Published in final edited form as:

J Chem Theory Comput. 2020 November 10; 16(11): 6763–6775. doi:10.1021/acs.jctc.0c00273.

Accelerated estimation of long-timescale kinetics from weighted ensemble simulation via non-Markovian “microbin” analysis

Jeremy Copperman, Daniel M. Zuckerman

Department of Biomedical Engineering, Oregon Health and Science University, Portland, OR

Abstract

The weighted ensemble (WE) simulation strategy provides unbiased sampling of non-equilibrium processes, such as molecular folding or binding, but the extraction of rate constants relies on characterizing steady state behavior. Unfortunately, WE simulations of sufficiently complex systems will not relax to steady state on observed simulation times. Here we show that a post-simulation clustering of molecular configurations into “microbins” using methods developed in the Markov State Model (MSM) community, can yield unbiased kinetics from WE data *before* steady-state convergence of the WE simulation itself. Because WE trajectories are directional and not equilibrium-distributed, the history-augmented MSM (haMSM) formulation can be used, which yields the mean first-passage time (MFPT) without bias for arbitrarily small lag times. Accurate kinetics can be obtained while bypassing the often prohibitive convergence requirements of the non-equilibrium weighted ensemble. We validate the method in a simple diffusive process on a 2D random energy landscape, and then analyze atomistic protein folding simulations using WE molecular dynamics. We report significant progress towards the unbiased estimation of protein folding times and pathways, though key challenges remain.

Graphical Abstract

Introduction

The weighted ensemble (WE) is a parallel path sampling strategy for rare events, proposed by Huber and Kim in 1996,¹ echoing earlier computational “splitting” strategies.² WE entails running a set of simulations, accompanied by replication (“splitting”) of promising trajectories and pruning (via “merge” events) of less important trajectories. As long as simple statistical rules regarding the splitting and merging of trajectories are followed, the

evolution of the trajectory ensemble is unbiased.^{3,4} By performing this splitting and merging based on a set of bins spanning the transition of interest, sampling can be greatly enhanced in rare or rate-limiting regions.⁵ Recent efforts have utilized WE simulations to determine the pathways and effective kinetic rates of increasingly challenging long-timescale biological processes such as cell and genetic switches,^{6–8} ion transport,⁹ protein-peptide binding,¹⁰ protein-ligand unbinding,^{11,12} protein-protein association,¹³ and protein folding.¹⁴

The Hill relation^{15,16} permits unbiased estimation of the mean first-passage time (MFPT), which is a proxy for the inverse rate constant, from steady-state probability flow. However, all applications of WE simulation for unbiased MFPT estimation are limited by the capability to obtain convergence to steady state (SS). In this study, we therefore address a question of considerable potential impact: Can SS kinetic information (i.e., the rate constant or MFPT) be extracted from unbiased transient data obtained via WE simulation *prior to reaching steady state*?

The data examined below will show that this is indeed the case, and that rate-constant estimates are most reliable when extracted using a non-Markov analysis^{16–18} of a very fine discretization of configuration space – typically much finer than was used to run the original WE simulations themselves. WE simulations typically use a discretization of configuration space into bins, each of which may contain several trajectories running in parallel^{1,19} – thus limiting the number of bins which can be used to tens or hundreds in practical cases.

We are considering the rate constant defined by $k_{AB} = 1/\text{MFPT}(A \rightarrow B)$, from one macrostate (A) to another (B), where A and B could be two non-overlapping conformational states, folded and unfolded states of a protein, or unbound and bound states of a complex. It has been established that k_{AB} can be obtained from WE simulations which have reached SS based on the Hill relation,^{15–17}

$$k_{AB} = \text{Flux}(A \rightarrow B; \text{SS}) \quad (1)$$

which uses the probability flux into state B – i.e., the probability arriving per unit time – based on trajectories initiated in A. Analytical and numerical investigation of the transient relaxation to SS in the Smoluchowski equation for the probability flux has shown the conditions in which to expect monotonic relaxation to SS, and generically supports the merit of sampling the one-way transition ensemble necessary to haMSM construction.²⁰ In simple 1D systems with a single energy barrier, the SS relaxation time can be arbitrarily faster than the MFPT. However, complex systems inevitably will require long relaxation times to reach SS, and during the transient relaxation period “direct” WE estimates obtained from the probability flux will typically underestimate the true SS flux.²⁰ A complementary approach utilizes the event duration distribution to boost the measured flux during the transient regime.²¹ Our concern is to extract accurate estimates of the SS flux and the entire SS distribution characterizing the transition, based on *transient* WE data – i.e., before steady state has been reached.

In prior work, we showed that a history-dependent non-Markov analysis of WE bins could be used to estimate the SS rate based on stationary solution of an appropriate transition matrix.^{16,22–24} Here, we term the non-Markov formulation a “history-augmented Markov state model” (haMSM) to emphasize both its relation to, and difference from, standard MSMs. In a haMSM, one constructs a separate transition matrix for each direction (A-to-B or B-to-A) based on the *subset* of trajectories which were most recently in macrostate A for the A-to-B direction (or B, for B-to-A). The history is equivalent to the directionality – i.e., the labeling of which macrostate was visited more recently. Once the history has been used to select the trajectory subset, rate estimation proceeds much as it would in a standard MSM calculation. Importantly, however, haMSMs provide unbiased estimates of the rates (inverse MFPTs), at arbitrary lag times.¹⁸ In this study we perform WE sampling of the one-way A-to-B ensemble, and the haMSM is itself just an MSM of the one-way WE trajectory ensemble; thus there is no direct comparison of an haMSM to a standard MSM in this context. We also emphasize that we are not computing “implied timescales” based on eigenvalues but the MFPT corresponding to a well-defined physical process.

Here we show that using finer (smaller) bins in a haMSM gives better performance for rate estimation from WE data, and we employ the clustering/binning processes which have been extensively developed in the MSM community.^{25–27} The motivation for using smaller bins is that larger bins are likely to possess internal free energy barriers leading to slower internal relaxation, which in turn will bias the transition probabilities (transition matrix elements) of the haMSM and ultimately the macroscopic transition rate estimate. Put another way, the distributions within smaller bins are more likely to be similar to the SS distribution, and the intra-bin distributions determine the transition probabilities.

Procedurally as shown in Fig. 1, we use fine “microbins” generated by analyzing uni-directional (A-to-B) data with the pyEMMMA software.²⁶ These microbins are used to generate a haMSM, whose stationary solution provides the desired rate constant k_{AB} .¹⁶ The stationary solution can also be used to initialize new WE simulations in the haMSM-estimated SS for validation. Our primary focus is showing that finer/smaller bins yield more accurate rate estimates, especially when compared to the relatively large bins used to run WE simulations. We note that practical WE simulations are limited in the number of bins which can be used because computing cost scales linearly with the number of bins.¹⁹

Methods

WE simulation of the $A \rightarrow B$ ensemble

We wish to describe the kinetics of macrostate (“state”) transitions in a molecular system using the Hill relation (1). We denote the initial/source state A and the target/sink state B, which are arbitrary non-overlapping regions of phase space. For the $A \rightarrow B$ transition we only need the “ α ” (A-to-B directed) subset of trajectories which were most recently in A; those most recently in B (B-to-A) are denoted β . We employ WE simulation to enable parallel sampling of the α trajectory ensemble: trajectories are initiated in A, and those which subsequently arrive at the absorbing sink at B, are restarted in A. The starting distribution in A is noted below for each system. WE provides an unbiased representation of

the α reactive trajectory ensemble,^{4,28} and the time-resolved probability flux into the sink state.

