Leveraging stochastic electronic structure methods at the exascale

E.J. Landinez [a], G. Abrusci [b], A. Alavi [c], V. G. Chilkuri [d], F. Coppens [e], A. Delval [e], C. Filippi [a], M. Hapka [f], M. Hoffer [e], W. Jalby [e], P. López Ríos [c], K. Nakano [g], P. Oliveira [e], R. Panades [a], K. Pernal [f], E. Posenitskiy [d], R. Shinde [a], A. Sokół [f], S. Sorella [g], A. Scemama [c]

[a] Computational Chemical Physics, MESA+ Institute for Nanotechnology, University of Twente,The Netherlands, [b] CINECA, Consorzio Interuniversitario, Casalecchio di Reno (Bologna), Italy, [c] Max-Planck Institute for Solid State Research, Stuttgart, Germany, [d] Laboratoire de Chimie et Physique Quantiques, Université de Toulouse (UPS), CNRS, France, [e] LI-PaRAD, UVSQ, Université Paris-Saclay, France, [f] Institute of Physics, Lodz University of Technology, Lodz, Poland, [g] International School for Advanced Studies (SISSA), Trieste, Italy

Abstract

We present here the software development strategy and the current achievements of the European Centre of Excellence “Targeting Real chemical accuracy at the EXascale” (TREX). The main objective of TREX is the development of a user-friendly and open-source software suite in the domain of stochastic electronic structure simulations, which integrates a set of flagship quantum Monte Carlo codes within an interoperable, high-performance platform. Core of our software efforts is the creation of two libraries, the TREXIO and the quantum Monte Carlo kernel library (QMCkl).

1. Quantum Monte Carlo (QMC)

Stochastic solution ot the interacting Schrödinger equation

Advantages of real-space QMC:

- Low memory requirements
- Distribute walkers on cores/nodes
- Often no blocking communication:
  → near-ideal scaling
- One QMC step is fast (~1ms)

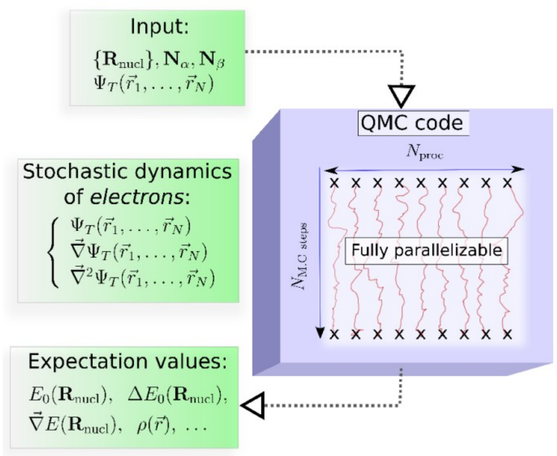

Figure 1. QMC workflow calculation

Challenges:

- Linear algebra with small matrices
- Difficult to parallelize within a QMC trajectory

2. Goals

- TREX (QMC and beyond) codes ready for exascale systems

NECI GAMMCOR TurboRVB

CHAMP QMC = Chem Quantum Package

Figure 2. EU exascale Leonardo and LUMI supercomputers

- Design ecosystem of efficient, scalable, and inter-operable software
- Robust management of complex workflows AiiDA

3. Strategy

Provide libraries with the common computational kernels for QMC instead of rewriting individual code

- Kernels developed by QMC experts in a human readable library
- Compile kernels into a HPC-library by HPC experts
- Scientist controls programming choices keeping the computational performance

4. TREXIO library

- Source code in pure C (C99) for the best performance and portability
- Bindings in Fortran (using ISO_C_BINDING) and Python (using SWIG)
- Official releases are available via GitHub, Conda, and Spack

