

Adaptive Markov state model estimation using short reseeding trajectories

Cite as: J. Chem. Phys. 152, 024103 (2020); doi: 10.1063/1.5142457

Submitted: 13 December 2019 • Accepted: 18 December 2019 •

Published Online: 8 January 2020



View Online



Export Citation



CrossMark

Hongbin Wan and Vincent A. Voelz^{a)} 

AFFILIATIONS

Department of Chemistry, Temple University, Philadelphia, Pennsylvania 19122, USA

^{a)}Electronic mail: voelz@temple.edu

ABSTRACT

In the last decade, advances in molecular dynamics (MD) and Markov State Model (MSM) methodologies have made possible accurate and efficient estimation of kinetic rates and reactive pathways for complex biomolecular dynamics occurring on slow time scales. A promising approach to enhanced sampling of MSMs is to use “adaptive” methods, in which new MD trajectories are “seeded” preferentially from previously identified states. Here, we investigate the performance of various MSM estimators applied to reseeding trajectory data, for both a simple 1D free energy landscape and mini-protein folding MSMs of WW domain and NTL9(1–39). Our results reveal the practical challenges of reseeding simulations and suggest a simple way to reweight seeding trajectory data to better estimate both thermodynamic and kinetic quantities.

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5142457>

I. INTRODUCTION

In the last decade, Markov State Model (MSM) methodologies have made possible accurate and efficient estimation of kinetic rates and reactive pathways for slow and complex biomolecular dynamics.^{1–5} One of the key advantages touted by MSM methods is the ability to use large ensembles of short-time scale trajectories for sampling events that occur on slow time scales. The main idea is that sufficient sampling using many short trajectories can circumvent the need to sample long trajectories.

With this in mind, many “adaptive” methods have been developed for the purpose of accelerating sampling of MSMs. The simplest of these can be called *adaptive seeding*, where one or more new rounds of unbiased simulations are performed by “seeding” swarms of trajectories throughout the landscape.⁶ The choice of seeds is based on some initial approximation of the free energy landscape, possibly from nonequilibrium or enhanced-sampling methods. Adaptive seeding can be performed by first identifying a set of metastable states and then initiating simulations from each state. If the seeding trajectories provide sufficient connectivity and statistical sampling of transition rates, an MSM can be constructed to accurately estimate both kinetics and thermodynamics.

Similarly, so-called *adaptive sampling* algorithms have been developed for MSMs in which successive rounds of targeted seeding simulations are performed, updating the MSM after each round.^{7,8} A simple adaptive sampling strategy is to start successive rounds of simulations from undersampled states, for instance, from the state with the least number of transition counts.⁹ A more sophisticated approach is the FAST algorithm, which is designed to discover states and reactive pathways of interest by choosing new states based on an objective function that balances undersampling with a reward for sampling desired structural observables.^{10,11} Other algorithms include REAP, which efficiently explores folding landscapes by using reinforcement learning to choose new states,¹² and surprisal-based sampling,⁷ which chooses new states that will minimize the uncertainty of the relative entropy between two or more MSMs.

A key problem with adaptive sampling of MSMs arises because we are often interested in *equilibrium* properties, while trajectory seeding is decidedly *nonequilibrium*. This may seem like a subtle point, because the dynamical trajectories themselves are unbiased, but of course, the ensemble of starting points for each trajectory are almost always *statistically* biased; i.e., the seeds are not drawn from the true equilibrium distribution. This can be problematic because most MSMs are constructed from transition rate estimators that

enforce detailed balance and assume trajectory data are obtained at equilibrium. The distribution of sampled transitions, however, will only reflect equilibrium conditions in the limit of long trajectory length.

One way around this problem is to focus mostly on the kinetic information obtained by adaptive sampling. A recent study of the ability of FAST to accurately describe reactive pathways concluded that the most reliable MSM estimator to use with adaptive sampling data is a row-normalized transition count matrix.¹¹ Indeed, weighted-ensemble path sampling algorithms focus mainly on sampling the kinetics of reactive pathways, information which can be used to recover global thermodynamic properties.^{13–17} A major disadvantage of this approach is that it ignores potentially valuable equilibrium information. As shown by Trendelkamp-Schroer and Noé,¹⁸ detailed balance is a powerful constraint to infer rare-event transition rates from equilibrium populations. Specifically, when faced with limited sampling, dedicating half of one's simulation samples toward enhanced thermodynamic sampling (e.g., umbrella sampling) can result in a significant reduction in the uncertainty of estimated rates, simply because the improved estimates of equilibrium state populations inform the rate estimates through detailed balance.

Another way around this problem, recently described by Nüske *et al.*, is to use an estimator based on observable operator model (OOM) theory, which utilizes information from transitions observed at lag times τ and 2τ to obtain estimates unbiased by the initial distribution of seeding trajectories.¹⁹ Although the OOM estimator is able to make better MSM estimates at shorter lag times, it requires the storage of transition count arrays that scale as the cube of the number states and dense-matrix singular-value decomposition, which can be impractical for MSMs with large numbers of states. Nüske *et al.* derived an expression quantifying the error incurred by nonequilibrium seeding, from which they conclude that such bias is difficult to remove without either increasing the lag time or improving the state discretization.

Here, we explore an alternative way to recover accurate MSM estimates from biased seeding trajectories, by reweighting sampled transition counts to better approximate counts that would be observed at equilibrium. Like the Trendelkamp-Schroer and Noé method,¹⁸ this requires some initial estimate of state populations, perhaps obtained from the previous rounds of adaptive sampling.

We are particularly interested in examining how the performance of this reweighting method compares with other estimators, in cases where it is impractical to generate long trajectories and instead one must rely on ensembles of short seeding trajectories. An example of a case like this is adaptive seeding of protein folding MSMs built from ultralong trajectories simulated on the Anton supercomputer.²⁰ Because such computers are not widely available, adaptive seeding using conventional computers may be one of the only practical ways to leverage MSMs to predict the effect of mutations, for example.

In this manuscript, we first perform adaptive seeding tests using 1D-potential energy models and compare how different estimators perform at accurately capturing kinetics and thermodynamics. We then perform similar tests for MSMs built from ultralong reversible folding trajectories of two mini-proteins, WW domain and NTL9(1–39). Our results, described below, suggest that reweighting

trajectory counts with estimates of equilibrium state populations can achieve a good balance of kinetic and thermodynamic estimation.

A. Estimators

We explored the accuracy and efficiency of several different transition probability estimators using adaptive seeding trajectory data as input: (1) a maximum-likelihood estimator (MLE), (2) an MLE estimator where the exact equilibrium populations π_i of each state i are known *a priori*, (3) an MLE estimator where each input trajectory is weighted by an *a priori* estimate of the equilibrium population of its starting state, (4) row-normalized transition counts, and (5) an observable operator model (OOM) estimator.

B. Maximum-likelihood estimator (MLE)

The MLE for a reversible MSM assumes that observed transition counts are independent and drawn from the equilibrium distribution so that reversibility (i.e., detailed balance) can be used as a constraint. The likelihood of observing a set of given transition counts, $L = \prod_i \prod_j p_{ij}^{c_{ij}}$, when maximized under the constraint that $\pi_i p_{ij} = \pi_j p_{ji}$, for all i, j , yields a self-consistent expression that can be iterated to find the equilibrium populations,^{3,21,22}

$$\pi_i = \sum_j \frac{c_{ij} + c_{ji}}{\frac{N_j}{\pi_j} + \frac{N_i}{\pi_i}}, \quad (1)$$

where $N_i = \sum_j c_{ij}$. The transition probabilities p_{ij} are given by

$$p_{ij} = \frac{(c_{ij} + c_{ji})\pi_j}{N_j\pi_i + N_i\pi_j}. \quad (2)$$

C. Maximum-likelihood estimator (MLE) with known populations π_i

Maximization of the likelihood function above, with the additional constraint of fixed populations π_i , yields a similar self-consistent equation that can be used to determine a set of Lagrange multipliers,¹⁸

$$\lambda_i = \sum_j \frac{(c_{ij} + c_{ji})\pi_j\lambda_i}{\lambda_j\pi_i + \lambda_i\pi_j}, \quad (3)$$

from which the transition probabilities p_{ij} can be obtained as

$$p_{ij} = \frac{(c_{ij} + c_{ji})\pi_j}{\lambda_j\pi_i + \lambda_i\pi_j}. \quad (4)$$

D. Maximum-likelihood estimator (MLE) with population-weighted trajectory counts

For this estimator, first a modified count matrix c'_{ij} is calculated,

$$c'_{ij} = \sum_k w^{(k)} c_{ij}^{(k)}, \quad (5)$$

where transition counts $c_{ij}^{(k)}$ from trajectory k are weighted in proportion to $w^{(k)} = \pi^{(k)}$, the estimated equilibrium population of the initial state of the trajectory. The idea behind this approach is to counteract the statistical bias from adaptive seeding by scaling the

observed transition counts proportional to their equilibrium fluxes. The modified counts are then used as input to the MLE.

E. Row-normalized counts

For this estimator, the transition probabilities are approximated as

$$p_{ij} = \frac{c_{ij}}{\sum_j c_{ij}}. \quad (6)$$

This approach does not guarantee reversible transition probabilities, which only occurs in the limit of large numbers of reversible transition counts. In practice, however, the largest eigenvectors of the transition probability matrix have very nearly real eigenvalues such that we can report relevant relaxation time scales and equilibrium populations.

