

Efficient exploration of discrete energy landscapes

Martin Mann

Bioinformatics Group, University of Freiburg, Georges-Köhler-Allee 106, D-79110 Freiburg, Germany

Konstantin Klemm

*Bioinformatics Group, Institute for Computer Science,
University of Leipzig, Härtelstraße 16-18, D-04107 Leipzig, Germany*

(Dated: November 12, 2010)

Many physical and chemical processes, such as folding of biopolymers, are best described as dynamics on large combinatorial energy landscapes. A concise approximate description of the dynamics is obtained by partitioning the micro-states of the landscape into macro-states. Since most landscapes of interest are not tractable analytically, the probabilities of transitions between macro-states need to be extracted numerically from the microscopic ones, typically by full enumeration of the state space or approximations using the Arrhenius law. Here we propose to approximate transition probabilities by a Markov chain Monte-Carlo method. For landscapes of the number partitioning problem and an RNA switch molecule we show that the method allows for accurate probability estimates with significantly reduced computational cost.

PACS numbers: 05.10.Ln, 87.15.H-, 02.50.Ga

I. INTRODUCTION

Energy landscapes [1–3] are a key concept for the description of complex physical and biological systems. In particular, the dynamics of structure formation (“folding”) of biopolymers, *e.g.* protein or ribonucleic acids, can be understood in terms of their energy landscapes [4, 5]. Formally, a landscape is determined by a set X of micro-states (or conformations), a neighborhood structure of X that encodes which conformations can be reached from which other ones, and an energy function $E : X \rightarrow \mathbb{R}$ which assigns an energy value to each state. In the case of ribonucleic acids (RNA) it has been demonstrated that the dynamics of the folding process can be captured in good approximation by merging large contiguous sets of micro-states into macro-states [6, 7]. A typical mapping is in terms of gradient basins: Each macro-state contains the micro-states from which a given local minimum is reached by steepest descent in energy, including the local minimum itself. The so-defined macro-states are also called inherent structures in the context of continuous disordered systems, see ref. [8] for a recent review.

Given a partitioning of the landscape, the dynamics is approximately described as a Markov chain on the set of macro-states. In order to obtain this description, the transition probabilities between macro-states in this Markov chain need to be extracted from the original energy landscape.

As a first approximation, the *Arrhenius* equation predicts that the transition probability is exponentially suppressed by the ratio between barrier height and temperature. The barrier height (also called activation energy) from minimum a to minimum b measures the minimal amount by which the system’s energy must increase along a path from a to b [6, 9, 10]. The accuracy of this approach is limited because it ignores the multiplicity of

low energy paths [11]. A more severe drawback is the complexity of computing barrier height itself. For landscapes of RNA secondary structure [12], the problem is NP-hard [13, 14].

Commonly used methods [6, 15–18] for precise transition rate estimation are based on enumeration of all micro-states. For landscapes of real combinatorial problems or long biopolymers with billions of micro-states, however, enumeration is impractical with the given time resources. Typically, limited storage capacity puts even more severe restrictions on the size of tractable problems because a large fraction of the enumerated micro-states needs to be kept in working memory. Some studies partially circumvent this problem by considering only the low-energy fraction of the landscape that is tractable with the available resources [9, 19–21]. Other heuristic approaches [22–26] restrict the landscape to the subset of states likely to be traversed by certain trajectories, *e.g.* folding from the open chain to the ground state of a biopolymer.

Here we make a contribution to the original challenge of capturing an arbitrary discrete landscape in terms of macro-states and transition probabilities. We suggest a Markov chain Monte-Carlo sampling method for transition matrix estimation. At difference with the earlier approaches, the memory requirement scales linearly with the number of non-zero transition probabilities to be determined. Other recent methods of stochastic landscape exploration [27, 28] use trajectories of the original dynamics for counting transitions between macro-states. In contrast, the idea behind the present method is to explicitly explore boundaries between macro-states. To this end, we confine the dynamics into a single macro-state b and find and count possible transitions from b to all adjacent macro-states. This strategy allows to select the regions of the landscape to be explored and the precision to be applied.

II. LANDSCAPE AND MICRO-STATE DYNAMICS

A discrete energy landscape is a triple (X, E, M) where

- X is a finite set of states,
- $E : X \rightarrow \mathbb{R}$ is an energy function on X , and
- $M : X \rightarrow \mathcal{P}(X)$ is a neighborhood function or “move set” that assigns to each state $x \in X$ the set of its directly accessible neighboring states. $\mathcal{P}(X)$ is the power set of X . Here we assume that M is symmetric, *i.e.* $x \in M(y) \Rightarrow y \in M(x)$. By Δ we denote the maximum number of neighbors, $\Delta = \max_{x \in X} |M(x)|$.

We consider a time-discrete stochastic dynamics on the state set X . Having the Markov property, the dynamics is defined by giving the transition probability $p_{x \rightarrow y}$ from each $x \in X$ to each $y \in M(x)$. Provided the system is in state x at time t , $p_{x \rightarrow y}$ is the probability that the system is in state y at time $t + 1$. With probability $p_{x \rightarrow x} = 1 - \sum_{y \in M(x)} p_{x \rightarrow y}$, the system remains at state x .

Specifically, the Metropolis probabilities at inverse temperature β ,

$$p_{x \rightarrow y} = \Delta^{-1} \min\{\exp(\beta[E(x) - E(y)]), 1\} \quad (1)$$

are used throughout this contribution. This choice, however, is not compulsory. All that follows, and in particular the estimation by sampling, applies to arbitrary choices of transition probabilities leading to *ergodic* Markov chains. The ergodicity is important because we need a unique stationary distribution $P(x)$ on X .

