Efficient Hybrid DFT Simulations of Solvated Biomolecules

J. Bernholc, E. L. Briggs, M. Hodak, W. Lu, and F. Rose
Center for High Performance Simulation and Department of Physics
NC State University, Raleigh, NC 27695-7518, USA

One of the most important challenges in quantum simulations of biomolecules is efficient and accurate inclusion of the solvent, because the solvent atoms usually far outnumber those of the solute. We have developed a hybrid method that allows for explicit quantum-mechanical treatment of the solvent at low computational cost. It combines Kohn-Sham (KS) density functional theory (DFT) with orbital-free (OF) DFT. KS DFT is used to describe the biomolecule and its first solvation shells, while the OF DFT is employed for the rest of the solvent. The OF method scales linearly with the number of atoms and is capable of handling 10^5 solvent molecules on current supercomputers, while taking only a small percentage of the total computational time. The compatibility between the KS and OF DFT methods enables seamless integration between the two. In particular, the flow of solvent molecules across the KS/OF interface is allowed and the total energy is conserved. This method is implemented in the RMG (real-space multigrid) code, which employs grids for efficient parallelization and multigrid preconditioning for convergence acceleration.

The hybrid method has been used to investigate the binding of copper ions to proteins involved in prion (PrP) and Parkinson's diseases, and to study an experimental drug for Alzheimer disease. Our results for PrP show how this protein binds multiple copper ions while undergoing complex structural rearrangements and becoming more resistant to misfolding and thus to initiation of the prion disease. For alpha-synuclein, the Parkinson's disease (PD) protein, we show that copper binding modifies the protein structurally, making it more susceptible to misfolding -- an initial step in the onset of PD. In Alzheimer disease (AD), a drug based on a copper chelator shows significant promise in clinical trials. We investigated the mechanism of its action and found a low-activation-energy pathway for copper removal from amyloid-β, the principal protein involved in AD.

Very recently, we have adapted the RMG method to the latest generation of supercomputers that use multicore CPUs and high-performance accelerators (GPUs). Our latest implementation uses one MPI process per node, rather than one per core. It achieves intra-node parallelization through POSIX threads and OpenMP, and efficiently distributes the computational load between CPUs and GPUs, while minimizing costly data transfers between CPU and GPU memories. The revamped code scales to over 100,000 CPUs and 10,000 of GPUs, and easily reaches multi-petaflop performance.