Rapid equilibrium sampling initiated from nonequilibrium data

Xuhui Huang[a], Gregory R. Bowman[b], Sergio Bacallado[c], and Vijay S. Pande[b,d,1]

[a]Department of Bioengineering, [b]Biophysics Program, [c]Department of Structural Biology, and [d]Department of Chemistry, Stanford University, Stanford, CA 94305

Communicated by Hans C. Andersen, Stanford University, Stanford, CA, August 12, 2009 (received for review January 27, 2009)

Simulating the conformational dynamics of biomolecules is extremely difficult due to the rugged nature of their free energy landscapes and multiple long-lived, or metastable, states. Generalized ensemble (GE) algorithms, which have become popular in recent years, attempt to facilitate crossing between states at low temperatures by inducing a random walk in temperature space. Enthalpic barriers may be crossed more easily at high temperatures; however, entropic barriers will become more significant. This poses a problem because the dominant barriers to conformational change are entropic for many biological systems, such as the short RNA hairpin studied here. We present a new efficient algorithm for conformational sampling, called the adaptive seeding method (ASM), which uses nonequilibrium GE simulations to identify the metastable states, and seeds short simulations at constant temperature from each of them to quantitatively determine their equilibrium populations. Thus, the ASM takes advantage of the broad sampling possible with GE algorithms but generally crosses entropic barriers more efficiently during the seeding simulations at low temperature. We show that only local equilibrium is necessary for ASM, so very short seeding simulations may be used. Moreover, the ASM may be used to recover equilibrium properties from existing datasets that failed to converge, and is well suited to running on modern computer clusters.

generalized ensemble methods | Markov state models | molecular dynamics simulations | RNA hairpin folding

The functions of biological macromolecules are in large part determined by their structure and dynamics. As such, many experimental techniques have been developed and applied to probe these properties, each of which has its strengths and weaknesses. Computational methods such as molecular dynamics (MD) and Monte Carlo (MC) simulations have the potential to complement such experiments by modeling the evolution of entire systems with atomic resolution. However, it is extremely difficult to obtain equilibrium sampling of even moderately sized systems in atomic simulations because of the rugged nature of the free energy landscapes that must be explored. Without adequate sampling, it is impossible to validate the parameters, or force fields, that determine the interactions between atoms or to address phenomena that occur on biologically relevant timescales.

Many methods have been developed in an attempt to address the sampling problem. Generalized ensemble (GE) algorithms, such as the replica exchange method (REM), or parallel tempering (1, 2) and simulated tempering (ST) (3, 4), are popular approaches for studying biomolecular folding (5–15). They attempt to overcome the sampling problem by inducing a random walk in temperature space while maintaining canonical sampling at each temperature. At high temperatures, energetic barriers may be crossed easily; at low temperatures, the system is generally constrained to local minima. However, recent studies have shown that GE simulations do not yield converged equilibrium sampling much faster than standard constant temperature MD if the phenomena of interest are non-Arrhenius (11, 12, 15–20).

For example, Zuckerman and Lyman (17) used the Arrhenius equation to argue that the maximum efficiency gain of GE simulations is no more than an order of magnitude at physiological temperatures, and Zheng et al. (18, 19) used a kinetic network model to show that there is an optimal temperature for non-Arrhenius folding kinetics, and any time spent above this temperature will decrease the efficiency of GE simulations. This lack of improvement is the result of the interplay between energy and entropy. Though high temperatures may facilitate the crossing of energetic barriers, entropic barriers will be more difficult to cross (12).

Thus, GE simulations will provide little improvement when the dominant barriers are entropic. Hansmann and coworkers (21, 22) have made some effort to improve the effectiveness of GE algorithms by optimizing the temperature spacing. However, these methods assume that diffusion in temperature space is the rate-limiting process, so crossing entropic barriers in the conformational space—which is the true rate-limiting process—is still a problem. A number of other methods also exist, such as umbrella sampling and milestoning (23). However, these methods require that the dominant reaction coordinate is known a priori, and this information is often unavailable.

