Parallelization of Monte Carlo Simulations using Prefetching

Harold Hatch [C, S] and Vincent Shen
Chemical Informatics Division, NIST, Gaithersburg, MD, U.S.A.
harold.hatch@nist.gov

In order to enable molecular simulations on a scale that is not currently practical, algorithms must efficiently utilize multicore processors that continue to increase in total core count over time with relatively stagnant clock speeds. Although parallelized molecular dynamics (MD) software has taken advantage of this trend in computer hardware, single-particle perturbations in Monte Carlo (MC) are more difficult to parallelize than system-wide perturbations in MD using domain decomposition. Instead, prefetching reconstructs the serial Markov chain after computing multiple MC trials in parallel. Canonical ensemble MC simulations of a Lennard-Jones fluid with prefetching resulted in up to a factor of 1.7 speedup using 2 threads, and a factor of 3 speedup using 4 threads. Flat histogram grand canonical ensemble simulations see further improvement with up to a factor of 3.5 speedup using 4 threads. Strategies for maximizing efficiency of prefetching simulations are discussed, including the potentially counterintuitive benefit of reduced acceptance probabilities. Determination of the optimal acceptance probability for a parallel simulation is simplified by theoretical prediction from serial simulation data. Finally, complete open-source code for parallel prefetch simulations was made available in the Free Energy and Advance Sampling Simulation Toolkit (FEASST).