2010PA5233 – Load Balancing of Molecular Properties Calculations In VeloxChem Program

Type: Preparatory Access Type C

Systems: SuperMUC-NG

Expected project start date: February, 2020

Code: VeloxChem

Description: The quantum chemical computations of molecular properties are one of the key components of various spectroscopies. They are extensively applied to study various molecular materials to gain a microscopic understanding of physical processes driving material response and to develop "structure-to-property" relationships used in the search for novel molecular materials. Among the quantum chemistry methods, the density functional response theory (DFRT) approaches are most frequently employed to investigate various linear and non-linear properties of molecular materials due to good ration between accuracy and computational cost. Unfortunately, DFRT approaches become computationally expensive for extended molecular systems, which contain over few hundreds of second-row atoms. To overcome this limitation of DFRT approaches, we recently developed a new quantum chemistry program VeloxChem (https://veloxchem.org), which is designed to effectively perform DFRT computations of various molecular properties in HPC environments. Currently, the VeloxChem program can effectively perform DFRT computations on systems up to 500 atoms using up to 16 000 CPU cores. In this project, we aim to enhance the scalability of DFTR computations in the VeloxChem program up to 100 000 CPU cores, and in this way, enable routine investigations of molecular properties in systems consisting of 800+ second-row atoms. This improvement will allow us to carry out investigations of the optical properties of realistic metal and metal oxide nanoparticles.

Technical Details: Two-electron integrals computation and exchange-correlation functional derivatives integration algorithms are paralleled using hybrid OpenMP/MPI scheme, where MPI used for communication across compute nodes and OpenMP used inside compute nodes (tasks based parallelization). Furthermore, computation is split over several domains of compute nodes using using split MPI communicators.

Current and target performance: Currently, VeloxChem scales up to 16000 CPU cores. We would like to achieve good scaling up to 100 000 CPU cores. This scaling/performance would make feasible modelling of optical properties of 1-4 nm nanoparticles.

Bottlenecks: MPI communication, load balancing between domains of compute nodes.

Level of collaboration: We would like to obtain help on profiling and at lesser degree in optimising MPI related code parts of VeloxChem program:

1. Profiling of realistic computations: estimated 1 person month.
2. Optimization of load balancing/communication code: estimates 2 person months.
3. Benchmarking of improved code on realistic computations: 1 person month.

Optimization Work: Optimization of MPI communication patterns at C++ layer of application. Help in implementing dynamic load balancing scheme for multiple MPI domains of compute nodes. Profiling and benchmarking of code in realistic computations scenarios.

PMs: 4