

RESEARCH ARTICLE | SEPTEMBER 23 2022

Configurational entropy, transition rates, and optimal interactions for rapid folding in coarse-grained model proteins

Margarita Colberg  ; Jeremy Schofield 



J. Chem. Phys. 157, 125101 (2022)

<https://doi.org/10.1063/5.0098612>





Lock-in Amplifier



Boxcar Averager

Boost Your Optics and Photonics Measurements

 Zurich Instruments

[Find out more](#)

Configurational entropy, transition rates, and optimal interactions for rapid folding in coarse-grained model proteins

Cite as: J. Chem. Phys. 157, 125101 (2022); doi: 10.1063/5.0098612

Submitted: 10 May 2022 • Accepted: 25 August 2022 •

Published Online: 23 September 2022



Margarita Colberg^{a)}  and Jeremy Schofield^{b)} 

AFFILIATIONS

Chemical Physics Theory Group, Department of Chemistry, University of Toronto, Toronto, Ontario M5S 3H6, Canada

^{a)} Author to whom correspondence should be addressed: margarita.gladikh@mail.utoronto.ca

^{b)} Electronic mail: jeremy.schofield@utoronto.ca

ABSTRACT

Under certain conditions, the dynamics of coarse-grained models of solvated proteins can be described using a Markov state model, which tracks the evolution of populations of configurations. The transition rates among states that appear in the Markov model can be determined by computing the relative entropy of states and their mean first passage times. In this paper, we present an adaptive method to evaluate the configurational entropy and the mean first passage times for linear chain models with discontinuous potentials. The approach is based on event-driven dynamical sampling in a massively parallel architecture. Using the fact that the transition rate matrix can be calculated for any choice of interaction energies at any temperature, it is demonstrated how each state's energy can be chosen such that the average time to transition between any two states is minimized. The methods are used to analyze the optimization of the folding process of two protein systems: the crambin protein and a model with frustration and misfolding. It is shown that the folding pathways for both systems are comprised of two regimes: first, the rapid establishment of local bonds, followed by the subsequent formation of more distant contacts. The state energies that lead to the most rapid folding encourage multiple pathways, and they either penalize folding pathways through kinetic traps by raising the energies of trapping states or establish an escape route from the trapping states by lowering free energy barriers to other states that rapidly reach the native state.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0098612>

I. INTRODUCTION

Coarse-grained models of polymers¹ and proteins^{2,3} are designed to bridge the gap in time scale between the motion of molecular components and slower, large-scale structural changes. A wide variety of such models exist, including Gō lattice models, in which monomers are restricted to lattice sites,^{4–9} elastic networks,^{10–12} and off-lattice linear chain models, which make use of a continuous force field that may include quantum effects in an approximate way.¹³ These models have been used to probe the mechanism of protein folding, the process by which a denatured protein reaches its experimentally determined native structure.¹⁴ The dynamics derived from coarse-grained models indicate that short, fast-folding proteins follow a hierarchical folding process. In this process, the backbone adopts secondary structural elements and a small number of nonlocal contacts early on, and then subsequently

folds in a directed fashion along a dominant pathway.^{15–17} Other simulation work has suggested that the folding is fast and efficient when the system is free of bottlenecks or kinetic traps, and when multiple pathways exist to the final state.^{18,19} Experimental evidence, primarily through studies of *cytochrome c*^{20–22} and RNase,²³ support the picture that the initial phase of the folding process consists of the formation of “foldon” subunits made up of secondary structures.^{23–25}

In recent years, there has been an interest in engineering polymer and protein systems that possess properties similar to those of naturally occurring proteins.^{26–29} Although a number of design considerations have been identified that are associated with secondary and tertiary structure, such as the use of residue sequences that have a propensity to form α -helices and other structural elements or that have hydrophobic side chains to facilitate packing, the relative importance of each design feature for a particular

folded structure is not easily determined.³⁰ Experimental and simulation studies suggest that secondary structure is important in providing the building blocks for foldons that nucleate the folding process,^{17,23–25} but other types of structures might provide a similar framework if they satisfy a set of physical characteristics. Understanding such requirements could provide insight into how to design synthetic polymer systems with protein-like structure and functionality.

Much recent work to connect the primary sequence to protein functionality has been data-driven, using machine learning methods trained on sequence data both predictively and generatively.^{31,32} Prominent among the predictive machine learning models is the AlphaFold³³ project of DeepMind, a subsidiary of Alphabet, Inc. Biophysical models of protein evolution frequently assume that evolution is determined by the sequences that optimize a structurally based “fitness landscape.”^{34,35} For example, lattice models of short proteins, for which an exact enumeration of configurations is feasible, have been used to study the connection between sequence and specific targeted folded structures.^{36–38} In the lattice models, where the dynamics of configurations itself is not well defined, a fitness landscape based on a target structure for a sequence is optimized via random mutations in the sequence space. Rather than being based on dynamical information, the fitness criterion is determined by the free energy of a Potts model,³⁹ defined in terms of the adjacency matrix of contacts in the protein.^{38,40} In contrast, while the microscopic dynamics in off-lattice models is clear, the definition of a fitness criterion for optimal folding is both conceptually and computationally challenging. Part of the difficulty in investigating the molecular features that determine protein structure and its connection to dynamics arises from the intractability of determining how the free energy landscape and dynamical folding pathways depend on sequence structure and external conditions such as the temperature. Unlike most off-lattice coarse-grained models of biomolecular systems, the structure and dynamics of discontinuous potential models based on distance constraints can be examined at any temperature for any choice of interaction energy once the entropy of the system’s states and the distribution of bond distances within each state are known.

The purpose of this paper is twofold: First, we introduce an efficient computational approach to evaluate both the configurational entropy and the mean first passage times for discontinuous potential models, based on adaptive event-driven sampling. We then present a variational optimization procedure in the context of a Markov state model to determine the state energies that minimize the first passage time, subject to a set of structural constraints. In this case, the first passage time is evaluated for a process in which an initial state with no bonds evolves to the fully bonded “native” state under a set of constraints determined by the thermodynamic requirement of a predominant native state population.

The outline for this paper is as follows: In Sec. II, the protein-like model is introduced.^{41–44} In Sec. III A, the configurational entropy is defined and related to the thermodynamic structure of the discontinuous potential model. In Sec. III B, a Markov state model for the simplified dynamics of the evolution of state populations is introduced. The explicit expressions are given for elements of the rate matrix that can be computed using only temperature-independent geometric information. Subsequently, we outline an adaptive procedure based on event-driven dynamics to evaluate the

configurational entropies, as well as the first passage times, that parameterize the rate matrix in the Markov state model. Adaptations to the method are discussed in the two sections that follow. We aim to apply the sampling approach in a massively parallel framework and introduce techniques to improve the rate of convergence in calculations involving states that differ substantially in their configurational entropy, for which the first passage times are large. This is followed by Sec. V, with the introduction of a variational principle to optimize the interactions that lead to rapid folding in the Markov state model, as well as several key-related measures that are useful to characterize the folding mechanism. The variational optimization of the folding time of two different model protein systems—a model of the crambin protein, which is rich in secondary structures, and a small model system that possesses a native state with a highly strained helical structure and frustrated intermediates—is discussed in Secs. V A and V B. Finally, concluding remarks are contained in Sec. VI.

II. THE COARSE-GRAINED MODEL

The model we consider here, similar to one introduced by Zhou and Karplus,^{41–43} is based on a coarse-grained approach in which each amino acid residue of a linear, protein-like chain is represented by a bead. The chain is immersed in a fluid in thermal equilibrium at a temperature T . We assume that the effect of the fluid is to alter the energy of the configurations of the chain and provide a stochastic environment for the motion of the beads diffusing in the fluid. In the chain, there are local and nonlocal bonds that connect the beads. Local bonds occur between nearest and next-nearest neighboring beads. These bonds can correspond to peptide bonds in the primary structure of a protein. Local bonds are modeled using an infinite square well potential,

$$U(r_{ij}) = \begin{cases} 0 & \text{if } \sigma_1 < r_{ij} < \sigma_2, \\ \infty, & \text{otherwise,} \end{cases} \quad (1)$$

where $U(r_{ij})$ is the potential energy of the local bond, r_{ij} is the distance between two nearest or next-nearest neighboring beads i and j , and σ_1 and σ_2 are the minimum and maximum bonding distances, respectively. For nearest neighbors, $\sigma_1 = 1$, which is taken as the unit of length in the model, and $\sigma_2 = 1.17$.⁴⁵ For next-nearest neighbors, $\sigma_1 = 1.4$ and $\sigma_2 = 1.67$ are chosen to restrict the bond angles to be between 75° and 112° to mimic the space that the side chains in amino acids would normally occupy in a protein.

Nonlocal bonds occur between beads that are not nearest or next-nearest neighbors. These bonds account for interactions between the side chains of amino acids in a protein to form its secondary structures. A nonlocal bond k , formed at a distance r_{ck} between beads i and j , is modeled using a step potential,

$$U_k(r_{ij}|\alpha) = \begin{cases} \infty & \text{if } r_{ij} < r_h, \\ \epsilon_k(\alpha) & \text{if } r_h \leq r_{ij} \leq r_{ck}, \\ 0 & \text{if } r_{ij} > r_{ck}, \end{cases} \quad (2)$$

where r_{ij} is the distance between the beads. The energy of the bond $\epsilon_k(\alpha)$ may depend conditionally on the overall configuration α of the system (i.e., the other nonlocal bonding distances). With this

flexible design of the bonding energy, the model can describe systems with nonlocal interactions that effectively allow for non-pairwise interactions in which the energy of a bond depends on the specific configuration involved. In this way, side chain interactions and temperature-dependent solvent effects, such as hydrophobicity, can be incorporated into the model in a mean-field way. At a distance of $r_{ij} = r_h = 1.25$, a hard-core repulsion accounts for excluded volume interactions. At a distance of r_{ck} , a bond forms between two nonlocal beads, which contributes a factor of $e_k(\alpha)$ to the total energy. Note that such a bond is either “on” or “off,” depending on the geometric distance between the beads forming the nonlocal bond. If two beads do not form a nonlocal bond, they will collide elastically at the hard-core repulsion distance r_h .

III. THE THERMODYNAMICS AND DYNAMICS OF THE COARSE-GRAINED MODEL

A. The configurational entropy

A configuration of a system with N monomer beads is specified by the $3N$ -dimensional vector of bead positions $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_N)$, where $\mathbf{r}_i$ is the position vector of bead i in the system. In a model with a step potential and infinite hard wall interactions, physically allowed configurations $\mathbf{R}$ of the system must satisfy distance constraints that force nearest and next-nearest beads in the chain to be within a short distance of one another determined by σ_1 and σ_2 . The entire configurational space of allowed configurations is geometrically partitioned into states of the system by the set of n_b nonlocal bonding distances $\{r_{ck}|k = 1, \dots, n_b\}$. A configurational state c can be represented as a binary string,

$$c = c_1 \dots c_{n_b}, \quad (3)$$

where each term c_i in the string c takes on a binary value of 1 if $x_i < r_{ci}$ and 0 if it is not. For example, for a model with three nonlocal bonds, the configuration 000 refers to an unfolded chain with no nonlocal bonds.

To develop the statistical mechanics of the model, we define the indicator function for a configurational state c ,

$$\mathbb{1}_c(\mathbf{R}) = \begin{cases} 1 & \text{if all constraints for } c \text{ are satisfied,} \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The partitioning of the $3N$ -dimensional space of microscopic configurations enables us to reduce the large number of allowed configurations (equal to the volume of the configurational space) to a finite and discrete set of $n_s = 2^{n_b}$ coarse-grained states for which structural and dynamical properties can be derived.

The coarse-grained model is unusual in that the dimensionless nonideal entropy, S_c , defined by

$$S_c = \ln \left(\frac{1}{V^N} \int \mathbb{1}_c(\mathbf{R}) d\mathbf{R} \right) \quad (5)$$

can be determined entirely by the distance constraints between the beads for any configuration c . The integral of $\mathbb{1}_c$ over the volume of configurations can be viewed as the volume of the subspace occupied by configuration c in the full configurational space.

