

Fast Equilibration of Water between Buried Sites and Bulk by MD with Parallel Monte Carlo Water Moves on GPUs

Ido Y. Ben-Shalom¹, Charles Lin², Brian K. Radak², Woody Sherman², and Michael K. Gilson^{1*}

¹Skaggs School of Pharmacy and Pharmaceutical Sciences, University of California, San Diego, 92093 La Jolla, California, USA

²Roivant Discovery, Boston, Massachusetts, 02110, USA

*To whom correspondence should be addressed, mgilson@health.ucsd.edu

1 ABSTRACT

Molecular dynamics (MD) simulations of proteins are commonly used to sample from the Boltzmann distribution of conformational states, with wide-ranging applications spanning chemistry, biophysics, and drug discovery. However, MD can be inefficient at equilibrating water occupancy for buried cavities in proteins that are inaccessible to the surrounding solvent. Indeed, the time needed for water molecules to equilibrate between bulk solvent and the binding site can be well beyond what is practical with standard MD, which typically ranges to hundreds of nanoseconds to low microseconds. We recently introduced a hybrid Monte Carlo/MD (MC/MD) method, which speeds up the equilibration of water between buried cavities and the surrounding solvent, while sampling from the thermodynamically correct distribution of states. While the initial implementation of the MC functionality led to considerable slowing of the overall simulations, here we address this problem with a parallel MC algorithm implemented on graphical processing units (GPUs). This results in speedups of 10-fold to 1000-fold over the original MC/MD algorithm, depending on the system and simulation parameters. The present method is available for use in the AMBER simulation software.

2 INTRODUCTION

Buried protein cavities may accommodate water molecules, and it is generally more favorable to fill a polar cavity with water than to leave it unoccupied.¹ Water exchange between bulk and cavities can occur on a very long time scale, depending on the nature of the free energy surface, and exchange can be particularly slow if these cavities are deeply buried.² Furthermore, large conformational changes may be required to allow the water to exchange with the bulk. Standard molecular dynamics (MD) simulations, which integrate Newton's equations of motion, are computationally demanding and can require a very long time to account for large conformational changes, adding to the challenge of efficiently sampling water exchange in buried cavities.³ Thus, a conformational ensemble from MD that does not efficiently sample the exchange of water between bulk and buried cavities may not be with concordance with the proper Boltzmann distribution associated with the conditions (e.g., solvent composition and temperature) of the simulation. Moreover, a simulation might appear to be converged, while in reality it is stuck in a local minimum and unable to sample important states for the buried water molecules.

As a consequence, properly equilibrating water molecules is essential for attaining accurate results in MD simulations, such as calculating protein–ligand binding free energies with MD-based methods.^{3,4} For example, alchemical methods that decouple the entire ligand from the binding site to compute absolute binding free energies,^{5–8} or which chemically modify the ligand to compute relative binding free energies,^{9–11} may not converge without accounting for resulting changes in the water occupancy of the binding site, especially for deeply buried cavities.¹² This is particularly problematic in drug discovery applications because the system may appear to be converged using standard MD analysis techniques, leading to false positive or false negative predictions.

Pioneering approaches to speed the equilibration of water molecules between bulk solvent and the protein interior have interleaved grand canonical Monte Carlo (GCMC) water moves⁸ with standard MD steps,^{3,9–12} and recent work¹³ has made this technology available in the OpenMM simulation toolkit.¹⁴ A GCMC move starts with an attempt (“try”) to insert/remove a water molecule to/from a random point in the simulated system. Each try is stochastically accepted or rejected with a probability based on its energy and the target value of the chemical potential of water in the system being simulated. Because the random locations at which water moves are tried and potentially accepted can be in the protein interior or in bulk solvent, such moves are able to equilibrate water between the protein interior and exteriors, as desired. However, the GCMC approach still poses some challenges. One results from the fact that the acceptance probability of a given try depends on the desired chemical potential of water, and additional work is typically needed to establish this value for the temperature, pressure, and composition, of the system being simulated. In addition, the GCMC/MD approach generates trajectories with a time-varying number of water molecules, a feature that can complicate bookkeeping. Finally, the probability of accepting a GC water insertion or removal move is typically low in condensed phase systems, so the method can be inefficient.