F. Observable operator model (OOM) theory

For an introduction to OOMs and their use in estimating MSMs, see Refs. 19, 23, and 24. OOMs are closely related to Hidden Markov Models (HMMs) but have the advantage, like MSMs, that they can be learned directly from observable-projected trajectory data.²⁴ Unlike MSMs, they require the collection of both one-step transition count matrices and a complete set of two-step transition count matrices. With enough training data and sufficient model rank, OOMs can recapitulate exact relaxation time scales, uncontaminated by MSM state discretization error.¹⁹

We used the OOM estimator as described by Nüske *et al.*¹⁹ and implemented in PyEMMA 2.3,²⁵ using the default method of choosing the OOM model rank by discarding singular values that contribute a bootstrap-estimated signal-to-noise ratio of less than 10. The settings used to instantiate the model in PyEMMA are as follows:

```
OOMReweightedMSM(reversible=True,  
count_mode='sliding', sparse=False,  
connectivity='largest',  
rank_Ct='bootstrap_counts', tol_rank=10.0,  
score_method='VAMP2', score_k=10,  
mincount_connectivity='1/n').
```

The OOM estimator returns two estimates: (1) a so-called *corrected* MSM, which derives from using the unbiased OOM equilibrium correlation matrices to construct a matrix of MSM transition probabilities, and (2) the *OOM time scales*, which (given sufficient training data) estimate the relaxation time scales uncontaminated by discretization error. Note that there is no corresponding MSM for the OOM time scales.

II. RESULTS

A. A simple example of adaptive seeding estimation error

To illustrate some of the estimation errors that can arise with adaptive seeding, we first consider a simple model of 1D diffusion over a potential energy surface $V(x)$. The continuous time evolution of a population distribution $p(x, t)$ is analytically described by the Fokker-Planck equation,

$$\frac{\partial}{\partial t} p(x, t) = D \frac{\partial^2}{\partial x^2} p(x, t) + \mu \frac{\partial}{\partial x} [p(x, t) \nabla V(x)], \quad (7)$$

where D is the diffusion constant and $\mu = D/k_B T$ is the mobility. Here, x and t are unitless values with D set to unity and the energy function $V(x)$ in units of $k_B T$.

Here, we consider a “flat-bottom” landscape, over a range $x \in (0, 2)$ with periodic boundaries, described by

$$V(x) = \frac{e^{-b(|x-1|-0.5)}}{1 + e^{-b(|x-1|-0.5)}}, \quad (8)$$

where $b = 100$. The landscape is partitioned into two states: x is assigned to state 0 if $x < 0.5$ or $x \geq 1.5$ and to state 1 otherwise [Fig. 1(a)]. In the limit of $b \rightarrow \infty$, the free energy difference between the states is exactly 1 ($k_B T$), with an equilibrium constant of $K = e \approx 2.718$. When $b = 100$, $K \approx 2.643$. This value of b was chosen to help the stability of numerical integration (performed at a resolution of $\Delta x = 2.5 \times 10^{-3}$ and $\Delta t = 5 \times 10^{-7}$).

To emulate adaptive seeding, we propagate the continuous dynamics starting from an initial distribution $p(x, 0)$ centered on state i and compute the transition probabilities p_{ij} from state i to j from the evolved density after some lag time τ [Fig. 1(b)].

We first consider a seeding strategy where simulations are started from the center of each state, which we model by propagating density from an initial distributions $p(x, 0) = \delta(x)$ to compute p_{00} and p_{01} as a function of τ and from $p(x, 0) = \delta(x - 1)$ to compute p_{10} and p_{11} [Fig. 1(c)]. By $t = 1$, the population has reached equilibrium. By detailed balance, the estimated equilibrium constant is $\hat{K} = p_{10}/p_{01}$, which systematically *underestimates* the true value [Fig. 1(d)]. This is because the initial distribution is nonequilibrium, in the direction of being too uniform. There are also errors in estimating the relaxation time scales. The different estimators described above (MLE, MLE with known equilibrium populations π_i , population-weighted MLE, and row-normalized counts) were applied to estimate the implied time scale τ_2 [Fig. 1(e)]. For this simple two-state system, all estimates are very similar to the analytical two-state result $\tau_2 = -\tau/\ln(1 - p_{10} - p_{01})$ (with the exception of MLE, which yields τ_2 estimates that are about 9% larger when $t = 0.1$), which approaches the true slowest relaxation time scale of the Fokker-Planck dynamics ($\tau_2 = 0.1013$) as the lag time increases. Until the population density becomes equilibrated, however, the computed value of τ_2 is *overestimated*. This is because at early times, outgoing population fluxes are underestimated.

We also consider a seeding strategy where simulations are started from randomly from inside each state, which we model by propagating density from initially uniform distributions across each state [Fig. 1(f)]. As before, we find that the equilibrium constant is underestimated, again, because the initial seeding is out of equilibrium [Fig. 1(g)]. The estimate is somewhat improved, however, by the “pre-equilibration” of each state. In this case, the implied time scale τ_2 is *underestimated* at early times [Fig. 1(h)]. This can be understood as arising from the *overestimation* of outgoing population fluxes, as not enough population has yet reached the ground state (state 0).

From these two scenarios, we can readily understand that the success of estimating both kinetics and thermodynamics from adaptive seeding greatly depends on how well the initial seeding distribution approximates the equilibrium distribution. Moreover, we can see that time scale estimates can be either underestimated or overestimated, depending on the initial seeding distribution. It is

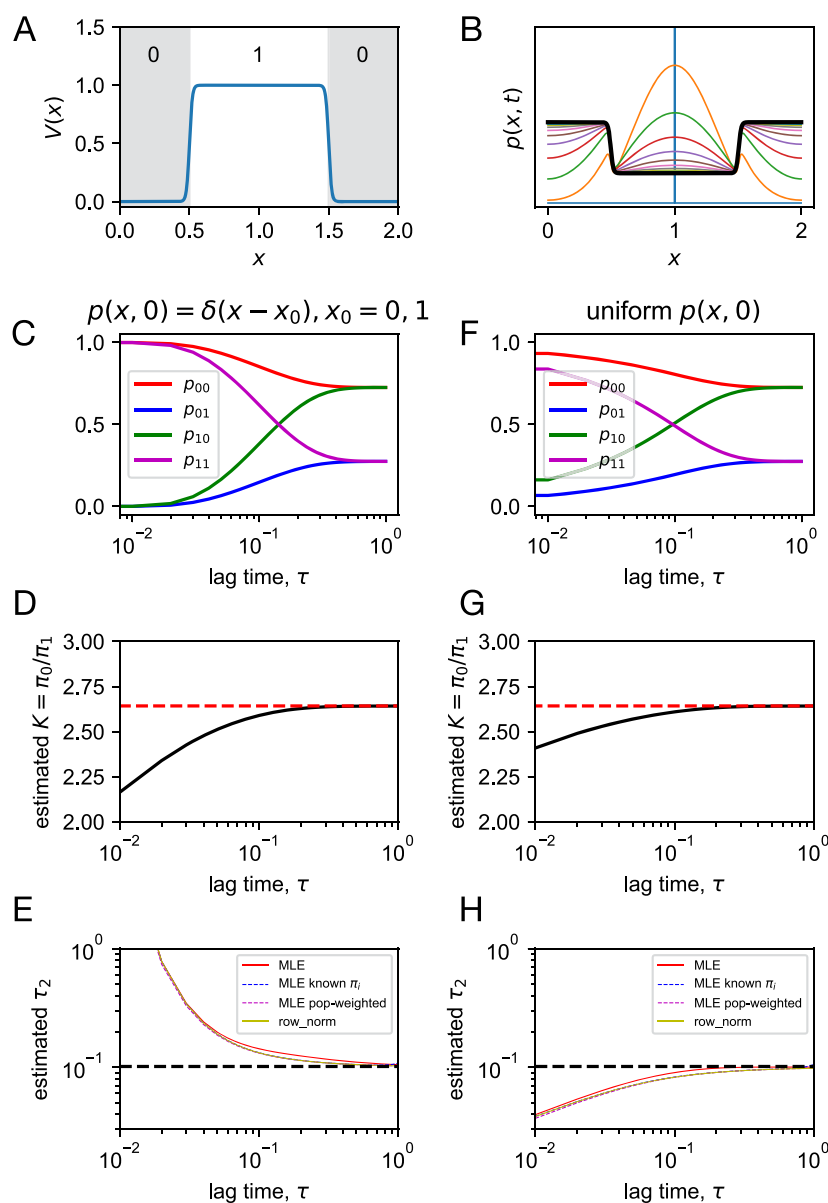


FIG. 1. Adaptive seeding in a simple two-state system. (a) The potential energy surface $V(x)$, annotated with state indices. (b) Propagation of an initial distribution $p(x, 0) = \delta(x - 1)$ (blue), plotted at increments $\Delta t = 0.1$, up to $t = 1$. Starting from an initial distribution $p(x, 0) = \delta(x - x_0)$ using $x_0 = 0, 1$, shown are estimated quantities as a function of lag time τ : (c) transition probabilities p_{ij} , (d) equilibrium constant $K = \pi_0/\pi_1$, and (e) relaxation time scale τ_2 , as computed by various estimators. Starting from uniform distributions $p(x, 0)$ over each state, shown are estimated quantities as a function of lag time τ : (f) transition probabilities p_{ij} , (g) equilibrium constant $K = \pi_0/\pi_1$, and (h) relaxation time scale τ_2 , as computed by various estimators.

important to note that these errors arise primarily because of the nonequilibrium sampling rather than finite sampling error or discretization error.³

B. Seeding of a 1-D potential energy surface

We next consider the following two-well potential energy surface, as used by Stelzl *et al.*:²⁶ $U(x) = -\frac{2k_B T}{0.596} \ln[e^{-2(x-2)^2-2} + e^{-2(x-5)^2}]$ for $x \in [1.5, 5.5]$ and $k_B T = 0.596$ kcal mol⁻¹. The state space is uniformly divided into 20 states of width 0.2 to calculate discrete-state quantities. Diffusion on the 1-D landscape is approximated by a Markov Chain Monte Carlo (MCMC) procedure in which new

moves are translations randomly chosen from $\delta \in [-0.05, +0.05]$ and accepted with probability $\min(1, \exp(-\beta[U(x+\delta) - U(x)]))$, i.e., the Metropolis criterion (Fig. 2).

