Non-equilibrium umbrella sampling on Blue Gene/P using one-sided communication

Jeff Hammond

Leadership Computing Facility
Argonne National Laboratory

17 March 2011

Motivation

Rare events pose a significant challenge for simulation since dynamics are simulated significantly slower than nature. Much compute time is spent just waiting for the conditions to initiate an event. Even under optimistic assumptions, some rare events will never be amenable to direct dynamics.

$$\tau_{MD} \approx 10^{-15}\ \mathrm{s} \qquad T_{event} \approx 10^{-6}\ \mathrm{s}$$
$$\tau_{comp} \approx 10^{-3}\ \mathrm{s} \qquad T_{sim} > 11\ \mathrm{days}$$

Non-equilibrium umbrella sampling (NEUS) replaces a sequential single dynamics trajectory with many trajectories which can be computed in parallel and is thus a massively-parallel algorithm for the time dimension of biomolecular simulation.

The Concept

The Concept

Mountain climbing would be easy if it was broken up into pieces and each climber just had to cover one section, like a relay race.

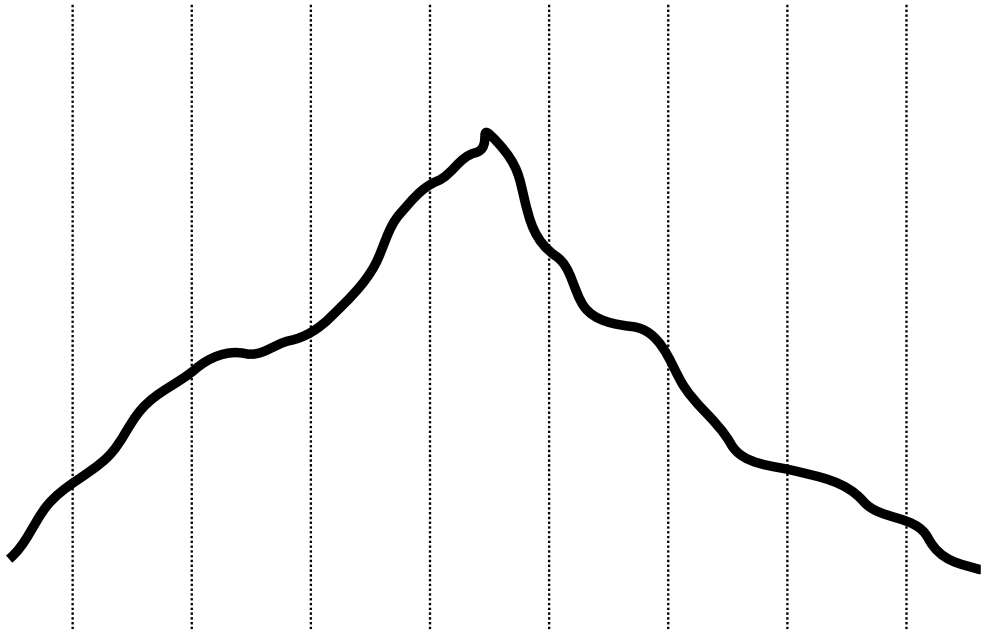

Non-equilibrium umbrella sampling is a way to make molecular mountain climbing computationally tractable.

Discretizing phase space

NEUS samples an order parameter space using one or more degrees of freedom from the system that describe the relevant dynamics. Many regions are sampled in parallel using independent simulations. NEUS enforces sampling in every region of phase space, thereby simulating all stages of the event simultaneously and eliminating the wait time required to see a rare event.