We recently introduced an alternative approach to equilibrating water in simulations, where translational MC moves allow water to jump between the protein interior and the bulk, while the number of waters in the simulation remains constant,¹⁵ and furthermore showed that this approach can improve both the precision and the accuracy of calculations of protein-ligand relative binding free energies.¹⁶ This MC/MD approach avoids the need of a GCMC/MD method to determine what chemical potential to use for each set of simulation conditions, and it also avoids any bookkeeping and software issues that might result from the time-varying number of water molecules. Although we took several algorithmic measures to speed the MC/MD method,¹⁵ integration of the MC steps with MD nonetheless markedly slowed the

simulations in our initial implementation. We have also collaboratively explored the use of nonequilibrium candidate MC (NMC) moves to increase the acceptance rate of the large translational MC water moves in this approach,¹⁷ but efficiency remains a concern.

Here, we present a different approach to improving the efficiency of the translational MC method. We describe the design, implementation, and validation of a new version of the standard MC/MD algorithm within the AMBER simulation code that utilizes a parallel MC algorithm implemented on graphical processing units (GPUs), resulting in speedups of roughly 10- to 1000-fold over the original MC/MD algorithm. The present method combines an MC method that allows multiple moves to be tried in parallel while still sampling configurations from the canonical distribution, with GPU acceleration of key bottleneck steps in the algorithm.

3 METHODS

3.1 DESCRIPTION OF THE ALGORITHM

3.1.1 WORKFLOW

The present algorithm iteratively alternates a block of N_{MD} standard MD steps and a multi-try MC step in which N_{MC} trial moves are processed in parallel, as detailed below. This is similar to our prior algorithm,¹⁵ where blocks of N_{MD} steps alternated with blocks of N_{MC} single-try MC steps. Each block of MD starts from the configuration present at the end of the prior MC step and *vice versa*. The velocities of all the atoms in the system at the start of each MD block are the same as those at the end of the prior MD block; that is, a moved water molecule carries its velocities. In addition, the orientation of a moved water is not modified, as the orientation of waters in the bulk is randomized by the MD blocks. Trajectory snapshots (atomic coordinates and velocities and system properties such as energy) are saved during the MD block, but not during MC.

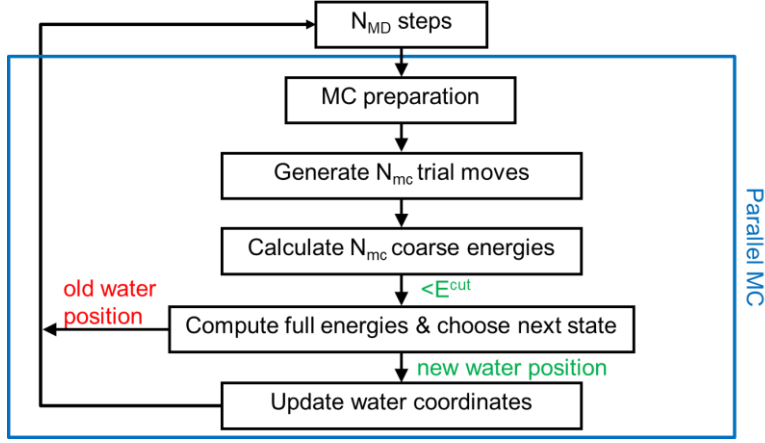


Figure 1: Flowchart of the parallel Monte Carlo/Molecular Dynamics procedure described in this work. Steps within the blue frame constitute the parallel MC part (see text for details). MC preparation comprises generation of the steric grid and the grid with the neighbor list for coarse energy calculations. Full energies are computed with the particle-mesh Ewald treatment of long-ranged electrostatics exactly as done during an MD step. Coarse and full energies are computed on the GPU, as are the steric grid and the coarse energy neighbor list.

3.1.2 PARALLELISM THROUGH MULTI-TRY MONTE CARLO STEPS

To facilitate an efficient GPU implementation, we moved from traditional serial MC to a parallel MC algorithm (previously described in the context of waste-recycling MC) in which multiple moves are tried in each step.^{18,19} The energies of all the move attempts are evaluated in parallel on the GPU, leading to dramatic wall-clock speed-up.

Each multi-try MC step starts with the system in an initial state indexed as $i = 0$ and with energy E_0 . From this state, additional trial states $i = 1 \dots N_{MC}$ are generated in which a water molecule, chosen with uniform probability from those within a region defined in the frame of reference of a residue, termed the MC region (Section 2.3), is moved to a sterically feasible location chosen with uniform probability within the same MC region. Thus, water moves can only occur within a single region, preserving detailed balance. Sterically feasible locations are identified via a steric grid within the MC region, as previously detailed.¹⁵ For each trial move, a fast but coarse approximation to the energy change, ΔE_i^{coarse} is then computed. If this energy difference is lower than a user-defined cutoff, i.e., $\Delta E_i^{coarse} < \Delta E^{cut}$ then a full PME energy calculation is done to provide energy E_i . Otherwise, i.e., if $\Delta E_i^{coarse} > \Delta E^{cut}$, the energy of this trial state is treated as if it were infinite.