We report the variability between individual WE simulations by calculating nominal 95% credibility regions (CR) utilizing a Bayesian bootstrapping procedure;²⁹ additional details about estimating the CR from haMSM estimates can be found in Adhikari et al.¹⁴

haMSM formulation

We wish to post-analyze molecular dynamics data harvested from WE simulation in a discretized configurational space using a transition matrix T . The matrix encodes (conditional) transition probabilities $T_{ij} = T_{i \rightarrow j}$ among bins or “microbins” i and j . Markov state models (MSMs) have been used to stitch together many independent simulations to approximate long-timescale processes^{30–32} but do not distinguish the α and β trajectory subsets.¹⁷

The haMSM is a transition matrix formulation containing history labels, namely α or β , so that transition matrix elements are calculated solely from the corresponding trajectory subsets.¹⁶ Compared to a standard MSM, a haMSM expands the transition matrix formulation for an N -state system from $N \times N$ elements into a $2N \times 2N$ labeled rate matrix. Here, we are concerned solely with the α (last in A) ensemble of the source/sink system, which limits our attention to the $N \times N$ transition (sub)matrix

$$T_{ij}^{\alpha} = P\{X_{t+\tau}^{\alpha} = j \mid X_t^{\alpha} = i\}, \quad (2)$$

where X_t^{α} is a trajectory in the α subset and τ is the lag time of the transition matrix.

The discretized version of the Hill relation (1) then becomes¹⁶

$$k_{AB} = \text{Flux}(A \rightarrow B; \text{SS}) = \sum_{i \notin B, j \in B} p_i^{\alpha} T_{ij}^{\alpha} \quad (\text{haMSM}) \quad (3)$$

with p_i^{α} the SS probability of microbin i based on SS solution of the transition matrix T^{α} .

This non-equilibrium α SS does not obey detailed balance because of the net flux into the B state originating in the A state. Remarkably, the haMSM formulation (3) yields the correct MFPT independent of the microbin set or lag-time used to construct the transition matrix^{16,17} so long as the matrix elements are calculated based on trajectories launched in every bin i according to the SS distribution. The lag-time independence is a powerful distinction from the traditional MSM because it means that all transitions collected in the α trajectory ensemble can be used to train the haMSM. In the atomistic protein folding example we use a 10ps lag time, which should be contrasted with the ~10–100ns lag times needed for accurate MSMs of molecular systems.^{25,32,33} In practice, then, training the haMSM should require significantly less trajectory data than for a standard MSM.

Because the haMSM transition matrix is built from the conditional transition probabilities between states, the SS distribution need only be reached locally *within* the defined bins, not globally over all configuration space. As noted above, we expect faster relaxation for smaller

bins – that is, for haMSMs with more microbins. This expectation is borne out by the data shown below.

haMSM construction

Construction of haMSMs proceeds in a highly similar manner to constructing a standard MSM, except via WE weights instead of simple transition counts. All weights from WE simulation are tracked¹⁹ and available for this analysis. Microbins which are not fully connected (e.g., having transitions in but not out) were removed from the analysis. To build the transition matrix, the transition weight w_{ij} from microbin i to j after lag τ was averaged over WE iterations and runs to construct the haMSM transition matrix using¹⁶

$$T_{ij} = \frac{\langle w_{ij} \rangle}{\langle w_i \rangle} \text{ with } w_i = \sum_j w_{ij}. \text{ Here all transitions are used – that is, transitions only observed a}$$

single time were included in the analysis. We used the same lag time τ for the transition matrix as the WE lag time, 10.0 ps in the protein folding systems and 50.0 ps in the 2D test system. Extracting the structural transition information at each step requires a large amount of file input/output, and can be slow depending on file access capability/concurrency, but is not computationally demanding.

In using the transition matrix to extract non-equilibrium SS kinetics via Eq. (3), we must apply suitable boundary conditions. The source/sink boundary conditions were enforced by employing the exact target state (sink, or macrostate B) and source state (A) definitions used in the WE simulation when constructing the haMSM, with the transition matrix row for the target state enforcing all probability transfers directly back to the source state. That is, with b the index of the sink state, and a the index of the source state,

$$T_{ba} = 1, \quad T_{b\{i \neq a\}} = 0, \quad (4)$$

The SS flux and mean first-passage time were estimated from the SS of the haMSM transition matrix via Eq. (3).

Microbin construction via clustering

History augmented Markov State Models (haMSMs) are constructed directly from the WE simulation trajectories. We employ “MSM-style” microbins, which is a new development in this work. Specifically, we apply unsupervised clustering methods to extract a set of microbins (Voronoi centers) representing the configuration space visited by the WE trajectories. For each haMSM based on clustering, the latest-occurring 100,000 structures from the training window are used as input to the clustering. The training window is chosen in different ways to address different questions (see below), and the windows used are always indicated in the results section. A new clustering was performed for each training window considered, so that only the information available inside the training window is ever used for the analysis. K-means clustering with a minimum RMSD metric based on all protein atom Cartesian coordinates (rotationally and translationally minimized) was performed using the pyEMMA software package,²⁶ requiring ~24 hours of computation over 12 CPUs for each clustering calculation. Given a set number of desired microbins and an

initial choice of bin centers, a k-means clustering is deterministic, but accordingly varies given a change in the desired number of microbins or initialization.

In the NTL9 low-friction protein folding system, we also explored the use of dimensionality-reduction methods, to extract a lower dimensional subspace spanning the slow conformational degrees of freedom. We first processed the atomic coordinates to extract the matrix of intra α -carbon distances. We tested Principal Component Analysis (PCA), as well as the Variational Approach for Markov Processes (VAMP)²⁷ dimensionality reduction method which leverages the time-lagged covariance matrix to extract a representation of the generator and the subspace of slow degrees of freedom, appropriate for non-equilibrium ensembles. Trajectories of every segment from the last iteration of the training window were traced back to the first step (unfolded structure). These trajectories were used as input to the VAMP clustering algorithm implemented in pyemma²⁶ with an input lag time of half the training window.

Initializing WE simulations in estimated SS

A direct approach to validating the stationary solution of the haMSM is to initialize new WE simulations in the haMSM estimated SS and see if SS properties are observed. Re-initializing WE simulations from the steady-state distribution estimated from prior WE runs requires extracting structures and weights representative of the estimated SS, presented in pseudocode below.

Algorithm 1 Extract representative SS structures and weights for WE simulation.

```

for WE bin  $i$  in  $\{1, \dots, N_{WE}\}$  do
  find the subset of microbins  $\Omega_i$  which map to each WE bin  $i$  via average progress coordinate
  for each new trajectory  $\xi$  allocated in WE bin  $i$  do
    choose microbin  $j$  for trajectory  $\xi$  from  $\Omega_i$  randomly by microbin haMSM SS weight
    extract library of structures and direct WE weights assigned to  $j$  from previous WE simulations
    choose structure for trajectory  $\xi$  from library randomly by direct WE weight
    assign relative weight of trajectory  $\xi$  proportional to haMSM SS weight of microbin  $j$ 
  end for
  enforce total weight in WE bin  $i$  to be total haMSM SS weights of all microbins in  $\Omega_i$ 
end for

```

Flux profiles

A stringent test of SS is to monitor the flux across collective iso-surfaces which separate source (A) and sink (B) states. At SS, the flux through any surface completely separating A from B will be identical, and constant—the SS flux J_{SS} . To calculate the flux profile from a WE simulation, we calculate values of the collective variable (specified below for each system) for WE trajectories as follows. The collective variable is split into a set of monotonically ordered bins, and we average the weight transitioning between these bins at each WE iteration. The flux passing bin i , J_i is then

$$J_i = \frac{1}{\tau_{WE}} \sum_{j > i, k \leq i} \langle w_{jk} - w_{kj} \rangle \quad (5)$$

where τ_{WE} is the WE lag time (i.e., interval between iterations) and bin ordering is set to define positive flux towards the sink (B) state.