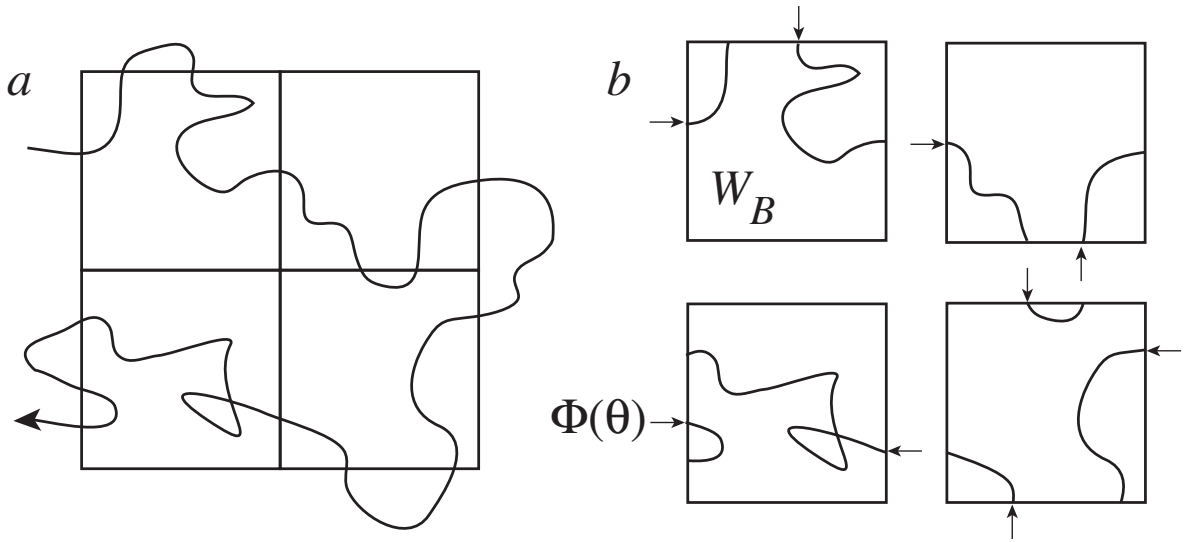

Combining probabilities

Each region is given a weight, that is determined using boundary crossing statistics. The weights are used to build the full steady-state distribution from the regional ones.

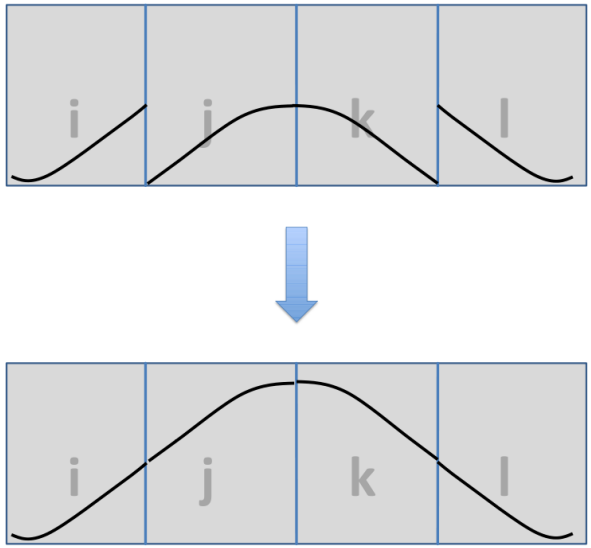

To conduct the "regional" simulations in nonequilibrium systems, one needs to be able to remove the bias in a physical way. The key insight is to run unbiased dynamics and restart walkers using a flux input distribution that is developed on-the-fly.