The sampling problem is exacerbated by the practice of viewing global equilibration of individual trajectories as a requirement for considering a simulation to have reached equilibrium. Global equilibration is most naturally obtained by running a simulation much longer than the longest relaxation time of the system, so that all degrees of freedom are equilibrated and many uncorrelated samples are generated from each metastable state. For example, the reversible folding metric holds that a simulation has reached equilibrium if there are multiple folding and unfolding events (24). Though this criterion is sufficient, it may not be necessary. Instead of requiring global equilibration of individual trajectories, we suggest that local equilibration may be sufficient. Local equilibration may be achieved by using multiple simulations, each of which visits only a subset of the metastable states with their correct Boltzmann probabilities but that together cover the entire accessible space. Local equilibration may require significantly less wall-clock time because shorter simulations (all of which may be run in parallel) are required. The main difficulty is to analyze multiple simulations appropriately.

Markov state models (MSMs) are a powerful tool that can be used to extract equilibrium properties from a dataset that satisfies the local equilibration criterion. MSMs partition phase space into metastable states such that intrastate transitions are fast, but interstate transitions are slow (25–29). Such separation of timescales ensures that the model is Markovian, that is, the

Author contributions: X.H., G.R.B., S.B., and V.S.P. designed research; X.H., G.R.B., and S.B. performed research; X.H., G.R.B., and S.B. contributed new reagents/analytic tools; X.H., G.R.B., and S.B. analyzed data; and X.H., G.R.B., S.B., and V.S.P. wrote the paper.

The authors declare no conflict of interest.

[1]To whom correspondence should be addressed. E-mail: pande@stanford.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0909088106/DCSupplemental.

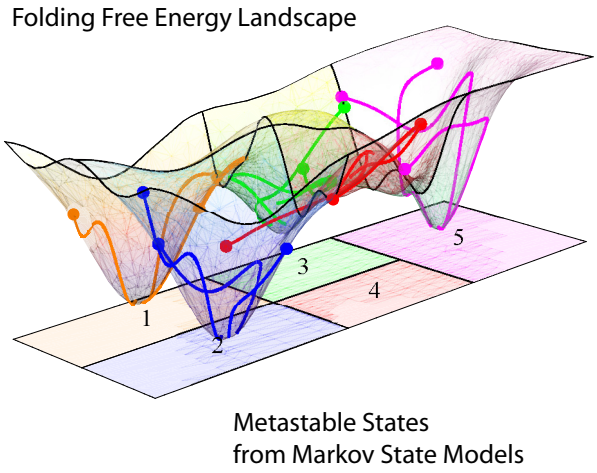


Fig. 1. A schematic free energy landscape with 3 representative seeding trajectories started from each basin and a projection of this free energy landscape onto a 2D plain showing the division into metastable states.

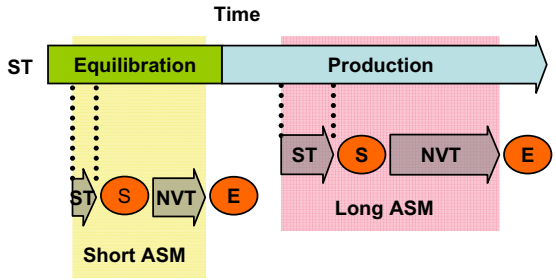


Fig. 2. Schematic of the adaptive seeding scheme. The top arrow represents our ST trajectories, which are split into equilibration (green) and production (light blue) phases. The light red and light yellow boxes encompass our long and short adaptive seeding schemes respectively. For each adaptive seeding scheme, the dotted lines demark the portion of the ST data used to identify the dominant thermodynamic, or metastable, states by building an MSM (S). Constant temperature (or canonical, NVT) simulations are then started from each state and used to build a new MSM (E) that captures the equilibrium distribution. Both the light yellow and pink boxes also encompass a portion of the original ST data that is equivalent to the amount of sampling used in the adaptive seeding scheme. These data are used as a baseline for judging the efficiency of the adaptive seeding scheme.

probability of being in a given state at time $t + \Delta t$, where Δt is called the lag time, depends only on the state at time t. The key point is to build a model with a lag time that is shorter than the timescale of the process of interest with few enough states that it may be understood easily. Usually, MSMs are used to study kinetics, but here we only derive thermodynamic information from them. In an MSM, the time evolution of a vector representing the population of each state may be calculated by repeatedly left-multiplying by the transition probability matrix

$$P(n\Delta t) = [T(\Delta t)]^n P(0) \qquad [1]$$

where P(nΔt) is a vector of state populations at time nΔt, and T is the column-stochastic transition probability matrix. The first left eigenvector of the transition matrix T corresponds to the equilibrium distribution (26). This can be an advantage and a useful opportunity, because obtaining kinetics from MSMs is challenging, but obtaining only the equilibrium thermodynamic properties might be a less demanding goal as less information is required. Indeed, we find that the populations of the dominant states are invariant with respect to the lag time so that very short simulations can be used.

