

Efficiency Comparison of Single- and Multiple-Macrostate Grand Canonical Ensemble Transition-Matrix Monte Carlo Simulations[†]

Harold W. Hatch,^{*,‡} Daniel W. Siderius,[‡] Jeffrey R. Errington,[¶] and Vincent K. Shen[§]

[‡]*Chemical Informatics Research Group, Chemical Sciences Division, National Institute of Standards and Technology, Gaithersburg, Maryland 20899-8380, United States*

[¶]*Department of Chemical and Biological Engineering, University at Buffalo, Buffalo, New York 14260, United States*

[§]*Chemical Informatics Research Group, Chemical Sciences Division, National Institute of Standards and Technology, Gaithersburg, Maryland 20899-8380, USA*

E-mail: harold.hatch@nist.gov

Abstract

Recent interest in parallelizing flat-histogram transition-matrix Monte Carlo simulations in the grand canonical ensemble, due to its demonstrated effectiveness in studying phase behavior, self-assembly and adsorption, has led to the most extreme case of single-macrostate simulations, where each macrostate is simulated independently

[†]Contribution of the National Institute of Standards and Technology, not subject to U.S. Copyright. Published as part of The Journal of Physical Chemistry virtual special issue "Pablo G. Debenedetti Festschrift"

with ghost particle insertions and deletions. Despite their use in several studies, no efficiency comparisons of these single-macrostate simulations have been made with multiple-macrostate simulations. We show that multiple-macrostate simulations are up to 3 orders of magnitude more efficient than single-macrostate simulations, which demonstrates the remarkable efficiency of flat-histogram biased insertions and deletions, even with low acceptance probabilities. Efficiency comparisons were made for supercritical fluids and vapor—liquid equilibrium of bulk Lennard-Jones and a three-site water model, self-assembling patchy trimer particles and adsorption of a Lennard-Jones fluid confined in a purely repulsive porous network, using the open source simulation toolkit FEASST. By directly comparing with a variety of Monte Carlo trial move sets, this efficiency loss in single-macrostate simulations is attributed to three related reasons. First, ghost particle insertions and deletions in single-macrostate simulations incur the same computational expense as grand canonical ensemble trials in multiple-macrostate simulations, yet ghost trials do not reap the sampling benefit from propagating the Markov chain to a new microstate. Second, single-macrostate simulations lack macrostate change trials that are biased by the self-consistently converging relative macrostate probability, which is a major component of flat histogram simulations. Third, limiting a Markov chain to a single-macrostate reduces sampling possibilities. Existing parallelization methods for multiple-macrostate flat-histogram simulations are shown to be more efficient than parallel single-macrostate simulations by approximately an order of magnitude or more in all systems investigated.

1 Introduction

Transition-matrix Monte Carlo (TMCMC) simulations in the grand canonical ensemble (GCE) are promising methods because they allow for the direct calculation of complex fluid phase behavior in a variety of applications. For example, in the study of liquid—vapor phase equilibrium, the pressure, density and free energy of the vapor and liquid phases may be calculated directly.¹ In addition to equilibrium states, metastable states² and the critical

point³ as well as other thermodynamic or structural properties of interest at a given temperature can also be obtained with these approaches. TMMC has also been successfully applied toward the study of adsorption on surfaces⁴⁻⁷ and inside porous materials.⁸⁻¹¹ TMMC may also be used in studies of self-assembly to calculate critical micelle concentrations, critical micelle temperatures and other structural transitions.^{12,13} Temperature extrapolation methods further improve the efficiency of TMMC by reducing the number of required simulations,¹⁴⁻¹⁹ and initialization with Wang—Landau can reduce convergence times.^{9,20}

TMMC can be parallelized in a variety of ways. In addition to all of the parallelization methods available in the canonical ensemble (CE), including domain decomposition,²¹ parallel configurational bias,²² and prefetching,²³ flat-histogram simulations in the GCE may also be parallelized by dividing the density range of interest into subranges assigned to different processors.^{13,24} However, the existing literature has little quantitative evidence of efficiency loss as more processors are used, or as the subranges become smaller.²⁵ While there is evidence that relatively small subranges may reduce efficiency,²⁵⁻²⁸ there are also legitimate reasons to utilize small density subranges. To our knowledge, there is no detailed study of the efficiency costs of this parallelization method in TMMC GCE simulations of vapor—liquid equilibrium, self-assembly and adsorption.

For TMMC in the GCE, the smallest possible density subinterval is a single macrostate, which means that the simulation has a fixed number of particles and the transitions to neighboring macrostates are computed using ghost particle insertions and deletions. This approach was utilized in a binary system in order to avoid a two dimensional biasing function.²⁹ More recently, single-macrostate simulations were also utilized in a study of adsorption of supercritical fluids¹⁷ and flexible³⁰ and rigid³¹ adsorbents. Cases such as these may be well suited for single-macrostate simulations because they avoid complications such as a two dimensional biasing function or the insertion and deletion of large or complex molecules with low acceptance probability.

A recent study of water adsorption utilized single-macrostate TMMC simulations and

reported large efficiency gains compared to conventional GCE Monte Carlo (MC) simulations.^{2,32,33} However, no efficiency comparison was made to multiple-macrostate TMMC simulations in the GCE. To our knowledge, such a comparison remains lacking.

Note that single-macrostate simulations are somewhat unique to TMMC, because other flat-histogram methods, such as Wang-Landau,³⁴ require at least two macrostates. A single-macrostate TMMC simulation does not allow transitions between macrostates, which means there is no self-consistent convergence of a bias function that seeks to “flatten” the sampling between macrostates. Instead, macrostate transitions in single-macrostate simulations are obtained by performing ghost insertions and deletions that do not actually change the microstate of the simulation or the Markov chain. In this manuscript, sampling refers to changes in the microstate. In the CE, sampling involves perturbations in particle positions and orientations. In the GCE, sampling also includes particle insertions and deletions, in addition to CE sampling. Because ghost insertions and deletions require nearly the same computational resources as GCE insertions and deletions but do not change the microstate, sampling efficiency with ghost particles is expected to be reduced compared to GCE. This sampling reduction could be negligible in situations where the CE sampling contributes more than GCE sampling, such as temperatures and densities where acceptance probabilities of GCE insertion and deletion trials are low. However, low GCE acceptance probabilities do not necessarily mean GCE trials do not contribute to sampling. In addition, low acceptance probabilities in GCE indicate high energy changes in ghost insertions and deletions, which means more samples are required for convergence. Thus, using ghost insertions and deletions does not solve the problem of low GCE acceptance probabilities, which is why configurational bias ghost insertions and deletions are used.³⁵ Currently, no studies to our knowledge have quantified the sampling efficiency of flat-histogram biased GCE trials relative to ghost insertions and deletions over a range of temperatures, densities and complex fluid behavior.

These open questions regarding the sampling efficiency of GCE trials are further complicated by the difficulty of quantifying efficiency in the GCE. While the efficiency of MC

simulations in the CE has been quantified for decades using energy fluctuation³⁶ and mean squared displacement,³⁷ these metrics are difficult to utilize in the GCE because they require tracking individual properties of particles over long CPU times. However, efficiency metrics that are amenable to the GCE as well as the CE have been utilized by observing the uncertainty in computed properties of interest, such as the GCE average energy or number of particles, as a function of CPU-time.^{38,39} Comparisons of different simulation methods are typically difficult to perform directly and can often involve a trade-off that depends on the model and properties of interest.⁴⁰ Because the only difference between single and double macrostate TMMC simulations is the sampling of biased GCE trials, we can directly quantify the effect of these trials on the efficiency of the simulation in a straightforward and easy to interpret comparison.

In this manuscript, the relative efficiency of single- and multiple-macrostate TMMC simulations in the GCE is quantified for a variety of complex fluids. This manuscript is organized as follows. Section 2 describes the simulation methods and the models used to investigate the relative efficiency in supercritical fluids, vapor—liquid equilibrium, self-assembly of micelles and adsorption in a repulsive porous network. Additional details concerning methods and models are also found in the Supporting Information. In Section 3, single- and double-macrostate simulations are rigorously compared using the time-dependent standard deviation of both the potential energy and the macrostate probability distribution. In addition, qualitative efficiency comparisons are made by quantifying the time-dependent convergence of benchmark simulations of large macrostate ranges. Finally, conclusions and future work are discussed in Section 4.