As a standard against which to compare estimates from seeding trajectories, we used a set of very long trajectories to construct an optimal MSM model using the above state definitions. To estimate the “true” relaxation time scales of the optimal model, we generated long MCMC trajectories of 10^9 steps, sampling from a series of scaled potentials $U^{(\lambda)}(x) = \lambda U(x)$ for $\lambda \in [0.5, 0.6, 0.7, 0.8, 0.9, 1.0]$. For each λ value, 20 trajectories were generated, with half of them starting from $x = 2.0$ and the other half starting from $x = 5.0$, resulting in a total of 120 trajectories. The DTRAM estimator of Wu

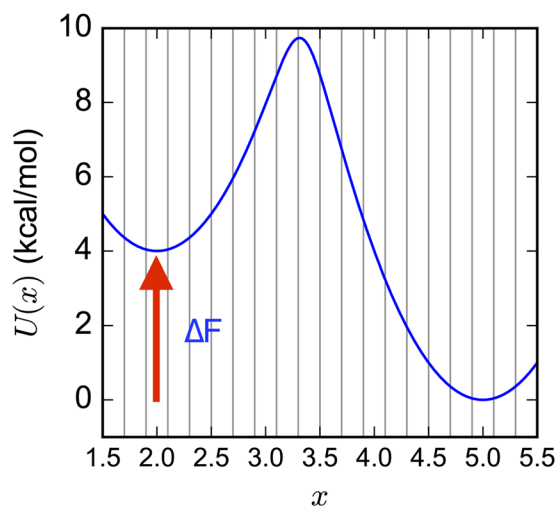


FIG. 2. A 1-D two-state potential was used to test the performance of MSM estimators on seeding trajectory data. Vertical lines denote the 20 uniformly spaced discrete states.

*et al.*²¹ was used to estimate the slowest relaxation time scale as $9.66 (\pm 1.37) \times 10^6$ steps, using a lag time of 1000 steps.

We next generated adaptive seeding trajectory data for this system. We limited the lag time to $\tau = 100$ steps and generated s trajectories ($5 \leq s \leq 1000$) of length $n\tau$ (for $n = 10, 100, 1000$) by initiating MCMC dynamics from the center positions of all 20 bins. The resulting data consisted of $20 \times s \times n$ transition counts between all states i and j in lag time τ , stored in a 20×20 count matrix of entries c_{ij} . From these counts, estimates of the transition probabilities, p_{ij} , were made.

Estimates of the slowest MSM implied time scales were computed as $\tau_2 = -\tau/\ln \mu_2$, where μ_2 is the largest nonstationary eigenvalue of the matrix of transition probabilities. Estimates of the free energy difference ΔF between the two wells were computed as $-k_B T \ln(\pi_L/(1 - \pi_L))$, where π_L is the estimated total equilibrium population of MSM states with $x < 3.3$. To estimate uncertainties in $\ln \tau_2$ and ΔF , 20 bootstraps were constructed by sampling from the s trajectories with replacement.

Results for the five different estimators are shown in Fig. 3. In all cases, time scale estimates improve with larger numbers of seeding trajectories and greater numbers of steps in each trajectory. From these results, one can clearly see that MLE (red lines in

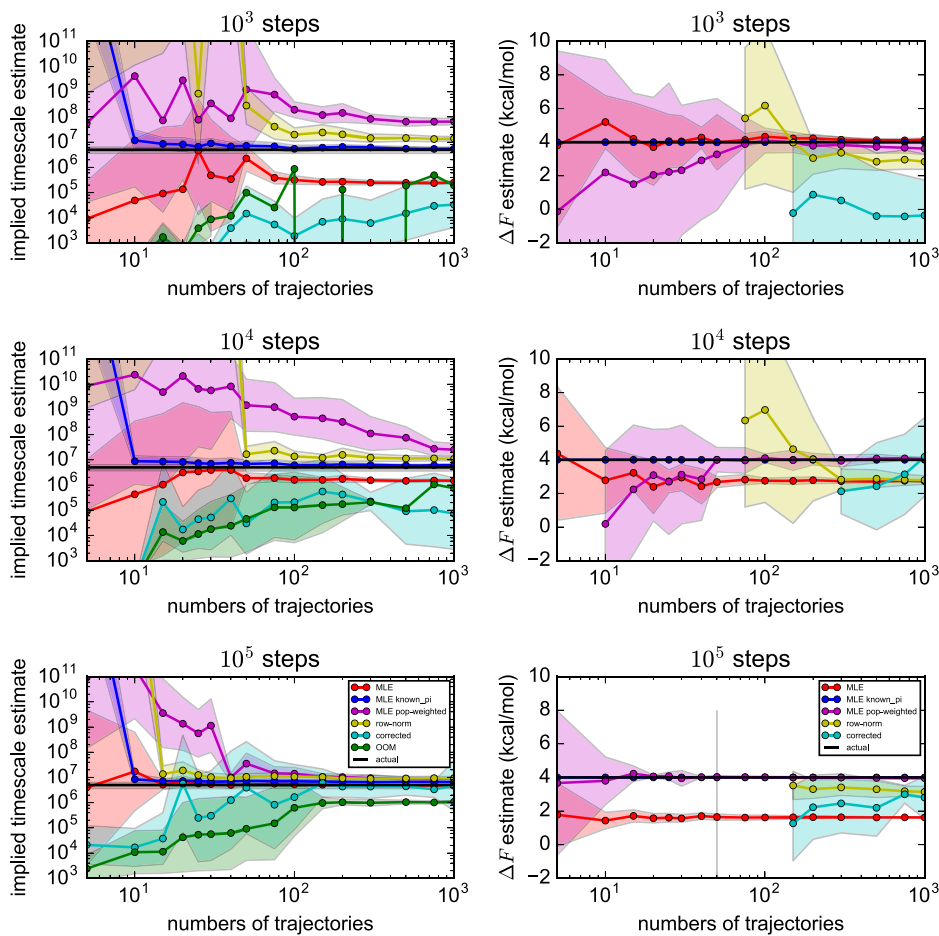


FIG. 3. Estimates of the slowest implied time scales (left) and two-well free energy differences (right), shown as a function of number of seeding trajectories of various lengths. Shown are results for (red) MLE, (blue) MLE with known equilibrium populations, (magenta) MLE with population-weighted trajectory counts, (yellow) row-normalized counts, (cyan) corrected MSMs, and (green) OOM time scales. Black horizontal lines denote the highly converged estimates from DTRAM. Shaded bounds represent uncertainties estimated from a bootstrap procedure.

Fig. 3), the default estimator in software packages like PyEMMA²⁵ and MSMBuilder,²⁷ consistently underestimates the implied time scale, although this artifact becomes less severe as more trajectory data is incorporated. At the same time, MLE estimates of equilibrium populations become systematically worse with more trajectory data. This is precisely because uniform seeding produces a biased, *nonequilibrium* distribution of sampled trajectories, but the MLE assumes that trajectory data are drawn from *equilibrium*. In contrast, the MLE with known equilibrium populations (blue lines in Fig. 3) accurately predicts implied time scales with much less trajectory data, underscoring the powerful constraint that detailed balance provides in determining rates. Of all the estimators, this method is the most accurate but relies on having very good estimates of equilibrium state populations *a priori*.

The row-normalized counts estimator (yellow lines in Fig. 3) estimates time scales well, provided that a sufficient number of observed transitions must be sampled to obtain proper estimates. This is necessary to achieve connectivity and to ensure that the transition matrix is at least approximately reversible. However, the free energy differences estimated by row-normalized counts are highly uncertain, and in the limit of large numbers of sufficiently long trajectories (10^5 steps), estimates of free energies become systematically incorrect due to the statistical bias from seeding. In comparison, the MLE with population-weighted trajectory counts (magenta lines in Fig. 3) is able to make much more accurate estimates of free energies with less trajectory data. We note that while here we are reweighting the transition counts with the known populations, the results are similar if reasonable approximations of the state populations are used (data not shown). Compared with the row-normalized counts estimator, implied time scales are slower to converge with more trajectory data but are comparable when the seeding trajectories are sufficiently long (10^5 steps). With sufficiently long trajectories (10^5 steps), the MLE time scales are better estimated than either the row-normalized counts estimator or the population-weighted trajectory counts, but MLE is flawed because the free energies are so biased from the seeding. The population-weighting trajectory counts estimator avoids this artifact to give good estimates of free energies.