Systems and Results

Our goal is to validate the use of the haMSM approach to extract SS kinetics from transient trajectory data in complex systems. To that end, we apply the haMSM analysis to WE

simulations of diffusion in a 2D random energy landscape, and to atomistic protein folding, described in detail in Ref. 14. We are interested in the extent to which haMSMs with many (thousands of) microbins can enable the estimation of the SS flux at molecular times less than the SS relaxation time. We will consider the ratio of the SS relaxation time (τ_{SS}) to the latest molecular time (WE simulation time) t_{mol} used in the haMSM training window where the haMSM can predict the SS flux, to be the computational acceleration. This is not a computational acceleration as compared to any other method but rather assesses our ability to exploit transient information.

Our validation has three stages. (i) When the haMSM is trained with a full set of data which approaches SS, then the predicted MFPT should be independent of the clustering. Even coarse clusterings produce exact results.^{16,18} (ii) When the haMSM is trained solely on transient data, we seek haMSMs which reliably predict the MFPT found from the more complete training – i.e., the SS value. (iii) The stationary solution of the haMSMs are used to re-initialize new WE simulations in the haMSM estimated SS, and these WE trajectories are analyzed for SS convergence.

Diffusion in a 2D Random Energy Landscape

We first tested our approach by simulating a particle diffusing in the 2D random energy landscape shown in Fig. 2. Parameters were chosen to emulate an amino acid in water at 300K ($m = 100$ daltons, $D = \frac{k_B T}{\gamma} = 6.08610^{-10} \frac{m^2}{s}$). Two stable states ($6k_B T$ deep Gaussian wells of 1.0nm width) were separated by 10nm with the addition of 40 randomly placed Gaussians (python code to generate the energy surface provided in the supplemental information). A confining radial potential was also placed at 9nm from the domain center. Particle trajectories were evolved according to the Langevin equation,

$$m \frac{\partial^2 \vec{x}}{\partial t^2} = -\gamma \frac{\partial \vec{x}}{\partial t} - \vec{\nabla} U(\vec{x}) + \vec{f}(t), \quad \langle f_i(t) \cdot f_j(t') \rangle = 2\gamma k_B T \delta(t - t') \delta_{ij} \quad (6)$$

with $\vec{x}(t)$ the particle position at time t , m the mass, γ the friction, $U(\vec{x})$ the potential energy surface, T the temperature (k_B Boltzmann's constant). Additionally, “recycling” was performed in the weighted ensemble implementation: trajectories which are found in the defined target state (B) are “recycled” back to the source state (A). State definitions were defined as all points within 0.1nm of the metastable states at A (x_A, y_A) = (4, -3)nm and B (x_B, y_B) = (-3, 4)nm, shown in Fig. 2.

Weighted ensemble requires several implementation choices. Bins were based upon the radial distance to the target state and placed every 0.2nm up to 10.0nm and a final bin edge at 12.0nm encompassing all trajectories with distances greater than 12.0nm. See Fig. 2 for state definitions and weighted ensemble binning. The system was evolved according to Eq. (6) using the OpenMM molecular dynamics Langevin integrator.³⁴ The weighted ensemble sampling was implemented with a 50.0ps lag time in WESTPA³⁵ with 4 trajectories per WE bin (200 trajectories for each WE simulation once full occupancy is reached). WE simulations, and numerical solution of the probability evolution were initialized from the source state.

We also computed the solution numerically in the Smoluchowski picture. Though inertial Langevin simulations were performed, the parameters chosen emulate an amino acid in water evolving on a potential with nanometer scale features and the system is well in the overdamped regime; inertial effects can be safely neglected as is borne out by the agreement with simulation data. Evolution of the probability distribution is thus approximated by the Smoluchowski equation with the addition of the recycling from target (sink, B) and initial (source, A) states,

$$\begin{aligned}\frac{\partial p}{\partial t} &= -\vec{\nabla} \cdot \vec{J} + \gamma_A(\vec{x}) \int_{d\Omega_B} \vec{J} \cdot \hat{n} d\vec{x} \\ \vec{J} &= \beta D \vec{f}(\vec{x}) p - D \vec{\nabla} p\end{aligned}\quad (7)$$

with $p(\vec{x}, t)$ the probability distribution in the domain at time t with absorbing boundary at the sink $p(\vec{x} \in B) = 0$, $\vec{J}$ the current. The source distribution $\gamma_A(\vec{x})$ is a Dirac delta function: trajectories are recycled to the center of the gaussian well at $[4.0, -3.0]\text{nm}$. Lastly, $\vec{f} = -\vec{\nabla} U$, $D = \frac{k_B T}{\gamma}$, and Ω_B is the boundary of B with inward-facing normal $\hat{n}$. Eq. (7) is the standard Smoluchowski equation with the addition of the source/sink boundary conditions for recycling,²⁰ and was solved numerically using the fipy³⁶ package. Slight variation was observed between different choices of grid sizes ($400 \times 400 - 800 \times 800$) and timestep (10.0–100.0ps) and this variation set the minimum and maximum range of the SS flux into the target state shown in Fig. 4.

Transient evolution between the WE trajectory ensemble and the deterministic probability evolution given from numerical solution of Eq. (7) are consistent (see supplementary Fig. S1), which is not surprising since WE is an unbiased path sampling procedure.³ A validation set of 50 long WE simulations of length 100ns were performed, showing relaxation at around $0.1\mu\text{s}$ to a SS indicating an MFPT of $\sim 1.0\mu\text{s}$ via Eq. (1). A second test set of 50 WE simulations of length 2.5ns were performed, and were used as input to build haMSM models and estimate the steady state distribution and flux into the target. Fig. 2C shows the transient probability distribution from the WE test set (no haMSM) at the end of the training window at 2.5ns, about 1/40 of the SS relaxation time and very far from the SS.

The capability to estimate the correct SS flux value from transient data before SS was dependent upon the size of the microbins, confirming our expectations. Fig. 3 shows that while the coarsest bins yielded haMSM SS flux estimation which were the same as the direct transient flux, the haMSM prediction increased monotonically up to the correct SS flux value, at about 1024 bins, and remained consistent for all finer microbins calculated (1024–16,384 microbins). In the limit of very many microbins, the haMSM microbins evidently become Markovian in the sense that the microbins are small compared to the size of the features on the 2D landscape. Since the very first trajectories in the WE simulation only reach the target state by 2.0ns, this represents about the earliest possible estimation of the SS flux possible using roughly 1/40 of the SS relaxation time. The aggregate simulation time in the training set was about $1.3\mu\text{s}$ which is roughly the MFPT itself. We note that it is likely that fewer WE simulations with fewer trajectories per WE simulation could have been used, although we do not explore this limit here.

The haMSM-predicted SS distribution does capture the overall scale and many important features of the SS distribution, shown in Fig. 2D. It is apparent that only approximate estimation of the entire SS distribution is necessary for accurate estimation of the SS target flux, and hence the MFPT. With more training data and a longer training window, the predicted haMSM distribution would converge upon the SS distribution (as would the direct WE distribution). With the limited training set at $t = 2.5\text{ns}$, the haMSM estimate is clearly not an exact reproduction of the SS distribution (from numerical solution of Eq. (7) shown in Fig. 2B).

WE simulations re-initialized in the haMSM-estimated SS demonstrated clear convergence to SS (Fig. 4). We used the haMSM SS distribution to initialize a set of fully independent 100 WE simulations in the estimated SS (see section “Initializing WE simulations in estimated SS” for details). The target flux remains steady, consistent with the numerical solution of the Smoluchowski equation, haMSM SS estimations, and the direct flux from WE simulations longer than the SS relaxation time, shown in Fig. 4.

For a more granular view to confirm SS, in Fig. 4B we plot the flux profile along the 1D reaction coordinate (distance to the target state). This flux profile should become flat at SS.^{37,38} While indeed the flux profile becomes much flatter after reweighting/restarting, wrong-way fluxes away from the target state are transiently observed after reweighting, indicative of the errors in the estimated SS distribution at large distances to the target. Within an additional 5ns these wrong-way fluxes relax to a nearly flat flux profile.