From the energies of the initial ($i = 0$) and trial ($i = 1, \dots, N_{MC}$) states, we obtain the following Boltzmann factors

$$q_i = e^{-E_i/RT}$$

where R is the gas constant and T is absolute temperature. Finally, one of the $N_{MC} + 1$ states (initial or trial) then is selected according to the following probabilities:

$$p_i = q_i / q_{sum}$$

where

$$q_{sum} = \sum_{i=0}^{N_{MC}} q_i$$

Note that the rationale for trying only sterically feasible locations and for spending the computer time to compute the full energy only for tries with a low approximate energy is that any very high energy state will almost certainly not be accepted as the new state of the system and so can be discarded without a costly full energy calculation.

3.1.3 DEFINITION OF THE MC REGION

The definition of the MC region is modified from the prior version of our code to a method that is compatible with all types of simulation boxes (e.g., rectangular prism, truncated octahedron). Here, the MC region in which waters are moved with MC steps is a rectangular prism with user-defined dimensions and centered at the center of coordinates of a user-defined set of atoms selected with the AMBER masking function. This selection is defined with the input keyword **mcwatmask**, and the keyword **mcgridshift** is used to define the length of the grid along each axis, in Å. For example, a value of 10 means that the grid will extend ± 10 Å from the center along each axis. Thus, one may enhance water exchange in a buried binding pocket by centering the MC region on a residue bordering the pocket and making the region large enough that it extends into the water outside the protein. If the grid is not defined, the grid dimensions will be the entire simulation box.

3.1.4 GPU IMPLEMENTATION

We transferred multiple components of the MC process from the CPU to the GPU. Most importantly, the values of ΔE_i^{coarse} for all trial moves in the MC step are evaluated in parallel on the GPU. The full AMBER energies are evaluated on the GPU with the PME treatment of long-ranged electrostatics for each trial configuration i with $\Delta E_i^{coarse} > \Delta E^{cut}$. As described in Section 3.2, this approach takes advantage of the multi-try MC method. In addition, two setup steps that are time consuming on the CPU are now executed on the GPU. Thus, for both the bit grid based steric grid as well as the subsequent grid-based neighbor list E^{coarse} calculations, hundreds of trial moves can be done in parallel by utilizing the GPU. As previously

detailed,¹⁵ the coarse calculation grid covers the entire simulation box and has a coarse spacing set to half the distance cutoff used for nonbonded interactions. Each coarse grid point is assigned a list of the atoms in its corresponding voxel, and E^{coarse} is computed by considering the interactions of a water molecule with the atoms in the same and the neighboring two coarse voxels in each direction.

3.2 EVALUATION AND OPTIMIZATION OF THE ALGORITHM

3.2.1 MAINTAINING THE BOLTZMANN DISTRIBUTION

We wished to confirm that the algorithm maintains detailed balance and thus yields a correct Boltzmann distribution. To do this, we tested the method on a hot water gas for which the acceptance rate for new trial water positions is high, so that any error introduced by the MC procedure would be apparent from a difference in the energy distribution relative to a plain MD simulation of the same system with a suitable thermostat. The test system is composed of 97 gas phase water molecules at 500 K in a constant volume box with a side length of about 56 Å along each axis. To maximize the effect of the MC moves, we used an MC region that covered the entire simulation box.

The MD component of the MC/MD calculations and the pure MD reference simulations used the SHAKE algorithm^{20,21} with the default tolerance of 0.00001 Å. Because these particular calculations were run at 500K, the atomic speeds are large, so the integration time step was set to 0.2 fs to minimize deviations from the Boltzmann ensemble. This maximizes compatibility with the MC steps, for which there are no finite time-step errors. (For the other simulations in this work, the MD time step was 2 fs.) The reference MD simulation was run for 20 ns, and the MD energy was recorded every 10 steps, which resulted in 10^7 energy calculations. The MC/MD calculations alternated runs of 10 MD steps with an MC step with $N_{MC} = 10$. As our MC/MD method is optimized to work for protein-ligand complexes and small cavities, using our method on a gas-phase system that is composed of mostly cavities is highly inefficient and is therefore considerably slower compared to MD. However, we think it is the most accurate way to examine the thermodynamic character of the system. We ran eight MC/MD calculations, each with 1.25×10^6 cycles, for a total of 10^7 energy calculations. All MC/MD and MD runs started with the same initial configuration of the equilibrated gas but used different random number seeds.