The “corrected” and OOM time scales from the OOM estimator consistently underestimate the slowest implied time scale for this test system, although these estimates improve with greater numbers and lengths of trajectories (cyan and green lines in Fig. 3). The OOM estimator is susceptible to statistical noise, as evidenced by the average OOM model rank selected by the PyEMMA implementation, which increases from 2 to 10 as the numbers of trajectories are increased (Fig. 4). Like the row-normalized counts estimator, the OOM estimator gives poorly converged estimates of two-well free energy differences.

C. Adaptive seeding of folding landscapes for WW domain and NTL9(1–39)

For studying protein folding, an adaptive seeding strategy may be useful for several reasons. For one, generating ensembles of short trajectories may simply be a practical necessity, due to the limited availability of special-purpose hardware to generate ultralong trajectories. Adaptive seeding may also be able to efficiently

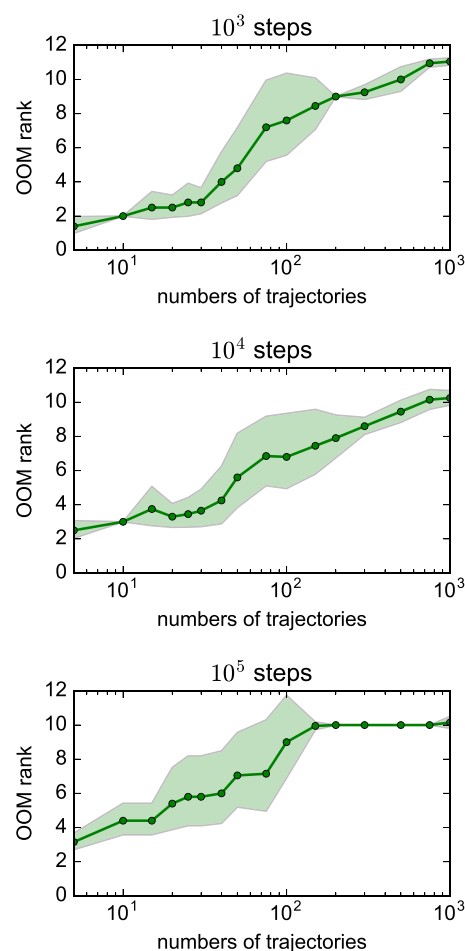


FIG. 4. The average OOM rank increases with the number of trajectories. Shaded bounds represent uncertainties estimated from a bootstrap procedure.

leverage an existing MSM (perhaps built from ultralong trajectory data) to predict perturbations to folding by mutations or to examine the effect of different simulation parameters like force field or temperature.

How can we best utilize adaptive seeding of MSMs built from ultralong trajectory data to make good estimates of both kinetics and thermodynamics? To address this question, here, we test the performance of various MSM estimators on the challenging task of estimating MSMs from adaptive seeding of protein folding landscapes. The folding landscapes come from MSMs built from ultralong reversible folding trajectories of fast-folding mini-proteins WW domain and NTL9(1–39).

WW domain is a 35-residue protein signaling domain with a three-stranded β -sheet structure that binds to proline-rich peptides. Its folding kinetics and thermodynamics have been studied extensively by time-resolved spectroscopy,^{28–32} and its folding mechanism has been probed extensively using molecular simulation studies.^{33–35} Moreover, an impressively large number of site-directed mutants of WW domain have been constructed to investigate folding

mechanism and discover fast-folding variants, including the Fip mutant (Fip35) of human pin1 WW domain, for which mutations in first hairpin loop region result in a fast folding time of $\sim 13.3 \mu\text{s}$ at 77.5°C .^{29,36} Further mutation of the second hairpin loop of Fip35 produced an even faster-folding variant, called GTT, in which the native sequence Asn-Ala-Ser (NAS) was replaced with Gly-Thr-Thr (GTT), resulting in a fast relaxation time $\sim 4.3 \mu\text{s}$ at 80°C .³⁵

NTL9(1-39) domain is a 39-residue truncation variant of the N-terminal domain of ribosomal protein L9, whose folded state consists of an α -helix and a three-strand β -sheet. NTL9(1-39), which has a folding relaxation time of $\sim 1.5 \text{ ms}$ at 25°C ,³⁷ has been extensively probed by both experimental and computational studies.³⁸⁻⁴¹ The K12M mutant of NTL9(1-39) is a fast-folding variant with a folding time scale of $\sim 700 \mu\text{s}$ at 25°C (Figs. 5 and 6).^{37,42}

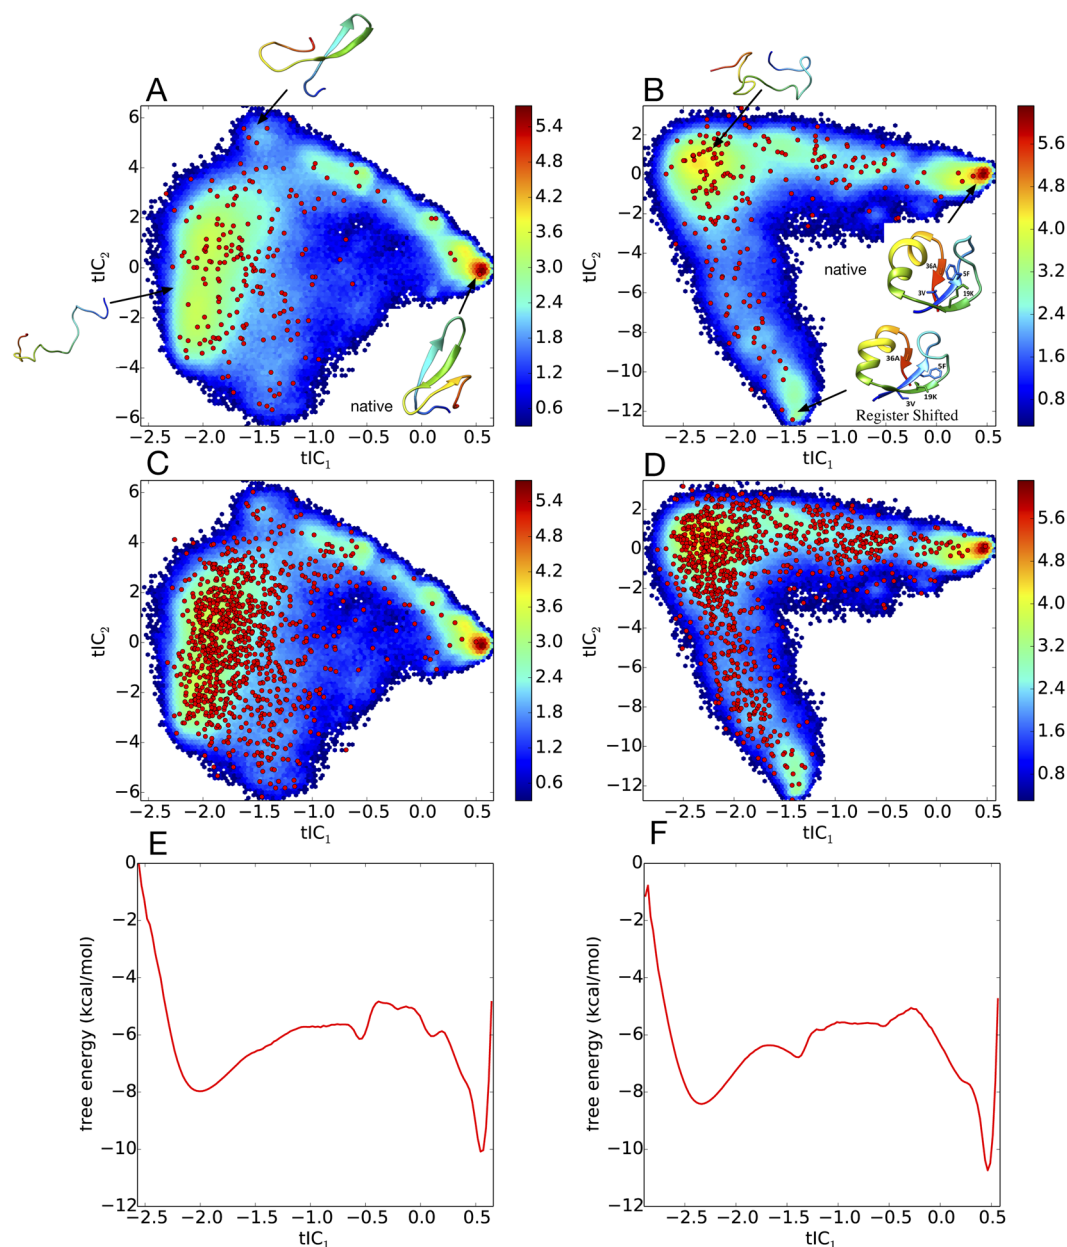


FIG. 5. 200-state MSMs of (a) GTT WW domain and (b) K12M NTL9(1-39), built from ultralong folding trajectory data, are shown projected onto the two largest tICA components. Red circles show the locations of the MSM microstates. The heat map shows a density plot of the raw trajectory data, with color bar values denoting the natural logarithm of histogram counts. 1000-state MSMs are shown for (c) GTT WW domain and (d) K12M NTL9(1-39). (e) and (f) Corresponding 1-D free energy profiles along the first tICA component, calculated from histogramming the trajectory data in bins of width 0.025.