When metastable intermediates exist along the transition pathway between macrostates of interest, we expect SS convergence to be slow and to approach the MFPT itself, and in this situation we expect MSM microbin clustering will have the most utility to accelerate the estimation of SS properties. The 2D diffusion of a particle on a random energy landscape explored in this section is a simple toy model, although it does capture some of the complexity expected for more challenging systems such as the atomistic protein folding systems. To investigate the dependence of the haMSM performance with energy landscape ruggedness, we also examined the A to B transition in the 2D random energy landscape at 150K, or $T/2$. As expected, the mean first-passage time and the steady-state relaxation time increase exponentially with an Arrhenius-like exponential dependence, and are 300x slower. Meanwhile, the haMSM with thousands of microstates can predict the SS flux with only 10x more WE simulation time than at 300K, so the acceleration is increased by a factor of 30, see Fig. S2 of the supplementary information. Generically then, we should expect a greater speedup when processes are slower, and more activated (energy-barrier dependent).

Atomistic Protein Folding

In protein folding the configurational space is very high dimensional, and dynamical motion spans timescales across many orders of magnitude. We show here that in these very challenging systems, haMSMs with thousands of microstates are able to accelerate estimation of SS and protein folding times. In the toy model system, it is possible to reach a point where microbins are truly Markovian in the sense that the microbins are so small there are no important intrabin features of the energy landscape. Here we construct protein-folding haMSMs with 10.0ps lag times, and do not expect this Markovian property to be satisfied for

any construction method or number of microbins.¹⁸ MSM models of NTL9 folding utilizing 100,000 microbins indicated that 10ns lag times were necessary before implied timescales leveled off, suggesting Markovian behavior.³² Variationally optimized MSM construction of protein folding trajectories utilized lag times of 50ns^{39,40} and 100ns.⁴¹ Despite the lack of truly Markovian microbins, we do find that haMSMs with many (thousands of) microbins accelerates SS estimation compared to the brute-force relaxation time to SS. While a full theoretical discussion of the SS relaxation time is beyond the scope of this work, we note that it should be very sensitive to the existence of metastable intermediates.⁴²

We study the folding of two proteins, the N-terminal domain of the ribosomal protein L9 (NTL9) and the IgG binding domain of streptococcal protein G. We analyze both weighted ensemble trajectories from Ref. 14 and from additional new WE simulations, following the same protocol. Computational wall time in these systems is discussed in more detail in Ref. 14 but we note that a single WE simulation with ~ 1000 trajectories accumulates roughly 0.5 ns/day of molecular time utilizing a single GPU/CPU (NVIDIA V100 gpu, Intel Xeon processor). The fast-relaxing low-friction NTL9 folding system serves as a well-converged system in which to validate the haMSM capability to accelerate SS estimation and explore the use of dimensionality reduction methods (PCA and VAMP) in the haMSM microbinning process. Utilization of dimensionality reduction methods significantly reduces the computational cost of the clustering but has a more subtle impact on the haMSM estimated SS, and in this work we explore their use only in the NTL9 folding system; we analyze the protein G folding system using haMSMs with microbins constructed using the all-atom based clustering with a minimum RMSD distance metric. We further attempt to validate the haMSM-estimated steady-state of the protein G folding system by reweighting and restarting a set of WE simulations in the haMSM estimated SS.

NTL9

NTL9 is a fast-folding globular protein^{32,43–45} without experimentally measurable folding intermediates.⁴⁶ The previously performed weighted ensemble simulations of the protein folding process,¹⁴ determined a protein folding time utilizing Eq. (1) consistent with the experimentally measured 1.2 – 1.4ms folding times.⁴⁶ In those implicit solvent atomistic simulations, flux profiles along the RMSD to the folded state indicated that the WE simulations effectively approached SS in both the low-friction ($\gamma = 1/5$ ps) and high-friction ($\gamma = \gamma_{\text{water}} = 1/80$ ps) systems within nanoseconds.¹⁴ Applying the haMSM analysis to both NTL9 systems (with a training window utilizing the final portion of the WE simulation), the haMSM estimated SS was independent of the method and number of microbins, confirming that the SS relaxation time was indeed short enough to approach the SS values via the brute force WE simulation.¹⁴

We now take advantage of this well validated system to test the capability of haMSMs to estimate SS from pre-SS training data. The previously reported WE simulations¹⁴ of 10 independent low-friction NTL9 folding simulations with a 2D progress coordinate (2D-WE) were run for $t_{\text{mol}} = 12$ ns and are here considered the validation set. Five newly run independent 2D-WE simulations are used as the test set for building haMSMs. We estimate the SS relaxation time to be ~ 5 ns in this system, which is the time needed for the average

flux of the validation set to reach the 95% confidence region (CR) for the SS flux. This is also the time needed for the upper bound of the test set direct flux to reach the lower bound of the validation set SS flux.

Fig. 5 shows that haMSMs can indeed accelerate the estimate of SS, dependent upon the microbin clustering and number of microbins in the haMSM. The haMSM estimated SS flux, from haMSM models utilizing WE simulation data from the test set only up to the training window indicated, increases with the number of microbins in the haMSM when the training window is significantly less than the SS relaxation time. As the SS relaxation time of $\tau_{SS} \sim 5$ ns is approached, the haMSM estimate becomes independent of the number of microbins, for all clustering methods studied here, shown in Fig. 5 for both PCA and VAMP dimensionality reduction. Results for all-atom clustering using a minimum RMSD distance metric are quantitatively similar to the results utilizing PCA dimensionality reduction: see Fig. 6.

In principle, with sufficient training data and sufficiently small microbins, it should be possible to extract SS estimates from arbitrarily short (small t_{mol}) trajectory data. In practice, the amount of trajectory and the ability to cluster configurational space into effective microbins will limit one's ability to leverage transient information. As shown in Fig. 6A, all clustering methods reach order-of-magnitude estimates of the SS flux within 1 ns of molecular time ($\sim 1/5$ of the transient, using aggregate simulation time $3.7 \mu s$) and a flux estimate within the validation set CR within 2.5 ns of molecular time ($\sim 1/2$ of the transient, using aggregate simulation time $11.8 \mu s$).

Flux profiles from the training set and the validation set (Fig. 6B) indicate continued relaxation towards SS even when the target flux and MFPT estimation have stabilized. The flux profile does not become completely flat, indicating the presence of a robust regime where flux into the sink (folded) state approaches SS before global convergence to SS. Steady-state estimation is most sensitive to the clustering procedure at early times in the transient, and becomes independent of the clustering method and number of microbins as SS is approached in the trajectory ensemble: see Figs. 5 and 6.

The haMSM steady-state probability distribution systematically redistributes weight to configurations similar to the folded structure (low RMSD to folded state) from unfolded configurations, as shown in Fig. 7, while a comparison of RMSD values and the fraction of native contacts on the landscapes are shown in the supplementary information Fig. S3. The PCA dimensionally-reduced landscapes show a landscape geometry where RMSD increases mostly monotonically. Meanwhile the VAMP landscape, which attempts to separate the subspace of "slow" degrees of freedom, shows a landscape where some unfolded large-RMSD structures are geometrically near small RMSD folded structures, with the largest geometrical distances between protein configurations of intermediate RMSD. It is not clear here which kinetic landscape is more correct.