For a given model with a prescribed set of interaction energies $\{\mathbf{E}\}$, the canonical probability, P_c , of a configuration c with a potential energy E_c at an inverse temperature β^* is

$$P_c = \langle \mathbb{1}_c \rangle = \frac{e^{-\beta^* E_c} e^{S_c}}{\sum_{\alpha=1}^{n_s} e^{-\beta^* E_\alpha} e^{S_\alpha}} = \frac{e^{-\beta^* F_c}}{\sum_{\alpha=1}^{n_s} e^{-\beta^* F_\alpha}}. \quad (6)$$

Here, $\langle \dots \rangle$ denotes the canonical ensemble average, $n_s = 2^{n_b}$ is the total number of configurations, and $F_c = E_c - T^* S_c$ is the free energy of configuration c .

As evident in Eq. (5), the entropy difference between two states is independent of both the temperature of the system and the set of interaction energies. As a result, once the configurational entropy for all states has been determined, the canonical probability of any state for *any* choice of interaction energies at any temperature can be evaluated. This flexibility permits us to examine not only how the morphology of the free energy landscape changes with temperature but also how changing the interaction energies of different states, which is similar to changing the molecular identity of each bead, influences the thermodynamics of the system. The generality of the model also allows for the effects of different interaction energies on aspects of the dynamics, such as the structural folding time, to be examined.

B. The transition rate matrix and the mean first passage times

In a viscous fluid environment, there is a separation of time scale between the typical time for a change of configuration of a protein and the time it takes to equilibrate locally in each state. Under such conditions, the evolution of populations of configurations at intermediate time scales that are long compared to the molecular time scale, but much shorter than the overall folding time, can be described by a Markov state model.⁴⁵ The dynamics can describe the folding process as a series of transitions between configurations, defined in Eq. (4), that differ by one bond; such transitions represent a structural change in the protein as bonds form or break.

In a Markov state model, a population of configurations, $\mathbf{P}(t) = \{P_1(t), \dots, P_{n_s}(t)\}$, evolves according to the continuous time Markovian dynamics,

$$\frac{d\mathbf{P}(t)}{dt} = \mathbf{K} \cdot \mathbf{P}(t), \quad (7)$$

where $\mathbf{K}$ is the transition rate matrix. The off-diagonal elements of $\mathbf{K}$ are the time-independent rates of transitioning from one state to another. Consider the case where states are ordered in index from fewest bonds to most bonds, and suppose $j > i$ is formed from i by the addition of a single bond. Then, due to diffusive barrier crossing, the inverse of K_{ji} is of the form⁴⁵

$$K_{ji}^{-1} = e^{-\beta^* F_{ij}} \tau_{(ij)}^- + \tau_{(ij)}^+ \\ = \frac{P_i}{P_j} \tau_{(ij)}^- + \tau_{(ij)}^+, \quad (8)$$

where $\tau_{(ij)}^-$ and $\tau_{(ij)}^+$ are the mean inner and outer equilibrium first passage times for the pair of states i and j : That is, $\tau_{(ij)}^+$ corresponds to the time required for a pair of beads, whose initial separation r_{ij} is greater than the transition state value r_c , to diffuse to r_c , averaged

over a (conditional) equilibrium distribution of initial separations. Correspondingly, $\tau_{(ij)}^-$ is the mean first passage time to r_c for beads averaged over an equilibrium distribution of initial distances $r_{ij} < r_c$. When the dynamics of beads in the solvent is diffusive, the mean first passage times for the transition from i to j can be estimated as⁴⁵

$$\tau_{(ij)}^+ = \frac{1}{D_{(ij)}} \int_{r_c}^{r_{\max}} \frac{(1 - C_{(ij)}^+(r))^2}{\rho_{(ij)}^+(r)} dr, \quad (9)$$

$$\tau_{(ij)}^- = \frac{1}{D_{(ij)}} \int_{r_{\min}}^{r_c} \frac{C_{(ij)}^-(r)^2}{\rho_{(ij)}^-(r)} dr, \quad (10)$$

where $D_{(ij)}$ is the self-diffusion coefficient for the relative distance r between beads involved in the bond that is formed or broken between states i and j in the solvent, ρ is the probability density of the bonding distance, and

$$C_{(ij)}^-(r) = \int_{r_{\min}}^r \rho_{(ij)}^-(x) dx,$$

$$C_{(ij)}^+(r) = \int_r^{r_{\max}} \rho_{(ij)}^+(x) dx$$

are the respective cumulative distributions of the distances r . The constants of integration $r_{\min}$ and $r_{\max}$ correspond to the minimum and maximum distances that can separate a pair of nonlocally bonding beads, and they can be taken to be zero and infinity, respectively, since the integrand vanishes in both limits. Generally speaking, the self-diffusion coefficients $D_{(ij)}$ depend on both the solvent friction as well as the internal friction that arises from the particular distance constraints determining the states i and j . By construction, the Markov state model obeys detailed balance,

$$K_{ij}P_j = K_{ji}P_i, \quad (11)$$

and Eq. (7) has a unique stationary equilibrium distribution of populations.

Note that the relative probability of states i and j ,

$$\frac{P_i}{P_j} = \frac{e^{-\beta^* E_i} e^{S_i}}{e^{-\beta^* E_j} e^{S_j}} = e^{-\beta^* (E_i - E_j)}, \quad (12)$$

plays an important role in determining the transition rates, but it does not affect the mean first passage times. In the low temperature limit, $P_j \gg P_i$, since state j has an additional bond relative to state i and, hence, $K_{ji} \approx 1/\tau_{(ij)}^+$. Under these conditions, the rate of back transitions, K_{ij} , will be exponentially small since

$$K_{ij} = \frac{P_i}{P_j} K_{ji} \approx e^{-\beta^* F_{ij}} \frac{1}{\tau_{(ij)}^+}.$$

The adaptive algorithms that we present in the following sections generate bond distances distributed according to the conditional equilibrium densities $\rho^+(r)$ and $\rho^-(r)$ during the iterative process. Using the set of recorded distances, the smooth fit of the probability density $\rho(r)$ and the cumulative distribution functions $C(r)$ are constructed in one of two ways: Either the empirical cumulative distribution is expanded in an orthonormal basis⁴⁶ or, alternatively, a maximum-likelihood estimate of an expansion of




FIG. 1. The integrand of the outer mean first passage time τ^+ for a 40-bead system for a nonlocal bond between beads separated by 20 monomers. The inset shows the probability density $\rho^+(r)$ of the bonding distance.

the logarithm of the density is constructed using splines.⁴⁷ Both approaches make use of goodness-of-fit statistical tests to judge the quality of the fit. In Fig. 1, the integrand for the outer first passage time for the transition between states in a 40-bead system is shown as a function of the bonding distance, as well as the probability density $\rho^+(r)$. The latter is constructed from a continuous spline fit of the logarithm of the density or the potential of mean force. With these continuous and smooth functions in hand, the mean first passage times in Eqs. (9) and (10) and the elements of the transition rate matrix $\mathbf{K}$ are easily evaluated numerically using Gaussian quadrature. Note that the integrand of the outer first passage time is determined primarily by the fit of the density in the transition region near r_c where the integrand is largest.

IV. NUMERICAL SOLUTION OF THE CONFIGURATIONAL ENTROPY AND FIRST PASSAGE TIME

In a complex system with a large number of beads, the shapes of the sub-volumes of different configurations in the high dimensional space are complicated, making Eq. (5) impossible to evaluate exactly for all but the simplest models. To compute the configurational entropies for larger model chains, we must resort to using Monte Carlo (MC) methods. In the next several sections, the algorithm used to compute the configurational entropy and first passage times is detailed. A brief overview of the procedure is as follows:

1. An ensemble of initial structures with no bonds is generated by selecting a set of bond distances, bond angles, and dihedral angles from the appropriate distribution of values. Configurations violating any distance constraints are rejected. The distribution of the distances of nonlocal interactions are used to estimate the first passage time to form each of the possible nonlocal bonds for the set of structures with a single bond [see Eq. (9)].

- The set of states defined by which nonlocal bonds are active is partitioned into layers based on the number of nonlocal bonds. The layer method is detailed in Sec. IV B.
- For states for which the estimate of the mean first passage time from earlier calculations is large ($\tau^+ > 10$), the attractive step potential depending on the active bond distance is replaced by a staircase potential in Sec. IV C to reduce the computational cost.
- The difference in entropy and the first passage times between pairs of states in adjacent layers that differ by only one bond are computed in parallel using an adaptive Monte Carlo sampling algorithm (see Sec. IV A) combined with replica-exchange between simulations at fixed intervals.
- In the adaptive procedure, rejection-free Monte Carlo updates are carried out using event-driven dynamics in a hybrid Monte Carlo algorithm.
 - The initial phase of the evaluation of the entropy difference consists of a fixed number (10^7) of adaptive adjustments of the entropy using a Wang–Landau algorithm [Eq. (18)].
 - Convergence of the entropy difference is assessed by applying the event-driven dynamics with fixed entropy values to generate a set of (independent) configurations and applying the G-test for uniformity given in Eq. (19).
 - If the G-test is not satisfied, the entropy values are adjusted using Eq. (20) and the previous step is repeated until the test is satisfied.
- The distances between beads forming nonlocal bonds, which are recorded at regular time intervals throughout the simulation, are used to obtain the mean first passage times given in Eqs. (9) and (10).
- The biased entropies and mean first passage times can, then, be used to construct the transition rate matrix in the Markov state model for a choice of interaction energies [see Eq. (8)].

A. Adaptive Monte Carlo sampling

For a molecular system suspended in a solvent in thermal equilibrium at inverse temperature β , the configurations of the molecule are canonically distributed. For a given model with a set of n_s energies $\{\beta^* \mathbf{E}\}$, the probability of a configuration in the ensemble is given by Eq. (6). The entropy difference between states i and j obeys

$$e^{S_i - S_j} = e^{S_{ij}} = \frac{e^{\beta^* E_i} P_i}{e^{\beta^* E_j} P_j},$$

and, hence,

$$S_{ij} = \ln\left(\frac{P_i}{P_j}\right) + \beta^* (E_i - E_j).$$

Suppose n_c samples of states are drawn *independently* with a canonical probability for a model with a set of energies $\{\beta^* \mathbf{E}\}$. The number of states of type i in the sample is denoted by n_i . Using

the empirical probability of state i , $\hat{P}_i = n_i/n_c$, an estimator of the entropy difference between states i and j can be defined as

$$\hat{S}_{ij} = \ln\left(\frac{n_i}{n_j}\right) + \beta^* (E_i - E_j). \quad (13)$$

If the states $\{1, 2, \dots, s\}$ are sampled independently, the set of counts $\{n_1, n_2, \dots, n_s\}$ is multinomially distributed with probability

$$P(\{n_1, n_2, \dots, n_s\}) = \frac{n_c!}{n_1! \dots n_s!} P_1^{n_1} P_2^{n_2} \dots P_s^{n_s}.$$

The mean and the variance of the entropy estimator in Eq. (13) are

$$\langle \hat{S}_{ij} \rangle = S_{ij} + \frac{1}{2n_c} \left(\frac{1}{P_j} - \frac{1}{P_i} \right) + O\left(\frac{1}{n_c^2}\right), \quad (14)$$

$$\sigma_S^2 = \frac{1}{n_c} \left(\frac{1}{P_i} + \frac{1}{P_j} \right) + O\left(\frac{1}{n_c^2}\right) \geq \frac{2n_s}{n_c}, \quad (15)$$

where all $P_k > 0$, which implies that $\langle \hat{S}_{ij} \rangle$ converges to S_{ij} as $n_c^{-1} \rightarrow \infty$. Note that the rate of convergence of the estimator is optimized when $P_i \approx P_j$, at which point the minimum value of the variance is $2n_s/n_c$. The minimum variance is achieved when $\beta^* E_i = S_j$. This choice of the set $\{\beta^* \mathbf{E}\}$ is not known *a priori* and must be determined self-consistently, as discussed below. Note that other choices of $\{\beta^* \mathbf{E}\}$ may result in $P_i \gg P_j$, in which case the empirical average of the entropy converges slowly due to a large standard error given that $1/(n_c P_j) \gg 1$.