Parallelization using Global Arrays

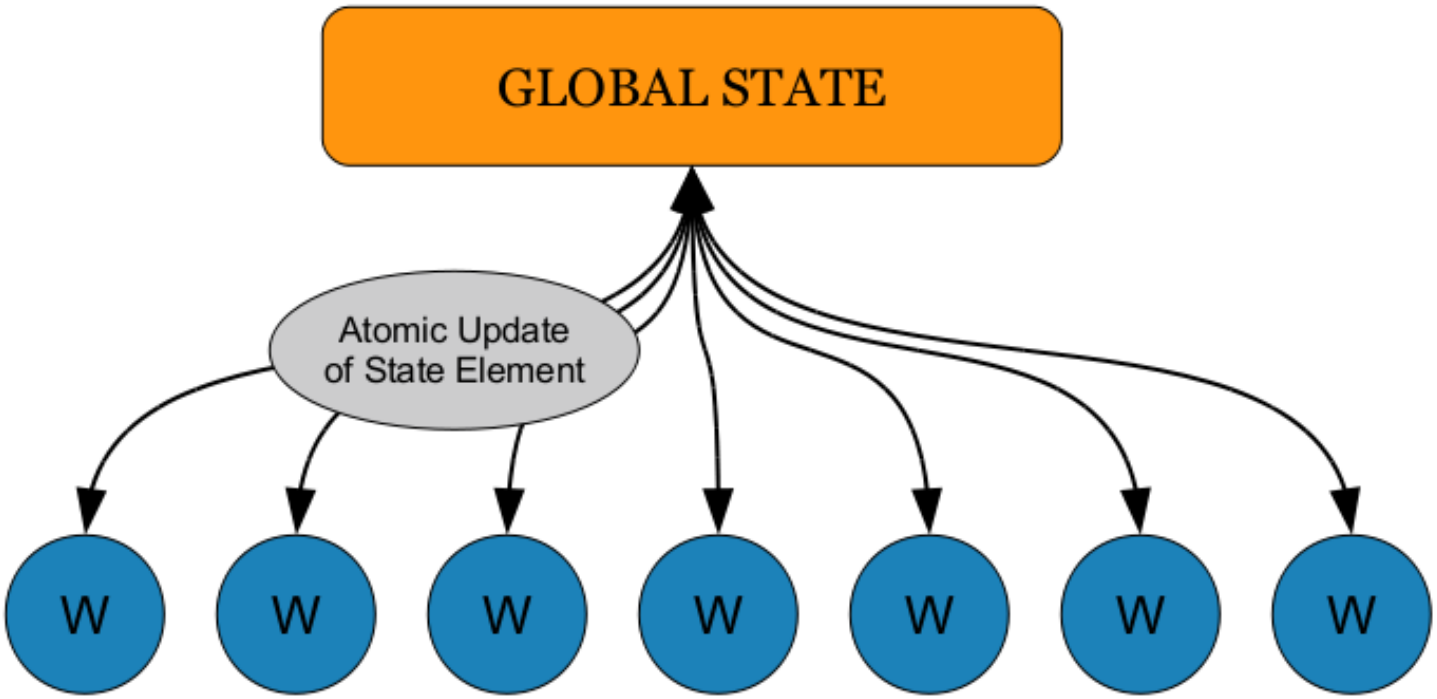

Parallelization using Global Arrays

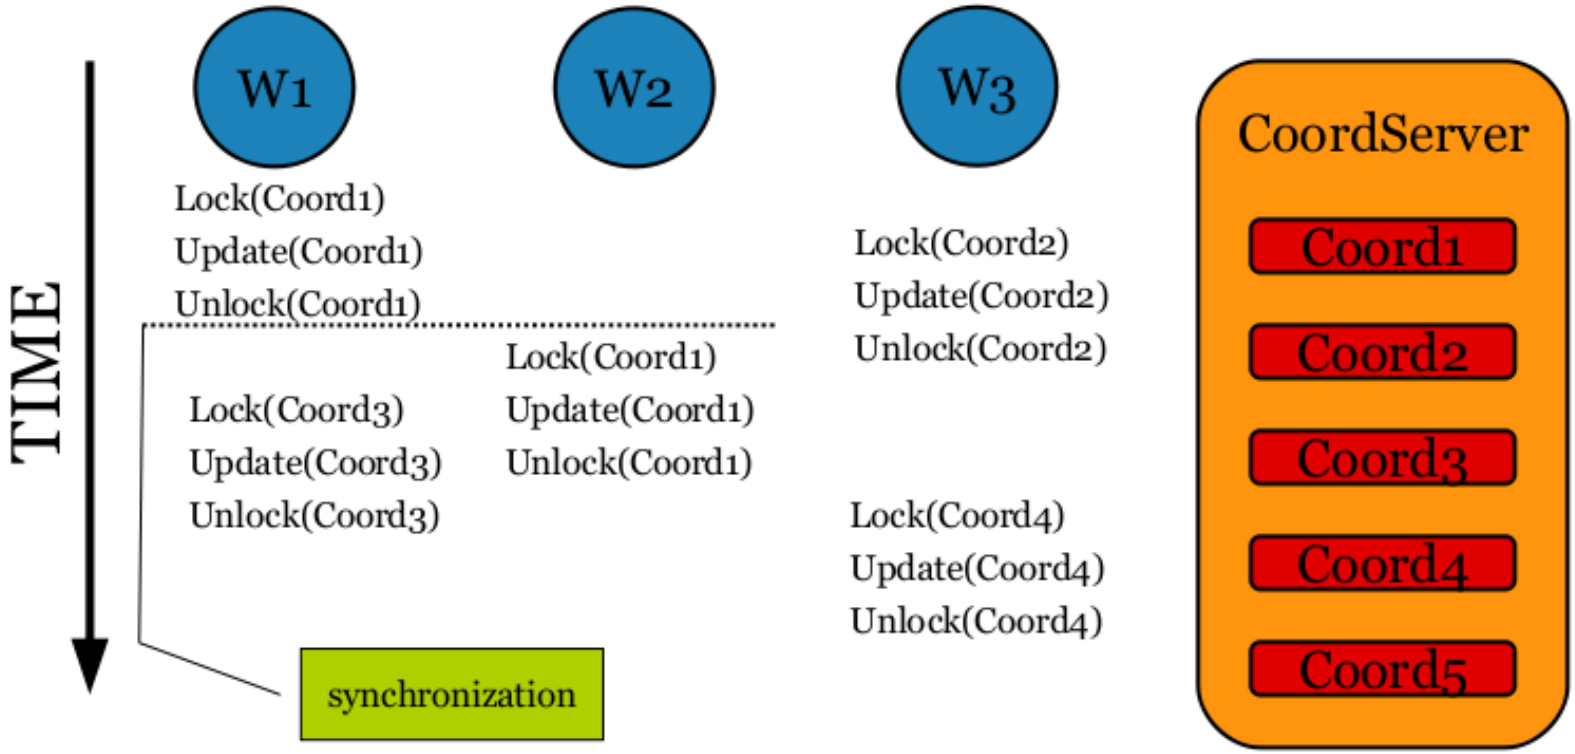

Algorithmic and Implementation Efficiency

The distributed nature of the NEUS global state allows