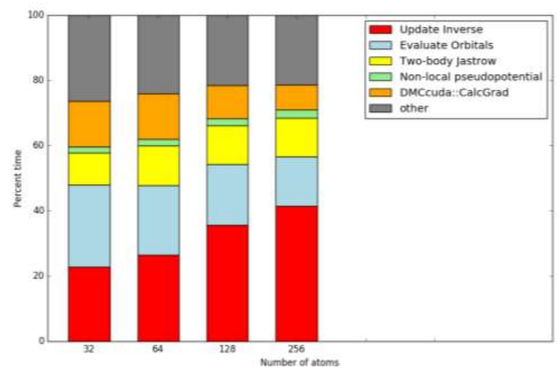

Profile of existing CUDA implementation. Only a handful of kernels need to be highly optimized and performance portable. Kernels are memory bandwidth &/or FLOP limited

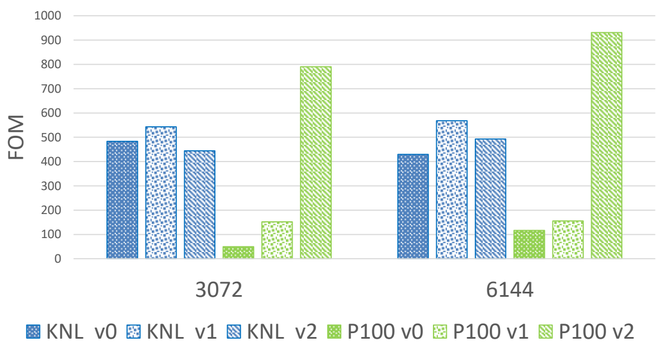

Initial investigation of OpenMP 4 for one kernel called for different problem sizes. Different versions of the kernel with different loop orderings are optimal on KNL and P100.

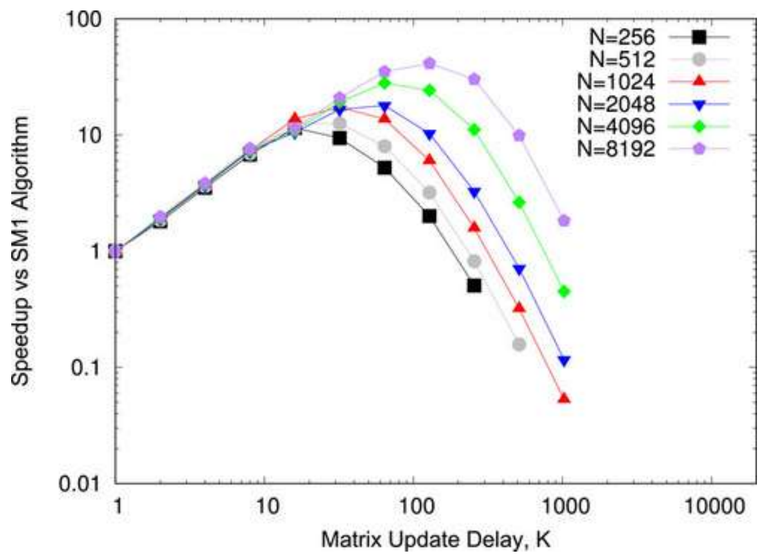

Performance of new OLCF developed delayed inverse update algorithm relative to existing algorithm. Matrix updates are delayed for K columns, then applied en bloc for higher efficiency. For N = 8192 and using K = 256, the measured speedup is 41.3× over the conventional algorithm.

*T. McDaniel et al J. Chem. Phys. **147** 174017 (2017)*

Initial Summit Results

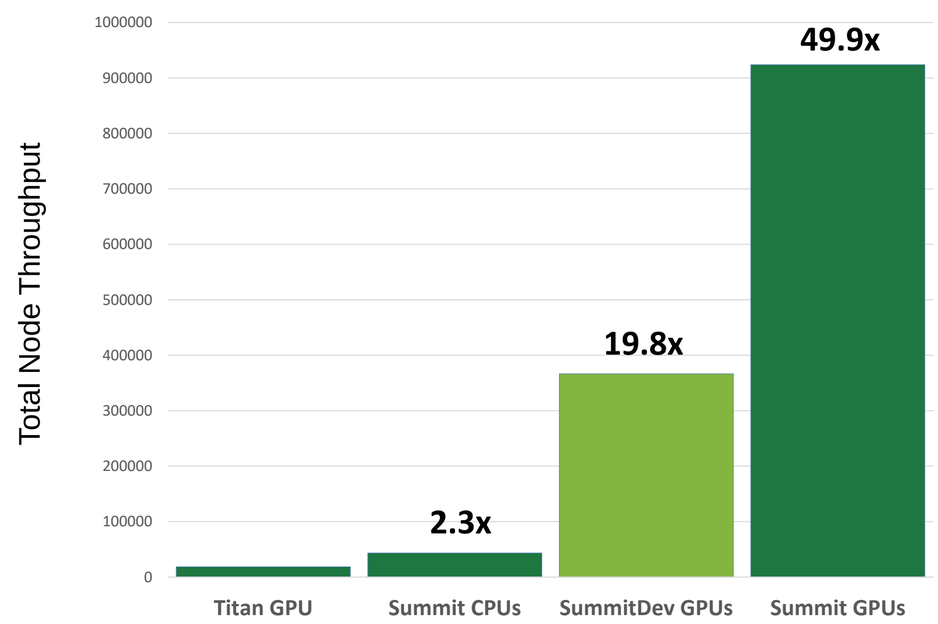

QMCPACK v3.4.0 NiO 128 atom cell. Power CPU reference uses optimal 2 MPI tasks, 42 OpenMP threads each and optimized "SoA" version.