3.2.2 EFFECTS OF PARAMETERS ON EFFICIENCY FOR A PROTEIN-LIGAND COMPLEX

Decreasing the value of E^{cut} speeds the MC step by reducing the number of costly full energy calculations, but it also risks perturbing the results due to excessive reliance on the less accurate coarse energy calculation. We explored this tradeoff, as well as the consequences of varying the number of MC tries per step, N_{MC} for a protein-ligand test system.

We used the Major Urinary Protein (MUP) structure with the PDB ID 1ZNK.²² The system was constructed using the TIP3P water model,²³ the ff14SB^{24,25} force field for the protein, and GAFF²⁶ for the ligand. The protein-ligand complex was solvated in rectangular boxes using tleap with initial buffer sizes of 11 Å. The net charge of the system was neutralized by addition of Na^+ . The simulations were performed on GPUs with AMBER20's GPU-accelerated PMEMD simulation code.²⁷⁻³⁰ During equilibration, Cartesian restraints to the starting structure were applied to all ligand and protein heavy atoms, with a force constant of $5 \text{ kcal mol}^{-1} \text{ Å}^{-2}$. After a brief minimization (50 steps of steepest descent plus 450 steps of conjugate gradient), the system was sequentially heated at a fixed volume with linear ramps joining the temperatures 5, 100, 200, and 298.15 K, at 20 ps per ramp, with each ramp followed by an additional 20 ps with pressure coupling. All simulations used a Langevin integrator with a 2 fs timestep with a friction coefficient of 2 ps^{-1} . Heating steps were run in NVT, while equilibration and production steps were run in NPT with pressure regulated at 1 atm with a Monte Carlo barostat.³¹ The SHAKE algorithm^{20,21} was used to constrain hydrogen bond lengths.³² The standard AMBER protocol for non-bonded interactions was followed. Thus, the particle mesh Ewald (PME) method³³ was used for periodic boundary conditions with an 8 Å cutoff for the short-ranged PME contribution as well as LJ interactions. A long-range continuum correction was used for the dispersive term.³⁴ Cartesian restraints with a force constant of $5 \text{ kcal mol}^{-1} \text{ Å}^{-2}$ were applied to the $\text{C}\alpha$ atoms and ligand atoms during the production run.

For this MUP system, we checked the effects of varying threshold of energy, E^{cut} , to send an MC try from a fast, coarse energy calculation to a slower, full pmemd calculation.¹⁵ The higher the value, more MC tries are sent to full pmemd calculation rather than being immediately rejected, and therefore the runtime is slower. Note that using a higher value can lead to an increase in the acceptance rate, but not a lower acceptance rate. This is because configurations that would otherwise have been rejected get a second chance to be accepted following a full energy calculation, as detailed in Section 4.2. We also investigated the effect of N_{MC} on the results and timings for this system. Thus, the number of MD steps per cycle was held at 10^3 , while N_{MC} was set to the following values: 10^3 , 10^4 , 2.5×10^4 , 5×10^4 , and 10^5 .

3.3 HARDWARE USED

All simulations were performed on Nvidia GeForce GTX 1080 Ti GPUs, with an Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GH.

3.4 APPLICATION TO A BURIED CAVITY IN T4 LYSOZYME

We used our MC/MD method to analyze the equilibrium water density in the moderately polar, buried, cavity³⁵ of the T4 Lysozyme L99A/M102Q double mutant. We used the structures of 2RBO¹³ and performed two MD simulations where the C α atoms were restrained. One simulation used only MD and the second used MC/MD; both included 200 ns of MD. In the MC/MD simulation, an MC step with $N_{MC} = 10^6$ was performed after every 10^3 MD steps. We then used the GIST software¹⁴ within cpptraj¹⁵ to compare the water densities in the buried cavities between the two simulations.