2 Methods


2.1 Single- and Double-Macrostate Simulations

Figure 1 illustrates the differences between a single-macrostate and a double-macrostate TMMC simulation in the GCE,^{1,41} where the relevant macrostate, N , is the number of fluid particles. In the single-macrostate simulation, transitions to the next highest and lowest macrostates are attempted in order to compute the transition probabilities, $P(N \rightarrow N + 1)$ and $P(N \rightarrow N - 1)$, respectively, but these transitions are always rejected, as shown by the red arrows in Figure 1, because they are outside the bounds of the allowed macrostate range. Note that transitions which are always rejected are also referred to as ghost particle insertions and deletions in this work. Similarly, a simulation with double macrostates also always rejects transitions which would change the macrostate outside of the bounds. However, the double macrostate simulation allows transitions between the allowed states, N and $N + 1$, shown in Figure 1. Thus, the attempted trials which transition between N and $N + 1$ may be accepted. These trials are referred to as GCE trials, or insertions and deletions. When accepted, these GCE trials change the microstate of the simulation, which improves sampling, and the trial acceptance is biased with the on-the-fly computed probability differences,⁴¹

$$\Delta \ln \Pi = \ln[P(N \rightarrow N + 1)/P(N + 1 \rightarrow N)]. \quad (1)$$

This improved sampling is expected in double macrostate simulations relative to single-macrostate simulations because GCE trials in double macrostate simulations can be accepted, will change the microstate, and are therefore part of the sampling algorithm. In contrast, ghost particle insertions and deletions will not change the microstate and thus are not a sampling algorithm, as opposed to GCE. Instead, ghost particle insertions and deletions are an analysis that could be computed from stored trajectories after CE sampling is performed. GCE insertions and deletions, when accepted, can lead to large changes in the microstate

two simulations with a single macrostate



one simulation with double macrostates



Figure 1: An illustration of the difference between two TMMC simulations with a single macrostate (top) and one TMMC simulation with double macrostates (bottom). In the GCE, the macrostates are the number of fluid particles, N . The arrows indicate trials which attempt to change the macrostate (i.e., insertions and deletions). The single red arrows indicate an attempt to change the macrostate outside the bounds of the simulation. The blue double arrow represents neighboring macrostates in a TMMC simulation which are in the allowed range, and thus trial attempts between these macrostates may be accepted.

relative to a small displacement typical in CE sampling. Thus, even relatively low acceptance probability GCE trials may contribute significantly to microstate sampling. Because both single and double macrostate simulations spend the majority of computational time computing the energy differences of the proposed trials, the only added computational cost of the double macrostate simulation is to apply the bias and possibly accept macrostate transitions. This additional computational cost is often negligible in comparison to the energy calculation. Thus, the major efficiency difference between the single and double macrostate simulations is the additional sampling benefit afforded by biased transitions in the double macrostate simulation, with negligible computational overhead. In other words, the parallelization cost of splitting a double macrostate simulation into two single-macrostate simulations is the automatic rejection of biased GCE trials, which is the cost required to completely decouple the macrostates and make them trivially parallelizable.

This process of decoupling the macrostates also introduces other related sources of efficiency loss. Confining the Markov chain to a single macrostate removes the self-consistent convergence of a bias function, which some may consider an essential component of flat histogram methods, while also limiting the possible sampling algorithms. In single-macrostate simulations, although the computational effort is spent to calculate $\Delta \ln \Pi$, it is not used

to bias macrostate changes, as in multiple-macrostate simulations. Macrostate changes can also allow efficient sampling, as will be discussed in subsequent sections for inhomogeneous structures.⁴²

2.2 Efficiency in Flat-Histogram Monte Carlo Simulations

For the results that follow, the CPU time-dependent standard deviation of a quantity of interest is calculated in order to compare the efficiency differences of single and double macrostate simulations. Standard deviations of the CE average potential energy, $\langle U(N) \rangle$, and the relative macrostate probability, $\Delta \ln \Pi$, given by eq 1, were obtained using the blocking method,^{43,44} as described in more detail in Section 1 of the Supporting Information. Over long times, t , the standard deviation, σ , is expected to decrease linearly on a log—log plot with a slope of negative one half,^{38,39}

$$\ln \sigma = b - 0.5 \ln t \tag{2}$$

where the intercept, b , was obtained by linear least squares regression. Rearrangement of eq 2 results in $\exp(b) = D = t^{1/2} \sigma$, where D is the difficulty as defined by Schultz and Kofke.³⁹

The relative efficiency of a simulation is defined as follows. For two simulations to obtain the same σ , the ratio of the CPU time required for simulation 1 relative to simulation 2 is

$$z_{12} = t_1/t_2, \tag{3}$$

where z_{12} is the efficiency of simulation 2 relative to simulation 1. For example, when simulation 2 is twice as fast, or obtains the same σ in half the time, as simulation 1, $z_{12} = t_1/t_2 = 2$. Thus, if $z_{12} > 1$, then simulation 2 is more efficient than simulation 1. In this work, we use the subscript 1 to denote a single-macrostate simulation and the subscript 2 to denote a double macrostate simulation. Thus, z_{12} is the efficiency of a double macrostate simulation relative to a single-macrostate simulation. In practice, it is difficult to obtain the

times at which two simulations have precisely the same σ . Instead, eq 2 may be substituted into eq 3, resulting in

$$z_{12} = \exp[2(b_1 - b_2)], \tag{4}$$

where the slope and intercept, b may be fit from one or multiple simulations to obtain z_{12} . Relative efficiency is related to the difficulty, $D = \exp(b)$, proposed by Schultz and Kofke³⁹ as $z_{12} = (D_1/D_2)^2$.

2.3 Simulation Models

Because flat-histogram (FH) TMMC simulations in the GCE have been shown in previous studies to be effective tools for studying vapor—liquid phase behavior,¹ self-assembly¹² and adsorption,⁴ the variety of models used in this study similarly span this range of phenomena to allow for a broad conclusion to be drawn about the efficiency of FH methods relative to ghost particle methods and to allow for a comparison of parallelization strategies. The models investigated here are the bulk Lennard-Jones (LJ) fluid, bulk extended simple point charge (SPC/E) water,⁴⁵ a micelle-forming patchy trimer particle^{12,13} and adsorption of an LJ fluid in a rigid and purely repulsive porous network shown in Figure 2. Here, we briefly describe relevant aspects of the simulated systems. Additional details and parameters for each fluid model can be found in Section 2 of the Supporting Information.

To study vapor—liquid phase behavior, bulk LJ and SPC/E fluids were simulated. Both of these bulk fluids undergo a series of structural transitions between the vapor and liquid phase.^{42,46} If these structural transitions are not well sampled, the relative accuracy of the free energy difference between the vapor and liquid phases will decrease. Bulk LJ was simulated at temperatures of $k_B T/\epsilon = 0.7$ and 1.5, while bulk SPC/E was simulated at temperatures of 300 and 525 K. The maximum number of particles simulated in bulk LJ and SPC/E, N_{max} , was 475 and 300, respectively, corresponding to a dense fluid. To examine the effect of system size, additional bulk LJ simulations were conducted for $N_{max} = 256$ and 2000.

To study self-assembly, patchy trimer particles with one attractive bead and two repulsive beads were simulated at the temperatures and densities known to form micelles.^{12,13} Unlike the bulk fluids, whose structures often span the periodic boundaries, many microscopic structural transitions occur in this fluid of micelles and free monomers. The maximum number of trimer particles was $N_{max} = 100$ and the temperature was $k_B T/\epsilon = 0.2$ or 0.275 .



Figure 2: Schematic illustrations of the patchy trimer particle with a blue, attractive bead and two red, repulsive beads (a) and the unit cell of the porous network (b). Black squares represent repulsive WCA sites, and white squares are open pores and channels. Starting from the left, the 2-dimensional slices are stacked on top of each other.

To study adsorption, we simulated an LJ fluid confined in a porous network with multiple, interconnected larger pores and smaller channels. In this case, the interaction of the porous network is designed to be purely repulsive, which is expected to lead to structural transitions that are associated with the successive filling of each pore with fluid. The maximum number of fluid particles was $N_{max} = 100$ and the temperature was $k_B T/\epsilon = 0.3$.