For flux (Fig. 6A), VAMP performed similarly to PCA when many independent components (ICs) were retained; when only the leading ICs were retained, accelerated order of magnitude estimation of SS rates were obtained but overshoot the true SS rate, and

approached the correct SS rate from above. This counter-intuitive non-monotonic behavior illustrates that different dimensionality reduction and clustering procedures can lead to systematic variation in the estimated steady-state when applied to transient training data, supporting testing multiple approaches. A related consideration is to understand how haMSM models fail under training data sparsity. We took representative haMSM models with 10,000 microbins at different training windows and systematically reduced the amount of training data (while maintaining the same final iteration of the training window). Generically, models are well-behaved and predict similar SS target flux until they fail through a loss of connectivity in the matrix, with some dependence on the dimensionality reduction method, see Fig. S4 in the supplementary information.

A concern in building any complex model from limited training data is the possibility of overfitting, and choosing optimal hyperparameters. The choice of dimensionality reduction method, and the optimal number of microbins in a haMSM model given a finite set of trajectory (training) data, requires attention in future work, and could perhaps be guided by cross-validation procedures inspired by those developed in the MSM community,^{41,47,48} but this will require further development for application in this context. In the NTL9 system, we have an independent steady-state WE validation set, and beyond the validation of the probability flux into the folded state, we present additional validation of the steady-state distribution Fig. S5 and the likelihood of the validation set trajectories in the haMSM models, see Fig. S6 of the supplementary information.

Protein G

We expect that in larger, slower-to-relax, and more computationally expensive systems the haMSM method will have the most utility. We know that protein G, sampled out to $t_{\text{mol}} = 15\text{ns}$, was not yet in SS by observing the flux profile along the WE progress coordinate shown in Fig. 8B. Experimental stopped flow kinetic measurements^{49,50} and coarse-grained structure-based simulations⁵¹ suggest that protein G has long-lived metastable on-pathway folding intermediates. Here we extend upon the protein G WE simulations and haMSM analysis we reported in Ref. 14, finding that in our atomistic WE folding simulations we observe slow relaxation to steady-state indicative of the presence of long-lived intermediates on the folding pathway. Using the haMSM estimated steady-state distribution projected along the fraction of native contacts, we find evidence for multiple metastable intermediates, see supplementary information Fig. S7.

Applying haMSM analysis to the original WE simulations (2D progress coordinate, friction $\gamma = \frac{5}{\text{ps}}$, 15 simulations of 15ns, aggregate simulation time 225 μs) haMSMs with 10^4 microbins constructed from the last 5ns of WE simulation predict a SS flux about 3 orders of magnitude higher than the direct flux, shown in Fig. 8. Constructing haMSMs from this trajectory ensemble, we observe that the predicted MFPT depends strongly upon the number of microbins, consistent with the use of training data in the transient regime seen above.

To validate haMSM estimates, WE simulations were re-initialized in the haMSM-estimated SS and seen to maintain a consistent probability flux into the folded state consistent with the experimentally derived folding rate: see Fig. 8A. We initialized a set of 15 independent WE

simulations (1D progress coordinate 1D-WE) in the haMSM-estimated SS; see section “Initializing WE simulations in estimated SS” for details. These reweighted simulations (rw1) were run for an additional 15ns. Aggregate simulation time of the reweighted simulations was 76 μ s. The lower bound of the direct flux CR remained inside the initial haMSM prediction CR, while the upper bound increased by a factor of ~ 10 . These flux values remained consistent for a second round of restarted WE simulations (rw2), as well as haMSM analysis of the reweighted data. Figure 8 shows that the haMSM estimated SS flux did not vary systematically with the number of microbins (10–10,000) and were within the direct flux CR, indicative of convergence to SS. The protein folding time estimated via the Hill relation Eq. (1) of 0.3–4.3ms is consistent with the experimentally derived folding time, which we take to be somewhere between the reported value of 3.1ms at pH 4.0⁴⁹ and 39ms at pH 11.2.⁵² Note that the lowered friction (1/16 of water) of the Langevin simulation should decrease the observed MFPT of the WE simulation. Also note that the experimental protein folding time stated in our prior work¹⁴ was erroneous.

Fig. 8B indicates that the convergence to SS improves with each reweighting, but even after an additional 15ns of 1D-WE simulation in the haMSM estimated SS (rw1), and a second round of reweighting and 2ns of 1D-WE simulation, the flux profile is not yet flat: true SS is not reached – see Discussion. We speculate that relaxation processes far from the folded state may be very slow but not strongly contributing to the folding process. However, a weaker property of SS is that the haMSM-estimated SS flux becomes independent of the binning, which is observed when analyzing the WE simulations (rw1) initialized in the haMSM estimated SS (see Fig. 8A).

The protein G folding system, where long-lived metastable intermediates make the relaxation to SS very slow, is a system where there is a large benefit of using a haMSM to accelerate SS estimation. The SS relaxation time in the in silico system is not known, but ultrarapid stopped flow kinetic measurements indicated the presence of an on-pathway metastable intermediate with a lifetime of 600–700 μ s⁵⁰ experimentally, which we take to be an estimate for the SS relaxation time. This $\tau_{SS} > 10\mu$ s would make WE simulation (without haMSM analysis) prohibitively costly, because hundreds of trajectories are integrated in parallel.

Discussion

The preceding data shows that the use of short-lag-time haMSMs with thousands of “microbins” can greatly accelerate the estimation of steady-state target fluxes, which translate directly into MFPT estimates. Not only does the haMSM analysis approach remove lag-time uncertainty, but we expect short lag times to be critical to analyzing mechanistic details on the ns scale and below. However, even though MFPT values evidently can be obtained prior to full SS relaxation, key challenges remain to determine the absolute convergence to the unique SS distribution.

Our investigation of flux profiles suggest that the determination of robust metrics to measure SS convergence remains an open problem which limits the ability to estimate the systematic error of rate estimation using the Hill relation. A flat flux profile along a reaction coordinate

which separates the macrostates (folded and unfolded) of interest is a requirement for SS convergence.^{20,37,38} However, this requires global convergence even in regions of configurational space which may not be important to the transition of interest, and our data indicates a robust temporal regime where accurate kinetic rates can be obtained before this global convergence is reached: see Figs. 6 and Fig. 8. Nevertheless, we believe it will be valuable to continue to examine the flux profile to understand whether steady state has truly been achieved.

We expect that iterative procedures which alternate between WE path sampling simulation and haMSM SS estimation steps (schematic shown in Fig. 1) will be the most efficient in determining converged haMSM transition rates. WE trajectory analysis using haMSMs at many levels of resolution (e.g. numbers of microbins) can be used as a practical measure of convergence, where rate estimates become insensitive to the details underlying the haMSM construction as SS is effectively approached, see Figs. 5 and 8. This property can be used as a less stringent indicator of effective SS convergence, similar to how the leveling off of implied timescales is an indicator of Markovian behavior in MSM model building. In contrast, haMSM estimation of SS during the initial transient regime is very sensitive to the number of microbins and the dimensionality reduction method performed preceding the clustering, as suggested by our data (Figs. 4, 5, 6, 8). The efficiency of convergence to SS, beyond the system-specific presence or lack of metastable intermediates, will depend upon how rapidly and effectively the configurational space can be mapped and sampled.

The procedure of launching new WE simulations in the haMSM estimated steady-state and probing for steady-state relaxation, as described and performed in the 2D random energy model and the protein G folding model, is a rigorous validation of haMSM models and steady-state probability flux prediction. However, since it is not always going to be possible to achieve a completely relaxed steady-state, and moreover since additional simulations are computationally expensive and probably not practical to guide haMSM hyperparameter selection, it will be important to develop and apply cross-validation procedures^{41,48,53,54} which rely on only the haMSM model and existing WE trajectories in the transient regime, without a steady-state validation set.