The estimator Eq. (13) requires a set of samples drawn from the canonical ensemble. Metropolis Monte Carlo (MMC) algorithms are an appealing sampling approach to generate a sample of states since they do not require computing normalizing factors to generate states with known probabilities. However, efficient implementations of the MMC algorithm require proposing trial configurations from the current state that are both statistically likely and yet differ significantly. For chain molecules, particularly those that have excluded volume constraints, this is a difficult task,⁴⁸ although methods using crankshaft rotations,^{49,50} configurational bias regrowth,^{51,52} and normalizing flows⁵³ exist to generate global changes to configurations.

The principal challenge in efficient MMC sampling in this context is the highly correlated way in which configurations must change to generate states of high probability. Dynamical sampling methods that evolve all degrees of freedom provide a viable solution for the rapid exploration of local structures. Here, we use the hybrid Monte Carlo (HMC) method to generate configurations with a canonical probability based on a dynamical updating scheme.⁵⁴ In this procedure, the dynamical updates must be time-reversible and must conserve phase space volume. In most applications, proposed configurations are generated by numerically solving the equations of motion for a given potential using symplectic split-operator integration schemes. For systems interacting via discontinuous step potentials and hard walls, the equations of motion are exactly solvable (within numerical precision) and the dynamics of the system is time-reversible and conserves phase space volume.

In our implementation of the HMC scheme, the current configuration $\mathbf{R}$ is augmented with momenta $\mathbf{P}$ drawn from a normal distribution with zero mean and unit variance so that the system acquires a kinetic energy $K(\mathbf{P}) = \mathbf{P}^2/2$. Then, the system is propagated forward for a time interval τ_p with Hamiltonian dynamics from an initial state $\mathbf{X} = (\mathbf{R}, \mathbf{P})$ to a final state $\mathbf{X}_{\tau_p} = (\mathbf{R}_{\tau_p}, \mathbf{P}_{\tau_p})$. The Hamiltonian $H(\mathbf{R}, \mathbf{P})$ is the sum of the kinetic energy $K(\mathbf{P})$ and a discontinuous potential $U(\mathbf{R})$. The final configuration $\mathbf{R}_{\tau_p}$ of the trajectory is then accepted or rejected as the next state in a Markov chain with acceptance probability given by

$$A(\mathbf{X} \rightarrow \mathbf{X}_{\tau_p}) = \min(1, e^{-\Delta H}), \quad (16)$$

where ΔH is the difference between the final and initial Hamiltonians. When event-driven dynamics generate trial configurations $\mathbf{R}_{\tau_p}$, the HMC algorithm proposes updates in a rejection-free manner since the Hamiltonian is exactly conserved so that $\Delta H = 0$, and the probability of acceptance of a trial configuration is unity. For discontinuous potential systems, the dynamic sampling trajectories are solved exactly (i.e., within numerical precision) using event-driven simulation methods, and efficient implementation of event-driven dynamics should make use of event trees, hybrid queues, and other cost-saving techniques.⁵⁵ The sampling procedure generates a set of states $\mathbf{R}$ asymptotically distributed with probability proportional to $e^{-U(\mathbf{R})}$. Instead of using an actual physical potential $U(\mathbf{R})$ to govern the dynamical updates in the Monte Carlo procedure, we use an estimate of the entropy $U(\mathbf{R}) = S_b(\mathbf{R})$ that approximates the true entropy $S(\mathbf{I}(\mathbf{R}))$, where $\mathbb{1}_c(\mathbf{R})$ is the indicator function for state c defined in Eq. (4). The HMC sampling procedure generates a Markov chain of states in which configuration i appears with probability

$$P_i(\mathbf{S}_b) = \frac{e^{-S_{i,b}} e^{S_i}}{\sum_{k=1}^{n_s} e^{-S_{k,b}} e^{S_k}} \approx \frac{1}{n_s}. \quad (17)$$

In order to ensure that the n_c samples are drawn independently, the time τ_s between recording configurations of the system should be larger than the largest outer mean first passage time between states when only local, dynamic updates are used to propose trial configurations. In this case, τ_s is set to be a multiple of the basic short propagation time τ_p of the dynamical updates.

However, in Eq. (17), the optimal values of the set of the biasing potential $\mathbf{S}_b^* = \{S_{i,b}^* | i = 1, \dots, n\}$ that lead to a uniform sampling are not known *a priori* and must be determined iteratively using an adaptive procedure. A number of adaptive methods that are effectively equivalent have been proposed in the literature to address this problem, including the Wang–Landau algorithm in its many flavors,⁵⁶ well-tempered metadynamics,^{57,58} and self-healing umbrella sampling.⁵⁹

The essential idea of the adaptive procedure is to construct a sequence of configurations $\{\mathbf{R}_i | \mathbf{S}_b^{(n)}\}$, in which each of the states $\mathbf{R}_{n+1}$ is obtained from the previous state $\mathbf{R}_n$ by applying an evolving transition matrix $\mathbf{T}(\mathbf{S}_b^{(n)})$. The parameters $\mathbf{S}_b^{(n)}$ are determined by a difference equation of the form

$$\begin{aligned} \mathbf{S}_b^{(n+1)} &= \mathbf{S}_b^{(n)} + \gamma_{n+1} \mathbf{f}(\mathbf{R}_{n+1} | \mathbf{S}_b^{(n)}) \\ &= \mathbf{S}_b^{(n)} + \gamma_{n+1} \mathbf{h}(\mathbf{S}_b^{(n)}) + \gamma_{n+1} \mathbf{H}(\mathbf{R}_{n+1} | \mathbf{S}_b^{(n)}), \end{aligned} \quad (18)$$

where γ_{n+1} is a decreasing function of n and the adaptive function $\mathbf{f}$ penalizes visits to the current state and encourages visits to other states. In Eq. (18), $\mathbf{h}$ is the mean drift in the difference equation at index n and $\mathbf{H}$ is the fluctuation around the mean.

The various algorithms differ in their choice of both the dependence of γ_n on the number of steps n and the form of the adaptive function $\mathbf{f}$. Here, we use the commonly chosen adaptive function $\mathbf{f}(\mathbf{R}) = \mathbf{I}(\mathbf{R})$ that penalizes future visits to the current state by increasing its entropy by γ_{n+1} . In general, the convergence of the sequence of biases $\{\mathbf{S}_b^{(n)}\}$ to a unique fixed-point solution $\mathbf{S}_b^*$ is difficult to establish for a particular choice of γ_n , $\mathbf{f}$, and the transition matrices $\mathbf{T}$, but it has been proved for the algebraic protocol $\gamma_n = \gamma^*/n^\alpha$, where $\alpha \in (1/2, 1]$, provided the transition matrices are sufficiently mixing.^{60,61} For example, for the parameter choice $\gamma^* = n_s$ and $\alpha = 1$, it has been shown⁶⁰ that the sequence $\{\mathbf{S}_b^{(n)}\}$ converges to $\mathbf{S}_b^*$ as n^{-1} , and that the set $\mathbf{S}_b^{(n)}$ has a multivariate normal distribution with a mean $\mathbf{S}_b^*$ and a covariance matrix proportional to $\mathbf{U}_t = n_s \gamma_n \mathbf{U}^* = n_s/t \mathbf{U}^*$, where $t = n/n_s$ is the state size-dependent scaled time between updates of γ_n . Here, $\mathbf{U}^*$ is a covariance matrix that depends on the fluctuations of H determined by the sequence of transition matrices $\mathbf{T}$. As a result, it is difficult to estimate $\mathbf{U}^*$ to determine the standard errors of the entropy values S_b in the adaptive procedure.

To assess the accuracy of the configurational entropy, we iterate Eq. (18) for a fixed number of total updates $t_f = mt$ with a large value of $m = 10^7$. At this point, a series of n_c independent trajectories are generated using the final set of biases and the number n_i of counts of uncorrelated states i recorded. The recorded empirical distribution of states is then checked for uniformity using a statistical test. Here, we use the G-test based on the statistic

$$G = -2 \sum_{i=1}^{n_s} n_i \ln \left(\frac{n_i}{e_i} \right), \quad (19)$$

where $e_i = P_i n_c = n_c/n_s$ is the expected number of counts of state i when P_i is uniform and the term⁶²

$$q_2 = 1 + \frac{n_s + 1}{6n_c} + \frac{n_s^2}{6n_c^2}$$

corrects for small sample sizes. When the sample counts $\{n_i\}$ are independent, the G-statistic is asymptotically χ^2 -distributed with $n_s - 1$ degrees of freedom, allowing the p-value of the computed statistic to be evaluated.

Convergence has been achieved when $p > p_c$ and the distribution of configurations is considered statistically consistent with a uniform distribution. If $p < p_c$, the biased entropy values in iteration n can be updated according to

$$S_{i,b}^{(n+1)} = S_{i,b}^{(n)} + \ln \left(\frac{n_i}{n_c} \right). \quad (20)$$

Strictly speaking, this additional iterative process is not required since the estimator for the configurational entropy is unbiased and has a variance that is close to optimal since the biased probabilities are already close to uniform, $P_i \approx P_j$. If desired, another iteration of the sampling can be performed with the updated bias values until convergence is obtained and $S_i = S_{i,b}$ within statistical resolution. More stringent statistical tests for convergence can be applied

if desired. For example, after the process has passed the condition $p > p_c$, the actual distribution of a set of G-statistics from independent runs can be tested against a χ^2 -distribution using a goodness-of-fit test, such as the Kolmogorov–Smirnov test.⁶³ It should be emphasized that the failure of the condition $p > p_c$ does not necessarily indicate that the data of counts are inconsistent with a multinomial distribution, since the statistical test also relies on the assumption that the samples are drawn independently.

The final configurational entropy difference between states i and j with counts n_i and n_j is

$$\Delta S_{ij} = \Delta S_{ij,b}^{(n)} + \ln(n_i/n_j) \pm \sqrt{\frac{A}{n_i + n_j}},$$

where samples are recorded at time intervals $\tau > n_b^2 \tau^+$ and A is the upper percentile value of the χ^2 -distribution with $n_s - 1$ degrees of freedom.⁶⁴ The length of the production run required for a given statistical resolution can be estimated using confidence intervals for multinomial proportions.^{64,65}

During each iteration, the bonding distances between all beads for each state i of the n_s states explored can be used to calculate the mean first passage times for state i . The bonding distances are distributed with the conditional equilibrium density for this state due to the fact that the bias $S_{i,b}$ is the same for all configurations in state i .

It is important to emphasize that the statistical analysis presented above assumes that each sampled configuration is drawn independently from the canonical probability density with a corresponding multinomial distribution of states. In practice, this will not be the case when local Monte Carlo proposals alone are used, since the proposed trial configurations are highly correlated with the current state. Correlations exist when the lengths of the trajectories τ_s are not long enough to generate independent configurations in the Markov chain. If the trajectory segments are only long enough to form or break a single bond, the overall dynamics in the state space is diffusive at best and the statistical tests for uniformity are inappropriate. Under these circumstances, the correlation time of the state counts $\{n_i\}$ must be analyzed to insure that successive states used in the convergence test are independent.

B. The layer simulation method

If the protein being modeled can form many bonds n_b , the number of possible states and the number of configurational entropy values $n_s = 2^{n_b}$ to be computed will be large. As n_s increases, the covariance matrix $\mathbf{U}_l$ of the adaptive procedure, which scales quadratically with n_s , becomes large. The convergence of the entropy to S_b^* will, therefore, be very slow, particularly when some of the transitions are infrequent due to long first passage times. The states generated via short trajectories remain correlated for increasingly long periods of time as the number of bonds increases. For example, when the full set of n_s states are sampled using local dynamical updates, the dynamics of the state space in the limit where the states are generated with uniform probability obeys a Master equation of the form

$$\frac{dP_i(t)}{dt} = \kappa(P_{i+1}(t) + P_{i-1}(t) - 2P_i(t)),$$

where $\kappa = 1/(\tau^- + \tau^+) \sim 1/\tau^+$ is the rate of transitions to neighboring states. These dynamics generate a uniform distribution of states $P_i \sim 1/n_s$ on time scales governed by the relaxation modes $\lambda_m = 2\kappa \sin^2(m\pi/(2(n_b + 1)))$. States remain correlated for time scales up to the overall equilibration time $\tau_{eq} \sim 1/\lambda_1 \sim (8/\pi)n_b^2 \tau^+$, where n_b is the total number of bonds that can be formed. Thus, to generate uncorrelated samples uniformly, the length of trajectories $\tau_s = s\tau_p$ should be scaled by n_b^2 , relative to two-state models for which trajectories of length τ^+ are adequate.