Here, we introduce the adaptive seeding method (ASM) and show that it rapidly yields converged thermodynamics, even when faced with entropic barriers, by exploiting many simulations at local equilibrium. This is achieved by (*i*) using nonequilibrium GE simulations to obtain broad sampling, (*ii*) building an MSM to identify all of the metastable states, as shown in Fig. 1, (*iii*) starting new constant temperature simulations at the temperature of interest from each metastable state in a process called seeding, and (*iv*) using MSMs to extract the correct equilibrium populations from the seeding simulations. Seeding short simulations from the known equilibrium distribution of alanine dipeptide has been shown to yield good models for its kinetics (30). A key advance in our new method is that it does not require that the initial sampling has reached equilibrium. We note that many nonequilibrium GE datasets have been generated due to the difficulty in reaching equilibrium, and that there is growing interest in recovering equilibrium properties from such datasets (31). Thus, one strength of the ASM is that steps 2–4 may be used to recover the correct equilibrium thermodynamic properties from a nonequilibrium dataset. Furthermore, this procedure may be iterated and combined with adaptive sampling (32) to most efficiently use one's computational resources, i.e., using the fewest and shortest trajectories necessary to achieve a good model, because minimizing wall-clock time is an important consideration for computer simulations.

To test the ASM, we apply it to a small biomolecular system with long time-scale dynamics: an 8-nucleotide RNA hairpin (5′-GCUUUUGC-3′) known as the UUUU tetraloop. Hairpins are a fundamental RNA secondary structure motif (33) and perform many biologically relevant functions, but our understanding of their folding is still incomplete (34–38). The folding of this hairpin is diffusion controlled (37, 39, 40), so despite its small size, the folding time is on the microsecond time scale, as measured by laser temperature jump experiments (35). Thus, capturing a single folding event with a single MD simulation with explicit solvent would likely take more than a year on a typical CPU. ASM, however, is able to reach converged equilibrium sampling within a week using many short parallel simulations, as judged by agreement on the populations of metastable states between distinct sets of simulations started from very different initial configurations. ASM is also found to yield converged thermodynamic properties with at least 6 times less sampling than GE simulations for this system. Finally, the fact that the most highly populated metastable state has a well-formed 2-base-pair stem, as in the NMR structure, provides some validation of the force field. Because there is no analytical solution for the equilibrium distribution of our RNA hairpin system, we also studied a 2D potential where the equilibrium populations can be computed analytically. Using this model, we confirm that ASM is much more efficient than ST, and also provide some guidelines for choosing the optimal number and length of the seeding simulations.

Results and Discussion

Comparison of ASM to ST. Here we compare the findings of our long ASM procedure with an equivalent amount of ST sampling, as depicted schematically in Fig. 2. We ran 2 distinct sets of simulations starting from a near-native and random-coil configuration, respectively, as shown in *SI Appendix*, Fig. S1. Thus, we are able to judge the convergence of our findings by comparing these 2 datasets.

The first step was to run an independent set of 1,000 18-ns ST simulations starting from each initial configuration to obtain broad sampling. During an initial equilibration phase (first 9 ns), the weights were updated using the simulated tempering equal acceptance ratio (STEAR) method (11, 41) described in *SI Appendix*. This procedure was found to give nearly equal sam-

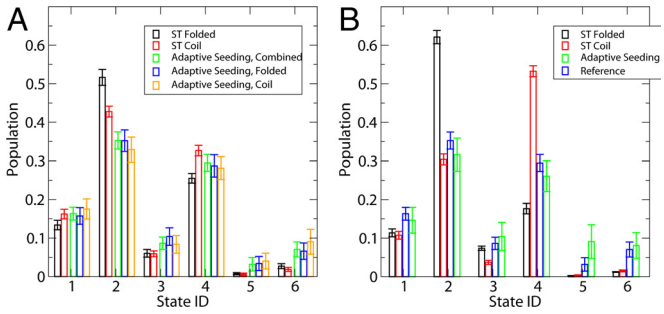


Fig. 3. Population of each state for (*A*) the long adaptive seeding scheme (lag time t = 4.5 ns) and (*B*) the short adaptive seeding scheme (lag time t = 4.5 ns). Bar graphs correspond to the mean values, and error bars represent standard deviations.