Although the force field does not appear to have biased our results to a significant degree, a few issues deserve mention. Here we have studied implicit solvent low-friction models of protein folding which have allowed for the proof of principle determination of folding times in qualitative agreement with experimental rates, while maintaining computational tractability. The validity of the rates and mechanisms observed in any biomolecular simulation are dependent on the quality of the system description (force-field, solvent, etc.). The protein folding times estimated in this work, rescaled by the ratio of the friction coefficient of water at room temperature to the friction used in the implicit solvent Langevin simulations (a factor of 16), are encouragingly within an order of magnitude of experimental folding times. There are many reasons why this comparison is qualitative only. For Brownian dynamics in simple systems, the friction simply rescales time, but the behavior is less straightforward in complex biomolecular systems.^{55,56} It is certainly physically reasonable to expect that the low-friction systems should have a shorter MFPT compared to high friction, but it is beyond the scope of this work to validate a time rescaling for the

implicit solvent simulations. Moreover, under any simulation protocol, the protein folding MFPT is sensitive to the exact definition of the folded/unfolded states.

While the original WE simulations¹⁴ were initialized from a single unfolded state, the ensemble of unfolded protein structures comprising the source state evolves toward a SS distribution in the WE bin encompassing the initial unfolded structure. Though trajectories which reach the folded state are fed back into the single initial unfolded state, the WE procedure merges trajectories randomly by weight, and trajectories near the folded state (sink) invariably have low weight due to the sink boundary condition (see supplementary Fig. S8). Hence, restarted trajectories in the single unfolded state immediately get pruned – i.e., merged with trajectories which have been allowed to evolve within the WE bin of the unfolded state; correspondingly, the nominal (and artificial) unfolded starting configuration is never observed in the SS re-initialized simulations. Thus it is more accurate to consider the source state of the folding simulations as being defined by the sub-ensemble of structures which map to the WE bin from which the single unfolded state belonged. Supplementary Fig. S9 shows the difference between the initial unfolded structure and the ensemble of unfolded structures from the same WE bin for the SS re-initialized protein G simulations.

In this work, we have focused on kinetics and the MFPT, but haMSMs can provide mechanistic insight in the same way that MSMs are often used to elucidate pathways and mechanisms in biomolecular processes.^{33,44,57,58} We expect that the short lag times and granular structural detail of the haMSMs developed here will allow more realistic determination of kinetic pathways and mechanisms⁵⁹ complementing the unbiased rate determination. We speculate that the identification of structural states/pathways which are low direct weight during the transient regime, but important for the SS folding process, are key to the capability of haMSMs with fine structural detail to estimate the SS distribution—see comparison of transient and SS distributions for the 2D system in Fig. 2, NTL9 folding in 7, and protein G folding in supplementary Fig. S8.

As a proof of principle, we focused on the application of the WE sampling and haMSM modeling method to two protein folding conformational transitions we had previously investigated. However, the framework applies to any dynamical process which can be framed as an A-to-B transition, notably protein-ligand binding/unbinding. We are actively developing well-validated protein+ligand WE simulations to develop the application of the haMSM modeling for the determination of protein-ligand binding/unbinding rates and mechanism.

Conclusions

By using fine bins or “microbins” generated from pyEMMA clustering, we show that WE trajectory data from the transient (pre-steady-state) regime can be used to construct unbiased haMSMs which reliably estimate steady state kinetics. The fine bins can sidestep potentially long relaxation times that will occur if there are internal barriers within large WE bins. If a haMSM is not used, standard or “direct” rate calculation from WE would otherwise require reaching the steady state which likely is impractical for many systems of interest. Therefore, the approach developed here could be of considerable practical importance.

We validated the fine-bin haMSM approach in a 2D toy model of a particle in a random energy landscape, obtaining accurate estimation of the MFPT using a training window only up to 1/40 of the SS relaxation time, and in atomistic WE simulation of NTL9 folding, using a training window 1/10 to 1/2 of the SS relaxation time. In the case of the more challenging protein G folding system, we initialized new WE simulations in the haMSM estimated steady-state which determined a folding MFPT consistent with experimentally measured values, in a tiny fraction of the experimentally measured lifetime of metastable intermediates. Nevertheless, key challenges remain to obtain global SS convergence as we have detailed.

Overall, the accelerated estimation of long-timescale kinetics using WE simulation combined with haMSM analysis, demonstrated here, supports the ongoing application to more accurate but computationally expensive explicit solvent models, and larger systems.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We are very appreciative of helpful discussions with David Aristoff, Gideon Simpson, Barmak Mostofian, Ernesto Suarez and Lillian Chong. Computational resources were provided by the University of Pittsburgh Center for Research Computing, and the Advanced Computing Center and Exacloud cluster at Oregon Health & Science University. This work was supported by NIH Grant R01GM115805.

References

- (1). Huber GA; Kim S Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophysical journal* 1996, 70, 97–110. [PubMed: 8770190]
- (2). Kahn H; Theodore H Estimation of particle transmission by random sampling. *National Bureau of Standards Applied Mathematics Series* 1951, 12, 27–30.
- (3). Zhang BW; Jasnow D; Zuckerman DM The “weighted ensemble” path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *The Journal of Chemical Physics* 2010, 132, 054107. [PubMed: 20136305]
- (4). Aristoff D Analysis and optimization of weighted ensemble sampling. *ESAIM: Mathematical Modelling & Numerical Analysis* 2018, 52.
- (5). Aristoff D; Zuckerman DM Optimizing weighted ensemble sampling of steady states. *arXiv preprint arXiv:1806.00860* 2018,
- (6). Donovan RM; Sedgewick AJ; Faeder JR; Zuckerman DM Efficient stochastic simulation of chemical kinetics networks using a weighted ensemble of trajectories. *Journal of Chemical Physics* 2013, 139.
- (7). Donovan RM; Tapia J-J; Sullivan DP; Faeder JR; Murphy RF; Dittrich M; Zuckerman DM Unbiased Rare Event Sampling in Spatial Stochastic Systems Biology Models Using a Weighted Ensemble of Trajectories. *PLOS Computational Biology* 2016, 12, e1004611. [PubMed: 26845334]
- (8). Tse MJ; Chu BK; Gallivan CP; Read EL Rare-event sampling of epigenetic landscapes and phenotype transitions. *PLoS Computational Biology* 2018, 14, e1006336. [PubMed: 30074987]
- (9). Adelman JL; Grabe M Simulating current-voltage relationships for a narrow ion channel using the weighted ensemble method. *Journal of Chemical Theory and Computation* 2015, 11, 1907–1918. [PubMed: 26392816]
- (10). Zwier MC; Pratt AJ; Adelman JL; Kaus JW; Zuckerman DM; Chong LT Efficient Atomistic Simulation of Pathways and Calculation of Rate Constants for a Protein-Peptide Binding Process:

Application to the MDM2 Protein and an Intrinsically Disordered p53 Peptide. *Journal of Physical Chemistry Letters* 2016, 7, 3440–3445.