To improve the rate of convergence of $S_b^{(n)}$ to S_b , we consider a layered simulation approach, in which short calculations are conducted in parallel to sample two states at a time differing by a single bond. If we define the layer ℓ to be the $\binom{n_b}{\ell}$ states in which there are ℓ nonlocal bonds that have formed and $n_b - \ell$ bonds that have not, each state in layer ℓ can lead to $n_b - \ell$ states in layer $\ell + 1$ by the formation of a single new bond. The pairing of all states connected in adjacent layers leads to a total set of $n_b n_s / 2$ pairs of connected states for which the difference in entropy is computed. If the entropy of a configuration is defined relative to the nonbonded state in layer 0, the entropy $\Delta S(\alpha_0, \alpha_\ell)$ of a particular state α_ℓ in layer ℓ can be estimated by the sum of the entropy differences between states in adjacent layers in a path that connects state α_0 to the state α_ℓ ,

$$\Delta S(\alpha_0, \alpha_\ell | \{\alpha_i\}) = \Delta S(\alpha_0, \alpha_1) + \Delta S(\alpha_1, \alpha_2) + \cdots + \Delta S(\alpha_{\ell-1}, \alpha_\ell),$$

where the path $\{\alpha_i\}$ used is $\alpha_0 \rightarrow \alpha_1 \rightarrow \alpha_2 \rightarrow \cdots \rightarrow \alpha_{\ell-1} \rightarrow \alpha_\ell$. However, when all states are dynamically connected and none are geometrically prohibited, there are a total of $\ell!$ unique paths that connect α_0 and α_ℓ , so a more precise estimate can be obtained by averaging over all paths that connect the same initial and final states,

$$\Delta S(\alpha_0, \alpha_\ell) = \frac{1}{\ell!} \sum_{\{\alpha_i\}} \Delta S(\alpha_0, \alpha_\ell | \{\alpha_i\}). \quad (21)$$

When each of the computations of $\Delta S(\alpha_i, \alpha_{i+1})$ has converged and the probability of the states α_i and α_{i+1} is the same, the mean of the estimator defined in Eq. (13) is zero with variance $4/n_c$, where n_c is the number of event-driven trajectories used to sample the states in the simulation. Hence, the variance of the estimator for a state in layer ℓ , Eq. (21), is

$$\sigma_{\Delta S_\ell}^2 \geq \frac{4}{(\ell - 1)! n_c}. \quad (22)$$

This estimate is useful to determine the number of independent configurations n_c chosen per iteration for a given level of precision. If the relative precision is set to 5%, then the number of sampled states should be chosen to be larger than $n_c \geq 1600/\Delta S_m$, where ΔS_m is an estimate of the minimum increase in entropy obtained by breaking one of the bonds. Typically, for the models considered here, this quantity is roughly unity (though often, it is three times larger).

To illustrate the layer method, consider a three-bond model for which $n_s = 8$. We present the layer approach in Fig. 2.

The initial state of the chain is linearly extended and entirely devoid of nonlocal bonds, which is represented by the binary string 000. The configuration 000 makes up layer 0 (in green). To obtain the configurations in layer 1 (in red) with a single nonlocal bond, we run three simulations: one in which we transition from 000 to 100 by forming the first bond, one from 000 to 010 by forming




FIG. 2. The layer approach for a 20-bead, three-bond model.

the second bond, and one from 000 to 001 where the third bond is formed. In the event-driven dynamical sampling of the two designated states, the active nonlocal bond that can be formed or broken is treated normally with dynamics governed by the step potential [see Eq. (2)], while the state of all other nonlocal bonds is fixed by an elastic collision at r_c (i.e., the step potential for these bonds is infinite and positive). To obtain the configurations in layer 2 (in blue), we start with each configuration in layer 1 and turn on one of each of the two remaining bonds. Thus, between the first and the second layer, we have a total of $\binom{3}{1} = 6$ simulations: 100 to 110, 100 to 101, 010 to 110, 010 to 011, 001 to 101, and 001 to 011. In each of the simulations, each existing bond in layer 1 is fixed and is not allowed to break. Overall, there are a total of $n_b n_s / 2 = 12$ computations of configurational entropy differences. To compute the entropy of a state in layer 2 relative to the nonbonded state, the average is taken over the paths connecting it to state 000. For example, $\Delta S_{110,000} = (\Delta S_{110,100} + \Delta S_{100,000} + \Delta S_{110,010} + \Delta S_{010,000}) / 2$.

To avoid quasi-ergodic sampling issues in which transitions between different types of structures for a given state are rare, the layer simulations are coupled together by replica-exchange Monte Carlo moves,^{66–68} in which configurations are exchanged between adjacent layers with unit probability when they satisfy the same bonding constraints. For example, a layer simulation connecting a state in layer $i - 1$ with a state in layer i that differs by a single bond can be coupled to a simulation between a pair of configurations in layers i and $i + 1$ that also differ by a single bond. The replica-exchange swaps between the Markov chains are accepted when both simulations are in states in layer i and, therefore, satisfy the same bonding pattern. The swap moves should be attempted frequently to optimize the efficiency of the replica-exchange sampling.^{69,70} Thus, a given set of configurations are exchanged with a frequency of 25% if each layer simulation consists of two states that only differ by a single bond. The exchange frequency can be increased by increasing the number of layers that are explored in a given chain. Similar replica-exchange algorithms have been proposed in the context of the Wang–Landau algorithm.^{71,72} An alternative parallel implementation⁷³ of the Wang–Landau algorithm that requires frequent

communication between stochastic trajectories uses an adaptive function $\mathbf{h}$. This function depends on a mean number of visits to update a shared set of biases $\mathbf{S}_b^{(n)}$. In a serial approach, population Monte Carlo algorithms, which generate pools of different structures for a pair of states and uniformly select a structure from the pool to be updated, can accomplish the same task.⁷⁴

As before, the bonding distances can be recorded and used to compute $2n_b$ inner or outer mean first passage times. The outer first passage time τ^+ between a source and destination state in the next layer can be also used to estimate the length of a trajectory $\tau_s \sim \tau^+$ to generate a statistically independent configuration in the next layer, where the destination state is a source state for the next layer. This information is useful in two ways. First, the computational cost of the procedure can be optimized by adapting the trajectory length τ_s to the pair of states. Second, problems can indicate when a pair of states either are not connected due to geometrical constraints that are impossible to satisfy or require unreasonably long trajectories due to a large value of τ^+ . In Sec. IV C, a biasing procedure is introduced to mitigate the problems associated with large first passage times.

Another issue that arises for models with a large number of bonds is that some bonding states have mutually exclusive distance constraints that cannot be satisfied simultaneously. In this event, the state is not allowed and must be removed from the model. The layer simulation approach provides a reliable method to eliminate states, since an estimate of the outer first passage time $\tau_{(ij)}^+$ between states i in layer ℓ and j in layer $\ell + 1$ is computed before the entropy of state j . If there are no outer collisions in all states i that are connected to state j , and $\tau_{(ij)}^+$ is infinite, state j can be eliminated from the layer $\ell + 1$.

The decomposition of the calculation of the configurational entropy into a number of independent calculations between adjacent layers reduces the computational demands of the task relative to a procedure in which all nonlocal bonds are active and $n_s = 2^{n_b}$ values of the entropy are evaluated simultaneously. This gain in serial efficiency is due to the reduction in the overall sampling time τ_s per sample required in the iterative procedure of verifying the uniform convergence of the sampled states that scales as the square of the number of bonds (see Sec. IV B). In Fig. 3, the relative efficiency of the layer method is demonstrated for a simple model with 20 beads and a variable number of bonds n_b . The two simulation approaches coincide for a model with a single bond where $n_b = 1$ and $n_s = 2$, and it is evident that the relative efficiency of the layer method increases roughly linearly with the number of bonds for a given choice of sampled configurations per state.

C. Biasing the entropy calculation: The staircase potential

As the mean first passage time between two states increases, the length of the trajectories τ_s required to sample independent configurations becomes prohibitively large, rendering the direct calculation of the configurational entropy difference between the states computationally inefficient. This situation arises when the probability density for the bond distance $\rho^+(r)$, which is the reaction coordinate for a change in state, is small in the vicinity of the transition state at $r = r_c$. In the vicinity of r_c , $C^+(r) \approx 0$ and the integrand in Eq. (9), which is proportional to $1/\rho^+(r)$, becomes large




FIG. 3. The ratio of the relative efficiency of the layers to the general procedure as a function of the number of bonds (or levels). In the general procedure, all states (rather than a pair) are accessible within a simulation. The relative efficiency is the ratio of the wall clock times needed to achieve convergence in each method. The data were generated by averaging the serial execution time to convergence of 100 instances of a 20-bead model with $n_i = 400$ configurations sampled per state i and a convergence level of $p = 0.25$.

(see Fig. 1). The probability density $\rho^+(r_c)$ at the transition distance can be small either because (i) in the unbonded state, the range of the bond distance allowed is broad and the configurational volume of the unbonded state is large or (ii) existing nonlocal bonds in the initial state introduce geometrical constraints in the chain that prevent the bonding distance from being reached unless beads in the rest of the chain are placed optimally. In both situations, the low probability of exploring the reaction coordinate values in the region of the transition state results in inefficient sampling. For systems with continuous potentials, sampling methods such as metadynamics^{57,58} or umbrella sampling⁵⁹ can be used to bias the stochastic sampling to visit improbable regions of the reaction coordinate. However, event-driven dynamical sampling is not amenable to the introduction of continuous force fields.

To encourage these bonds to form more readily while maintaining the discontinuous nature of the model, we introduce modifications to the potential described in Eq. (2) to bias the calculation of the configurational entropy by reducing the required trajectory time τ_s . This bias is a computational device to calculate the biased entropy and mean first passage times for the original model. Like an adaptive binning strategy,⁷³ we subdivide the outer region of the constant energy potential to create a discontinuous potential that resembles a staircase, as in Fig. 4.

To simulate dynamics in the staircase potential, we use a layer approach in which each step of the staircase defines a new state that is implemented in a separate layer. As a result of the restrictions to the bond distance, the mean first passage time between the staircase regions is small, and only short trajectories are required to evaluate the ratio of the relative number of states.

When the initial state is divided into ℓ sub-states (such as those defined by the $\ell = 3$ regions 1, 2 and 3 in Fig. 4), the configurational entropy for the i to j transition in the original model is not the sum of the entropy differences between the regions due to the fact that the total number of states in configuration i is the sum of the number of




FIG. 4. Example of a staircase with three steps. Note that $\epsilon_3 > \epsilon_2 > \epsilon_{ij}$ and $r_3 > r_2 > r_c$.

states in each of the regions. Rather, the entropy difference between states i and j is

$$e^{\Delta S_{ij}} = \frac{n_i}{n_j}, \quad (23)$$

$$n_i = \sum_{k=1}^{\ell} n_{ik},$$

where n_i is the total volume of state i and n_{ik} is the volume of region k . If each of the regions is treated as a separate layer, $\tilde{S}_{k,k-1} = \ln(n_{ik}/n_{i,k-1})$ corresponds to the entropy difference between the sub-states defined by adjacent regions k and $k-1$. If the volume of region j is taken as n_{j0} , then we find

$$\Delta S_{ij} = \ln \sum_{k=1}^{\ell} \exp \left\{ \sum_{m=1}^k \tilde{S}_{m,m-1} \right\}. \quad (24)$$

Each value of $\tilde{S}_{m,m-1}$ is readily computed using the dynamical sampling procedure in Sec. IV B. Note that the estimate of the entropy difference in Eq. (24) has an approximate asymptotic variance $4\ell/n_c$ when $\tilde{S}_{m,m-1} > 1$. Consequently, the number of trajectories per iteration should be scaled appropriately for a given level of precision.

The number of steps ℓ in the staircase potential and the location of each of the steps r_i can be estimated from the expected difference in entropy ΔS_{ij} and the distribution of the relevant bond distance from the simulation in the previous layer. If the drop in entropy in each of the regions is constant, $\tilde{S}_{m,m-1} \approx \tilde{S}$, which means $\Delta S_{ij} \approx \ell \tilde{S}$, and hence $\ell \approx \Delta S_{ij}/\tilde{S}$. The location r_i of staircase i is determined from the cumulative bond distance distribution $C^+(r)$ by the

condition that $C^+(r_i) = e^{-\tilde{S}} C^+(r_{i+1})$, where the outermost region satisfies $C^+(r_\ell) = e^{-\tilde{S}}$.