for efficient strong-scaling due to the majority of asynchronous updates being nonconflicting.

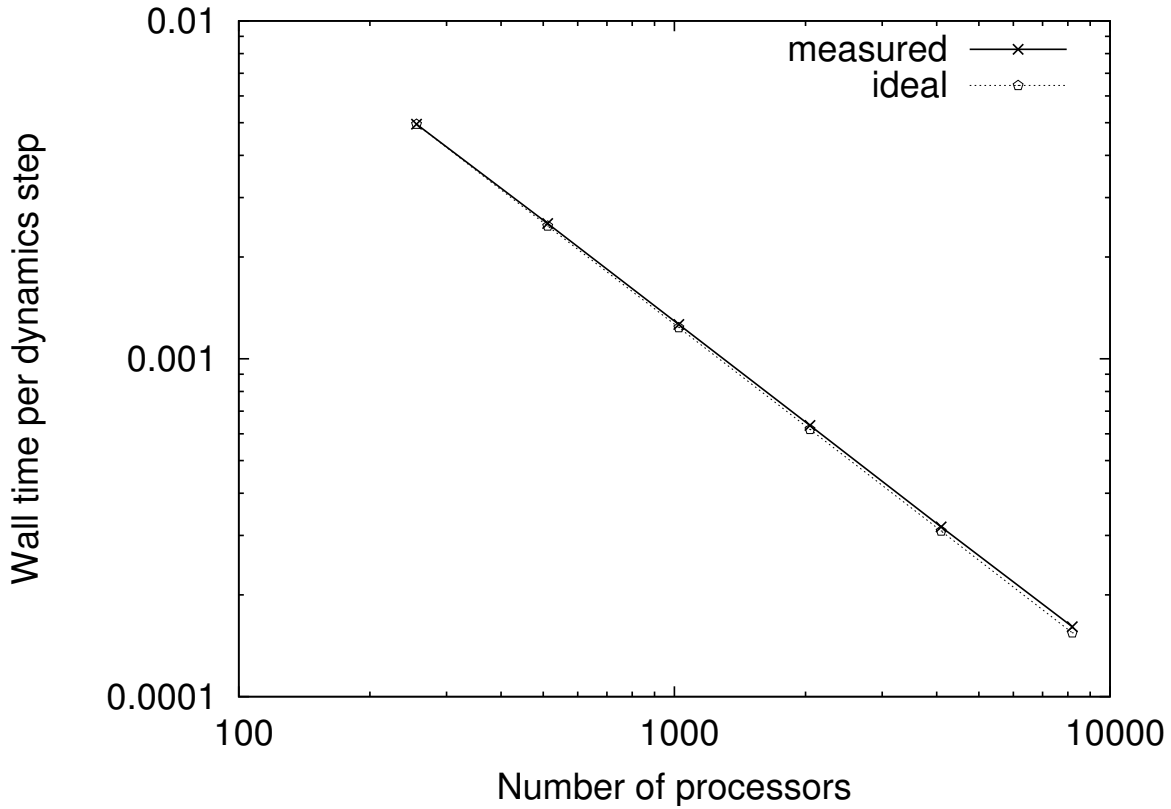

Figure: Strong scaling to 8192 cores with 96% efficiency.

Acknowledgments

UChicago: Alex Dickson, Aaron Dinner