4 RESULTS

4.1 THERMODYNAMIC TEST SYSTEM

Our thermodynamic test system, gas-phase water at 500K, allows us to check that we maintain the correct Boltzmann distribution for a system that has a large number of accepted MC tries; here ~57% of the MC steps resulted in a new water configuration. The potential energy distribution provided by the MC/MD method agrees closely with that provided by a pure MD simulation of the same system (Figure 2A). In addition, the results are highly consistent across all eight MC/MD runs, as shown in Figure 2B, which shows all simulations starting from the same configuration energy, diverging as sampling proceeds, and converging to very similar mean energies. For pure MD, the mean potential energy over 10^7 energy evaluations was -126.8 kcal mol⁻¹, while the mean potential energy for MC/MD for 10^7 energy evaluations divided into 8 separate 1.25×10^6 runs was -127.0 kcal mol⁻¹ (Figure 2B). The good agreement between pure MD and MC/MD supports the validity of the parallel MC method and its correct maintenance of detailed balance. Hence, it can be used to compute experimental observables, such as binding free energies.




Figure 2: A) Probability distribution of potential energies for MD (red, 10^7 energy calculations) and MC/MD (blue, 10^7 energy calculations). The energies were divided to bins of 0.1 kcal mol⁻¹. B) Cumulative average of the potential energy. Red MD simulation. Gray eight replicates of the MC/MD simulations.

4.2 PARAMETER OPTIMIZATION AND EFFICIENCY

We examined the effects of E^{cut} and N_{MC} for the MUP protein-ligand test system. Setting E^{cut} to 15, 12, and 10 kcal mol⁻¹ resulted in 0.002, 0.0006 and 0.0002, of MC tries, respectively, being sent for full energy calculations. However, the fraction of tries accepted does not change significantly between these E^{cut} thresholds (Figure 3A). This indicates that using the lower threshold tried here, 10 kcal mol⁻¹ is sufficient and that there is no need for a higher, and thus more rigorous, energy threshold, which results in a slowdown of the program, as shown in Figure 3B. In addition, as expected, increasing the number of MC tries (N_{MC}) increases the number of acceptances, though with somewhat diminishing acceptance rates for higher values of this parameter (Figure 3A). Presumably, there are some MC cycles where the water molecule chosen is particularly stable and/or there are few trial sites available where the water could make favorable interactions with its surrounding atoms.




Figure 3: Results for the MUP test system. A) Mean number of accepted tries over three replicates with $N_{MD} = 10^3$ and several values of N_{MC} , given in the legend, over a total of 10^6 cycles. B) Mean runtime for the same three MC/MD replicates with the present code and for our prior implementation with serial MC on the CPU (AMBER20). The value for MD without any MC is indicated on the y axis.

The present parallel, GPU-based MC implementation runs dramatically faster than our prior serial, CPU-based MC code, which is reported as AMBER 20 in Figure 3B. Thus, the new algorithm is ~ 10 times faster for $N_{MC} = N_{MD} = 10^3$ and ~ 270 times faster for $N_{MC} = 10^5$ with $N_{MD} = 10^3$. Furthermore, in the AMBER20 version, increasing N_{MC} past 10^3 results in an approximately proportional increase in computer time, indicating that the time comes to be dominated by the time required for the MC steps. For the current version, in contrast, the cost of the MC steps remains considerably less than that for the MD steps.

This speed enhancement has important implications for the integration of the present technology into absolute and relative binding free energy calculations. Thus, for one of our prior relative binding free energy calculations (the $W1 \rightarrow W2$ perturbation, decharge windows), using MC/MD to equilibrate a single λ window took ~ 9 hours,¹⁶ while in the current version, this only requires ~ 30 seconds, a 1000x speedup.

4.3 ILLUSTRATIVE APPLICATION TO A CAVITY IN T4 LYSOZYME

The present MC/MD method may be applied to study the hydration of buried cavities. Here, we demonstrate this use case for the engineered L99A/M102Q mutant of T4 lysozyme cavity.³⁵ Simulations with MC/MD and pure MD were initiated with no water molecules in cavity (Figure 4A). The MC/MD simulations provided a smooth, physically plausible water

distribution within the cavity (Figure 4B), while water did not penetrate to the cavity in a pure MD simulation (Figure 4C).




Figure 4: Buried cavity in the T4 Lysozyme double mutant L99A/M102Q, also showing residue GLN 102, which was used to define the MC region. A) empty cavity before simulations. B) cavity with water oxygen atom isodensity contours (blue: 0.1; red: 2.0, as multiples of bulk density), based on MC/MD simulations and computed with GIST in cpptraj. C) empty cavity observed from a pure MD simulation of roughly equal time to the MC/MD simulation shown in panel B.