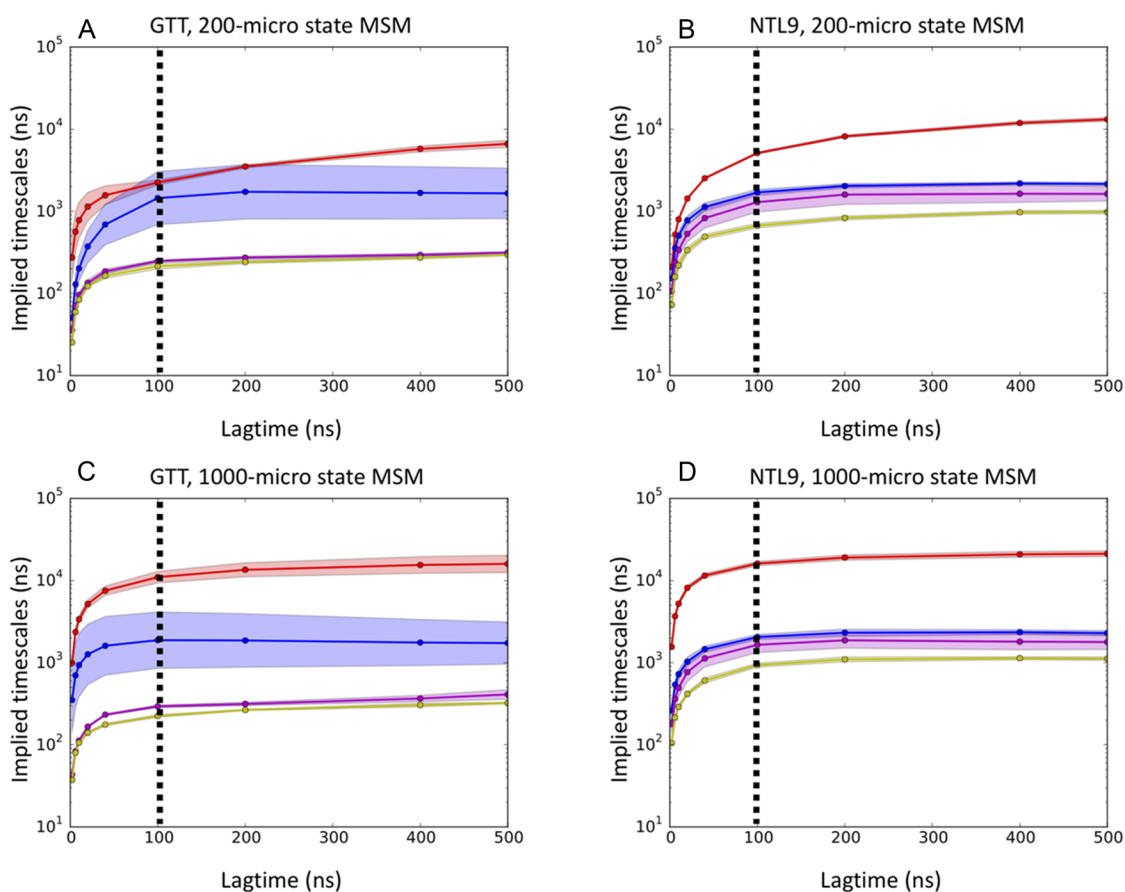


FIG. 6. Implied time scale plots for 200-state MSMs of (a) GTT WW domain and (b) K12M NTL9(1–39) and for 1000-state MSMs of (c) GTT WW domain and (d) K12M NTL9(1–39). Uncertainties (shaded areas) for GTT and NTL9 were estimated using 5-fold and 8-fold leave-one-out bootstraps, respectively. The dotted vertical line marks the lag time of 100 ns chosen to construct the MSMs.

D. Markov state model construction

Molecular simulation trajectory data for GTT WW domain and K12M NTL9(1–39) were provided by Shaw *et al.*,²⁰ Our first objective was to construct high-quality reference MSMs from the available trajectory data, to be used as benchmarks against which we could compare adaptive seeding results.

Two independent ultralong trajectories of 651 and 486 μ s, performed at 360K using the CHARMM22* force field, were used to construct MSMs of GTT WW domain, as previously described by Wan and Voelz.⁴³ The MSM was constructed by first projecting the trajectory data to all pairwise distances between C_α and C_β atoms, performing dimensionality reduction via tICA (time-lagged independent component analysis),^{44,45} and then clustering in this lower-dimensional space using the k -centers algorithm to define a set of metastable states. A 1000-state MSM for GTT WW domain was constructed using a lag time of 100 ns. The generalized matrix Raleigh quotient (GMRQ) method⁴⁶ was used to choose the optimal number of states (1000) and the number of independent tICA components

(eight). The MSM gives an estimated folding relaxation time scale of 10.2 μ s, which is comparable with the folding time of $21 \pm 6 \mu$ s estimated from the analysis of the trajectory data by Lindorff-Larsen *et al.*²⁰ It is also comparable with the 8 μ s time scale estimate from a three-state MSM model built using sliding constraint rate estimation by Beauchamp *et al.*⁴⁷ As described below, since seeding trajectory data for a 1000-state MSM were too expensive to analyze using the OOM estimator, we additionally built a 200-state model, using the same methods, which gives an estimated folding time scale of 2.3 μ s.

Four trajectories of the K12M mutant of NTL9(1–39), of lengths 1052 μ s, 990 μ s, 389 μ s, and 377 μ s, simulated at 355 K using the CHARMM22* force field, were provided by Shaw *et al.*²⁰ Choosing a lag time of 200 ns, we constructed an MSM using the same methodology as GTT WW domain described above, resulting in a 1000-state model utilizing 6 tICA components, which gives a folding time scale estimate of $\sim 18 \mu$ s. Although this is a faster time scale than the experimentally measured value at 25 °C, it accurately reflects the time scale of folding events observed in the trajectory

data (at 355 K) and is similar to folding time scales measured by T-jump 2-D IR spectroscopy at 80 °C ($\sim 26.4 \mu\text{s}$).⁴¹ We also note that the previous analysis of the K12M NTL9(1–39) trajectory data has given similar time scale estimates. Analysis of the trajectory data by Lindorff-Larsen *et al.* yielded an estimate of $\sim 29 \mu\text{s}$,²⁰ a three-state MSM model built using sliding constraint rate estimation by Beauchamp *et al.* gives an estimate of $\sim 16 \mu\text{s}$,⁴⁶ and an MSM by Baiz *et al.* gives a time scale estimate of $18 \mu\text{s}$.⁴¹ As we did for GTT WW domain, and using the same methods, we additionally built a 200-state MSM, which gives an estimated folding time scale of $5.2 \mu\text{s}$.

E. Estimated folding rates and equilibrium populations from reseeded simulations

To emulate data sets that would be obtained by adaptive reseeded, we randomly sampled segments of the reference trajectory data that start from each metastable state. An advantage of emulating reseeded trajectory data in this way is the ability to recapitulate the original model in the limit of long trajectories. The full set of reference trajectory data (~ 1.1 ms) for GTT WW domain composed of 12 reversible folding events, while the full set for K12M NTL9(1–39) (~ 2.8 ms) composed of 14 reversible folding events.

To test the performance of different estimators on reseeded trajectory data, we sampled either 5 or 10 seeding trajectories of various lengths from each MSM state. By design, the length of the reseeded trajectories was limited to 500 ns for the 200-state MSMs and to 250 ns for the 1000-state MSMs, because trajectories longer than this length result in data sets that exceed the total amount of reference trajectory data. For reseeded 1000-state MSMs with 10 trajectories, the trajectory lengths were limited to 100 ns for GTT and to 200 ns for NTL9 so as not to exceed the total amount of original trajectory data.

For each set of emulated reseeded trajectory data, MSMs were constructed using various estimators and compared with the reference MSM. Implied time scale estimates were computed using a lag time of 40 ns to facilitate the analysis of seeding trajectories as short as 50 ns. For the MLE estimator with population reweighting and the MLE estimator with known populations, we used the equilibrium populations estimated from the reference MSMs. In cases where the row-normalized counts estimator fails at large lag times due to the disconnection of states, a single pseudocount was added to the elements of the transition count matrix that are nonzero in the reference MSM.