We have found that for all models considered here, the largest change in entropy $\Delta S_{ij} \approx 12$ so that a typical choice of $\tilde{S} = 4$ requires the introduction of no more than three staircase regions. Larger choices of $\tilde{S}$ result in less efficient sampling since the first passage time between regions increases exponentially with $\tilde{S}$. For the special case of $\ell = 3$, we have

$$\Delta S_{ij} = \tilde{S}_{1,j} + \tilde{S}_{2,1} + \tilde{S}_{3,2} + \ln \left[1 + e^{-\tilde{S}_{3,2}} (1 + e^{-\tilde{S}_{2,1}}) \right]. \quad (25)$$

For the crambin and frustrated models considered in Sec. V, the typical values of the step location were $r_2 \approx 1.8$ and $r_3 \approx 2.5$.

The introduction of the staircase potential greatly reduces the workload of computing the entropy difference between a pair of states that infrequently interconvert, and it also improves the accuracy of the outer first passage times. The staircase bias increases the rate of convergence of both the initial Wang–Landau estimates of the entropy and the subsequent procedure for the verification of convergence. To demonstrate this explicitly, we consider a 40-bead chain with a single nonlocal bond between bead pair [10,30]. The entropy difference between the nonbonded and bonded states is $S_0 - S_1 = 9.0 \pm 0.1$, and the mean first passage time for this model is $\sim \tau^+ = 627 \pm 29$ due to the average separation between the bonding beads. The convergence rate of the Wang–Landau procedure depends on the magnitude of the standard deviation σ_t of the drift term in the adaptive adjustments. As is clear in Fig. 5, the introduction of additional staircase states into the system at a fixed computational cost reduces the standard error of the Wang–Landau procedure. This error decreases as $t^{-1/2}$ with time step $t = n_s \tau$, even though the cost-per-iteration of the algorithm increases linearly with




FIG. 5. The standard deviation σ_t of the drift term vs $1/\sqrt{t}$ for a fixed scaled time step $t = n_s \tau$ averaged over 50 realizations of the adaptive procedure. The model is a 40-bead system in which distant beads 10 and 30 form a nonlocal bond at $r_c = 1.5$. The black line denotes a system with no staircase, and the red, green, and blue lines denote systems with one, two, and three additional steps. The inset shows the value of $S_b(t)$ vs $1/\sqrt{t}$ for the corresponding systems.

the number of steps in the staircase. For this model, the reduction saturates after the inclusion of two steps in the staircase (green line in Fig. 5). The ratio of the standard errors $\sigma_i(t) \sim \sigma_i/\sqrt{t}$ for a simulation without a staircase to one with i steps is roughly $\sigma_0/\sigma_3 \approx 23/3$, indicating that the inclusion of the steps reduces the computational time needed at a given level of statistical resolution by a factor of roughly 60. At the same time, the efficiency of the validation procedure to establish uniformity is also improved, since the trajectory length required for each independent sample decreases from $\tau_s \sim 8\tau^+/\pi \sim 6400$ trajectories of unit length to less than $\tau_s \sim 50$.

The bias introduced by the stairs also improves the calculation of the outer first passage time. The sampling of the reaction coordinate is enhanced in the vicinity of the transition state at r_c , where the integrand of the first passage time is the largest, as is apparent in Fig. 1 for the 40-bead model. The density of the reaction coordinate $\rho^+(r)$ is constructed by stitching together continuous fits of the densities in each of the staircase regions. To improve the quality of the fit of the integrand in the region near r_c , a larger number of sampling points in the staircase region containing r_c should be used. Without enhancing the sampling, the standard error of the estimated outer first passage time is large when transitions are rare.

V. FOLDING DYNAMICS, PATHWAYS, AND EVOLUTION

The simplicity of the discontinuous potential model allows both the free energy and the transition rate matrix $\mathbf{K}$ in a Markovian description of the dynamics to be determined analytically for any choice of state energies at any temperature. These features enable the study of how folding pathways from a nonbonded initial state to the fully bonded “folded” state change with these parameters.

The utility of Markov chains in describing the dynamics of chemical and biophysical systems has long been recognized, and vast literature exists on the subject (for example, see Refs. 75 and 76). A number of properties are of interest in a Markov state model of protein dynamics. Since $\mathbf{K}$ is a regular, square matrix satisfying detailed balance, it has a unique zero eigenvector that corresponds to the equilibrium populations. The transition rate matrix can be written in terms of a diagonal matrix as $\mathbf{K} = \mathbf{U}\lambda\mathbf{U}^{-1}$, where $\mathbf{U}$ is a matrix with eigenvectors of $\mathbf{K}$ as the columns, and λ is a diagonal matrix with eigenvalues $\lambda_i \leq 0$ on the diagonal. The spectrum of eigenvalues $\{\lambda_i\}$ can be useful to determine if a small number of states dominate the long-time dynamics of the system. When this is the case, reduction techniques such as stochastic complementation may be profitably applied to reduce the dimensionality of the Markov model.⁷⁷ In addition, the probability of particular paths starting from an initial distribution of states to the folded state can be analyzed to find dominant folding pathways and potential bottlenecks in the nonequilibrium first passage path ensemble.^{78,79}

Functional proteins have evolved to carry out specific tasks under stressful environmental conditions. Since their function is intimately linked to their three-dimensional structure, their structure must be resilient to thermal stress. This suggests that a fast-folding, single-domain protein should not only exhibit a strong preference for its active structure over a range of temperatures, but it should also rapidly equilibrate or refold to this “native” structure if perturbed. Naturally evolved proteins of this type have optimized

sequences and energies of configurations that result in such characteristics.

The evolution of sequences selected to optimize thermal stability can be examined in the Markov state model by considering the variation of the folding time with respect to the set of interactions $\{\beta^* \mathbf{E}\} = \{\mathbf{E}^*\}$ in the model. The folding time can be analyzed by considering the probability density $S(t)$ of the system in a non-native configuration at time t in the presence of an absorbing state f , defined as

$$S(t) = \sum_{i \neq f} P_i(t).$$

Here, f is taken to be the index of the folded (native) state and $P_i(t)$ is the population of state i at time t . Assuming the initial state of the system is the fully unfolded state of index u , $S(t)$ can be written for the Markov state model as

$$S(t) = \sum_{i=1}^{n_s-1} \left(e^{\tilde{\mathbf{K}}t} \right)_{iu},$$

where $\tilde{\mathbf{K}}$ is the square matrix of rank $n_s - 1$ obtained by removing the row and column from the transition matrix $\mathbf{K}$ corresponding to the native state f . This matrix is invertible and has negative real eigenvalues.

The first passage time density $f(t)$ to the folded state is

$$f(t) = -\frac{dS(t)}{dt} = -\sum_{i \neq f} \frac{dP_i(t)}{dt},$$

and hence the mean and variance of the folding time are given by

$$\mu_t(\{\mathbf{E}^*\}) = \int_0^\infty t f(t) dt = -\sum_{i=1}^{n_s-1} \tilde{\mathbf{K}}_{iu}^{-1}, \quad (26)$$

$$\sigma_t^2(\{\mathbf{E}^*\}) = \sum_{i,j=1}^{n_s-1} (2 \tilde{\mathbf{K}}_{ij}^{-1} \tilde{\mathbf{K}}_{ju}^{-1} - \tilde{\mathbf{K}}_{iu}^{-1} \tilde{\mathbf{K}}_{ju}^{-1}), \quad (27)$$

which depend on the choice of the set $\{\mathbf{E}^*\}$ of dimensionless interaction energies $E_i^* = \beta^* E_i$.

There are a number of dynamical measures that are helpful to understand the characteristic behavior of a Markov state model. We consider an ensemble of “reactive” trajectories defined as the set of trajectories initiated from the unfolded state u that reach the folded state f without revisiting the initial state.⁷⁸ The definition of the ensemble makes use of the committor probability q_i^+ that a trajectory from a given state i reaches the folded target state f before reaching the unfolded state u ,

$$q_i^+ = -\sum_{j \neq (u,f)} \mathbf{K}_{fj} \tilde{\mathbf{K}}_{ji}^{-1}, \quad (28)$$

where the matrix $\tilde{\mathbf{K}}$ is obtained from $\mathbf{K}$ by removing the rows and columns of the u and f states. We assume that set of macrostates defining the originating set in the reactive ensemble consists only of the unfolded state u and that the committors q_i^+ are nonzero for

$i \neq u$. One defines the transition probability matrix $\mathbf{T}$ of passing from state i to state j from the transition rate matrix $\mathbf{K}$ as^{78,79}

$$T_{ji} = \frac{K_{ji}}{\sum_{k \neq i} K_{ki}}, \quad (29)$$

and the reactive transition matrix $\tilde{\mathbf{T}}$ with elements $\tilde{T}_{jf} = \delta_{j,f}$, $\tilde{T}_{u,i} = 0$, as

$$\tilde{T}_{ji} = \begin{cases} \frac{q_j^+ T_{ji}}{q_i^+ T_{ju}} & \text{if } i \neq u, i \neq f, \\ \frac{q_j^+ T_{ji}}{\sum_{j \neq u} q_j^+ T_{ju}} & \text{if } i = u. \end{cases} \quad (30)$$

The elements $\tilde{N}_{ji}$ of the fundamental matrix $\tilde{\mathbf{N}} = (\mathbf{I} - \tilde{\mathbf{T}})^{-1}$ are the expected number of visits to state j from state i in the reactive ensemble. The expected number of visits $\tilde{\theta}_j$ for any state j from an ensemble of reactive trajectories initiated from the unfolded state u is given by $\tilde{\theta}_j = \tilde{N}_{ju}$. Similarly, the visitation probability matrix is denoted as $\tilde{\mathbf{H}}$, and it satisfies $\tilde{\mathbf{N}} = \mathbf{I} + \tilde{\mathbf{H}} \cdot \tilde{\mathbf{N}}$. Its elements $\tilde{H}_{ji}$ correspond to the probability that a reactive trajectory initiated at state i will reach state j , where the reactive probability

$$r_j^+ = \tilde{H}_{ju} \quad (31)$$

is the probability that state j will be visited along the reactive path, starting from the unfolded state.⁷⁹

The reactive flux, $\tilde{J}_{ji}$, measuring the reactive rate from state i to state j , is defined as⁷⁸

$$\tilde{J}_{ji} = \left(\frac{q_j^+ T_{ji}}{\sum_k q_k^+ T_{ki}} \right) \tilde{\theta}_i. \quad (32)$$

The reactive probability r_j^+ and the reactive fluxes $\tilde{\mathbf{J}}$ provide useful measures of the probability of different pathways, and the importance of a particular state in the folding process. We make use of these quantities in Secs. V A and V B.

The evolutionary process for the model system can be simulated by defining a set of beneficial physical characteristics that the system should have. For real biological systems, the selection pressures vary according to their environment and the required physical function of the biomolecule. The relevant physical characteristics, which depend on the set of state energies $\{E_i^* | i = 1, \dots, n_s\}$, could include the requirement that the native state is the most probable state of the system over a large range of temperatures. Additionally, interactions can be selected that make both the mean folding time μ_t , given in Eq. (26), and the variance σ_t^2 in Eq. (27), as small as possible assuming a fixed ratio of the probabilities of the unfolded state to folded state, P_u/P_f . These constraints on the selection of energies ensure that the protein not only folds and refolds quickly, but is also unlikely to have folding pathways that trap intermediate structures for extended periods of time.