5 DISCUSSION

The present MC/MD method achieves speedups of 10-1000x over our prior version by using a parallel MC algorithm that is effectively ported to the GPU. A further update allows it to support all simulation box types by defining the steric grid not in terms of the boundary but in terms of the distances from a molecular component of interest, such as a ligand or a residue at the edge of a binding pocket. The ability to equilibrate buried water molecules with minimal slowdown means this can be integrated with alchemical free energy calculations and performed as part of any system setup to establish preferred locations of water molecules. This is important as it solves the need for *a priori* knowledge regarding the hydration state of the cavity and allows for water sampling in the context of a flexible protein. It may also find application for structural optimization of water molecules in crystallographic or cryo-electron microscopy structures, and to accelerate sampling of structural processes that involve the penetration of water into the protein interior, potentially including protein conformational changes or full denaturation. It may be thought that the acceptance rate would be higher for GCMC than for the present MC method, because a trial move for GCMC does not require paying the penalty of removing water from its initial location. In fact, this should not be the case, as the acceptance probability for a move to a given location with GCMC should be the same as the mean

acceptance probability to the same location from a simulation of bulk water having the same chemical potential as that used in the GCMC calculation.

A typical MC simulation of water makes small incremental water moves, with correspondingly small energy changes, and thus can have a large acceptance probability for these moves. In contrast, here the whole point is to make large translational MC moves that allow jumps between the protein cavity and the surrounding water. This leads to a rather low acceptance rate and motivates efforts to maximize efficiency. Following the marked advance in speed described here, possible future directions may include omitting moves that would be only within the surrounding water or within the cavity, as these are not very interesting and are accounted for by the MD steps, while being careful to avoid any possible violations of detailed balance. Another direction is to allow small clusters of water to move together in order to increase the rate of acceptance for moves into large nonpolar cavities. However, the present method is already highly efficient, and is currently available for use in the AMBER simulation software.

6 ACKNOWLEDGEMENTS

MKG acknowledges funding from National Institute of General Medical Sciences (GM061300 and GM100946). These findings are solely of the authors and do not necessarily represent the views of the NIH. MKG has an equity interest in and is a cofounder and scientific advisor of VeraChem LLC. CL, BKR, and WS are employees of Roivant Discovery.