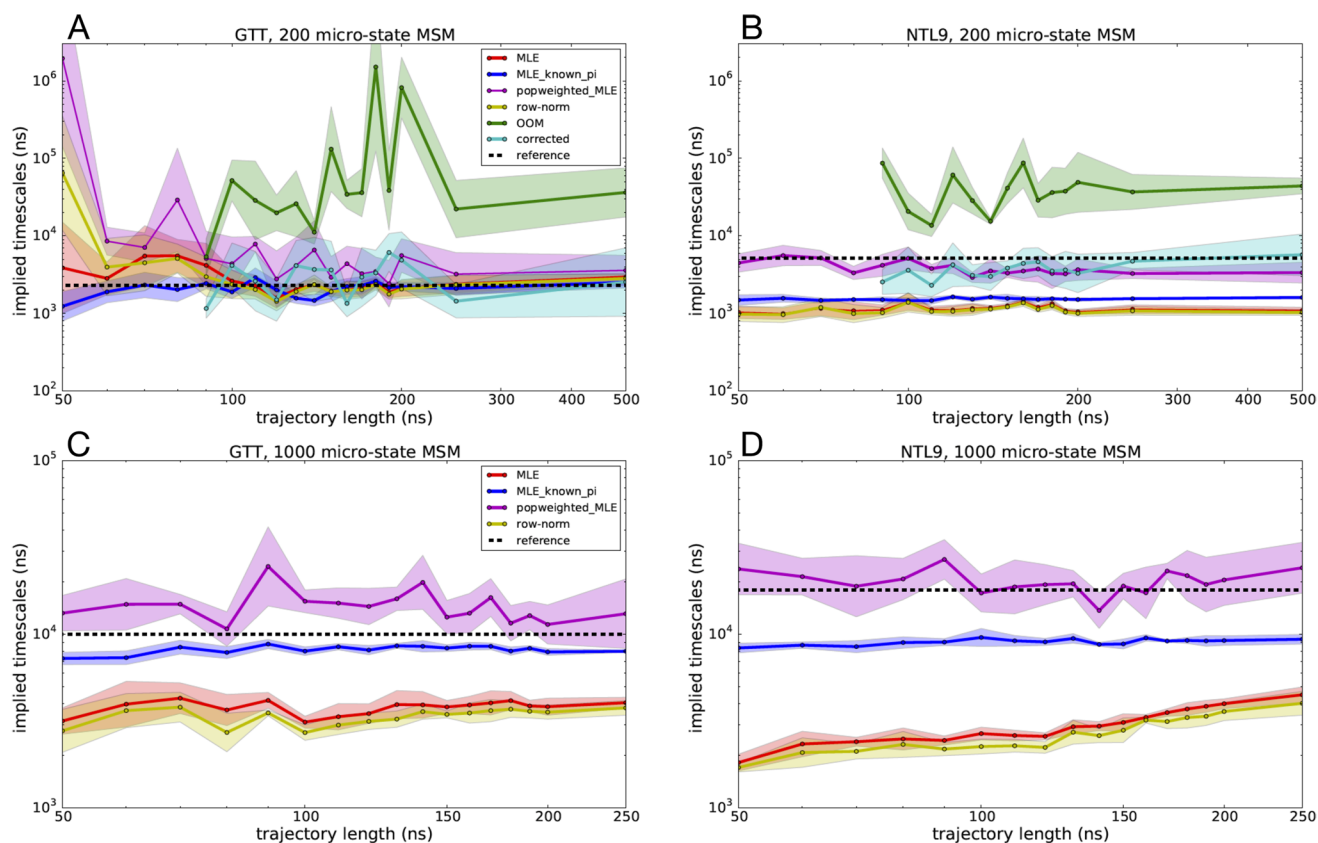


FIG. 7. Slowest implied time scales predicted by various MSM estimators as a function of seeding trajectory length, where five independent trajectories were started from each microstate. Estimated slowest time scales are shown for 200-microstate MSMs of (a) GTT WW domain and (b) K12M NTL9(1–39) and for 1000-microstate MSMs of (c) GTT WW domain and (d) K12M NTL9(1–39). Uncertainties were computed using a bootstrap procedure.

1. Estimates of slowest implied time scales

We applied different estimators to the reseeded trajectory data generated for 200-state and 1000-state MSMs of GTT WW domain and K12M NTL9(1–39) and compared the predicted slowest implied time scales to the reference MSMs. Similar implied time scale estimates were obtained using five independent reseeded trajectories initiated from each state (Fig. 7) and ten independent reseeded trajectories (Fig. 8).

Across most of the estimators tested, implied time scale estimates are statistically well converged using trajectories of 100 ns or more. For the MLE-based estimators in particular (MLE, MLE with known populations, MLE with population reweighting), the estimates have small uncertainties and are not strongly sensitive to trajectory lengths tested. This suggests that while there is enough seeding data to combat statistical sampling error, there remain systematic errors in estimated time scales, mainly due to the estimation method used.

The MLE performs well for the 200-state MSMs of GTT but consistently underestimates the slowest implied time scale in all other MSMs. As seen with our tests above with the 1-D two-state potential, this error is likely due to using a detailed balance

constraint with statistically biased seeding trajectory data. The MLE with known populations is much more accurate at estimating the slowest implied time scale and with less uncertainty. For the NTL9 MSMs, however, the MLE with known populations systematically underestimates the slowest implied time scale, again suggesting that the biased seeding trajectory data have an influence despite the strong constraints enforced by the estimator. MLE with population reweighting has more uncertainty than the other MLE-based estimators but overall is arguably the most accurate estimator across all the MSM reseeded tests.

The OOM-based estimates of the slowest implied time scale (which could only be performed for the computationally tractable 200-state MSMs) required seeding trajectory data of least 90 ns, in order to have sufficient trajectory data to perform the bootstrapped signal-to-noise estimation, from which the OOM rank is selected. While the predicted OOM time scales overestimate the relaxation time scales of the 200-state MSMs of both GTT and NTL9 systems, the time scale estimates from the corrected MSMs more closely recapitulate the reference MSMs time scales, with accuracy comparable with the MLE with population reweighting, especially for seeding trajectories over 200 ns.

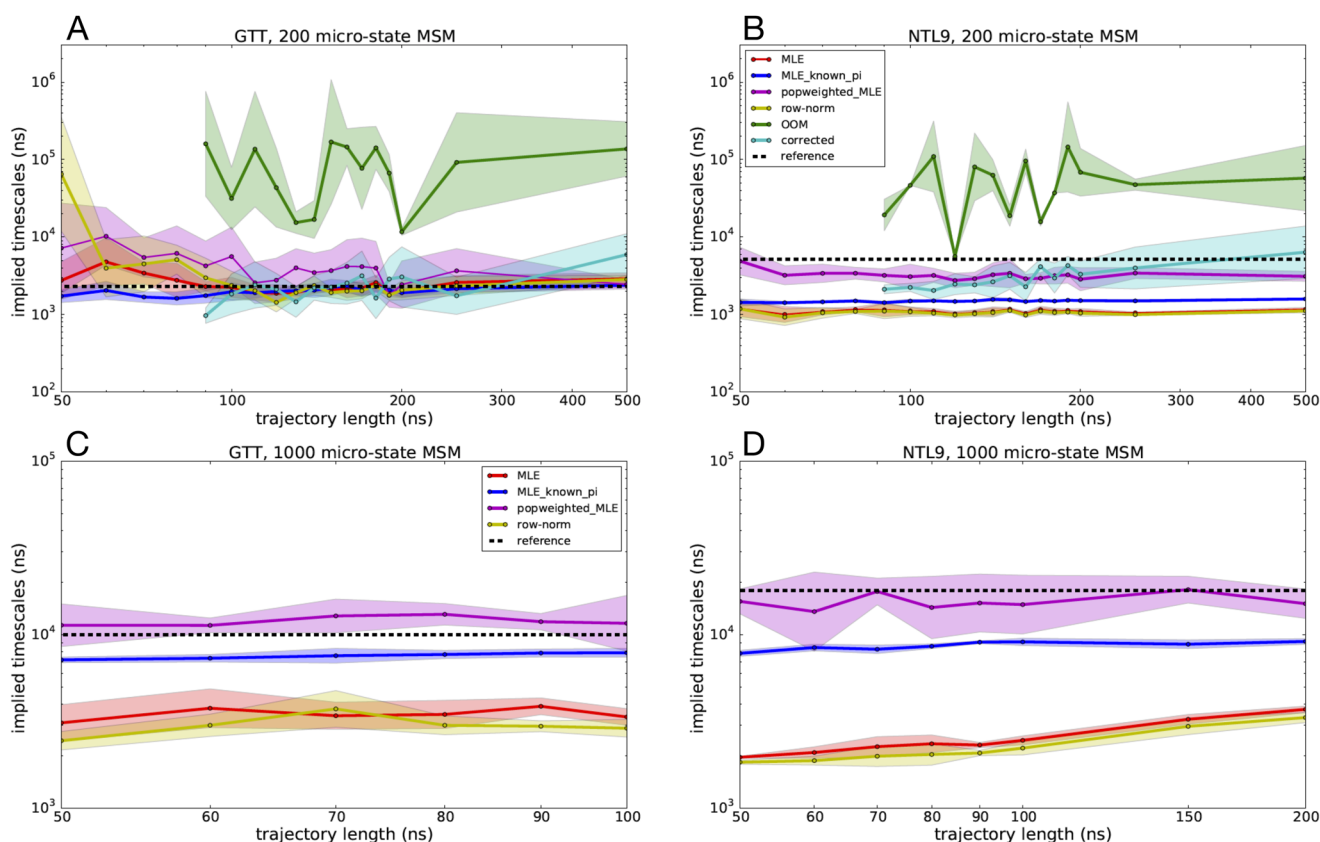


FIG. 8. Slowest implied time scales predicted by various MSM estimators as a function of seeding trajectory length, where ten independent trajectories were started from each microstate. Estimated slowest time scales are shown for 200-microstate MSMs of (a) GTT WW domain and (b) K12M NTL9(1–39) and for 1000-microstate MSMs of (c) GTT WW domain and (d) K12M NTL9(1–39). Trajectory lengths longer than 100 ns for GTT WW domain and 200 ns for K12M NTL9(1–39) result in trajectory data sets larger than the reference model and are not shown. Uncertainties were computed using a bootstrap procedure.

The row-normalized counts estimator yields slowest implied time scales estimates accelerated compared with the reference MSM, similar to the MLE results. This acceleration is greater than might be expected given our tests with the 1-D two-state potential and may in part be due to the necessity of adding pseudocounts to the transition count matrix to ensure connectivity.

2. Estimates of native-state populations

We applied different estimators to the reseeded trajectory data generated for 200-state and 1000-state MSMs of GTT WW domain and K12M NTL9(1–39) and compared the predicted native-state populations with the reference MSMs. The native-state population was calculated as the equilibrium population of the MSM state corresponding to the native conformation. Similar native-state populations were estimated using five independent reseeded trajectories initiated from each state (Fig. 9) and ten independent reseeded trajectories (Fig. 10). These estimates are compared with the native-state populations estimated from the

reference MSMs: 75.3% and 74.8% for the 200-state and 1000-state MSMs of GTT WW domain, respectively, and 75.9% and 79.4% for the 200-state and 1000-state MSMs of K12M NTL9(1–39), respectively.

As we saw in our tests using a 1-D two-state potential, obtaining accurate estimates of the equilibrium populations from very short seeding trajectories is a challenging task. In those tests, we found that the most accurate estimates of native-state populations come from estimators that utilize some *a priori* knowledge of native-state populations. Here, the MLE with known populations recapitulates the native population exactly (by definition), while the MLE with population reweighting successfully captures the native-state population in a 200 microstate model of both GTT and NTL9 and slightly overestimates the native state population for the 1000-state MSMs. In contrast, the MLE, row-normalized counts, and OOM estimators all underestimate the folded populations. With the exception of some slight improvement that begins to occur as the seeding trajectory length increases, these underestimates are independent of the trajectory length, indicating that the nonequilibrium seeding is

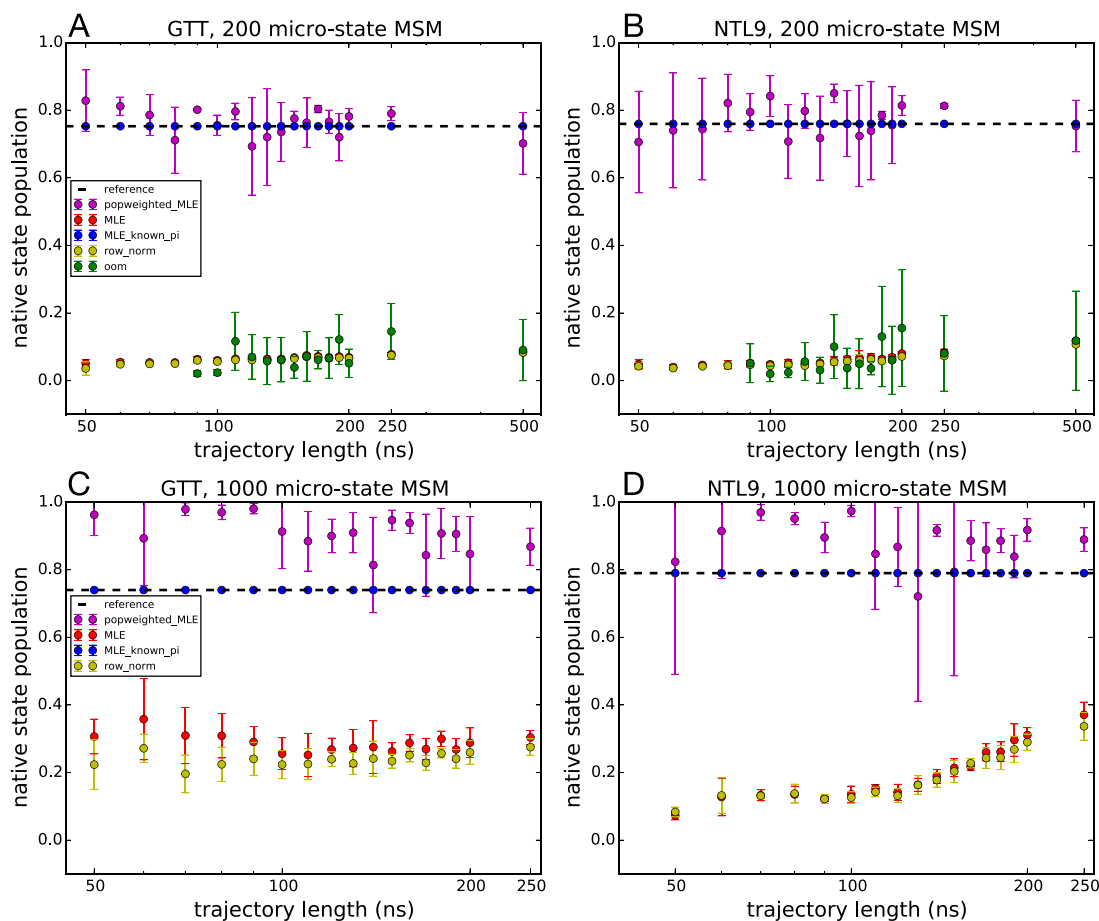


FIG. 9. Native state populations predicted by various MSM estimators as a function of seeding trajectory length. Estimated populations are shown for 200-microstate MSMs of (a) GTT WW domain and (b) K12M NTL9(1–39) and for 1000-microstate MSMs of (c) GTT WW domain and (d) K12M NTL9(1–39). Adaptive seeding data were generated using five independent trajectories from each microstate. Uncertainties were computed using a bootstrap procedure.

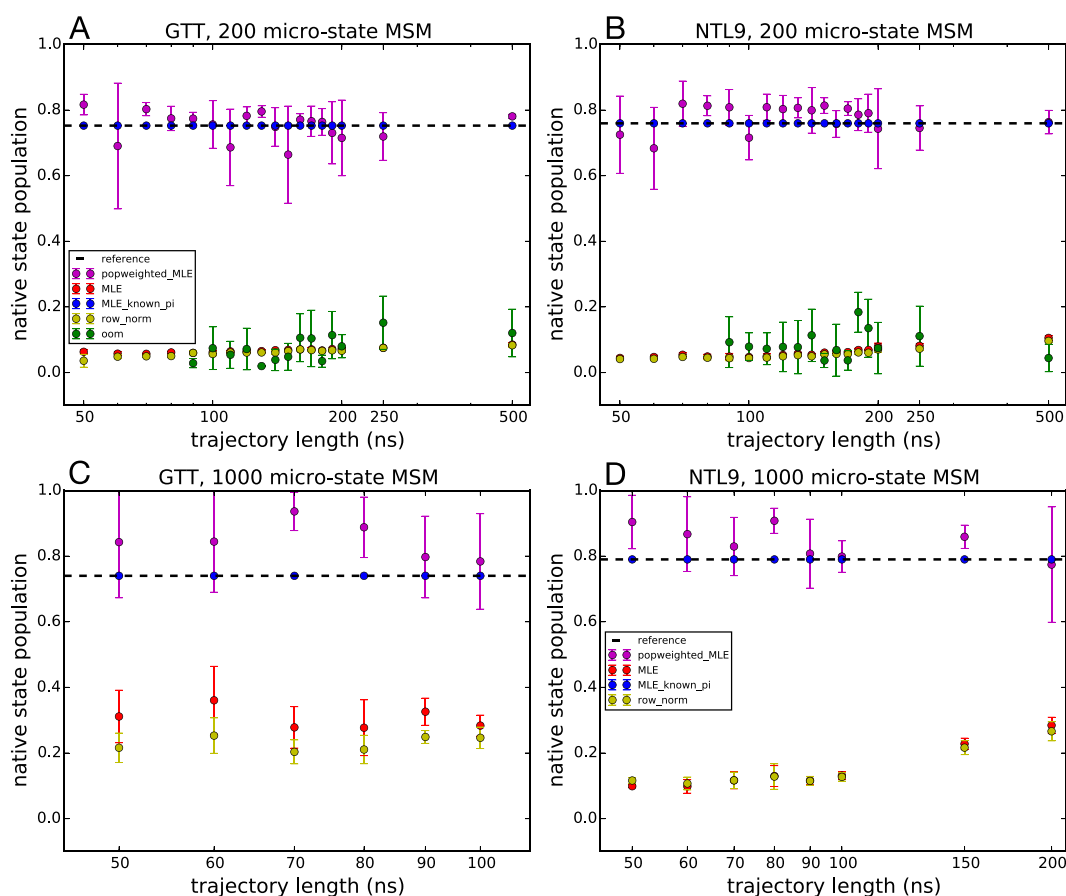


FIG. 10. Native state populations predicted by various MSM estimators as a function of seeding trajectory length. Estimated populations are shown for 200-microstate MSMs of (a) GTT WW domain and (b) K12M NTL9(1–39) and for 1000-microstate MSMs of (c) GTT WW domain and (d) K12M NTL9(1–39). Adaptive seeding data were generated using ten independent trajectories from each microstate. Trajectory lengths longer than 100 ns for GTT WW domain and 200 ns for K12M NTL9(1–39) result in trajectory data sets larger than the reference model and are not shown. Uncertainties were computed using a bootstrap procedure.

to blame. The native state populations are underestimated in these cases because there are many more MSM microstates belonging to the unfolded-state basin of the folding landscape.

III. DISCUSSION

Markov state model approaches have enjoyed great success in the last decade due to their ability to integrate ensembles of short trajectories. In light of this success, the allure of adaptive sampling methods, which promises to leverage the focused sampling of short trajectories, is understandable. Recent studies have shown that a judiciously chosen adaptive sampling strategy can accelerate barrier crossing, state discovery, and pathway sampling.^{11,47}

Our work in this manuscript underscores that a key problem with the practical use of adaptive sampling algorithms is the statistical bias introduced by focused sampling, which makes difficult the unbiased estimation of both kinetics and thermodynamics. We have illustrated these problems in simple 1D diffusion

models and in large-scale all-atom simulations of protein folding by testing the performance of various MSM estimators on adaptive seeding data. In all cases, we find that estimators that incorporate some form of *a priori* knowledge about equilibrium populations are able to correct for sampling bias to some extent, resulting in accurate estimates of both slowest implied time scales and equilibrium free energies. Moreover, we have shown that an MLE estimator with population-weighted transition counts is a very simple way to counteract this bias and achieves reasonably accurate results.

As pointed out by others previously,¹⁸ the success of such estimators is further demonstration of the importance of incorporating thermodynamic information into MSM estimates, given the powerful constraint that detailed balance provides. With this in mind, it is no surprise that new multiensemble MSM estimators like TRAM⁴⁸ and DHAMed²⁶ have been able to achieve more accurate results than previous estimators. Thus, while many existing adaptive sampling strategies focus sampling to refine estimated transition rates between states, we expect that improved adaptive sampling

estimators may come from similar “on-the-fly” focusing that seeks to refine thermodynamic estimates as well.

As adaptive sampling strategies and estimators continue to improve, we expect the more efficient use of ultralong simulation trajectories combined with ensembles of short trajectories for adaptive sampling, especially to probe the effects of protein mutations, or different binding partners for molecular recognition.

ACKNOWLEDGMENTS

The authors thank the participants of Folding@home, without whom this work would not be possible. We thank D. E. Shaw Research for providing access to folding trajectory data. This research was supported in part by the National Science Foundation through major research instrumentation Grant No. CNS-09-58854 and the National Institutes of Health Grant No. 1R01GM123296 and NIH Research Resource Computer Cluster Grant No. S10-OD020095.

REFERENCES

- 1 F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, *Proc. Natl. Acad. Sci. U. S. A.* **106**, 19011 (2009).
- 2 V. A. Voelz, G. R. Bowman, K. Beauchamp, and V. S. Pande, *J. Am. Chem. Soc.* **132**, 1526 (2010).
- 3 J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, *J. Chem. Phys.* **134**, 174105 (2011).
- 4 J. D. Chodera and F. Noé, *Curr. Opin. Struct. Biol.* **25**, 135 (2014).
- 5 F. Noé, J. Chodera, G. Bowman, V. Pande, and F. Noé, *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*, Advances in Experimental Medicine and Biology Vol. 797 (Springer, 2014).
- 6 X. Huang, G. R. Bowman, S. Bacallado, and V. S. Pande, *Proc. Natl. Acad. Sci. U. S. A.* **106**, 19765 (2009).
- 7 V. A. Voelz, B. Elman, A. M. Razavi, and G. Zhou, *J. Chem. Theory Comput.* **10**, 5716 (2014).
- 8 Z. Shamsi, A. S. Moffett, and D. Shukla, *Sci. Rep.* **7**, 12700 (2017).
- 9 S. Doerr and G. De Fabritiis, *J. Chem. Theory Comput.* **10**, 2064 (2014).
- 10 M. I. Zimmerman and G. R. Bowman, *J. Chem. Theory Comput.* **11**, 5747 (2015).
- 11 M. I. Zimmerman, J. R. Porter, X. Sun, R. R. Silva, and G. R. Bowman, *J. Chem. Theory Comput.* **14**, 5459 (2018).
- 12 Z. Shamsi, K. J. Cheng, and D. Shukla, *J. Phys. Chem. B* **122**, 8386 (2018).
- 13 B. W. Zhang, D. Jasnow, and D. M. Zuckerman, *J. Chem. Phys.* **132**, 054107 (2010).
- 14 M. C. Zwier, J. L. Adelman, J. W. Kaus, A. J. Pratt, K. F. Wong, N. B. Rego, E. Suárez, S. Lettieri, D. W. Wang, M. Grabe, D. M. Zuckerman, and L. T. Chong, *J. Chem. Theory Comput.* **11**, 800 (2015).
- 15 A. Dickson and S. D. Lotz, *J. Phys. Chem. B* **120**, 5377 (2016).
- 16 S. D. Lotz and A. Dickson, *J. Am. Chem. Soc.* **140**, 618 (2018).
- 17 T. Dixon, S. D. Lotz, and A. Dickson, *J. Comput.-Aided Mol. Des.* **32**, 1001–1012 (2018).
- 18 B. Trendelkamp-Schroer and F. Noé, *Phys. Rev. X* **6**, 011009 (2016).
- 19 F. Nüske, H. Wu, J.-H. Prinz, C. Wehmeyer, C. Clementi, and F. Noé, *J. Chem. Phys.* **146**, 094104 (2017).
- 20 K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, *Science* **334**, 517 (2011).
- 21 H. Wu, A. S. J. S. Mey, E. Rosta, and F. Noé, *J. Chem. Phys.* **141**, 214106 (2014).
- 22 G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande, *J. Chem. Phys.* **131**, 124101 (2009).
- 23 H. Jaeger, *Neural Comput.* **12**, 1371 (2000).
- 24 H. Wu, J.-H. Prinz, and F. Noé, *J. Chem. Phys.* **143**, 144101 (2015).
- 25 M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Perez-Hernandez, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz, and F. Noé, *J. Chem. Theory Comput.* **11**, 5525 (2015).
- 26 L. S. Stelzl, A. Kells, E. Rosta, and G. Hummer, *J. Chem. Theory Comput.* **13**, 6328 (2017).
- 27 M. P. Harrigan, M. M. Sultan, C. X. Hernández, B. E. Husic, P. Eastman, C. R. Schwantes, K. A. Beauchamp, R. T. McGibbon, and V. S. Pande, *Biophys. J.* **112**, 10 (2017).
- 28 H. Nguyen, M. Jäger, A. Moretto, M. Gruebele, and J. W. Kelly, *Proc. Natl. Acad. Sci. U. S. A.* **100**, 3948 (2003).
- 29 M. Jäger, Y. Zhang, J. Bieschke, H. Nguyen, M. Dendle, M. E. Bowman, J. P. Noel, M. Gruebele, and J. W. Kelly, *Proc. Natl. Acad. Sci. U. S. A.* **103**, 10648 (2006).
- 30 F. Liu and M. Gruebele, *Chem. Phys. Lett.* **461**, 1 (2008).
- 31 K. Dave, M. Jäger, H. Nguyen, J. W. Kelly, and M. Gruebele, *J. Mol. Biol.* **428**, 1617 (2016).
- 32 F. Liu, D. Du, A. A. Fuller, J. E. Davoren, P. Wipf, J. W. Kelly, and M. Gruebele, *Proc. Natl. Acad. Sci. U. S. A.* **105**, 2369 (2008).
- 33 Y. Gao, C. Zhang, X. Wang, and T. Zhu, *Chem. Phys. Lett.* **679**, 112 (2017).
- 34 R. B. Best, G. Hummer, and W. A. Eaton, *Proc. Natl. Acad. Sci. U. S. A.* **110**, 17874 (2013).
- 35 S. Piana, K. Sarkar, K. Lindorff-Larsen, M. Guo, M. Gruebele, and D. E. Shaw, *J. Mol. Biol.* **405**, 43 (2011).
- 36 H. Nguyen, M. Jäger, J. W. Kelly, and M. Gruebele, *J. Phys. Chem. B* **109**, 15182 (2005).
- 37 J.-C. Horng, V. Moroz, and D. P. Raleigh, *J. Mol. Biol.* **326**, 1261 (2003).
- 38 J.-H. Cho and D. P. Raleigh, *J. Mol. Biol.* **353**, 174 (2005).
- 39 C. R. Schwantes and V. S. Pande, *J. Chem. Theory Comput.* **9**, 2000 (2013).
- 40 J.-H. Cho, W. Meng, S. Sato, E. Y. Kim, H. Schindelin, and D. P. Raleigh, *Proc. Natl. Acad. Sci. U. S. A.* **111**, 12079 (2014).
- 41 C. R. Baiz, Y.-S. Lin, C. S. Peng, K. A. Beauchamp, V. A. Voelz, V. S. Pande, and A. Tokmakoff, *Biophys. J.* **106**, 1359 (2014).
- 42 J.-H. Cho, S. Sato, and D. P. Raleigh, *J. Mol. Biol.* **338**, 827 (2004).
- 43 H. Wan, G. Zhou, and V. A. Voelz, *J. Chem. Theory Comput.* **12**, 5768 (2016).
- 44 G. Perez-Hernandez, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *J. Chem. Phys.* **139**, 015102 (2013).
- 45 R. T. McGibbon and V. S. Pande, *J. Chem. Phys.* **142**, 124105 (2015).
- 46 K. A. Beauchamp, R. McGibbon, Y.-S. Lin, and V. S. Pande, *Proc. Natl. Acad. Sci. U. S. A.* **109**, 17807 (2012).
- 47 E. Hruska, J. R. Abella, F. Nüske, L. E. Kavrakci, and C. Clementi, *J. Chem. Phys.* **149**, 244119 (2018).
- 48 H. Wu, F. Paul, C. Wehmeyer, and F. Noé, *Proc. Natl. Acad. Sci. U. S. A.* **113**, E3221 (2016).