Here, we consider the simple loss function for the constrained variational optimization of the set of energies $\{\mathbf{E}^*\}$,

$$\mathcal{L} = \mu_t(\{\mathbf{E}^*\}), \quad (33)$$

to be the average folding time given in Eq. (26) from the nonbonded state u to the folded state f . The most favorable choice of interaction energies for a fixed ratio of P_u/P_f is determined by minimizing

the loss function $\mathcal{L}$ with respect to the $n_s - 2$ adjustable interaction energies. Additionally, the optimization is constrained such that no intermediate state i is substantially populated by including the inequality condition $P_i/P_f \leq 0.005$ to maintain the dominance of the native population. Many choices of selective pressure, defined by the loss function and constraints, are possible and relevant for other types of proteins with different functionality. Note that the gradients of the loss function can also be computed analytically from $\tilde{\mathbf{K}}$, and the probabilities of the configurations can be used to accelerate the minimization of the loss function. For the simple loss function in Eq. (33), the gradients are given by

$$\frac{\partial \mu_i}{\partial E_k^*} = \sum_{\ell, m, n=1}^{n_s-1} \tilde{K}_{\ell m}^{-1} \frac{\partial \tilde{K}_{mn}}{\partial E_k^*} \tilde{K}_{ni}^{-1}, \quad (34)$$

and the derivatives of the $\mathbf{K}$ matrix are

$$\frac{\partial K_{ji}}{\partial E_k^*} = \begin{cases} K_{ji}^2 \tau_{(ij)}^- \frac{P_i}{P_j} (\delta_{i,k} - \delta_{j,k}), & j > i, \\ K_{ji}^2 \tau_{(ij)}^+ \frac{P_i}{P_j} (\delta_{i,k} - \delta_{j,k}), & i > j, \\ -\sum_{l \neq i} \frac{\partial K_{li}}{\partial E_k^*}, & i = j, \end{cases} \quad (35)$$

where states are ordered by their number of bonds from fewest to most. The derivatives of the matrix $\tilde{\mathbf{K}}$ or $\tilde{\mathbf{K}}^*$ can be obtained from the derivatives of $\mathbf{K}$ by the appropriate removal of the rows and columns at the index of the absorbing and source states. The numerical minimization of Eq. (33) is complicated by the high dimension of the parameter search. Standard minimization algorithms, such as the Nelder–Mead or Broyden–Fletcher–Goldfarb–Shanno (BFGS) methods, make use of gradients, but this can lead to difficulties in locating the global minimum. For such situations, methods that combine local gradient search algorithms with multiple trajectory sampling are suitable.^{80–82}

A. The structure and dynamics of crambin

We now examine a coarse-grained model of crambin (PDB ID: 1EJG),^{83,84} a 46-residue protein of unknown function found naturally in cabbage. The three-dimensional crystal structure of the protein, shown in Fig. 6(a), has been measured with 0.48 Å resolution with x-ray crystallography.⁸⁵ The structure of crambin is interesting, since it contains three important structural

motifs: α -helices (in red), antiparallel β -sheets (in yellow), and disulfide bridges (in blue). The discontinuous model of crambin was constructed from the crystal structure in the Protein Data Bank. To model α -helices, attractive interactions representing hydrogen bonds occur only between nonlocal beads whose indices are $a = 2 + 4k$ and $b = a + 4l$, where $l = 1, 4, 5, \dots$ and k is any positive integer.⁴⁴ The omission of bonds between monomers separated by eight or twelve beads is done to discourage the formation of turns and introduce rigidity along the protein's primary structure to prevent it from collapsing in on itself over short distances. For other nonlocal interactions, bonds were identified based on distances in the crystal structure, an idea used in the construction of elastic network models.¹⁰ In particular, crambin was assumed to have bonds formed at a distance $r_c = 1.5$ between a set of beads separated by four residues, viz., $\{[6, 10], [10, 14], [14, 18]\}$, and $\{[22, 26], [26, 30]\}$, that form two short α -helices. The relative orientation of the helices is restricted by an antiparallel β -sheet formed by bonds $\{[2, 34], [3, 33]\}$. The β -sheets are linked to the terminal section of the protein, which has a random coil nature, by disulfide bridges $\{[3, 40], [4, 32]\}$, and an additional bridge bond $[16, 26]$ links the helices. The resulting “native” structure when all bonds are formed is shown in Fig. 6(b).

The evaluation of the drop in entropy and the outer first passage times for the formation of the long-range disulfide bonds benefits from the use of the staircase bias, given that the decrease in entropy for such bonds is roughly 10 and $\tau^+ \sim 10^3$. The input parameters, mean first passage times, and biased entropies for the simulation of crambin are available on GitHub (see Acknowledgments) in the hybridmc/examples folder. To visualize the free energy landscape and the kinetics for a pairwise-additive model in which each bond formed lowers the potential energy by a fixed amount ϵ_b , we use disconnectivity graphs.^{86,87} The node levels in the graphs are determined by the dimensionless free energies, and the transition state free energies are set by adding the negative logarithm of the rate to the state's free energy. Changes in the morphology of the free energy landscape as the temperature is modified can be tracked by the structure of the disconnectivity graphs. Unsurprisingly, the disconnectivity graph for the model crambin system shown in Fig. 7(a) exhibits a “funnel-shape” at low temperatures ($\beta^* = 12$), in which the fully bonded structure corresponds to a deep-lying node in the graph, centrally flanked by local minima over a wide range of free energy values.⁸⁶ The folding dynamics for the model, which exhibits no kinetic traps, is particularly simple. The average folding




FIG. 6. The model crambin system. Panel (a) is a cartoon representation of the structure of the crystallized protein. Panel (b) is the fully folded minimum entropy state of the 46-bead, 10-bond model. The structure in panel (c) contains five bonds leading to a helical structure but none of the four nonlocal bonds between distant monomers that culminate the folding process. In all three figures, the beads participating in nonlocal bonds in the α -helices are in red, the β -sheets are in yellow, and the disulfide bridges are in blue and cyan for beads 16 and 40, respectively.




FIG. 7. The disconnectivity graphs for the model crambin system in the low temperature regime, with $\beta = 12$. (a) Tree diagram: crambin model. (b) Tree diagram: optimized crambin model.

time needed to pass from an initial state with no bonds to the target native state, shown in Fig. 8 in dimensionless units inversely proportional to the self-diffusion coefficient D , decreases monotonically as β increases and approaches a constant value. Note that this does not imply that the folding rate is fastest at low temperatures, since from kinetic theory,⁸⁸ the diffusion coefficient is expected to scale as $D \sim \beta^{-1/2}$ so that the physical folding time increases at low temperatures.

The most probable pathways of transitioning from the non-bonded state to the folded target state can be visualized using




FIG. 8. The average folding times in units of $1/D$ as a function of inverse temperature β for the crambin model.

network diagrams, in which each state appears as a node whose size is represented by the probability r^+ [see Eq. (31)] that the state is visited in the reactive ensemble. The connecting arrows represent the reactive flux [see Eq. (32)]. For the equal bond energy model at low temperatures where the rate of breaking a bond is small and the committer probability $q_i^+ \sim 1$ for all bonded states, the folding pathways primarily consist of two distinct parts: Five local bonds are formed first, leading to a helical intermediate state 617 [colored red in Fig. 6(c), with β -sheet [2, 34], disulfide bridge [4, 32], and α -helix bonds [6, 10], [14, 18], [26, 30] turned on], followed by the formation of the disulfide bridge [16, 26], the β -sheet [2, 34], [3, 33], and the [4, 32] bond. In 70% of the folding pathways, the most probable final transition to the folded structure involves the formation of the disulfide bridge [3, 40], linking the β -sheet to the random coil end of the chain, denoted as state 1021 to state 1024.

Heretofore, we have assumed that the formation of a bond changes the energy of a configuration by an amount ϵ_b . Suppose we are interested in determining the optimal set of interactions that lead to a given structure, while maintaining a set of physical requirements, i.e., the fully bonded structure has a free energy that is well-separated from other structures so that it is thermodynamically preferred over a range of temperatures yet reached quickly from a fully unbonded configuration. In principle, since the coarse-grained models are allowed a nonadditive (i.e., not pairwise) decomposition of the potential energy to permit hidden effects such as hydrophobicity not directly incorporated into the model [see Eq. (2)], arbitrary choices of the energies of states are possible provided they are physical. To mimic evolutionary behavior, we minimize the mean folding time in Eq. (33) with respect to the set of state energies $\{E^*\}$, subject to the constraints that (1) the ratio of the probability of the unfolded state to the folded state is fixed (i.e., the state energies

of the folded and unfolded states are constant), (2) the probability that each partially folded state cannot be too large, enforced by a constraint $0.005 < P_i/P_f$, and (3) the maximum energy of a given state is restricted to a finite value (taken here to be less than 10, well above the zero energy of the unbonded state). For the crambin model with ten bonds, there are 1022 intermediate states whose energies are varied to minimize the mean folding time. To carry out the minimization procedure of a loss function with many possible local minima, we use methods that combine the BFGS search algorithms with multiple trajectory sampling.^{80–82}

The result of the minimization procedure with $P_u/P_f = 10^{-52}$, a value of the relative probability corresponding to $\beta^* = 12$ when the bond energy is fixed at $\epsilon_b = 1$, lowers the mean folding time by a factor of roughly 2 over a range of temperature values, as shown in Fig. 8. Nonetheless, the smooth funnel morphology of the

disconnectivity graph is maintained [see Fig. 7(b)]. The disconnectivity graph of the optimized model is more segmented, particularly in the last level of states, with most states at a given level having similar probability and hence roughly the same free energy. From the network diagram of the optimized model shown in Fig. 9(b), it is apparent that the folding mechanism is significantly altered. The optimization yields energies of states that make the pathways leading to the helical transition state equally likely (similar values of r^+), and the effect of bottleneck state 617 is mitigated by substantially facilitating the β -sheet formation by decreasing the energies of states with long-range bonds (such as the β -sheet [2, 34] and [3, 33] bonds) to allow additional connecting pathways at level 5. These findings are consistent with the view that structure grows locally and models with local stabilizing interactions that compensate the conformational entropy loss as local structure forms result in faster folding.⁸⁹ The change in folding mechanism and the increase in the folding rate correlate with the increase in the “contact order” in which the mean separation in sequence between bonding beads^{90,91} is weighted by the reactive probability r^+ for that bond. There are also rapid transitions between the state 830 and 943 that both have the [3, 33] bond and then form or break the adjacent β -sheet [2, 34] bond. The energies of states in the penultimate level are similarly adjusted to create three equally likely pathways to the final state. These results imply that even a system with a smooth funnel will fold more quickly when the state energies allow for a multitude of pathways rather than passing through a fixed sequence of states, in agreement with studies of fast-folding proteins.¹⁸

B. Eliminating frustration and misfolding

The multiplicity of competing interactions in real systems can give rise to free energy landscapes with many local minima, resulting in long-lived metastable structures. Small proteins that fold quickly




FIG. 9. Simplified network diagrams of the most probable network of folding pathways for the model crambin system at low temperatures, where $P_u/P_f = 10^{-52}$, and in which only states with reactive probability $r^+ > 0.1$ are shown [see Eq. (31)]. The size of each node is representative of the probability of visiting the state in the reactive ensemble, and the size of the arrows between nodes represents the reactive flux between them [see Eq. (32)]. On the left is the network for a system with bonds of equal energy with bottleneck state 617 colored red, and on the right is the network following the optimization of the folding time. Note that the two models have very different networks and folding pathways. (a) Network diagram: equal bond energy crambin model. (b) Network diagram: optimized energy crambin model.




FIG. 10. The reactive and average folding times (in units of $1/D$) as a function of inverse temperature β for the frustrated model. At high temperatures (low values of β), the equilibrium probability of the folded state is small, leading to a large folding time due to trajectories that return to the unfolded state. At low temperatures (large β), the slow transitions out of trapping states lead to folding times that increase rapidly with β .




FIG. 11. Example structures of the frustrated model system. On the left is the fully folded state (structure 37), in which all the bonding constraints between bonding atoms (yellow) are satisfied. In the center is the transition state (structure 36), satisfying all local helical bonds, with the largest flux between it and the final state. On the right is a kinetically trapped state (structure 29), in which the bond between bead 7 (blue) and bead 11 (red) cannot be formed without breaking existing bonds. There are three low-lying trapping configurations in the model, identified as states 29, 30, and 34. (a) Fully-bonded state: state 37. (b) Maximum flux state: state 36. (c) State 29 lacking a [7, 11] bond.

have amino acid sequences that lead to thermodynamically stable configurations and avoid kinetically trapped metastable states.