7 REFERENCES

- (1) Bissantz, C.; Kuhn, B.; Stahl, M. A Medicinal Chemist's Guide to Molecular Interactions. *J. Med. Chem.* **2010**, *53* (14), 5061–5084. <https://doi.org/10.1021/jm100112j>.
- (2) Poornima, C. S.; Dean, P. M. Hydration in Drug Design. 2. Influence of Local Site Surface Shape on Water Binding. *J. Comput. Aided Mol. Des.* **1995**, *9* (6), 513–520.
- (3) Deng, Y.; Roux, B. Computation of Binding Free Energy with Molecular Dynamics and Grand Canonical Monte Carlo Simulations. *J. Chem. Phys.* **2008**, *128* (11), 115103. <https://doi.org/10.1063/1.2842080>.
- (4) Ross, G. A.; Bruce Macdonald, H. E.; Cave-Ayland, C.; Cabedo Martinez, A. I.; Essex, J. W. Replica-Exchange and Standard State Binding Free Energies with Grand Canonical Monte Carlo. *J. Chem. Theory Comput.* **2017**, *13* (12), 6373–6381. <https://doi.org/10.1021/acs.jctc.7b00738>.
- (5) Tembe, B. L.; Mc Cammon, J. A. Ligand-Receptor Interactions. *Comput. Chem.* **1984**, *8* (4), 281–283. [https://doi.org/10.1016/0097-8485\(84\)85020-2](https://doi.org/10.1016/0097-8485(84)85020-2).
- (6) Gilson, M. K.; Given, J. A.; Bush, B. L.; McCammon, J. A. The Statistical-Thermodynamic Basis for Computation of Binding Affinities: A Critical Review. *Biophys. J.* **1997**, *72* (3), 1047–1069. [https://doi.org/10.1016/S0006-3495\(97\)78756-3](https://doi.org/10.1016/S0006-3495(97)78756-3).
- (7) Mobley, D. L.; Gilson, M. K. Predicting Binding Free Energies: Frontiers and Benchmarks. *Annu. Rev. Biophys.* **2017**, *46*, 531–558. <https://doi.org/10.1146/annurev-biophys-070816-033654>.
- (8) Mezei, M. Grand-Canonical Ensemble Monte Carlo Study of Dense Liquid: Lennard-Jones, Soft Spheres and Water. *Molecular Physics* **1987**, *61* (3), 565–582. <https://doi.org/10.1080/00268978700101321>.
- (9) Woo, H.-J.; Dinner, A. R.; Roux, B. Grand Canonical Monte Carlo Simulations of Water in Protein Environments. *J. Chem. Phys.* **2004**, *121* (13), 6392–6400. <https://doi.org/10.1063/1.1784436>.
- (10) Ross, G. A.; Russell, E.; Deng, Y.; Lu, C.; Harder, E. D.; Abel, R.; Wang, L. Enhancing Water Sampling in Free Energy Calculations with Grand Canonical Monte Carlo. *J. Chem. Theory Comput.* **2020**, *16* (10), 6061–6076. <https://doi.org/10.1021/acs.jctc.0c00660>.
- (11) Bruce Macdonald, H. E.; Cave-Ayland, C.; Ross, G. A.; Essex, J. W. Ligand Binding Free Energies with Adaptive Water Networks: Two-Dimensional Grand Canonical Alchemical Perturbations. *J. Chem. Theory Comput.* **2018**, *14* (12), 6586–6597. <https://doi.org/10.1021/acs.jctc.8b00614>.
- (12) Ross, G. A.; Bodnarchuk, M. S.; Essex, J. W. Water Sites, Networks, And Free Energies with Grand Canonical Monte Carlo. *J. Am. Chem. Soc.* **2015**, *137* (47), 14930–14943. <https://doi.org/10.1021/jacs.5b07940>.
- (13) Samways, M. L.; Bruce Macdonald, H. E.; Essex, J. W. Grand: A Python Module for Grand Canonical Water Sampling in OpenMM. *J. Chem. Inf. Model.* **2020**, *60* (10), 4436–4441. <https://doi.org/10.1021/acs.jcim.0c00648>.
- (14) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. OpenMM 7: Rapid Development of High Performance Algorithms for Molecular Dynamics. *PLoS Comput Biol* **2017**, *13* (7), e1005659. <https://doi.org/10.1371/journal.pcbi.1005659>.
- (15) Ben-Shalom, I. Y.; Lin, C.; Kurtzman, T.; Walker, R. C.; Gilson, M. K. Simulating Water Exchange to Buried Binding Sites. *J. Chem. Theory Comput.* **2019**, *15* (4), 2684–2691. <https://doi.org/10.1021/acs.jctc.8b01284>.
- (16) Ben-Shalom, I. Y.; Lin, Z.; Radak, B. K.; Lin, C.; Sherman, W.; Gilson, M. K. Accounting for the Central Role of Interfacial Water in Protein–Ligand Binding Free Energy