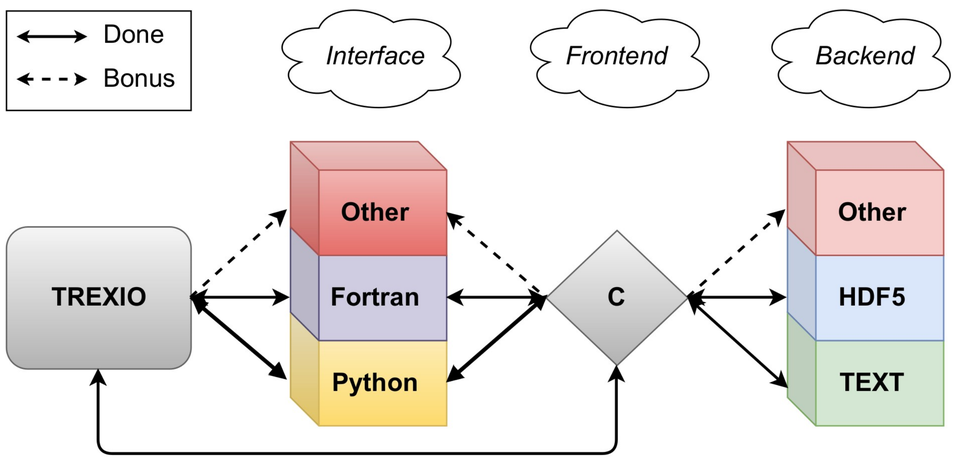

Figure 3. Schematic design of the TREXIO library

Backend → HDF5: Default, HPC I/O based on HDF

Backend → TXT: Debug, fallback when HDF5 not available

5. TREXIO FORMAT

Stores electronic wave functions data

- Single-reference Ψ
- Multi-reference $\Psi = \sum_i^N c_i \Psi_i$ → 001 110

- Portable and unified format
- HPC I/O

- **Data is organized into groups**

• Metadata	• basis	• ao	• mo	• cell	• pbc
• Electron	• ecp	• ao_1e_int	• mo_1e_int	• qmc	• Etc..
• Nucleus	• determinant	• ao_2e_int	• mo_2e_int	• rdm	

6. QMCkl library

- QMCkl is a C library with C-compatible API
- System functions in C (memory allocation, thread safety,etc)
- Computational kernels in Fortran for readability
- Documentation library: Literate programming with org-mode

7. Demonstrations

QMC ecosystem

Interoperability between codes through unified TREXIO format

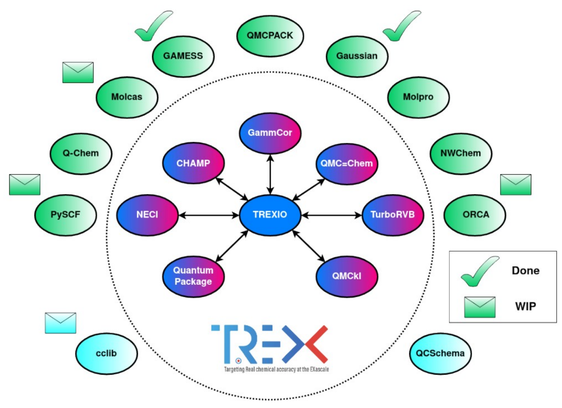

	HF/VMC Energy(Ha)	VMC 2-body Jastrow Energy (Ha)
Quantum Package	-36.5055	-
CHAMP	-36.5051±0.0004	-36.6419±0.0005
TurboRVB	-36.5060±0.0004	-36.6414±0.0004
QMC=Chem	-36.5061±0.0005	-36.6410±0.0004

Table 1. Benchmark benzene calculation using same TREXIO input

Performance improvement

QMCkl Integration with codes and kernels customization

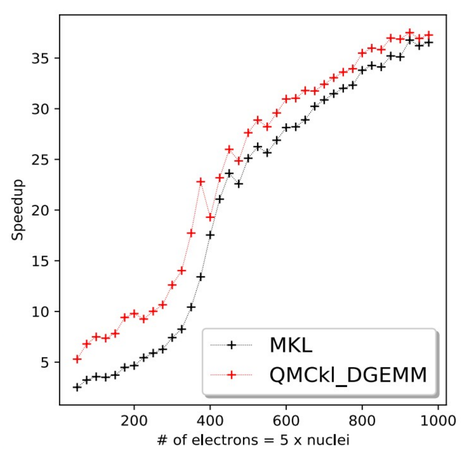

Figure 4. Speedup from kernel customization for CHAMP 3-body Jastrow factor and DGEMM in QMCkl

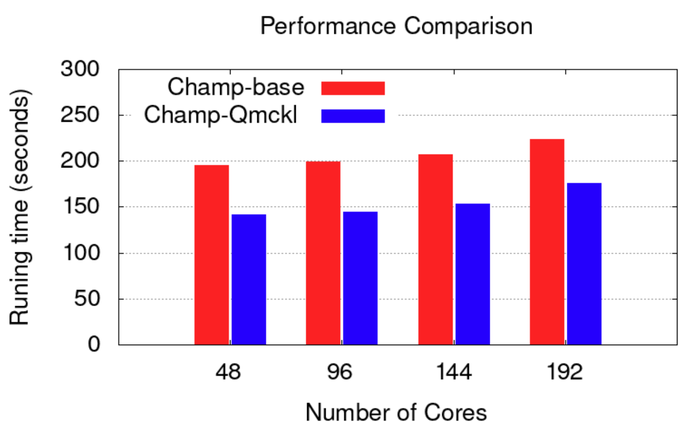

Figure 5. Current speedup of ~26% for energy calculation in CHAMP after integrating QMCkl in MO-subroutines

Acknowledgments

Targeting Real chemical accuracy at the EXascale

The European Centre of Excellence in Exascale Computing TREX - Targeting Real Chemical Accuracy at the Exascale – has received funding from the European Union’s Horizon 2020 - Research and Innovation program - under grant agreement no. 952165. The content of this document does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.

Get libraries here!