To examine the role of kinetic traps and their elimination through a selection process, we consider a 14-bead model with a set of bonding interactions $\{[4, 12], [3, 7], [5, 9], [7, 11], [9, 13]\}$, all formed at $r_c = 1.5$. Highly bonded structures in this model, shown in Fig. 11, resemble a short α -helix that folds over due to the long-range interaction between beads 4 and 12. The model mimics misfolding due to the existence of kinetic traps: For the trapping states, the manner in which the structure satisfies a set of bonding constraints geometrically prohibits the formation of the additional

bonds required to reach the fully bonded structure. An example of such a structure is shown in Fig. 11(c). The identification of trapping states and the calculation of their entropies can be difficult. Within the layer approach, the trapping states are found by identifying which and how many states in a pool of possible structures are incapable of reaching a target state the next level down in a short trajectory. Each of the pools is iteratively constructed in parallel, starting from the fully unfolded state with no bonds. For a given state, its pool of structures is generated by using the pools of all structures in the previous layer that can reach the target state by the formation of a single bond.




FIG. 12. Disconnectivity graphs of the 14-bead, 5-bond model. Note the difference in locations of the trapping states 29, 30, and 34 as well as state 17 in the respective graphs. (a) Frustrated model. (b) Optimized model.




FIG. 13. The first passage times in units of $1/D$ from the trapping states to the folded state. The inset in both figures shows the reactive probability of passing through the trapping states. Note that the similarities of the first passage times to the mean first passage time in Fig. 10 indicates that the first passage times out of the three trapping states dictate the overall folding behavior of both models. (a) Frustrated model. (b) Optimized model.

For this model system, we find a total of five trapping states: Two of the states are in the third level and have three bonds, and three are found in the fourth level and cannot form the $[3, 7]$ bond (structure 29), the $[5, 9]$ bond (structure 30), or the $[9, 13]$ bond (structure 34) due to the preexisting long-range bond between beads 4 and 12. In Fig. 12(a), the trapping states appear as a separate fork in the disconnectivity graph of the model system at low temperatures, since dynamical events with high free energy barriers that break bonds must occur to reach the fully folded state. The effect of the trapping states on the dynamics is significant and leads to a qualitatively different temperature dependence from the fast-folding model of crambin. At high temperatures, as in the crambin model, the fully folded structure is thermodynamically unfavorable and reactive trajectories have low probability. Instead, the main contribution to the average first passage time to the target state comes from the nonreactive paths that repeatedly revisit the unfolded state. The probability of nonreactive trajectories rapidly decreases with temperature, and the minimum folding time is reached at intermediate values of β near $\beta = 6$. At low temperatures, many folding trajectories become kinetically trapped, and the folding time increases exponentially as the free energy barrier increases. This kinetic trapping, whose inverse temperature dependence is plotted in Fig. 13(a), follows the same trends as the mean folding time [see Fig. 10].

As is evident in Fig. 14, the constrained minimization of the folding time with respect to the state energies eliminates the effect of the trapping states by raising the energy of the trapping states so that either they have negligible reactive probability r^+ at all temperatures (state 34) or they are in resonance with state 17 in the previous level with fewer bonds (states 29 and 30). In the optimized model, these states have a low activation barrier and rapidly break the bond connecting them to the less bonded state 17, which appears with enhanced reactive probability. At the same time, the flux of non-trapping states in the final layer is optimized by lowering the energies of those states to allow multiple pathways of similar probability to pass to the target state. These effects are evident in the disconnectivity graph of the optimized model in Fig. 12(b) by the equal barrier

heights of the layer of states and the shift of trapping states 29, 30, and 34 to higher points. Additionally, the trap outlet state 17, which contains the $[4, 12]$, $[7, 11]$, and $[9, 13]$ bonds, is repositioned in the tree-like structure.

The changes in the network diagrams of the folding [see Fig. 15(b)] highlight the disappearance of node 34 and the increased flux of transitions among states in the middle levels of the network. In both the equal bond energy and optimized energy models, the main final transition occurring in 70% of the folding trajectories to the target state consists of the formation of the $[4, 12]$ bond. As a result, the qualitative nature of the folding pathways is similar to




FIG. 14. Energy adjustments to minimize the average folding time when the free energy difference between the folded and unfolded states is constrained to be 60 and $P_u/P_f = 10^{-26}$. Note that the optimization raises the energies of the trapping states 29, 30, and 34 to either minimize their path probability and flux or facilitate escape by lowering backward transition barriers to outlet state 17, which now lies below the energies of the trapping states.




FIG. 15. Network diagrams of the folding process for the frustrated model. The size of a node i is representative of the probability r_i^+ that the state is visited in the reactive ensemble, and the size of the arrows between nodes represents the reactive flux between them [see Eq. (32)]. On the left is the network for a system with bonds of equal energy, and on the right is the network following the optimization of the folding time. Note the disappearance of trapping state 34 and the flux between trapping states 29 and 30 (all colored red) and state 17 (colored green) that appears in the optimized network. (a) Network diagram: frustrated model. (b) Network diagram: optimized model.

that of the crambin model: The system first forms a helical element that subsequently folds into the final structure. After optimization and removal of the trapping kinetics, the temperature dependence of the mean first passage time from the unfolded to folded states in Fig. 13(b) approaches a constant as β increases, the same qualitative behavior observed in the rapid-folding crambin system.

VI. DISCUSSION AND CONCLUSIONS

In this work, we introduced methodology to address the computational challenges of computing the entropy and mean first passage times for a linear chain model of proteins in which monomers interact discontinuously. These quantities appear as parameters in Markov state models of the population dynamics. The methods combine adaptive sampling algorithms with statistical tests to compute reliable interval estimates for all quantities. Given the exponential growth of the number of states with the number of bonding interactions included in the model, parallel algorithms are a critical requirement to investigate large, complex models. The level-based calculations in which individual pairs of linked states are conducted in parallel with the inclusion of either replica-exchange or population Monte Carlo components improve the rate of convergence for the entire set of state entropies and allow the computation to be carried out in massively parallel platforms with coarse-grained parallelism. The numerical sampling difficulties associated with large first passage times are lessened by the introduction of intermediate staircase states, which was shown to significantly reduce the sampling required for a given statistical resolution.

There are possible improvements to the sampling that are relatively simple to implement. The methods presented here rely on sampling states using a model with a discontinuous potential with event-driven dynamics that prevent the implementation of continuous adaptive biases frequently used in the molecular simulation community. However, auxiliary sampling chains based on dynamical trajectories governed by continuous potential approximations to Heaviside and infinite square well functions can be applied to generate trial Monte Carlo updates, provided that the acceptance criterion is suitably adjusted.⁹² The continuous potentials can be adaptively adjusted along bonding distances using methods such as well-tempered metadynamics^{37,58} and self-healing umbrella sampling.⁵⁹ However, some care is required to ensure that the continuous potential system does not frequently allow configurations that violate the strict geometrical constraints of the model. Current studies along these lines are underway.

Given the computational cost of models with many nonlocal bonds, an interesting open question is whether or not the methods of machine learning on small systems can be used to accurately infer the entropies and first passage times of more complicated models. Nonetheless, machine learning algorithms frequently require large sets of training data to be useful. The sampling algorithms introduced here can help with the task of generating the necessary training datasets.

Machine learning methods may also prove useful in the classification of trapping states that subdivide configurations determined by their bonding patterns alone. Geometrical descriptors that generalize the state indicator functions will allow for more accurate evaluation of the probability of the kinetic traps as well as their first passage times.

The main appeal of the Markov state model of the discontinuous potential lies in the possibility to evaluate both the structure and dynamics for an infinite number of choices of interactions at arbitrary temperatures, once the density of states and the first passage times have been calculated for a given choice of bonding pattern. In Secs. V A and V B, we demonstrated how this flexibility may be exploited to select interaction energies that enhance physical properties or desired functional characteristics. In these sections, we analyzed the mean passage time from the unbonded state to the fully bonded state for a model of the crambin protein, which folds quickly and has a free energy landscape with a funnel-shaped morphology, and for a short 14-bead helical protein designed to exhibit a more complex free energy landscape and trapping kinetics.

For the crambin model, a choice of pairwise-additive bond energies for states led to a simple mechanistic folding pathway, in which the helical portions of the model protein formed first with no clear preference of order, followed by the passage with near unit probability through a helical bottleneck state. In the second step of the folding, distant bonds linking regions of the helix to one another lead to a penultimate state with all bonds present except for those linking the most distant edges of the chain. The optimization of the state energies with a fixed relative probability of the unfolded to folded states resulted in a different folding mechanism and folding rates that were twice as large. Interestingly, the initial stage of the folding process in the optimization involved the rapid formation of the local helical bonds. However, the passage through the restrictive bottleneck state was discouraged by adjustments that lowered the energy of states with long-range bonds. Similarly, multiple pathways to the final state were found due to energy adjustments of the long-range bonds to compensate for their different entropic values.

The folding process for the 14-bead model system with frustration was also initiated by the formation of local interactions and a helical precursor to the final folded state. However, the inclusion of a high density of interactions in the model introduces a number of kinetic traps that are reflected in an exponentially increasing folding time as the temperature decreases, and new distinct branches appear in the disconnectivity graphs of the free energy landscape at low temperatures. The deepest-lying trapping states determine the folding time at low temperatures. In this case, the optimization of energies destabilized the trapping states so that they either had a negligible probability in the folding pathway or were positioned in resonance with states with fewer bonds to enable the rapid breaking of a bond.

The optimal energies for rapid folding depend on the choice of constraints employed in the optimization procedure, and these constraints should reflect conditions that are realizable for molecular systems. If the bonds formed are intended to represent weak electrostatic or hydrogen bonding interactions between segments of the chain, the maximum drop in the state's energy should be restricted in magnitude. Large increases in the state energies can easily be achieved through steric repulsion or torsional strain. It is likely chain stiffness along the peptide backbone effectively limits the density of bonds in the chain to avoid this type of kinetic trap. It would be interesting to explore the inclusion of other information in the loss function. For example, a target electrostatic map for the folded

structure could be included as a penalty in the loss function and monomer-specific partial charges used to determine optimal residue sequences.

The linear chain model can be generalized to include side chain beads interacting with other beads to allow for the inclusion of steric effects of bulky residues as well as attractive nonlocal bonding. Such features are important in determining the overall three-dimensional structure of real proteins. The sampling methods and optimization procedure of the Markov state model introduced here can be applied without modification.

The folding mechanism and optimization of the state energies in the models of fast-folding proteins analyzed in this work are indicative of the type of issues that can be explored with the discontinuous potential model. Its simplicity opens the door to explore general questions that are difficult to address by other means. Other avenues to be explored include the following: Given a particular three-dimensional structure, what state energies lead to rapid folding and thermodynamic stability? To what extent is the optimization of the native state of a protein for fast refolding dictated by its topology? How additive are the energies of biopolymers? How do biomolecular systems avoid kinetic traps? Why do certain motifs of secondary structure appear and not others, and what role do secondary structures play in the folding pathways? Does the optimization of the folding time confirm well-established principles of fast-folding proteins, such as the importance of Ramachandran angles, the existence of foldons, and the statistical correlation between contact order^{90,91} and folding rate? What are the differences in the interaction patterns of fast-folding vs intrinsically disordered proteins? Work along these lines is underway.