- Calculations. *J. Chem. Theory Comput.* **2020**, *16* (12), 7883–7894. <https://doi.org/10.1021/acs.jctc.0c00785>.
- (17) Bergazin, T. D.; Ben-Shalom, I. Y.; Lim, N. M.; Gill, S. C.; Gilson, M. K.; Mobley, D. L. Enhancing Water Sampling of Buried Binding Sites Using Nonequilibrium Candidate Monte Carlo. *J. Comput. Aided Mol. Des.* **2021**, *35* (2), 167–177. <https://doi.org/10.1007/s10822-020-00344-8>.
 - (18) Frenkel, D. Speed-up of Monte Carlo Simulations by Sampling of Rejected States. *Proceedings of the National Academy of Sciences* **2004**, *101* (51), 17571–17575. <https://doi.org/10.1073/pnas.0407950101>.
 - (19) Kim, J.; Rodgers, J. M.; Athènes, M.; Smit, B. Molecular Monte Carlo Simulations Using Graphics Processing Units: To Waste Recycle or Not? *J. Chem. Theory Comput.* **2011**, *7* (10), 3208–3222. <https://doi.org/10.1021/ct200474j>.
 - (20) Ryckaert, J.-P.; Ciccotti, G.; Berendsen, H. J. C. Numerical Integration of the Cartesian Equations of Motion of a System with Constraints: Molecular Dynamics of n-Alkanes. *J. Comput. Phys.* **1977**, *23* (3), 327–341. [https://doi.org/10.1016/0021-9991\(77\)90098-5](https://doi.org/10.1016/0021-9991(77)90098-5).
 - (21) Miyamoto, S.; Kollman, P. A. Settle: An Analytical Version of the SHAKE and RATTLE Algorithm for Rigid Water Models. *J. Comput. Chem.* **1992**, *13* (8), 952–962. <https://doi.org/10.1002/jcc.540130805>.
 - (22) Malham, R.; Johnstone, S.; Bingham, R. J.; Barratt, E.; Phillips, S. E. V.; Laughton, C. A.; Homans, S. W. Strong Solute–Solute Dispersive Interactions in a Protein–Ligand Complex. *J. Am. Chem. Soc.* **2005**, *127* (48), 17061–17067. <https://doi.org/10.1021/ja055454g>.
 - (23) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. *J. Chem. Phys.* **1983**, *79* (2), 926–935. <https://doi.org/10.1063/1.445869>.
 - (24) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. Comparison of Multiple Amber Force Fields and Development of Improved Protein Backbone Parameters. *Proteins* **2006**, *65* (3), 712–725. <https://doi.org/10.1002/prot.21123>.
 - (25) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. Ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from Ff99SB. *J. Chem. Theory Comput.* **2015**, *11* (8), 3696–3713. <https://doi.org/10.1021/acs.jctc.5b00255>.
 - (26) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and Testing of a General Amber Force Field. *J. Comput. Chem.* **2004**, *25* (9), 1157–1174. <https://doi.org/10.1002/jcc.20035>.
 - (27) Götz, A. W.; Williamson, M. J.; Xu, D.; Poole, D.; Le Grand, S.; Walker, R. C. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born. *J. Chem. Theory Comput.* **2012**, *8* (5), 1542–1555. <https://doi.org/10.1021/ct200909j>.
 - (28) Salomon-Ferrer, R.; Götz, A. W.; Poole, D.; Le Grand, S.; Walker, R. C. Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 2. Explicit Solvent Particle Mesh Ewald. *J. Chem. Theory Comput.* **2013**, *9* (9), 3878–3888. <https://doi.org/10.1021/ct400314y>.
 - (29) Le Grand, S.; Götz, A. W.; Walker, R. C. SPFP: Speed without Compromise—A Mixed Precision Model for GPU Accelerated Molecular Dynamics Simulations. *Comput. Phys. Commun.* **2013**, *184* (2), 374–380. <https://doi.org/10.1016/j.cpc.2012.09.022>.
 - (30) Mermelstein, D. J.; Lin, C.; Nelson, G.; Kretsch, R.; McCammon, J. A.; Walker, R. C. Fast and Flexible Gpu Accelerated Binding Free Energy Calculations within the Amber Molecular Dynamics Package. *J. Comput. Chem.* **2018**, *39* (19), 1354–1358. <https://doi.org/10.1002/jcc.25187>.

- (31) Åqvist, J.; Wennerström, P.; Nervall, M.; Bjelic, S.; Brandsdal, B. O. Molecular Dynamics Simulations of Water and Biomolecules with a Monte Carlo Constant Pressure Algorithm. *Chemical Physics Letters* **2004**, *384* (4–6), 288–294. <https://doi.org/10.1016/j.cplett.2003.12.039>.
- (32) Steinbrecher, T.; Mobley, D. L.; Case, D. A. Nonlinear Scaling Schemes for Lennard-Jones Interactions in Free Energy Calculations. *The Journal of Chemical Physics* **2007**, *127* (21), 214108. <https://doi.org/10.1063/1.2799191>.
- (33) Darden, T.; York, D.; Pedersen, L. Particle Mesh Ewald: An N·log(N) Method for Ewald Sums in Large Systems. *J. Chem. Phys.* **1993**, *98* (12), 10089–10092. <https://doi.org/10.1063/1.464397>.
- (34) Allen M. P. and Tildesley D. J. *Computer Simulation of Liquids*; Clarendon Press: Oxford, 1987.
- (35) Boyce, S. E.; Mobley, D. L.; Rocklin, G. J.; Graves, A. P.; Dill, K. A.; Shoichet, B. K. Predicting Ligand Binding Affinity with Alchemical Free Energy Methods in a Polar Model Binding Site. *J. Mol. Biol.* **2009**, *394* (4), 747–763. <https://doi.org/10.1016/j.jmb.2009.09.049>.

for Table of Contents use only

Fast Equilibration of Water between Buried Sites and Bulk by MD with Parallel Monte Carlo Water Moves on GPUs

Ido Y. Ben-Shalom, Charles Lin, Brian K. Radak, Woody Sherman, and Michael K. Gilson