- (11). Dixon T; Lotz SD; Dickson A Predicting ligand binding affinity using on - and off - rates for the SAMPL6 SAMPLing challenge. *Journal of Computer-Aided Molecular Design* 2018, 0, 0.
- (12). Dickson A Mapping the Ligand Binding Landscape. *Biophysical Journal* 2018, 115, 1707–1719. [PubMed: 30327139]
- (13). Saglam AS; Chong LT Protein-protein binding pathways and calculations of rate constants using fully-continuous, explicit-solvent simulations. *Chemical Science* 2019, 10, 2360–2372. [PubMed: 30881664]
- (14). Adhikari U; Mostofian B; Copperman J; Subramanian SR; Petersen AA; Zuckerman DM Computational Estimation of Microsecond to Second Atomistic Folding Times. *Journal of the American Chemical Society* 2019,
- (15). Hill TL Free energy transduction and biochemical cycle kinetics; Courier Corporation, 2005.
- (16). Suarez E; Lettieri S; Zwier MC; Stringer CA; Subramanian SR; Chong LT; Zuckerman DM Simultaneous computation of dynamical and equilibrium information using a weighted ensemble of trajectories. *Journal of Chemical Theory and Computation* 2014, 10, 2658–2667. [PubMed: 25246856]
- (17). Suárez E; Pratt AJ; Chong LT; Zuckerman DM Estimating first-passage time distributions from weighted ensemble simulations and non-Markovian analyses. *Protein Science* 2016, 25, 67–78. [PubMed: 26131764]
- (18). Suarez E; Adelman JL; Zuckerman DM Accurate estimation of protein folding and unfolding times: beyond Markov state models. *Journal of Chemical Theory and Computation* 2016, 12, 3473–3481. [PubMed: 27340835]
- (19). Zuckerman DM; Chong LT Weighted ensemble simulation: review of methodology, applications, and software. *Annual review of biophysics* 2017, 46, 43–57.
- (20). Copperman J; Aristoff D; Makarov DE; Simpson G; Zuckerman DM Transient probability currents provide upper and lower bounds on non-equilibrium steady-state currents in the Smoluchowski picture. *Journal of Chemical Physics* 2019,
- (21). DeGrave AJ; Chong LT Reducing the impact of transient effects in rate-constant estimation using the weighted ensemble strategy. *bioRxiv* 2018, 453647.
- (22). Warmflash A; Bhimalapuram P; Dinner AR Umbrella sampling for nonequilibrium processes. *The Journal of Chemical Physics* 2007, 127, 114109. [PubMed: 17887830]
- (23). Dickson A; Warmflash A; Dinner AR Nonequilibrium umbrella sampling in spaces of many order parameters. *The Journal of Chemical Physics* 2009, 130, 02B605.
- (24). Vanden-Eijnden E; Venturoli M Exact rate calculations by trajectory parallelization and tilting. *Journal of Chemical Physics* 2009, 131, 044120.
- (25). Chodera JD; Noé F Markov state models of biomolecular conformational dynamics. *Current Opinion in Structural Biology* 2014, 25, 135–144. [PubMed: 24836551]
- (26). Scherer MK; Trendelkamp-Schroer B; Paul F; Pérez-Hernández G; Hoffmann M; Plattner N; Wehmeyer C; Prinz J-H; Noé F PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *Journal of Chemical Theory and Computation* 2015, 11, 5525–5542. [PubMed: 26574340]
- (27). Paul F; Wu H; Vossel M; De Groot BL; Noé F Identification of kinetic order parameters for non-equilibrium dynamics. *Journal of Chemical Physics* 2019, 150, 164120.
- (28). Bhatt D; Zhang BW; Zuckerman DM Steady-state simulations using weighted ensemble path sampling. *The Journal of Chemical Physics* 2010, 133, 014110. [PubMed: 20614962]
- (29). Mostofian B; Zuckerman DM Statistical Uncertainty Analysis for Small-Sample, High Log-Variance Data: Cautions for Bootstrapping and Bayesian Bootstrapping. *Journal of Chemical Theory and Computation* 2019, 15, 3499–3509. [PubMed: 31002504]
- (30). Singhal N; Snow CD; Pande VS Using path sampling to build better Markovian state models: predicting the folding rate and mechanism of a tryptophan zipper beta hairpin. *The Journal of Chemical Physics* 2004, 121, 415–425. [PubMed: 15260562]

- (31). Noé F; Horenko I; Schütte C; Smith JC Hierarchical analysis of conformational dynamics in biomolecules: transition networks of metastable states. *The Journal of Chemical Physics* 2007, 126, 04B617.
- (32). Voelz VA; Bowman GR; Beauchamp K; Pande VS Molecular simulation of ab initio protein folding for a millisecond folder NTL9 (1– 39). *Journal of the American Chemical Society* 2010, 132, 1526–1528. [PubMed: 20070076]
- (33). Plattner N; Noé F Protein conformational plasticity and complex ligand-binding kinetics explored by atomistic simulations and Markov models. *Nature Communications* 2015, 6, 7653.
- (34). Eastman P; Swails J; Chodera JD; McGibbon RT; Zhao Y; Beauchamp KA; Wang LP; Simonnet AC; Harrigan MP; Stern CD; Wiewiora RP; Brooks BR; Pande VS OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Computational Biology* 2017, 13, e1005659. [PubMed: 28746339]
- (35). Zwier MC; Adelman JL; Kaus JW; Pratt AJ; Wong KF; Rego NB; Suárez E; Lettieri S; Wang DW; Grabe M; Zuckerman DM; Chong LT WESTPA: An interoperable, highly scalable software package for weighted ensemble simulation and analysis. *Journal of Chemical Theory and Computation* 2015, 11, 800–809. [PubMed: 26392815]
- (36). Guyer JE; Wheeler D; Warren JA FiPy: Partial differential equations with Python. *Computing in Science & Engineering* 2009, 11.
- (37). Gardiner C Stochastic methods; Springer Berlin, 2009; Vol. 4.
- (38). Risken H; Frank T The Fokker-Planck Equation: Methods of Solution and Applications; Springer Science & Business Media, 1996; Vol. 18.
- (39). Husic BE; McGibbon RT; Sultan MM; Pande VS Optimized parameter selection reveals trends in Markov state models for protein folding. *Journal of Chemical Physics* 2016, 145.
- (40). Husic BE; Pande VS Note: MSM lag time cannot be used for variational model selection. *Journal of Chemical Physics* 2017, 147, 176101.
- (41). Scherer MK; Husic BE; Hoffmann M; Paul F; Wu H; Noé F Variational selection of features for molecular kinetics. *Journal of Chemical Physics* 2019, 150, 194108.
- (42). Bolhuis PG; Chandler D; Dellago C; Geissler PL Transition Path Sampling: Throwing ropes over rough mountain passes, in the dark. *Annual Review of Physical Chemistry* 2002, 53, 291–318.
- (43). Lindorff-Larsen K; Piana S; Dror RO; Shaw DE How Fast-Folding Proteins Fold. *Science* 2011, 334, 517–520. [PubMed: 22034434]
- (44). Schwantes CR; Pande VS Improvements in Markov State Model construction reveal many non-native interactions in the folding of NTL9. *Journal of Chemical Theory and Computation* 2013, 9, 2000–2009. [PubMed: 23750122]
- (45). Nguyen H; Maier J; Huang H; Perrone V; Simmerling C Folding simulations for proteins with diverse topologies are accessible in days with a physics-based force field and implicit solvent. *Journal of the American Chemical Society* 2014, 136, 13959–13962. [PubMed: 25255057]
- (46). Horng JC; Moroz V; Raleigh DP Rapid cooperative two-state folding of a miniature α - β protein and design of a thermostable variant. *Journal of Molecular Biology* 2003, 326, 1261–1270. [PubMed: 12589767]
- (47). Kellogg EH; Lange OF; Baker D Evaluation and optimization of discrete state models of protein folding. *Journal of Physical Chemistry B* 2012,
- (48). McGibbon RT; Pande VS Variational cross-validation of slow dynamical modes in molecular kinetics. *Journal of Chemical Physics* 2015, 142, 124105.
- (49). Park SH; O’Neil KT; Roder H An early intermediate in the folding reaction of the B1 domain of protein G contains a native-like core. *Biochemistry* 1997, 36, 14277–14283. [PubMed: 9400366]
- (50). Park SH; Shastry MC; Roder H Folding dynamics of the B1 domain of protein G explored by ultrarapid mixing. *Nature Structural Biology* 1999, 6, 943–947. [PubMed: 10504729]
- (51). Shimada J; Shakhnovich EI The ensemble folding kinetics of protein G from an all-atom Monte Carlo simulation. *Proceedings of the National Academy of Sciences of the United States of America* 2002, 99, 11175–11180. [PubMed: 12165568]
- (52). Alexander P; Orban J; Bryan P Kinetic Analysis of Folding and Unfolding the 56 Amino Acid IgG-Binding Domain of Streptococcal Protein G. *Biochemistry* 1992, 31, 7243–7248. [PubMed: 1510916]