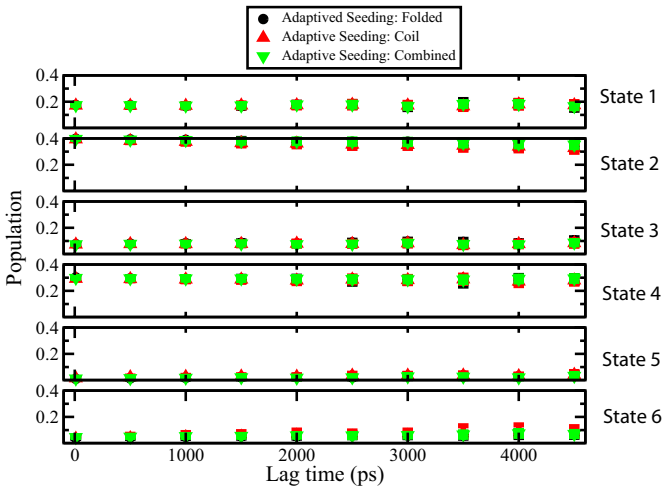


Fig. 4. Population of each state for the long adaptive seeding scheme as the lag time is varied.

pling of each temperature and converged weights for each dataset (*SI Appendix*). During the subsequent 9-ns production phase, the weights were held constant. These 2 sets of ST simulations do not reach converged sampling because of their short length, but they should be able to reach all of the metastable states.

To identify the metastable states, we built an independent MSM from the production phase of each dataset. First, all of the conformations from every temperature were divided into a large number of small sets of very structurally similar, and therefore likely kinetically similar, conformations called microstates using a hierarchical K-medoids clustering algorithm as described in *SI Appendix*. We then used spectral clustering (42) (Perron-Cluster Cluster Analysis [PCCA]; refs. 26, 43, and 44) refined with simulated annealing to lump microstates that can interconvert rapidly into larger states called metastable states, whereas conformations separated by large, free energy barriers are grouped into different states, as depicted in Fig. 1. This algorithm was developed by Chodera et al. (26) and is also described in *SI Appendix*. This procedure yielded 6 states for each dataset.

To obtain equilibrium sampling we then seeded simulations from each metastable state. Specifically, 100 random conformations were chosen from each metastable state and used as starting points for 10-ns constant temperature MD simulations at 300 K. The equilibrium distribution was extracted by building a new MSM. A common state definition is necessary to compare different datasets, so this MSM was built using all of the seeding data. Populations with error bars for each independent dataset were then determined under the same state definition using a Bayesian method developed by Noe (45). Fig. 3*A* shows that the populations from each seeding dataset, as well as the combined data, are in strong agreement and are therefore converged to the equilibrium distribution. Populations for an equivalent amount of folded and coil ST data (19 ns) were also calculated by considering only those conformations at 300 K. These 2 ST datasets have not converged yet. In particular, there is a relatively obvious difference in the populations of states 2 and 4 (about 10% and 7% respectively).

MSMs are usually used to study kinetics (25). To get a reasonable number of states and ensure that the model is Markovian, a relatively long lag time must be used, although it generally ought to be shorter than the timescale of the process of interest. Furthermore, to ensure accurate kinetics, each simulation must be at least a few times longer than the lag time so that multiple crossings of each barrier may be observed. For example, Chodera et al. (26) show that a 20-state MSM for the folding of the helical Fs peptide (which occurs on a timescale of tens of nanoseconds) requires a lag time of 5 ns. Thus, obtaining accurate kinetics for the UUUU tetraloop, which folds on a microsecond timescale, should likely require orders of magnitude longer simulations than for the Fs peptide. However, obtaining accurate thermodynamics may require significantly less sampling. In particular, short lag times where the system is not Markovian may still be sufficient to estimate thermodynamic properties. In fact, Fig. 4 shows that the equilibrium populations of each state are identical within statistical error regardless of the lag time. Similar observations have been made by Hummer and coworkers (29), who found that the free energy profile for a water dewetting transition can be predicted using a very short lag time at which the kinetics are not reproduced well. In addition to the error due to non-Markovian effects, the statistical error due to insufficient sampling of transition events will also be smaller for thermodynamic properties. In a model with N states there are only N thermodynamic parameters to determine, whereas getting accurate kinetics requires determining all N^2 pairwise transition probabilities. Sampling all possible transitions overdetermines the free energy differences between states. Thus, obtaining accurate thermodynamics may require significantly less sampling.

Minimizing the Simulation Length. To push the limits of the ASM, we repeated the previous procedure using drastically less data (see short ASM in Fig. 2). Ten times less data were used for equilibration, 6 times less ST data were used to identify the states, and the seeding simulations were half as long. To maximize our use of this minimal data, we combined the folded and coil ST data to identify the metastable states used for seeding. Fig. 3*B* shows the populations obtained from this procedure compared with a reference distribution from our long ASM runs and an equivalent amount of ST data started from both folded and coil states. All of these populations were determined using the previous state definition. The populations from these short ASM runs were found to be in agreement with the previously determined equilibrium distribution, whereas the ST data deviated significantly from equilibrium. To determine the limits of ASM, we also performed the same analysis using fewer and shorter trajectories. First we held the seeding trajectory length constant at 5 ns and varied the number of trajectories initiated from each state, finding that as few as 70 trajectories from each state gave reasonable agreement with the reference distribution. We also held the number of seeding trajectories started from each state constant at 100 and varied the trajectory length, finding that as little as 2-ns-long seeding simulations gave reasonable agreement with the reference. Thus, our ASM method reaches equilibrium at least 6 times faster than ST. These

findings suggest that the ASM is significantly more efficient than GE simulations for sampling conformational changes that are diffusion controlled, as in hairpin folding.

To address any concerns about the validity of our reference distribution, we also studied a simple model where the equilibrium populations can be computed analytically. The model is based on a discrete-state system introduced by Zwanzig (46) as a simple model for protein folding (see *SI Appendix* for details). There are 4 metastable states in the system (folded, unfolded, and 2 intermediate states), among which the folded state is favored energetically, and the unfolded state is favored entropically (see *SI Appendix*, Fig. S5). This is an attractive system for testing ASM because it has non-Arrhenius folding kinetics (i.e., the folding rate decreases with temperature; see *SI Appendix*, Fig. S7).

We compared the efficiency with which ST and ASM reach the equilibrium state populations as a function of length and number of trajectories. As shown in *SI Appendix*, Fig. S12, ASM converges to the correct distribution with 4–7 times shorter simulations than ST. We suggest that seeding simulations longer than the slowest intramacrostate equilibration time should always be sufficient for convergence. In practice, however, much shorter simulations may be used, as discussed previously. When using shorter simulations, one should test that independent sets of simulations started from different configurations converge to the same distribution, and that the equilibrium distribution is invariant with respect to the lag time. We also found that using more than 200 trajectories does not increase the efficiency of ST, whereas ASM continues to scale favorably with the number of trajectories—up to 600 trajectories in this example. The optimal number of simulations to run depends on one's tolerance for statistical error. Currently, an equal number of simulations are seeded from each state. In the future, however, adaptive sampling (30) could be used to start an optimal number of simulations from each metastable state to further optimize the efficiency of this method.

There are many factors contributing to the improved efficiency of ASM. By using short GE simulations to identify the metastable states, the ASM is able to exploit the ability of GE simulations to rapidly cross energetic barriers, and avoid the penalty incurred at high temperatures for entropic barriers by using seeding simulations at low temperatures. Furthermore, only short seeding simulations are necessary, because the use of MSMs requires only local, rather than global, equilibration. Global equilibration metrics, such as reversible folding, require that each simulation is long enough to cross every barrier multiple times. Local equilibration, however, may be obtained with many short simulations run in parallel because each run need only be long enough to cross a single barrier. By using MSMs to identify the metastable states, we can initiate seeding simulations from uncorrelated conformations within every metastable state and thereby ensure every barrier is crossed.

The ASM also has limitations. For example, the initial sampling has to be broad enough to identify all of the metastable states. Failure to do so will quickly become apparent because some states will be populated in one dataset but not in another. This situation may be remedied by iterating the ASM—that is, seeding ST simulations from each state to obtain broader sampling, building an MSM to identify the metastable states, and performing new constant-temperature seeding simulations. In addition, seeding simulations at physiological temperatures are only able to cross barriers on the order of a few kT. However, this should be sufficient for most biological systems. Finally, we note that the random selection of initial configurations from each state may lead to some error if the seeding simulations are not long enough. In the future, this method might be improved by choosing initial configurations from an equilibrium distribution prepared within each metastable state (47).

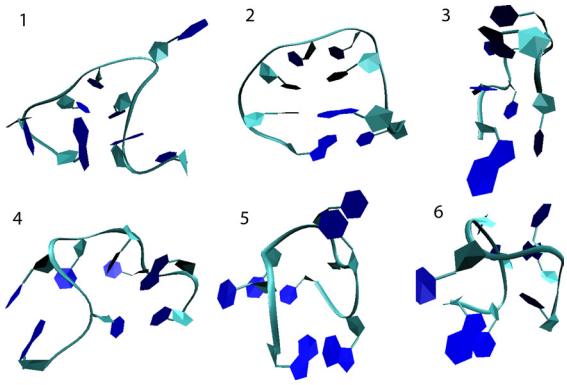


Fig. 5. Representative structure for each of the 6 metastable states. The numbering is the same as in Figs. 3 and 4.

Examining the States. Fig. 3*B* shows that the short folded ST data spent a disproportionate amount of time in state 2, whereas the coil ST data spent a disproportionate amount of time in state 4. Based on this finding, we hypothesized that state 2 is the native state, and that state 4 is a random coil state. To test this hypothesis, we extracted representative structures for each state. The representative structure for each state is the configuration with the greatest density of nearby conformations (mathematically this is the conformation with the minimal rmsd to every other conformation in the state).

Fig. 5 shows the representative structures for each state. In fact, state 2 is the folded state, having a well-formed 2-base-pair stem. Our ability to identify this state without including any knowledge of the native state and the fact that it is the most populated state lends credibility to the force field used, AMBER99 (48). Furthermore, state 4 is a random coil. The other states represent various collapsed nonnative states. For example, state 1 has native-like base stacking interactions but no clear base pairing between the 2 sides of the stem. State 3 has interactions between bases 1 and 8 as well as 2 and 7, but they are stacking interactions instead of base pairing interactions. These findings are consistent with both experimental and computational work showing that small RNA hairpins have folding intermediates with contacting end residues but without well-formed stems (34, 36). Fully validating the force field will require longer simulations to get accurate kinetic predictions and more extensive comparisons with experimental observables.

Conclusions

We have introduced the adaptive seeding method (ASM) and shown that it samples significantly more efficiently than GE simulations, which have found widespread use in studying biological systems (7, 12, 14, 49), for a 2D simple potential and RNA hairpins. The ASM takes advantage of the broad sampling possible with GE methods but can more effectively cross entropic barriers using constant temperature simulations. Moreover, by requiring local equilibration rather than global equilibration, only relatively short simulations are necessary, and these simulations may be run in parallel, rendering the calculation particularly well suited to modern computing clusters. MSMs are then used to extract global equilibrium populations from these short simulations. Besides serving as an efficient sampling algorithm, the ASM may also be used to recover equilibrium properties from nonequilibrium datasets. Thus, the ASM holds great promise for validating force fields and bridging the gap between experimental and computational timescales.

In the future, we plan to apply the adaptive seeding method to larger systems. We also hope to explore alternative sampling methods for identifying initial states. For example, coarse-

grained simulations could be used to identify the dominant states of a system and seed all-atom MD simulations that would elucidate the atomic details of the free energy surface. Alternatively, implicit solvent simulations run at low viscosity could be used to rapidly identify the dominant states and seed explicit solvent simulations to provide more accuracy. Finally, adaptive sampling (32) with longer simulations may be used to obtain accurate kinetics from MSMs.

Materials and Methods

Two distinct sets of ST simulations were run: one started from a folded state and the other from a random coil. An independent MSM was then built for each dataset to identify the dominant metastable states. We use MSMBuilder (50) to build an MSM. At first, conformations were first split into a large number of microstates using a hierarchical K-medoids clustering algorithm with the all-heavy-atom rmsd as the distance metric (e.g., we generated 1,597 microstates for long ASM seeding runs). Kinetically related microstates were then lumped together using PCCA (26, 43, 44). One hundred random conformations were then chosen from each state and used as starting points for constant temperature 300-K MD simulations, still maintaining 2 distinct sets of simulations. New MSMs were built from these constant temperature datasets. A Bayesian method (46) was used to calculate the populations of each state with error bars, and the models were compared based on these values. The original ST simulations were also extended to match the sampling of the constant temperature simulations. State populations with error bars for these long ST runs were computed using bootstrapping and compared to the populations from the constant temperature simulations. More details are available in *SI Appendix*.

ACKNOWLEDGMENTS. We thank Dr. John Chodera, Dr. Vince Voelz, and Daniel Ensign for their helpful insights on sampling and MSM building. X.H. acknowledges the support of Professor Michael Levitt. Support was received from the National Institutes of Health Roadmap Grant U54 GM072970 (to X.H.), the National Science Foundation Graduate Research Fellowship Program (to G.R.B.), and a Stanford Graduate Fellowship (to S.B.). This work was also funded by National Institutes of Health Grants R01-GM062868 and P01-GM066275. Computer resources were provided by National Science Foundation Award CNS-0619926 and Folding@Home volunteers. MSMBuilder is available on SIMTK (https://simtk.org/home/msmbuilder).

1. Hansmann UH, Okamoto Y (1999) New Monte Carlo algorithms for protein folding. *Curr Opin Struct Biol* 9:177–183.
2. Sugita Y, Okamoto Y (1999) Replica-exchange molecular dynamics method for protein folding. *Chem Phys Lett* 141–151.
3. Lyubartsev AP, et al. (1992) New approach to Monte Carlo calculation of the free energy: Method of expanded ensembles. *J Chem Phys* 96:1776–1783.
4. Marinari E, Parisi G (1992) Simulated tempering: A new Monte Carlo scheme. *Europhys Lett* 19:451–458.
5. Zhou R, Berne BJ, Germain R (2001) The free energy landscape for beta hairpin folding in explicit water. *Proc Natl Acad Sci USA* 98:14931–14936.
6. Mohanty S, Hansmann UH (2006) Folding of proteins with diverse folds. *Biophys J* 91:3573–3578.
7. Nymeyer H, Garcia AE (2003) Simulation of the folding equilibrium of alpha-helical peptides: A comparison of the generalized Born approximation with explicit solvent. *Proc Natl Acad Sci USA* 100:13934–13939.
8. Liu P, et al. (2006) Hydrophobic aided replica exchange: An efficient algorithm for protein folding in explicit solvent. *J Phys Chem B* 110:19018–19022.
9. Im W, Brooks CL III (2004) De novo folding of membrane proteins: An exploration of the structure and NMR properties of the fd coat protein. *J Mol Biol* 337:513–519.
10. Roitberg AE, Okur A, Simmerling C (2007) Coupling of replica exchange simulations to a non-Boltzmann structure reservoir. *J Phys Chem B* 111:2415–2418.
11. Huang X, Bowman GR, Pande VS (2008) Convergence of folding free energy landscapes via application of enhanced sampling methods in a distributed computing environment. *J Chem Phys* 128:205106.
12. Rhee YM, Pande VS (2003) Multiplexed-replica exchange molecular dynamics method for protein folding simulation. *Biophys J* 84:775–786.
13. Pitera JW, Swope WC, Abraham FF (2008) Observation of noncooperative folding thermodynamics in simulations of 1BBL. *Biophys J* 94:4837–4846.
14. Zhang W, Wu C, Duan Y (2005) Convergence of replica exchange molecular dynamics. *J Chem Phys* 123:154105.
15. Periole X, Mark AE (2007) Convergence and sampling efficiency in replica exchange simulations of peptide folding in explicit solvent. *J Chem Phys* 126:014903.
16. Nymeyer H (2008) How efficient is replica exchange molecular dynamics? An analytic approach. *J Chem Theory Comput* 4:626–636.
17. Zuckerman DM, Lyman E (2006) A second look at canonical sampling of biomolecules using replica exchange simulation. *J Chem Theory Comput* 2:1200–1202, and erratum (2006) 2:1693.
18. Zheng W, et al. (2008) Simple continuous and discrete models for simulating replica exchange simulations of protein folding. *J Phys Chem B* 112:6083–6093.
19. Zheng W, et al. (2007) Simulating replica exchange simulations of protein folding with a kinetic network model. *Proc Natl Acad Sci USA* 104:15340–15345.
20. Sanbonmatsu KY, Garcia AE (2002) Structure of Met-enkephalin in explicit aqueous solution using replica exchange molecular dynamics. *Proteins* 46:225–234.
21. Nadler W, Hansmann UH (2007) Dynamics and optimal number of replicas in parallel tempering simulations. *Phys Rev E Stat Nonlin Soft Matter Phys* 76:065701.
22. Nadler W, Hansmann UH (2007) Optimizing replica exchange moves for molecular dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys* 76:057102.
23. Faradjian AK, Elber R (2004) Computing time scales from reaction coordinates by milestoning. *J Chem Phys* 120:10880–10889.
24. Rao F, Caflisch A (2003) Replica exchange molecular dynamics simulations of reversible folding. *J Chem Phys* 119:4035–4042.
25. Noe F, Fischer S (2008) Transition networks for modeling the kinetics of conformational change in macromolecules. *Curr Opin Struct Biol* 18:154–162.
26. Chodera JD, et al. (2007) Automatic discovery of metastable states for the construction of Markov models of macromolecular conformational dynamics. *J Chem Phys* 126:155101.
27. Buchete NV, Hummer G (2008) Coarse master equations for peptide folding dynamics. *J Phys Chem B* 112:6057–6069.
28. Swope WC, Pitera JW, Suits F (2004) Describing protein folding kinetics by molecular dynamics simulations. 1. Theory. *J Phys Chem B* 108:6571–6581.
29. Sriraman S, Kevrekidis IG, Hummer G (2005) Coarse master equation from Bayesian analysis of replica molecular dynamics simulations. *J Phys Chem B* 109:6479–6484.
30. Gerhard H, Ioannis GK (2003) Coarse molecular dynamics of a peptide fragment: Free energy, kinetics, and long-time dynamics computations. *J Chem Phys* 118:10762–10773.
31. Ytreberg FM, Zuckerman DM (2008) A black-box re-weighting analysis can correct flawed simulation data. *Proc Natl Acad Sci USA* 105:7982–7987.
32. Singhal N, Pande VS (2005) Error analysis and efficient sampling in Markovian state models for molecular dynamics. *J Chem Phys* 123:204909.
33. Uhlenbeck OC (1990) Tetraloops and RNA folding. *Nature* 346:613–614.
34. Bowman GR, et al. (2008) Structural insight into RNA hairpin folding intermediates. *J Am Chem Soc* 130:9676–9678.
35. Ma H, et al. (2006) Exploring the energy landscape of a small RNA hairpin. *J Am Chem Soc* 128:1523–1530.
36. Ma H, et al. (2007) DNA folding and melting observed in real time redefine the energy landscape. *Proc Natl Acad Sci USA* 104:712–716.
37. Ansari A, Kuznetsov SV, Shen Y (2001) Configurational diffusion down a folding funnel describes the dynamics of DNA hairpins. *Proc Natl Acad Sci USA* 98:7771–7776.
38. Sorin EJ, Rhee YM, Pande VS (2005) Does water play a structural role in the folding of small nucleic acids? *Biophys J* 88:2516–2524.
39. Ansari A, Kuznetsov SV (2005) Is hairpin formation in single-stranded polynucleotide diffusion-controlled? *J Phys Chem B* 109:12982–12989.
40. Levitt M (1972) Folding of nucleic acids. *Polymerization in Biological Systems*, Ciba Foundation Symposium 7, eds Wolstenholme GEW, O'Connor M (Elsevier, Amsterdam), pp 146–171.
41. Bowman GR, Pande VS (2009) Simulated tempering yields insight into the low-resolution Rosetta scoring functions. *Proteins* 74:777–778.
42. Schütte C, Huisinga W (2003) Biomolecular conformations can be identified as metastable sets of molecular dynamics. *Handbook of Numerical Analysis: Special Volume: Computational Chemistry, Vol. X*, ed Ciarlet PG (Elsevier, Amsterdam), pp 699–744.
43. Deuflhard P, et al. (2000) Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains. *Lin Alg Appl* 315:39–59.
44. Deuflhard P, Weber M (2005) Robust Perron cluster analysis in conformation dynamics. *Lin Alg Appl* 398:161–184.
45. Noe F (2008) Probability distributions of molecular observables computed from Markov models. *J Chem Phys* 128:244103.
46. Zwanzig R (1995) Simple model of protein folding kinetics. *Proc Natl Acad Sci USA* 92:9801–9804.
47. Chodera JD, et al. (2006) Long-time protein folding dynamics from short-time molecular dynamics simulations. *Multiscale Model Simul* 5:1214–1226.
48. Wang J, Cieplak P, Kollman PA (2000) How well does a restrained electrostatic potential (RESP) model perform in calculating conformational energies of organic and biological molecules? *J Comp Chem* 21:1049–1074.
49. Garcia AE, Paschek D (2008) Simulation of the pressure and temperature folding/unfolding equilibrium of a small RNA hairpin. *J Am Chem Soc* 130:815–817.
50. Bowman GR, Huang X, Pande VS (2009) Using generalized ensemble simulations and Markov state models to identify conformational states. *Methods*, 10.1016/j.ymeth.2009.04.013.