ACKNOWLEDGMENTS

Financial support from the Natural Sciences and Engineering Research Council of Canada is gratefully acknowledged. Computations were performed on the Cedar supercomputer at Simon Fraser University, which is funded by the Canada Foundation for Innovation under the auspices of Compute Canada, WestGrid, and Simon Fraser University. Code for this project is available at <https://github.com/margaritacolberg/hybridmc>.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Margarita Colberg: Writing – original draft (lead); Writing – review & editing (equal). **Jeremy Schofield:** Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹S. Y. Joshi and S. A. Deshmukh, *Mol. Simul.* **47**, 786 (2021).
- ²S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, and A. Kolinski, *Chem. Rev.* **116**, 7898 (2016).
- ³S. Kmiecik, M. Kouza, A. Badaczewska-Dawid, A. Kloczkowski, and A. Kolinski, *Int. J. Mol. Sci.* **19**, 3496 (2018).
- ⁴H. Taketomi, Y. Ueda, and N. Gö, *Int. J. Pept. Protein Res.* **7**, 445 (1975).
- ⁵Y. Ueda, H. Taketomi, and N. Gö, *Biopolymers* **17**, 1531 (1978).
- ⁶N. Gö and H. Taketomi, *Int. J. Pept. Protein Res.* **13**, 235 (1979).
- ⁷N. Gö and H. Abe, *Biopolymers* **20**, 991 (1981).
- ⁸H. Abe and N. Gö, *Biopolymers* **20**, 1013 (1981).
- ⁹S. Takada, *Biophys. Physicobiol.* **16**, 248 (2019).
- ¹⁰M. M. Tirion, *Phys. Rev. Lett.* **77**, 1905 (1996).
- ¹¹L. Yang, G. Song, and R. L. Jernigan, *Proc. Natl. Acad. Sci. U. S. A.* **106**, 12347 (2009).
- ¹²J. Schofield, P. Inder, and R. Kapral, *J. Chem. Phys.* **136**, 205101 (2012).
- ¹³M. Orozco, *Chem. Soc. Rev.* **43**, 5051 (2014).
- ¹⁴K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, *Science* **334**, 517 (2011).
- ¹⁵M. Karplus and D. L. Weaver, *Protein Sci.* **3**, 650 (1994).
- ¹⁶A. R. Fersht, *Curr. Opin. Struct. Biol.* **7**, 3 (1997).
- ¹⁷K. A. Dill, S. B. Ozkan, M. S. Shell, and T. R. Weikl, *Annu. Rev. Biophys.* **37**, 289 (2008).
- ¹⁸P. G. Wolynes, J. N. Onuchic, and D. Thirumalai, *Science* **267**, 1619 (1995).
- ¹⁹P. E. Leopold, M. Montal, and J. N. Onuchic, *Proc. Natl. Acad. Sci. U. S. A.* **89**, 8721 (1992).
- ²⁰M. Krishna, L. Hoang, Y. Lin, and S. W. Englander, *Methods* **34**, 51 (2004).
- ²¹M. M. G. Krishna, H. Maity, J. N. Rumbley, Y. Lin, and S. W. Englander, *J. Mol. Biol.* **359**, 1410 (2006).
- ²²H. Maity, M. Maity, M. M. G. Krishna, L. Mayne, and S. W. Englander, *Proc. Natl. Acad. Sci. U. S. A.* **102**, 4741 (2005).
- ²³W. Hu, Z.-Y. Kan, L. Mayne, and S. W. Englander, *Proc. Natl. Acad. Sci. U. S. A.* **113**, 3809 (2016).
- ²⁴Y. Bai, T. R. Sosnick, L. Mayne, and S. W. Englander, *Science* **269**, 192 (1995).
- ²⁵W. Hu, B. T. Walters, Z.-Y. Kan, L. Mayne, L. E. Rosen, S. Marqusee, and S. W. Englander, *Proc. Natl. Acad. Sci. U. S. A.* **110**, 7684 (2013).
- ²⁶B. Kuhlman, G. Dantas, G. C. Ireton, G. Varani, B. L. Stoddard, and D. Baker, *Science* **302**, 1364 (2003).
- ²⁷F. Richter and D. Baker, in *Synthetic Biology*, edited by H. Zhao (Academic Press, Boston, 2013), pp. 101–122.
- ²⁸Z. Chen, M. C. Johnson, J. Chen, M. J. Bick, S. E. Boyken, B. Lin, J. J. De Yoreo, J. M. Kollman, D. Baker, and F. DiMaio, *J. Am. Chem. Soc.* **141**, 8891 (2019).
- ²⁹W. Zhou, T. Šmidlehner, and R. Jerala, *FEBS Lett.* **594**, 2199 (2020).
- ³⁰M. H. Cordes, A. R. Davidson, and R. T. Sauer, *Curr. Opin. Struct. Biol.* **6**, 3 (1996).
- ³¹W. P. Russ, M. Figliuzzi, C. Stocker, P. Barrat-Charlaix, M. Socolich, P. Kast, D. Hilvert, R. Monasson, S. Cocco, M. Weigt, and R. Ranganathan, *Science* **369**, 440 (2020).
- ³²A. Hawkins-Hooker, F. Depardieu, S. Baur, G. Couairon, A. Chen, and D. Bikard, *PLoS Comput. Biol.* **17**, e1008736 (2021).
- ³³J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis, *Nature* **596**, 583 (2021).
- ³⁴J. Echave and C. O. Wilke, *Annu. Rev. Biophys.* **46**, 85 (2017).
- ³⁵X. Ding, Z. Zou, and C. L. Brooks III, *Nat. Commun.* **10**, 5644 (2019).
- ³⁶K. A. Dill, *Curr. Opin. Struct. Biol.* **3**, 99 (1993).
- ³⁷E. I. Shakhnovich and A. M. Gutin, *Proc. Natl. Acad. Sci. U. S. A.* **90**, 7195 (1993).
- ³⁸H. Jacquin, A. Gilson, E. Shakhnovich, S. Cocco, and R. Monasson, *PLoS Comput. Biol.* **12**, e1004889 (2016).
- ³⁹F. Y. Wu, *Rev. Mod. Phys.* **54**, 235 (1982).
- ⁴⁰F. Morcos, N. P. Schafer, R. R. Cheng, J. N. Onuchic, and P. G. Wolynes, *Proc. Natl. Acad. Sci. U. S. A.* **111**, 12408 (2014).
- ⁴¹Y. Zhou and M. Karplus, *Proc. Natl. Acad. Sci. U. S. A.* **94**, 14429 (1997).
- ⁴²Y. Zhou, M. Karplus, J. M. Wichert, and C. K. Hall, *J. Chem. Phys.* **107**, 10691 (1997).
- ⁴³Y. Zhou and M. Karplus, *J. Mol. Biol.* **293**, 917 (1999).
- ⁴⁴H. B. Movahed, R. van Zon, and J. Schofield, *J. Chem. Phys.* **136**, 245103 (2012).
- ⁴⁵J. Schofield and H. Bayat, *J. Chem. Phys.* **141**, 095101 (2014).
- ⁴⁶R. van Zon and J. Schofield, *J. Chem. Phys.* **132**, 154110 (2010).
- ⁴⁷J. Schofield, *J. Phys. Chem. B* **121**, 6847 (2017).
- ⁴⁸J. Schofield and M. A. Ratner, *J. Chem. Phys.* **109**, 9177 (1998).
- ⁴⁹P. H. Verdier and W. H. Stockmayer, *J. Chem. Phys.* **36**, 227 (1962).
- ⁵⁰S. K. Kumar, M. Vacatello, and D. Y. Yoon, *J. Chem. Phys.* **89**, 5206 (1988).
- ⁵¹J. I. Siepmann and D. Frenkel, *Mol. Phys.* **75**, 59 (1992).
- ⁵²D. Frenkel, G. C. A. M. Mooij, and B. Smit, *J. Phys.: Condens. Matter* **4**, 3053 (1992).
- ⁵³D. Rezende and S. Mohamed, in *Proceedings of the 32nd International Conference on Machine Learning, PMLR '15*, edited by F. Bach and D. Blei (Proceedings of Machine Learning Research, Lille, France, 2015), Vol. 37, p. 1530.
- ⁵⁴S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, *Phys. Lett. B* **195**, 216 (1987).
- ⁵⁵D. C. Rapaport, *The Art of Molecular Dynamics Simulation*, 2nd ed. (Cambridge University Press, Cambridge, 2004).
- ⁵⁶F. Wang and D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001).
- ⁵⁷A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12562 (2002).
- ⁵⁸A. Barducci, G. Bussi, and M. Parrinello, *Phys. Rev. Lett.* **100**, 020603 (2008).
- ⁵⁹S. Marsili, A. Barducci, R. Chelli, P. Procacci, and V. Schettino, *J. Phys. Chem. B* **110**, 14011 (2006).
- ⁶⁰G. Fort, E. Moulines, and P. Priouret, *Ann. Stat.* **39**, 3262 (2012).
- ⁶¹G. Fort, B. Jourdain, T. Lelièvre, and G. Stoltz, *Stat. Comput.* **27**, 147 (2017).
- ⁶²P. J. Smith, D. S. Rae, R. W. Manderscheid, and S. Silbergeld, *J. Am. Stat. Assoc.* **76**, 737 (1981).
- ⁶³G. Marsaglia, W. W. Tsang, and J. Wang, *J. Stat. Software* **8**, 1 (2003).
- ⁶⁴L. A. Goodman, *Technometrics* **7**, 247 (1965).
- ⁶⁵W. L. May and W. D. Johnson, *J. Stat. Software* **5**, 1 (2000).
- ⁶⁶C. J. Geyer, “Markov chain Monte Carlo maximum likelihood,” in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface* (Interface Foundation of North America, 1991), pp. 156–163.
- ⁶⁷C. J. Geyer and E. A. Thompson, *J. Am. Stat. Assoc.* **90**, 909 (1995).
- ⁶⁸R. M. Neal, *Stat. Comput.* **6**, 353 (1996).
- ⁶⁹S. B. Opps and J. Schofield, *Phys. Rev. E* **63**, 056701 (2001).
- ⁷⁰D. Sindhikara, Y. Meng, and A. E. Roitberg, *J. Chem. Phys.* **128**, 024103 (2008).
- ⁷¹T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, *Phys. Rev. Lett.* **110**, 210603 (2013).
- ⁷²F. Moreno, S. Davis, and J. Peralta, *Comput. Phys. Commun.* **274**, 108283 (2022).
- ⁷³L. Bornn, P. E. Jacob, P. Del Moral, and A. Doucet, *J. Comput. Graph. Stat.* **22**, 749 (2013).
- ⁷⁴V. Elvira, L. Martino, D. Luengo, and M. F. Bugallo, *Signal Process.* **131**, 77 (2017).
- ⁷⁵N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (North Holland, 2007).
- ⁷⁶L. J. S. Allen, *An Introduction to Stochastic Processes With Applications to Biology* (Prentice-Hall, Upper Saddle River, NJ, 2003).
- ⁷⁷C. D. Meyer, *SIAM Rev.* **31**, 240 (1989).
- ⁷⁸M. von Kleist, C. Schütte, and W. Zhang, *J. Stat. Phys.* **170**, 809 (2018).
- ⁷⁹D. J. Sharpe and D. J. Wales, *Phys. Rev. E* **104**, 015301 (2021).
- ⁸⁰L.-Y. Tseng and C. Chen, in *2008 IEEE Congress on Evolutionary Computation* (IEEE World Congress on Computational Intelligence, Hong Kong, 2008), p. 3052.

- ⁸¹T. Liao, M. Montes de Oca, D. Aydin, T. Stützle, and M. Dorigo, in *Genetic and Evolutionary Computation Conference, GECCO '11* (Association for Computing Machinery, 2011), p. 125.
- ⁸²U. Kumar, Jayadeva, and S. Somit, in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation, GECCO '15* (Association for Computing Machinery, 2015), p. 33.
- ⁸³M. Ormó, A. B. Cubitt, K. Kallio, L. A. Gross, R. Y. Tsien, and S. J. Remington, "Crystal structure of the *Aequorea victoria* green fluorescent protein," *Science* **273**, 1392–1395 (1996).
- ⁸⁴C. Jelsch, M. M. Teeter, V. Lamzin, V. Pichon-Pesme, R. H. Blessing, and C. Lecomte, *Proc. Natl. Acad. Sci. U. S. A.* **97**, 3171 (2000).
- ⁸⁵A. Schmidt, M. Teeter, E. Weckert, and V. S. Lamzin, *Acta Crystallogr., Sect. F: Struct. Biol. Cryst. Commun.* **67**, 424 (2011).
- ⁸⁶O. M. Becker and M. Karplus, *J. Chem. Phys.* **106**, 1495 (1997).
- ⁸⁷D. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses*, Cambridge Molecular Science (Cambridge University Press, Cambridge, 2004).
- ⁸⁸S. Chapman and T. G. Cowling, *The Mathematical Theory of Non-Uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction, and Diffusion in Gases* (Cambridge University Press, 1990).
- ⁸⁹V. Muñoz and W. A. Eaton, *Proc. Natl. Acad. Sci. U. S. A.* **96**, 11311 (1999).
- ⁹⁰K. W. Plaxco, K. T. Simons, and D. Baker, *J. Mol. Biol.* **277**, 985 (1998).
- ⁹¹A. R. Dinner and M. Karplus, *Nat. Struct. Biol.* **8**, 21 (2001).
- ⁹²R. Iftimie, D. Salahub, D. Wei, and J. Schofield, *J. Chem. Phys.* **113**, 4852 (2000).