2.4 Flat-Histogram Monte Carlo Simulation Details

All simulations utilized the open source FEASST version 0.20.1 simulation software.⁴⁷ To aid in reproducibility, an example of the simulations described in this Section is provided as a tutorial in FEASST. Additional simulation details are included in Sections 2 and 3 of the Supporting Information. Each condition contained 32 independent simulations initiated with different random number seeds. MC trials included rigid displacements, rotations, single particle insertions and deletions. Patchy trimer simulations also had rigid cluster translations and rotations, as described in ref 12. Simulations of the LJ fluid in the porous network utilized dual-cut configurational bias (DCCB) for GCE trials.⁴⁸

The effect of specialized trials on efficiency was also studied by performing simulations with and without the following trials. For one density series of simulations of the bulk LJ fluid at $k_B T/\epsilon = 0.7$, DCCB with 8 first bead insertions and deletions was used with the WCA reference potential. In addition, aggregation volume bias (AVB) was used in both the CE⁴⁹⁻⁵¹ and GCE^{52,53} for the bulk and confined LJ fluid with inner and outer radii of 0.9σ and 1.375σ respectively. The CE AVB used AVBMC2 and a variant of AVBMC3 of equal probability to attempt that was ten times lower than traditional displacement trials.^{47,51} The AVBMC3 variant is described in Section 3 of the Supporting Information. The GCE AVB insertions and deletions were also attempted with equal probability that was ten times lower than traditional particle insertion and deletion trials.^{52,53} Note that the GCE AVB trials in the porous network simulations also utilized 4 first bead insertions and deletions with DCCB.^{13,48} Finally, an additional bulk LJ simulation also utilized trial displacements with the maximum set to a constant value of half the periodic boundary length in each dimension, and these whole-domain displacements were attempted with ten times lower probability than the usual, tuned displacements.

3 Results and Discussion

3.1 Single Versus Double Macrostate Simulations

To compare the efficiency of the single and double macrostate simulations, consider Figure 1. If the macrostate changes indicated by the blue double arrow do not improve sampling, then the two types of simulations should give the same error estimate if run for the same total amount of CPU time. However, if the change in the macrostate from the blue double arrow does improve sampling, the single-macrostate simulations will spend CPU time computing those macrostate transition probabilities without the additional sampling and will thus become less efficient than the double macrostate simulation.

The block standard deviation as a function of CPU time is expected to decrease on a

log—log plot with a linear slope of $-1/2$, as shown in eq 2 and demonstrated in Figure 3 for a supercritical LJ fluid with $k_B T/\epsilon = 1.5$ and $N = 200$. While both simulations were run for 24 hours, the single-macrostate simulations were parallelized with two simulations, and the CPU time reported is the aggregate CPU time utilized by both cores. The relative efficiencies of the two simulations are then obtained for 32 independent pairs of simulations, with the linear fits shown by the often overlapping black lines in Figure 3. The linear fits were truncated in this case to only include time beyond 1 h to ensure a fit to a linear region, which appears as 2 h for the single-macrostate simulations due to the cumulative time of both parallel processes. For the potential energy, the efficiency of the double macrostate simulation relative to the single-macrostate simulations is $z_{12} = 6.6$ with a standard deviation of the mean of 0.3; for the change in the natural logarithm of the probability, the relative efficiency, z_{12} , is 3 with a standard deviation of the mean of 0.1. Thus, two simulations with a single macrostate each consume about 3 to 6.6 times more CPU resources than one simulation with double macrostates in order to obtain the same standard deviation for the potential energy and $\Delta \ln \Pi$, respectively. All of the remaining results in this manuscript report on the relative efficiency with respect to the standard deviation of the change in the macrostate probability distribution, $\sigma_{\Delta \ln \Pi}$, because every GCE averaged property depends on it. However, in this case and many others, the efficiency metrics based on potential energy gave results either similar to or favoring multiple macrostates, but always within the same order of magnitude range as metrics based on the macrostate probability distribution.

The computation of the efficiency of double macrostates relative to single-macrostate simulations, z_{12} , was illustrated in detail in Figure 3 at one condition. In Figure 4, we show z_{12} over the entire range of models, densities and temperatures simulated. The first sanity check is the $N = 0$ and 1 results for all the bulk simulations (e.g., not the porous network). At $N = 0$ for bulk fluids, $z_{12} \approx 1$ is expected because the sampling of 0 or 1 particles is trivial and the efficiency is not affected by the sampling of biased GCE trials. This sanity check also demonstrates that any overhead due to multiple-macrostate TMMC or parallelization

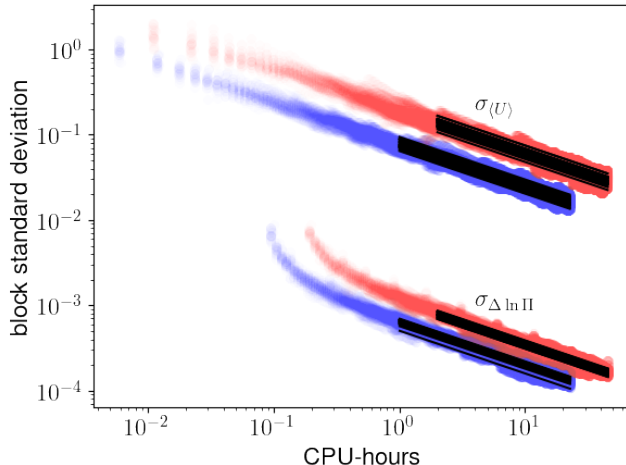


Figure 3: Block standard deviation as a function of time with bulk LJ, $N = 200$, $k_B T/\epsilon = 1.5$ for one simulation with double macrostates (blue) and two simulations with a single macrostate each (red). The block standard deviation of the energy and $\Delta \ln \Pi$ is shown by the top pair and bottom pair, respectively, including data and linear fits from 32 independent simulations.

is minimal, or else it would affect the $N = 0$ and 1 results where the energy calculation does not take up the majority of the CPU time. On the other hand, the $N = 1$ case in a porous network depends on sampling with a non trivial energy calculation.⁵⁴ In this case, $z_{12} \approx 4$, which already shows the single-macrostate simulations to be highly inefficient.

It is also important to note that all measures of relative efficiency compared simulations with the same trial move sets, operational parameters, computer hardware and simulation software. Thus, any optimization or inefficiency introduced by these arbitrary choices is expected to affect both simulations to the same extent. The validity of this assumption is demonstrated by the sanity check at $N = 0$ and 1 for bulk fluids in Figure 4.

The supercritical LJ fluid at $k_B T/\epsilon = 1.5$ is the least affected by sampling differences of single and double macrostates, shown by the dashed black line in Figure 4. This is likely because the supercritical configurations are most often homogeneous, as opposed to inhomogeneous, where distinct structural motifs can be observed in simulations of vapor—liquid equilibrium,^{42,46} self-assembly¹² and adsorption.⁴ Despite the high temperature and lack of structure, the single-macrostate simulations require up to three times more CPU time




Figure 4: Efficiency of double macrostate relative to single-macrostate simulations, z_{12} , as a function of the number of fluid particles, N , divided by the maximum number of particles for a high density fluid, N_{max} , for the following simulations: bulk LJ (black circle) with $k_B T/\epsilon = 1.5$ (dashed lines) and $k_B T/\epsilon = 0.7$ (solid lines), bulk SPC/E (red bar) with $T = 525$ K (dashed lines) and $T = 300$ K (solid lines), trimer patchy particles (blue x) with $k_B T/\epsilon = 0.275$ (dashed lines) and $k_B T/\epsilon = 0.2$ (solid lines), and the confined LJ fluid with $k_B T/\epsilon = 0.3$ (green triangle). Lines between the points are guides for the eye. The gray rectangle highlights $z_{12} = 1$ to 2 . Error bars, often smaller than symbols, are the standard deviation of the mean from 32 independent simulations.

to obtain the same standard deviation than the double macrostate simulations, as reflected by $z_{12} = 3$ in this density range. The relative efficiency, z_{12} , then decreases to a value of 1.47 at $N_{max} = 475$ or $\rho = 0.928$. Note that the parallelization methods should ideally target $z_{12} = 1$, and that $z_{12} = 2$ is the absolute minimum efficiency for any reasonable parallelization algorithm over two processors. For $z_{12} = 2$, the same standard deviation in $\Delta \ln \Pi$ is obtained in one simulation of double macrostates run for half the total CPU time as both single-macrostate simulations. This means that when $z_{12} = 2$, the parallelization does not speed up the calculation, but requires twice as many processors. Moreover, $z_{12} = 3$ means that the parallelization over two processors requires 1.5 times as long as the double macrostate simulation while using twice as many processors, so the parallelization not only consumes more resources but actually increases the time it takes to obtain a result to the desired precision. The gray rectangle in Figure 4 shows this region between $z_{12} = 1$ and 2 that represents the efficiency limit of parallelization algorithms. Even for the supercritical LJ fluid, there are only limited cases at high and low density where single-macrostate simulations are a reasonable parallelization strategy.