- (53). Kellogg EH; Lange OF; Baker D Evaluation and optimization of discrete state models of protein folding. *Journal of Physical Chemistry B* 2012, 116, 11405–11413.
- (54). Wu H; Noé F Variational Approach for Learning Markov Processes from Time Series Data. *Journal of Nonlinear Science* 2020, 30, 23–66.
- (55). Shen MY; Freed KF Long time dynamics of Met-enkephalin: Comparison of explicit and implicit solvent models. *Biophysical Journal* 2002, 82, 1791–1808. [PubMed: 11916839]
- (56). Anandakrishnan R; Drozdetski A; Walker RC; Onufriev AV Speed of conformational change: Comparing explicit and implicit solvent molecular dynamics simulations. *Biophysical Journal* 2015, 108, 1153–1164. [PubMed: 25762327]
- (57). Noé F; Fischer S Transition networks for modeling the kinetics of conformational change in macromolecules. *Current Opinion in Structural Biology* 2008, 18, 154–162. [PubMed: 18378442]
- (58). Voelz VA; Bowman GR; Beauchamp K; Pande VS Molecular simulation of ab initio protein folding for a millisecond folder NTL9(1–39). *Journal of the American Chemical Society* 2010, 132, 1526–1528. [PubMed: 20070076]
- (59). Suárez E; Zuckerman DM Pathway Histogram Analysis of Trajectories: A general strategy for quantification of molecular mechanisms. *arXiv* 2018,

Figure 1: Using haMSM trajectory analysis of WE simulations to estimate steady state (SS) behavior. WE simulation trajectories are used to construct haMSM models used to estimate the SS distribution and the effective protein folding rate ($k_{\text{fold}} = \text{Flux}(\text{unfolded} \rightarrow \text{folded}; \text{SS})$) and check for convergence. New WE simulations can then be re-initialized from the haMSM estimated SS for validation or iteratively to continue convergence to SS.

Figure 2:

2D random energy landscape and distributions. **A** Potential energy in the domain (units of $k_B T$ at $T = 300 K$) with source (A) and sink (B) states labeled, and weighted ensemble binning of the distance to the sink (black lines). **B** SS distribution $-\log p_{SS}$ of the one-way feedback process from numerical solution of the Smoluchowski equation. **C** Transient distribution $-\log p(t)$ from WE simulation at roughly 1/50 of the SS relaxation time, $t = 2.5 ns$ (MFPT $\sim 1.0 \mu s$). **D** haMSM estimated SS distribution $-\log p_{SS}$ using WE simulation training set up to $t = 2.5 ns$. 2D images from WE distributions have been smoothed with a 1 pixel gaussian kernel for visual clarity.

Figure 3:

Microbin dependence of the haMSM model in the 2D random energy system. Top: haMSM-estimated SS flux from the WE simulation training set using a training window from $t = 0$ – 2.5 ns, as a function of the number of microbins used in the haMSM model (black squares). Estimated SS flux saturates within the min/max region from numerical solution of the Smoluchowski equation (shaded blue) with $\sim 10^3$ microbins. Bottom: 2D random energy landscape (black contours) with haMSM microbins (4, 64, 1024) overlaid (colors).

Figure 4:

Validation of steady state behavior obtained from the WE-haMSM pipeline for the 2D random energy system. **A** Direct flux into the target from weighted ensemble simulation validation set (gray lines) and 95% confidence region (shaded gray) as a function of simulation (molecular) time, and WE training set out to $t = 2.5\text{ns}$ (black lines). Min/max region from numerical solution of the Fokker-Planck equation (shaded blue), and 95% confidence region from haMSM SS flux estimates from models with $nC > 10^3$ microbins (orange). Direct flux from reweighted and restarted WE simulations (red lines) and confidence region (shaded red). **B** Flux profile along the distance to the target (sink) state before haMSM reweighting (black triangles) and after reweighting (red triangles). Filled

left-pointing triangles depict flux directed towards the target, and empty right-pointing triangles depict flux directed away from the target.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 5:

Steady state estimation for NTL9 protein-folding using haMSMs with PCA and VAMP microbins. Shown is the haMSM-estimated SS flux for NTL9 (low-friction) protein folding at $t_{mol} = 0.5, 0.63, 1.0, 1.5, 2.5, 4.0$ ns (colored squares) as a function of the number of haMSM microbins, and SS flux from WE validation set (shaded blue horizontal bar). **A** Principal Component Analysis (PCA) dimensionality reduction. **B** VAMP dimensionality reduction.

Figure 6:

NTL9 (low-friction) folding target flux and current profile. **A** Effect of varying training haMSM training window. Predicted SS flux from 3 separately calculated haMSMs with 10,000 microbins clustered with k-means using an all-atom minimum RMSD metric (red diamonds, filled), PCA retaining 10 dimensions (green squares, unfilled), VAMP retaining 10 dimensions (gold circles, filled) and 100 dimensions (orange circles, unfilled) plotted at the final iteration of the training window used (in every case, the training window starts at the first iteration). The spread in values reflects stochastic variation in the clustering process based on identical training data. The haMSM prediction CR (shaded gray) is from all haMSM models at each training window, the validation CR (shaded blue) is from the 10 WE simulations reported in Ref. 14, and the direct flux confidence region (shaded red) is from

the independent test set of 5 WE simulations. **B** Flux profile along the RMSD to the folded state from the training set (red triangles) and from the validation set (black triangles). Filled left-pointing triangles depict flux directed towards the target, and empty right-pointing triangles depict flux directed away from the target.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Figure 7:

NTL9 folding landscapes based on different coordinates. **A** PCA landscapes (x-axis PC1, y-axis PC2) at $t_{\text{mol}} = 0.5, 1.0, 3.0$ ns constructed from 2D-WE protein folding simulations. Left: Scatter plot of RMSD to folded structure. Middle: haMSM estimated SS distribution $-\log p_{\text{SS}}$ Right: Difference in haMSM estimated SS distribution and direct transient distribution (at time t_{mol}) from WE $-\log p_{\text{SS}} + \log p_{\text{direct}}$. Red (blue) shows an increase (decrease) in probability of the SS compared to the direct transient. **B** VAMP landscapes (x-axis VAMP1, y-axis VAMP2) at $t_{\text{mol}} = 0.5, 1.0, 3.0$ ns constructed from 2D-WE protein folding simulations. Left: RMSD Middle: haMSM estimated SS. Right: Difference in haMSM SS and direct transient distribution. 2D images of distributions have been smoothed with a 1 pixel gaussian kernel for visual clarity.

Figure 8:

Protein G (low-friction) folding target flux and current profile. **A** Flux into the folded state as a function of t_{mol} . Direct flux CR from a set of 15 2D-WE simulations initialized from the unfolded state (shaded gray) and the direct flux CR from a set of 15 reweighted/restarted (rw1) 1D-WE simulations initialized in the haMSM estimated SS (shaded red). A second set of 15 reweighted/restarted (rw2) 1D-WE simulations initialized in the haMSM estimated SS (shaded green) with rw1 as training data. haMSM estimated SS flux with 100, 1000, 5000, and 10,000 microbins (blue - light green dashed lines) and CR reflecting the variation between individual WE runs and haMSM analysis (with 10,000 microbins) plotted within their respective training windows, and experimental folding rate at pH 4.0 (gold dashed line). **B** Flux profile along the RMSD to the folded state before haMSM reweighting (black

triangles) and after reweighting (red triangles). Flux profile at the end of rw1 simulations (maroon triangles) and beginning of rw2 simulations (green triangles). Filled left-pointing triangles depict flux directed towards the target, and empty right-pointing triangles depict flux directed away from the target.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript