

Alex Dickson¹, Jeff R. Hammond² and Aaron R. Dinner¹

¹ The University of Chicago (adickson@uchicago.edu, dinner@uchicago.edu)

² Argonne National Laboratory (jhammond@alcf.anl.gov)

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.

$$\tau_{MD} \approx 10^{-15} \text{ s}$$

$$\tau_{event} \approx 10^{-6} \text{ s}$$

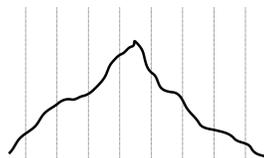
$$\tau_{comp} \approx 10^{-3} \text{ s}$$

$$\tau_{sim} > 11 \text{ days}$$

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

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.

How it works

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.

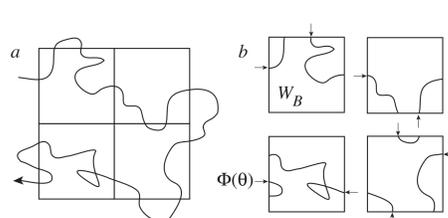


Figure: Discretizing phase space.

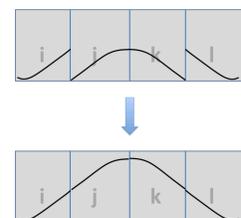


Figure: Combining probabilities.

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.

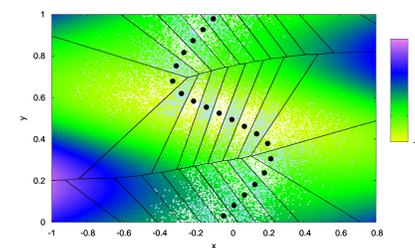
To conduct the "regional" simulations in nonequilibrium systems, one needs to be able to remove the bias in a physical way. The key insight is to run unbiased dynamics and restart walkers using a flux input distribution that is developed on-the-fly.

Nonequilibrium Umbrella Sampling in Detail

NEUS is an umbrella sampling algorithm that is applicable to nonequilibrium systems. It divides the phase space of a system into different regions and conducts restricted simulations within each region.

Strings:

Irregular regions can be formed by Voronoi polyhedra defined using a set of phase space points ("images") that together form a string that winds its way from reactants to products. The string is updated during sampling by periodically moving the images towards the average of the walker position in the last time interval.

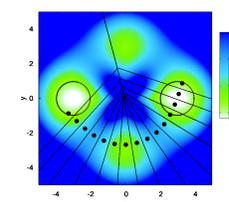
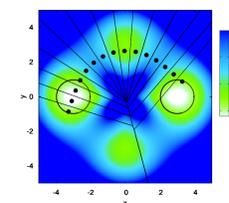


Separating the transition path ensemble:

We can also separate trajectories based on their basin of origin by defining two ensembles: S_A for trajectories originating in basin A , and S_B for trajectories originating in basin B . This allows for a better description of transition paths that are separate due to nonequilibrium path splitting, and as a bonus, it allows for the **easy calculation of transition rates!**

$$k_{AB} = \frac{\bar{\Phi}_{B|S_A}}{\bar{h}_A}$$

where $\bar{\Phi}_{B|S_A}$ is the flux into basin B from the S_A ensemble, and \bar{h}_A is the total weight of all regions in S_A .



Algorithmic and Implementation Efficiency

NEUS outperforms FFS when long-lived intermediate states are present. Both methods outperform conventional sampling when the transitions are sufficiently rare.

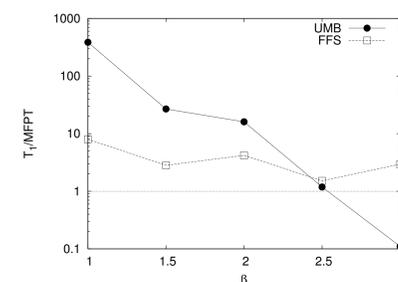


Figure: Efficiency results from the circular model potential, where β is the inverse temperature. Transitions become more rare as β increases.

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

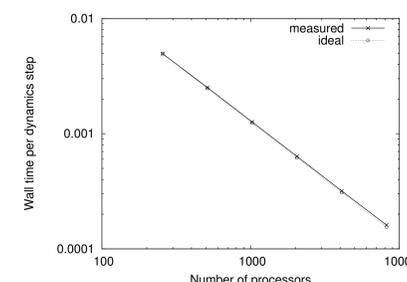
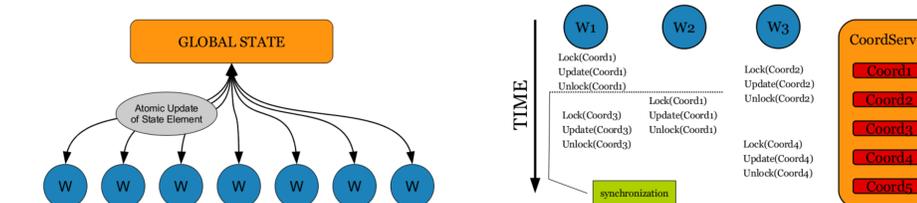


Figure: Strong scaling to 8192 cores with 96% efficiency.

Parallelization using Global Arrays

NEUS is implemented using a global state updated by atomic (complex) operations asynchronously. Global Arrays (GA) Put/Get operations in conjunction with Lock/Unlock provide a means to implement the necessary tools in a straightforward way. The implementation of NEUS on top of GA is called CoordServer.



Because Global Arrays provides a *distributed* global state, atomic updates can occur simultaneously except when locked. A master-worker approach would force unnecessary serialization of updates.

Results for RNA

Evidence for long-lived stable-states: The competition between the flow and the native contacts results in a rich dynamics that includes competing metastable states. We are interested in the intermediate case, where contracted and extended states coexist, and transitions between them occur on long, macroscopic timescales.

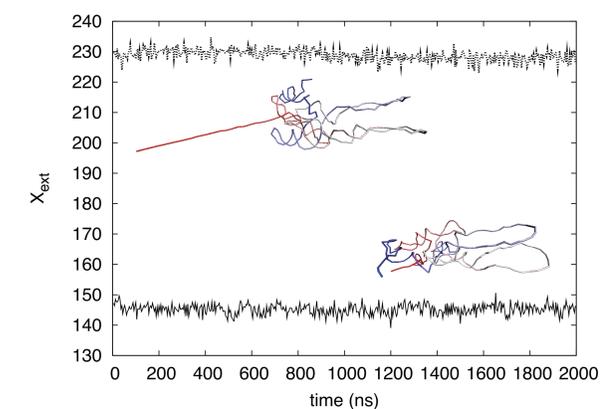


Figure: Extension vs. time for two trajectories with different initial conditions.

Problem Description: Each of the 262 nucleotides is represented by a single bead. Adjacent beads in the chain are connected with a FENE potential, and secondary and tertiary interactions are modelled with a Lennard-Jones potential. Repulsive terms keep the chain local straight to mimic steric repulsion. The flow is simulated with the stochastic rotation dynamics method.

Flow is simulated with the stochastic rotation dynamics wherein solvent is represented by a large number of infinitesimal particles that are grouped into cells of a lattice. Each step of the algorithm is comprised of a free streaming step and a "collision" step in which the velocity of each particle in the cell is rotated around the cell's average velocity vector by a random rotation matrix. The RNA beads are included in the collisions, through which the solvent will influence the polymer.