The relative efficiency, z_{12} , increases by an order of magnitude when conditions with vapor—liquid equilibrium are simulated. For example, when the temperature of the LJ fluid is decreased to $k_B T/\epsilon = 0.7$ within the density range of vapor—liquid equilibrium, the efficiency of the double macrostate simulation increases up to 20 times more than the single-macrostate simulations. In addition, for SPC/E water, which also phase separates at 300 K and 525 K, double macrostate simulations are up to 10 to 20 times more efficient than single-macrostate simulations. Even though insertion and deletion acceptance probabilities decrease from high to low temperature, z_{12} increases. For example, the GCE insertion and deletion acceptance probability of bulk LJ at $N/N_{max} = 400/475$ is approximately 1.6×10^{-3} at $k_B T/\epsilon = 1.5$ and 5×10^{-4} at $k_B T/\epsilon = 0.7$. This demonstrates that GCE trials can greatly impact the sampling efficiency even at low acceptance probabilities, because free-energy biased GCE sampling is the only difference between the single and double macrostate

simulations.

One explanation for the remarkably important contributions of biased GCE sampling to simulation efficiency is that the shape of an inhomogeneous structure is likely to evolve quicker with insertions and deletions than with simple small-scale displacements. A series of structures form between the vapor and liquid, specifically the droplet, cylinder and slab.^{42,46} GCE insertions and deletions allow particles to be efficiently added to or removed from the structures, the surrounding vapor, or interfaces. It is important for these metastable transition-state structures to be sampled well or else the relative probability of the vapor and liquid, which affects all equilibrium properties, may not be computed accurately. Later, we will present results examining the efficiency impact of trial moves that are intended to facilitate the sampling of these structures.

At conditions where patchy trimer particles self-assemble into micelles, z_{12} , reaches 2 orders of magnitude, as shown by the blue lines in Figure 4. While the lower temperature of $k_B T/\epsilon = 0.2$ formed micelles of a relatively precise size of 13 particles, the higher temperature $k_B T/\epsilon = 0.275$ simulations formed clusters that were less well-defined and were near the approximately defined critical micelle temperature. In both cases, $N_{max} = 100$, allowing for over half a dozen micelles. Because the only difference in the simulations is the biased GCE sampling, these results show their strikingly important contributions to the simulation efficiency. Note that rigid cluster translations and rotations are present in the simulation, which are CE trials. Thus, it is possible for two partially formed micelles to approach each another in the single-macrostate simulations by moving the clusters together as two wholes. However, due to detailed balance, rigid cluster moves which result in two clusters becoming one must be rejected, so coalescence must occur via single-particle translations and rotations. On the other hand, biased GCE sampling allows particles to be added directly to the more stable cluster and removed from the less stable cluster.

For the confined LJ fluid studied in this work, the relative efficiency, z_{12} , reaches over 3 orders of magnitude, as shown by the green line in Figure 4. At these low temperatures of

$k_B T/\epsilon = 0.3$, the fluid particles would coalesce and freeze or vitrify in bulk, but the fluid is frustrated by the size and shape of the pores and the small channels connecting them. Thus, the most stable configurations are those associated with successive pore filling. However, the CE trials clearly struggle to sample these pores without biased GCE sampling. If two pores were half filled, the CE trials would require each particle to traverse the channels in order to fill a single pore. On the other hand, biased GCE sampling allows particles to be removed from one pore and added to another. Thus, we will now consider different trial move sets which improve the sampling in either the CE, or the GCE, to investigate their effect on the relative efficiency.

The dependence of z_{12} on different Monte Carlo trial move sets is shown in Figure 5 for the bulk and confined LJ fluid. For reference, the same data presented in Figure 4 for bulk LJ fluid at vapor—liquid equilibrium for $k_B T/\epsilon = 0.7$ with only single particle translations, insertions and deletions (i.e., no configurational bias) is shown by the black circles with a solid line in both Figures 4 and 5. When the sampling of GCE insertions and deletions is improved by using DCCB with 8 attempts of a WCA reference potential and a cell list, the efficiency of the double macrostate simulation increases relative to the single-macrostate simulations. This increase in z_{12} is likely because improved GCE sampling is wasted in single-macrostate simulations that will reject the insertions and deletions regardless. The effect of DCCB is especially pronounced at the higher densities, and takes z_{12} above the critical value of parallel feasibility, $z_{12} = 2$, even at the highest macrostate, $N_{max} = 475$, $\rho = 0.928$. This is likely because DCCB increases the GCE acceptance probability at high density without an overly burdensome calculation of the energy of the reference potential due to a judicious use of a cell list. If the GCE trial acceptance probability is prohibitively small, to the point that the GCE sampling contributes little to the overall sampling, then z_{12} may approach unity. However, additional effort spent to improve the GCE sampling will disproportionately help the multiple-macrostate simulations relative to the single-macrostate simulations. Additionally, it is typically worth the effort to improve the GCE sampling using

better trial moves if the GCE acceptance probabilities become prohibitively small with the existing set of trials. Prohibitively large energies in ghost particle insertions and deletions also reduce the efficiency of single-macrostate simulations, which is why configurational-bias has also been implemented for ghost particles.³⁵ The addition of GCE AVB moves leads to modest increases in z_{12} relative to DCCB, as shown by the blue squares with dashed lines in Figure 5. This result still reinforces the notion that increased effort put into improving GCE sampling serves only to increase z_{12} in favor of multiple-macrostate simulations.




Figure 5: Same as Figure 4, except the bulk LJ fluid at $k_B T/\epsilon = 0.7$ is shown in black or blue and the confined LJ fluid at $k_B T/\epsilon = 0.3$ is shown in green with the following trial move sets: single particle translations, insertions and deletions (circles with solid lines) CE AVB (x with dotted lines), and GCE AVB (square with dashed lines). The porous network always used GCE DCCB, while the bulk LJ fluid also includes GCE DCCB (black triangles with dotted lines) and whole-domain CE translations (black stars with dash dot lines). The bulk LJ fluid with a smaller and larger system size, $N_{max} = 256$ and 2000, are shown with orange pentagons and purple horizontal bars, respectively.

While improving GCE sampling increases z_{12} , improving CE sampling decreases z_{12} . This is demonstrated in the LJ fluid at $k_B T/\epsilon = 0.7$ by introducing either whole-domain translations or CE AVB trials, shown by the black star and blue x, respectively, in Figure 5. Although increasing CE sampling may decrease z_{12} by up to a factor of 2, the single-macrostate simulations are still up to an order of magnitude less efficient than the double macrostate simulations. Thus, the improvement in the CE sampling for the trials investigated

is not effective enough to completely overshadow the efficiency gains from GCE sampling.

To gain some understanding of the difficulty of replicating GCE efficiency in the CE, consider the relationship between GCE insertions and deletions to whole-domain displacements. A whole-domain displacement is essentially a combined deletion and insertion trial. However, the GCE trials are more efficient by approximately the inverse of the acceptance probability. As an example, if the probability of insertion and deletion is 0.1, then 100 attempts of each type of trial would result in about 10 accepted insertions and 10 accepted deletions. In the same CPU time, instead 100 whole-domain displacements could have been attempted. However, the acceptance probability of the whole-domain displacement is the product of the probabilities of insertion and deletion, which leads to only 1 accepted whole-domain displacement in this example. Thus, GCE insertions or deletions are approximately 10 times more efficient in this example, which is the inverse of the acceptance probability. Furthermore, the whole-domain trials will not benefit from GCE FH bias because they do not change the macrostate.

The effect of system size on the efficiency is also shown in Figure 5 by comparing the black circles with $N_{max} = 475$ against the purple horizontal bars with $N_{max} = 2000$ and orange pentagons with $N_{max} = 256$. The volume of the domain was scaled proportionally with the number of particles so that the densities of the N/N_{max} points are nearly equivalent within the constraint that the number of particles is an integer. Other than the system size difference, all other aspects of the simulations were identical. Increasing the system size increases z_{12} for $N/N_{max} \leq 0.42$, while z_{12} is unchanged within statistical uncertainty for $N/N_{max} \geq 0.63$. Decreasing the system size decreases z_{12} for $N/N_{max} \leq 0.21$, increases z_{12} for $N/N_{max} = 0.42$, while z_{12} is unchanged within statistical uncertainty for $N/N_{max} \geq 0.63$. These differences in z_{12} with system size may be attributed to the system-size dependence of the inhomogeneous structures.

Similar effects of increased CE or GCE sampling seen in the bulk LJ fluid with vapor equilibrium is also apparent in the confined LJ fluid. The green circles with the solid line in

Figure 5 show z_{12} for the LJ fluid in a porous network at $k_B T/\epsilon = 0.3$ with tunable single-particle translations and DCCB insertions and deletions. DCCB was used for all simulations in the porous network because it is an efficient way to identify pores and channels, even for $N = 1$. While the addition of CE AVB trials reduced z_{12} by over one order of magnitude in some cases, this was still not enough to overcome the up to three orders of magnitude in z_{12} in the original trial set. This order of magnitude reduction in z_{12} is likely due to the ability of AVB CE moves to transfer particles between pores without needing to traverse the channels. The addition of GCE AVB trials also led to modest differences in z_{12} compared to the original or reference set of trial moves. Also note that AVB does not affect the efficiency at $N = 0$ and 1 because AVB trials cannot be attempted with less than 2 particles.

3.2 Parallel Flat-Histogram Monte Carlo Simulations

While the previous section considered only single or double macrostate simulations, many applications require simulations with hundreds or more macrostates (i.e., number of particles) to span the relevant density range for studying phase behavior, self-assembly or adsorption. In this section, the relative convergence of single-macrostate simulations are compared with multiple-macrostate simulations when parallelized on high-performance computer (HPC) nodes.




Figure 6: Illustration of nine simulations with single macrostates of 0 to 8 particles load balanced on three processors, p_1 , p_2 and p_3 (a) and multiple-macrostate simulations on each processor (b). The red and blue vertical lines are as described in Figure 1.

An illustration of two possible parallelization strategies, single-macrostate simulations or multiple-macrostate simulations, is shown in Figure 6. In the first strategy shown in Figure

6a, nine simulations of single macrostates with a particle number of 0 to 8 are simulated in parallel with a pool of 3 processors, p_1 , p_2 and p_3 . For this illustration, assume the time to complete the simulation scales with the number of particles squared, N^2 . The goal is to balance the load of each processor equally to reduce wasted CPU resources. For load balancing, each processor simulates the highest particle number that has not already been simulated. As an example, the first processor, p_1 begins with the largest number of particles, 8. Then, the second processor, p_2 , simulates 7 particles and the third processor simulates 6, all with ghost insertions and deletions, as indicated by the red vertical lines in Figure 6. Because the third processor, p_3 , has the least particles, it finishes first and moves to 5 particles, and the second processor begins on 4 particles. The third processor then picks up its third macrostate, $N = 3$, because the first processor is still processing 8 particles (e.g., $8^2 > 5^2 + 6^2$). This ultimately results in p_1 simulating only 8 and 2 particles, while p_2 simulates 7, 4, 1, and 0 particles and p_3 simulates 6, 5, and 3. The expected total CPU time of each processor is expected to vary by no more than approximately 5 %.

The second parallelization strategy shown in Figure 6b sets the number of multiple-macrostate simulations equal to the number of processors. Contrary to the single-macrostate simulations, each processor is fixed to a multiple-macrostate simulation with a contiguous subset of macrostates and run for the same CPU time as all of the other processors. The range is contiguous so transitions are allowed between macrostates, as shown by the blue vertical dashed lines in Figure 1 and 6. In this case, load balancing is represented by how subsets of macrostates are divided among processors. This load balancing choice is as equally arbitrary as the choice of the number of trials per macrostate in single-macrostate simulations and the order in which the macrostates are assigned to the processors.

Benchmark simulations over a large density range were performed to compare parallelization performance. These benchmark simulations include simulations with one macrostate each and simulations with multiple macrostates. For simulations with only a single macrostate per processor, ghost insertions and deletions were performed, as described above and illus-

trated in Figure 6a. These single-macrostate TMMC simulations are parallelized by pooling the resources of all 32-cores of a node, e.g., running each macrostate until all macrostates are finished. To improve load-balancing, the highest density simulations were run first because they take the longest to complete. Each macrostate, N , was equilibrated with $Nn_s/10$ trials and then run in production with Nn_s trials, where $n_s = 10^3, 10^4, 10^5, 10^6$ for bulk LJ and SPC/E and $n_s = 10^4, 10^5, 10^6, 10^7$ for trimers and confined LJ.¹⁷ The simulation protocols were otherwise the same as those described in Section 2.4, using the same trial move sets as described for each model in Figure 4 with no AVB or whole-domain displacement trials. The porous network simulations were the only model that used DCCB in these parallelized simulations.

Benchmark TMMC simulations with multiple macrostates were also parallelized to compare with the single-macrostate simulations. The trial move sets in these multiple-macrostate simulations were identical to the single-macrostate simulations. The parallel scheme used in these multiple-macrostate simulations follow the general approach illustrated in Figure 6b and found in a number previous studies,^{9,12} without communication between processors.^{13,24} First, macrostate ranges for each processor are divided according to an arbitrarily chosen exponential factor, γ . For bulk LJ, bulk SPC/E, patchy trimers, and LJ in a porous network, $\gamma = 2.5, 1.75, 1.5$ and 1.5 , respectively. For example, in bulk LJ, $\gamma = 2.5$ and $N_{max} = 475$ resulted in the following macrostate minimums of the 32 processors: 0, 120, 158, ..., 458, 464, 470. The minimum number of macrostates in the highest density simulations was 5 for LJ with $\gamma = 2.5$ and $N_{max} = 475$, while the minimum macrostate range was 2 for patchy trimers with $\gamma = 1.5$ and $N_{max} = 100$. These values of γ were arbitrarily chosen, with a higher γ spending more CPU time on higher densities relative to lower densities. Section 3.3 discusses how to account for this arbitrary choice when comparing with single-macrostate simulations.

For initialization, the minimum number of particles is obtained as described for the single and double macrostate TMMC simulations. The TMMC benchmark simulations were

terminated after 3.86, 8.06, 2.57 and 2.06 h, for bulk LJ, SPC/E, trimers and the porous network, respectively, which are equivalent to the elapsed real time of the second highest n_s for single-macrostate simulations (i.e., $n_s = 10^5$ for LJ and SPC/E and 10^6 for trimers and confined LJ). The $\ln \Pi$ over all macrostates is then obtained by splicing together all of the collection matrices in order, in exactly the same fashion as the single-macrostate simulations. The $\ln \Pi$ from these benchmark simulations were then compared with reference simulations that were run for a very long time. For the bulk LJ and SPC/E fluids, these reference simulations were taken from the SRSW.⁵⁵ For the patchy trimer particles and the LJ fluid in a porous network, reference TMMC simulations were run in the same way as these benchmark simulations, but for 50 days.

3.3 Parallel Benchmarks

To compare the efficiency of parallelized single-macrostate simulations with parallelized flat-histogram simulations, we will examine the relative convergence of these two methods with short benchmark simulations that are compared to a reference simulation that was run for much longer. First, consider a reference simulation of a bulk LJ fluid at $k_B T/\epsilon = 0.7$ reweighted to the μ that gives equal probability of the vapor and liquid phases, as shown at the top of Figure 7. The red dashed line is data taken from the SRSW⁵⁵ for $k_B T/\epsilon = 0.7$. The vapor peak on the left occurs at very small densities, while the liquid peak on the right occurs around 432 particles for $V = (8\sigma)^3$, and then the probability drops until truncation at $\Pi(N = 475) < e^{-50}$. The transition region becomes flat around 200 particles, where a slab configuration is the minimum free energy structure and grows in width at constant free energy at equilibrium^{42,46} (e.g., $\Delta \ln \Pi \approx 0$ at $N \approx 200$). We will now use this reference simulation to compare the relative convergence of shorter benchmark simulations.

Single macrostate simulations were parallelized with a pool of 32 processors and run for $n_s = 10^3, 10^4, 10^5$ and 10^6 , shown by yellow, green, blue and purple, respectively, in Figure 7. Because the best converged purple and blue simulations are not distinguishable on the




Figure 7: Natural logarithm of the macrostate probability, $\ln \Pi$, for the Lennard-Jones fluid at $k_B T/\epsilon = 0.7$, reweighted to equilibrium (top). Results from the SRSW are shown by a red dashed line. Parallelized single-macrostate non flat-histogram simulations are shown for $n_s = 10^3$ (yellow), $n_s = 10^4$ (green), $n_s = 10^5$ (blue) and $n_s = 10^6$ (purple). The black dots show multiple-macrostate TMMC results obtained with the same CPU time as the $n_s = 10^5$ single-macrostate TMMC. The change in the natural logarithm of the macrostate probability, $\Delta \ln \Pi(N)$ (middle). The absolute value of the difference between the $\Delta \ln \Pi$ of the benchmark simulation and reference results, $|\Delta \Delta \ln \Pi|$ (bottom).

scale of $\ln \Pi$ variations over the entire macrostate range, it is instructive to consider instead the change in $\ln \Pi$ from $N - 1$ to N , $\Delta \ln \Pi(N) = \ln[\Pi(N)/\Pi(N - 1)]$. Also, $\Delta \ln \Pi$ is a more convenient quantity than $\ln \Pi$ to consider for the following reasons: $\Delta \ln \Pi$ is independent of normalization, $\sum \Pi = 1$, $\Delta \ln \Pi$ shifts by only a constant when reweighting with respect to μ , and $\Delta \ln \Pi$ varies over N on a smaller scale than $\ln \Pi$. Even so, the best converged purple and blue simulations are not distinguishable from the reference simulation on the scale used for the middle of Figure 7. However, they are distinguishable when plotted as $|\Delta \Delta \ln \Pi|$, the absolute value of the difference in the $\Delta \ln \Pi$ of the benchmark simulation relative to the $\Delta \ln \Pi$ of the reference simulation, as shown at the bottom of Figure 7. A smaller $|\Delta \Delta \ln \Pi|$ indicates better agreement between the benchmark and reference simulations. As expected, the purple simulations have a smaller $|\Delta \Delta \ln \Pi|$ than the blue, on average, because the purple was run for 10 times longer.

To investigate how the efficiency of these parallelized single-macrostate simulations com-

pare with parallelized multiple-macrostate simulation, shown by the black dots in Figure 7, were run for the same CPU time on the same HPC node as the single-macrostate simulations with $n_s = 10^5$, shown by the blue line in Figure 7. Note that the multiple-macrostate parallel simulations were divided into the same number of subsets as available processors, 32, using an exponential factor, $\gamma = 2.5$, which for example led to a subset of 0 to 120 on the first processor, and 470 to 475 on the last processor, as described in more detail in Section 3.2. For most macrostates, the multiple-macrostate simulations have a smaller $|\Delta\Delta \ln \Pi|$ than not only the blue single-macrostate simulations with $n_s = 10^5$, but even smaller than the purple single-macrostate simulations with $n_s = 10^6$, even though they were run for 10 times longer than the multiple-macrostate simulations. However, at higher values of N , the multiple-macrostate simulations have a similar $|\Delta\Delta \ln \Pi|$ as the single-macrostate simulations with $n_s = 10^5$.

In order to compare the relative efficiency of single and multiple-macrostate simulations, we look at the distribution of $|\Delta\Delta \ln \Pi|$ values for each type of simulation. More efficient simulations will have distributions favoring more negative values of $\ln |\Delta\Delta \ln \Pi|$. However, a proper comparison must account for the computational effort that increases with N . For example, obtaining $\ln \Pi$ at high values of N requires more CPU time than low values of N , as discussed in Section 3.2. The arbitrary choice of γ to decompose the multiple-macrostate simulation range into subsets is related to the arbitrary choice of Nn_s trials per single-macrostate simulation. For example, $N^\delta n_s$ trials could instead be performed per single macrostate. If $\delta = 0.5$, a greater portion of computational effort would be spent on the lower N values, relative to $\delta = 1$. Because it would be difficult to choose a γ and δ that are guaranteed to divide computational effort among macrostates in the same proportion, the difference between the benchmark and long-time $\Delta \ln \Pi$ must be weighted by the effort required to compute that difference as a function of macrostate. The amount of computational effort to compute a change in energy at a given macrostate scales as approximately N^2 for the models considered in this manuscript because the potential cutoff is roughly half of the

periodic domain length. For this reason, probability distributions of $|\Delta\Delta \ln \Pi|$ are weighted by N^2 as follows.

$$P(|\ln \Delta\Delta \ln \Pi|) = \frac{\sum_i N_i^2 H(|\ln \Delta\Delta \ln \Pi_i|)}{\sum_i N_i^2} \quad (5)$$

where the sum over i is for each unique value of N and $H(x)$ is 1 when x is within the histogram subinterval or bin.

In Figure 8, we show $P(\ln |\Delta\Delta \ln \Pi|)$ for the lowest temperatures in the following four models considered: bulk LJ, bulk SPC/E, self-assembly of patchy trimers, and confined LJ. A smaller $\Delta\Delta \ln \Pi$ means better convergence of a short benchmark simulation, and the black dotted results shown in Figure 8 were multiple-macrostate simulations run for the same CPU time on the same node as the blue single-macrostate simulations. For the bulk LJ fluid at $k_B T/\epsilon = 0.7$, the multiple-macrostate simulations ran as well as the $n_s = 10^5$ at the highest densities and better than $n_s = 10^6$ at the lowest densities. This qualitatively agrees with the results of Figures 4 and 5. For the bulk SPC/E fluid at $T = 300$ K, the multiple-macrostate simulations have $\Delta\Delta \ln \Pi$ similar to the single-macrostate simulations with $n_s = 10^6$, even though the multiple-macrostate simulations were run for 10 times less total CPU time. Similarly, the micelle-forming patchy trimers at $k_B T/\epsilon = 0.2$ have similar $\Delta\Delta \ln \Pi$ to the single-macrostate simulations with $n_s = 10^7$, even though the multiple-macrostate simulation was run for 10 times less total CPU time. The LJ fluid in a porous network at $k_B T/\epsilon = 0.3$ had much smaller $\Delta\Delta \ln \Pi$ than the single-macrostate simulations with $n_s = 10^7$, even though the multiple-macrostate simulation was run for 10 times less total CPU time. Thus, parallelization with multiple macrostates is an order of magnitude or more efficient than parallelization with single macrostates in these practical applications.

4 Conclusions

Single macrostate simulations are shown to be inefficient relative to multiple-macrostate TMMC for a variety of models, thermodynamic conditions and Monte Carlo trial move

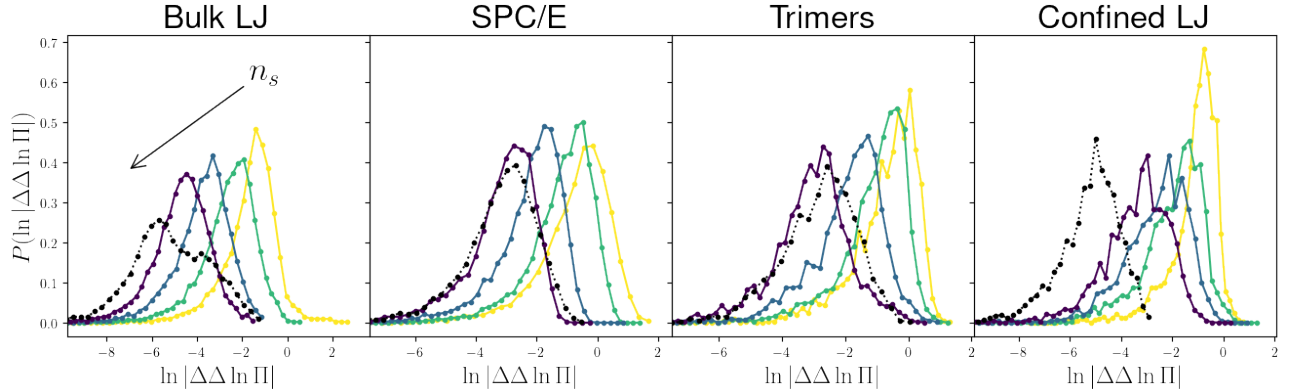


Figure 8: Probability distributions, $P(\ln |\Delta\Delta \ln \Pi|)$, averaged over 32 parallelized simulations with 50 bins each. Starting from the left panel, simulations of bulk LJ at $k_B T/\epsilon = 0.7$, SPC/E at $T = 300$ K, trimers at $k_B T/\epsilon = 0.2$, and confined LJ at $k_B T/\epsilon = 0.3$ are shown for parallelized single-macrostate simulations, starting with $n_s = 10^3$ for bulk LJ and SPC/E, and $n_s = 10^4$ for trimers and confined LJ (yellow). Each darker color has 10 times more n_s than the preceding one, from yellow to green to blue to purple. The black dotted line is multiple-macrostate TMMC simulations run for the same CPU time as the blue single-macrostate simulation.

sets. Single macrostate simulations can be 1 order of magnitude less efficient than double macrostate simulations for bulk vapor—liquid equilibrium calculations, 2 orders of magnitude less efficient for self-assembly and 3 orders of magnitude less efficient for fluid adsorption in a repulsive porous network. These results highlight the strikingly effective sampling of grand canonical ensemble insertion and deletion trials in these cases. Improving the grand canonical ensemble sampling with specialized trials makes single-macrostate simulations even less efficient than multiple-macrostate simulations, while improved CE sampling reduces the efficiency gap between single and double macrostate simulations. However, this reduction in the efficiency gap was not enough to make single-macrostate simulations a reasonable choice in the models, thermodynamic conditions and trial move sets investigated.

This study does not address the selection of optimal macrostate ranges in parallel simulations.²⁷ Future work could also include the simulation of larger or more complex molecules than three site Lennard-Jones and water models. For complex molecules, specialized techniques which improve the acceptance probabilities of insertions and deletions, such as configurational bias⁴⁸ and growth expanded ensembles,⁵⁶ are expected to work in favor of multiple

macrostates relative to single macrostates, based on the comparisons made between different trial move sets. In this study, the simulation efficiency is impacted by both the flat-histogram bias and GCE microstate sampling in a coupled manner. Decoupling their efficiency impact could be an area of interest in future work.

5 Acknowledgements and Disclaimer

This manuscript was funded by the National Institute of Standards and Technology and is not subject to U.S. Copyright. Certain commercial firms and trade names are identified in order to specify the usage procedures adequately for reproducibility. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that related products are necessarily the best available for the purpose. JRE gratefully acknowledges the financial support of the National Science Foundation (grant no. CHE-1900344).

Supporting Information Available

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jpcc.3c00613>.

Block averages in flat-histogram simulations describing the procedure for computing block averages in flat-histogram Monte Carlo simulations, additional simulation details regarding simulation box sizes, cutoff parameters, particle topology and confinement structure, additional Monte Carlo simulation details describing run-time parameters that enable comparison between different macrostate division schemes, specialized Monte Carlo trial moves and simulation initialization.

References

- (1) Errington, J. R. Direct calculation of liquid-vapor phase equilibria from transition matrix Monte Carlo simulation. *J. Chem. Phys.* **2003**, *118*, 9915–9925.
- (2) Siderius, D. W.; Hatch, H. W.; Errington, J. R.; Shen, V. K. Comments on “Monte Carlo simulations for water adsorption in porous materials: Best practices and new insights”. *AIChE J* **2022**, *68*, e17686.
- (3) Panagiotopoulos, A. Z. Critical parameters of the restricted primitive model. *J. Chem. Phys.* **2002**, *116*, 3007–3011.
- (4) Rane, K. S.; Kumar, V.; Errington, J. R. Monte Carlo simulation methods for computing the wetting and drying properties of model systems. *J. Chem. Phys.* **2011**, *135*, 234102.
- (5) Kumar, V.; Sridhar, S.; Errington, J. R. Monte Carlo simulation strategies for computing the wetting properties of fluids at geometrically rough surfaces. *J. Chem. Phys.* **2011**, *135*, 184702.
- (6) Guo, W.; Bali, P.; Errington, J. R. Calculation of the Saturation Properties of a Model Octane–Water System Using Monte Carlo Simulation. *J. Phys. Chem. B* **2018**, *122*, 6260–6271.
- (7) Guo, W.; Errington, J. R. Effect of Carboxylic Acid on the Wetting Properties of a Model Water–Octane–Silica System. *Langmuir* **2019**, *35*, 6540–6549.
- (8) Siderius, D. W.; Shen, V. K. Use of the Grand Canonical Transition-Matrix Monte Carlo Method to Model Gas Adsorption in Porous Materials. *J. Phys. Chem. C* **2013**, *117*, 5861–5872.
- (9) Shen, V. K.; Siderius, D. W. Elucidating the effects of adsorbent flexibility on fluid

- adsorption using simple models and flat-histogram sampling methods. *J. Chem. Phys.* **2014**, *140*, 244106.
- (10) Siderius, D. W.; Mahynski, N. A.; Shen, V. K. Relationship between pore-size distribution and flexibility of adsorbent materials: statistical mechanics and future material characterization techniques. *Adsorption* **2017**, *23*, 593–602.
- (11) Shen, V. K.; Siderius, D. W.; Mahynski, N. A. Molecular simulation of capillary phase transitions in flexible porous materials. *J. Chem. Phys.* **2018**, *148*, 124115.
- (12) Hatch, H. W.; Mittal, J.; Shen, V. K. Computational study of trimer self-assembly and fluid phase behavior. *J. Chem. Phys.* **2015**, *142*, 164901.
- (13) Hatch, H. W.; Yang, S.-Y.; Mittal, J.; Shen, V. K. Self-assembly of trimer colloids: effect of shape and interaction range. *Soft Matter* **2016**, *12*, 4170–4179.
- (14) Mahynski, N. A.; Blanco, M. A.; Errington, J. R.; Shen, V. K. Predicting low-temperature free energy landscapes with flat-histogram Monte Carlo methods. *J. Chem. Phys.* **2017**, *146*, 074101.
- (15) Mahynski, N. A.; Errington, J. R.; Shen, V. K. Multivariable extrapolation of grand canonical free energy landscapes. *J. Chem. Phys.* **2017**, *147*, 234111.
- (16) Mahynski, N. A.; Jiao, S.; Hatch, H. W.; Blanco, M. A.; Shen, V. K. Predicting structural properties of fluids by thermodynamic extrapolation. *J. Chem. Phys.* **2018**, *148*, 194105.
- (17) Witman, M.; Mahynski, N. A.; Smit, B. Flat-Histogram Monte Carlo as an Efficient Tool To Evaluate Adsorption Processes Involving Rigid and Deformable Molecules. *J. Chem. Theory Comput* **2018**, *14*, 6149–6158.
- (18) Mahynski, N. A.; Hatch, H. W.; Witman, M.; Sheen, D. A.; Errington, J. R.; Shen, V. K.

- Flat-histogram extrapolation as a useful tool in the age of big data. *Mol. Simul* **2020**, *0*, 1–13.
- (19) Monroe, J. I.; Hatch, H. W.; Mahynski, N. A.; Shell, M. S.; Shen, V. K. Extrapolation and interpolation strategies for efficiently estimating structural observables as a function of temperature and density. *J. Chem. Phys.* **2020**, *153*, 144101.
- (20) Shell, M. S.; Debenedetti, P. G.; Panagiotopoulos, A. Z. An improved Monte Carlo method for direct calculation of the density of states. *J. Chem. Phys.* **2003**, *119*, 9406–9411.
- (21) Anderson, J. A.; Jankowski, E.; Grubb, T. L.; Engel, M.; Glotzer, S. C. Massively parallel Monte Carlo for many-particle simulations on GPUs. *J. Comput. Phys* **2013**, *254*, 27–38.
- (22) Vlugt, T. J. H. Efficiency of Parallel CBMC Simulations. *Mol. Simul* **1999**, *23*, 63–78.
- (23) Hatch, H. W. Parallel Prefetching for Canonical Ensemble Monte Carlo Simulations. *J. Phys. Chem. A* **2020**, *124*, 7191–7198.
- (24) Rane, K. S.; Murali, S.; Errington, J. R. Monte Carlo Simulation Methods for Computing Liquid-Vapor Saturation Properties of Model Systems. *J. Chem. Theory Comput* **2013**, *9*, 2552–2566.
- (25) Virnau, P.; Müller, M. Calculation of free energy through successive umbrella sampling. *J. Chem. Phys.* **2004**, *120*, 10925–10930.
- (26) Valleau, J. P. In *Advances in Chemical Physics*; Prigogine, I., Rice, S., Eds.; John Wiley & Sons, Ltd, 1999; pp 369–404.
- (27) Escobedo, F. A.; Abreu, C. R. A. On the use of transition matrix methods with extended ensembles. *J. Chem. Phys.* **2006**, *124*, 104110.

- (28) Escobedo, F. A.; Martinez-Veracoechea, F. J. Optimization of expanded ensemble methods. *J. Chem. Phys.* **2008**, *129*, 154107.
- (29) Shen, V. K.; Errington, J. R. Determination of fluid-phase behavior using transition-matrix Monte Carlo: Binary Lennard-Jones mixtures. *J. Chem. Phys.* **2005**, *122*, 064508–064508–17.
- (30) Witman, M.; Wright, B.; Smit, B. Simulating Enhanced Methane Deliverable Capacity of Guest Responsive Pores in Intrinsically Flexible MOFs. *J. Phys. Chem. Lett* **2019**, *10*, 5929–5934.
- (31) Mazur, B.; Formalik, F.; Roztocki, K.; Bon, V.; Kaskel, S.; Neimark, A. V.; Firlej, L.; Kuchta, B. Quasicontinuous Cooperative Adsorption Mechanism in Crystalline Nanoporous Materials. *J. Phys. Chem. Lett* **2022**, 6961–6965.
- (32) Datar, A.; Witman, M.; Lin, L.-C. Improving Computational Assessment of Porous Materials for Water Adsorption Applications via Flat Histogram Methods. *J. Phys. Chem. C* **2021**, *125*, 4253–4266.
- (33) Datar, A.; Witman, M.; Lin, L.-C. Monte Carlo simulations for water adsorption in porous materials: Best practices and new insights. *AIChE J* **2021**, *67*, e17447.
- (34) Wang, F.; Landau, D. P. Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States. *Phys. Rev. Lett* **2001**, *86*, 2050–2053.
- (35) Shah, J. K.; Marin-Rimoldi, E.; Mullen, R. G.; Keene, B. P.; Khan, S.; Paluch, A. S.; Rai, N.; Romanielo, L. L.; Rosch, T. W.; Yoo, B.; Maginn, E. J. Cassandra: An open source Monte Carlo package for molecular simulation. *Journal of Computational Chemistry* **2017**, *38*, 1727–1739.
- (36) Mountain, R. D.; Thirumalai, D. Quantative measure of efficiency of Monte Carlo simulations. *Phys. A: Stat. Mech. Appl.* **1994**, *210*, 453–460.

- (37) Kolafa, J. On optimization of Monte Carlo simulations. *Mol. Phys.* **1987**, *63*, 559–579.
- (38) Cortés Morales, A. D.; Economou, I. G.; Peters, C. J.; Ilja Siepmann, J. Influence of simulation protocols on the efficiency of Gibbs ensemble Monte Carlo simulations. *Mol. Simul* **2013**, *39*, 1135–1142.
- (39) Schultz, A. J.; Kofke, D. A. Quantifying Computational Effort Required for Stochastic Averages. *J. Chem. Theory Comput* **2014**, *10*, 5229–5234.
- (40) Paluch, A. S.; Shen, V. K.; Errington, J. R. Comparing the Use of Gibbs Ensemble and Grand-Canonical Transition-Matrix Monte Carlo Methods to Determine Phase Equilibria. *Ind. Eng. Chem. Res* **2008**, *47*, 4533–4541.
- (41) Errington, J. R.; Shen, V. K. Direct evaluation of multicomponent phase equilibria using flat-histogram methods. *J. Chem. Phys.* **2005**, *123*, 164103–164103–9.
- (42) MacDowell, L. G.; Shen, V. K.; Errington, J. R. Nucleation and cavitation of spherical, cylindrical, and slablike droplets and bubbles in small systems. *J. Chem. Phys.* **2006**, *125*, 034705.
- (43) Flyvbjerg, H.; Petersen, H. G. Error Estimates on Averages of Correlated Data. *J. Chem. Phys.* **1989**, *91*, 461–466.
- (44) Grossfield, A.; Patrone, P. N.; Roe, D. R.; Schultz, A. J.; Siderius, D. W.; Zuckerman, D. M. Best Practices for Quantification of Uncertainty and Sampling Quality in Molecular Simulations [Article v1.0]. *Living J. Comput. Mol. Sci.* **2018**, *1*, 5067.
- (45) Berendsen, H. J. C.; Grigera, J. R.; Straatsma, T. P. The missing term in effective pair potentials. *J. Phys. Chem.* **1987**, *91*, 6269–6271.
- (46) MacDowell, L. G.; Virnau, P.; Müller, M.; Binder, K. The evaporation/condensation transition of liquid droplets. *J. Chem. Phys.* **2004**, *120*, 5293–5308.

- (47) Hatch, H. W.; Mahynski, N. A.; Shen, V. K. FEASST: Free Energy and Advanced Sampling Simulation Toolkit. *J. Res. Natl. Inst. Stan* **2018**, *123*, 123004.
- (48) Vlugt, T. J. H.; Martin, M. G.; Smit, B.; Siepmann, J. I.; Krishna, R. Improving the efficiency of the configurational-bias Monte Carlo algorithm. *Mol. Phys* **1998**, *94*, 727–733.
- (49) Chen, B.; Siepmann, J. I. A Novel Monte Carlo Algorithm for Simulating Strongly Associating Fluids: Applications to Water, Hydrogen Fluoride, and Acetic Acid. *J. Phys. Chem. B* **2000**, *104*, 8725–8734.
- (50) Wierchowski, S.; Kofke, D. A. A general-purpose biasing scheme for Monte Carlo simulation of associating fluids. *J. Chem. Phys.* **2001**, *114*, 8752–8762.
- (51) Chen, B.; Siepmann, J. I. Improving the Efficiency of the Aggregation-Volume-Bias Monte Carlo Algorithm. *J. Phys. Chem. B* **2001**, *105*, 11275–11282.
- (52) Chen, B.; Siepmann, J. I.; Oh, K. J.; Klein, M. L. Aggregation-volume-bias Monte Carlo simulations of vapor-liquid nucleation barriers for Lennard-Jonesium. *J. Chem. Phys.* **2001**, *115*, 10903–10913.
- (53) Chen, B.; Siepmann, J. I.; Oh, K. J.; Klein, M. L. Simulating vapor–liquid nucleation of n-alkanes. *J. Chem. Phys.* **2002**, *116*, 4317–4329.
- (54) Siderius, D. W.; Hatch, H. W.; Shen, V. K. Temperature Extrapolation of Henry’s Law Constants and the Isotheric Heat of Adsorption. *J. Phys. Chem. B* **2022**, *126*, 7999–8009.
- (55) Shen, V. K., Siderius, D. W., Kregelberg, W. P., Hatch, H. W., Eds. *NIST Standard Reference Simulation Website, NIST Standard Reference Database Number 173*; National Institute of Standards and Technology: Gaithersburg, MD, 20899, 2006.

- (56) Escobedo, F. A.; Pablo, J. J. d. Expanded grand canonical and Gibbs ensemble Monte Carlo simulation of polymers. *J. Chem. Phys.* **1996**, *105*, 4391–4394.

TOC Graphic