III. PARTITIONING AND MACRO-STATE DYNAMICS

A partitioning of the landscape is a mapping F from the set of micro-states X into a set of macro-states B . Our goal here is to find a dynamics on B that does have the Markov property while following the original micro-state dynamics as closely as possible. In general, however, a Markov chain is not obtained as the direct mapping $(F(x_t))_{t=0}^{\infty}$ of a Markov chain $(x_t)_{t=0}^{\infty}$ generated by the dynamics on X . The reason can be sketched as follows. When the system is in a macro-state $b \in B$, the probability of exiting to a macro-state c depends on where exactly (in which micro-state) the system is inside b . The micro-state assumed inside b , however, depends on how the system entered b , which is again influenced by the macro-state a assumed before entering b .

Thus, the following simplifying assumption is made [29]. Given that the system is found in macro-state $b \in B$, the micro-state $x \in X$ is distributed as

$$P_b(x) = \begin{cases} P(x) / \sum_{y \in F^{-1}(b)} P(y) & \text{if } x \in F^{-1}(b) \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

This is the stationary distribution P of the whole system restricted to micro-states in b and normalized appropriately. Under this assumption, the probability of a transition to macro-state c , when being in macro-state $b \neq c$ is

$$q_{b \rightarrow c} = \sum_{x \in F^{-1}(b)} \left(P_b(x) \sum_{y \in M(x) \cap F^{-1}(c)} p_{x \rightarrow y} \right). \quad (3)$$

The inner sum is the probability of going to a micro-state y belonging to macro-state c and being a neighbor of x , given that the system is in state x . The outer sum represents the equilibrium weighting of the micro-states x inside the given macro-state b . A straight-forward method determines the exact transition probabilities by performing the sums in Eq. (3), *i.e.* exhaustive enumeration of all micro-states and all neighbors [6, 17].

Throughout this contribution, we consider the usual partitioning of X with respect to gradient basins but the method is not restricted to this choice. Two micro-states $x, y \in X$ lie in the same macro-state $F(x) = F(y)$ if and only if the steepest descent walks starting in x and y terminate in the same local minimum. A state $u \in X$ is called local minimum, if $E(v) > E(u)$ for all $v \in M(u)$. For a given landscape and partitioning, the macro-state transition probabilities can be estimated by the sampling algorithm presented in the next section.

IV. SAMPLING METHOD

The method we introduce computes an estimate of the transition probabilities q in Eq. (3) by a standard importance sampling restricted to a macro-state b using the micro-state probabilities $P_b(x)$ defined in Eq. (2). Being in state $x_t \in F^{-1}(b)$ at time t , a neighbor $z \in M(x_t)$ is drawn at random with equal probabilities. The suggestion is accepted as the next state, $x_{t+1} = z$, with probability $\min\{1, P_b(z)/P_b(x_t)\}$. Otherwise the state remains the same, $x_{t+1} = x_t$. This choice guarantees that the relative frequency of state x tends towards the relative frequency $P_b(x)$ for increasing chain length $t \rightarrow \infty$ [30]. For a realization of a Markov chain of length t_{\max} , transition probabilities are estimated as

$$q'_{b \rightarrow c} = \frac{1}{t_{\max}} \sum_{t=1}^{t_{\max}} \sum_{y \in M(x_t) \cap F^{-1}(c)} p_{x_t \rightarrow y}. \quad (4)$$

In practice, the inner summation is performed only once at each time t , because each neighbor y of x_t contributes to the transition probability to exactly one macro-state $F(y)$.

Computation time is saved by storing visited micro-states of basin b and their sets of neighbors with transition probabilities in a data structure with fast search access, *e.g.* in a hash table. This is particularly advantageous in cases with broadly distributed micro-state probabilities such as Boltzmann weights at low temperature.

Here the Markov chain will encounter the highly probable (low energy) micro-states many times but neighbor sets and transition probabilities are computed only once per state. In the usual cases where macro-states are defined as basins of local minima, memory of visited states also saves time in evaluating the macro-state assignment function F : When the gradient walk starting at state x reaches a micro-state known to be in basin b , x itself is known to belong to b . Thus in many cases the walk may be terminated before reaching the ground state. Keeping previously visited micro-states in memory, however, is not necessary for the method to work. It may be handled according to the available resources. One may simply stop storing micro-states when the designated memory has been filled.

So far we have described how to estimate probabilities of transitions from *one* macro-state b to others. The result is the b -th column vector $(q'_{b \rightarrow c})_{c \in B}$ of the estimated transition matrix q' as given in Eq. 4. By applying the procedure separately to each macro-state, the full matrix q' is obtained. This can be implemented as an iterative exploration of the energy landscape without initial knowledge of the set of macro-states. Whenever a neighbor y of a state x in the Markov chain belongs to a macro-state $F(y)$ not previously seen, we add the pair $(F(y), y)$ to a queue Q of macro-states yet to work on. Initially, Q may contain only one particular pair (b, x_0) , *e.g.* the completely unfolded state x_0 of a polymer and the corresponding macro-state $b = F(x_0)$. The iterative exploration of the landscape is implemented in the following loop. (i) Extract a pair (b, x_0) from Q ; (ii) generate Markov chain inside b , starting at x_0 ; (iii) obtain estimates according to Eq. (4) and add newly discovered macro-states to Q ; (iv) If Q is not empty, resume at (i). Note, this method is directly parallelizable and will easily profit from distributed computing. Several independent realizations of Markov chains with respect to different macro-states can be run simultaneously, extracting from and feeding to the same queue. An implementation of the method is part of the Energy Landscape Library [31].

V. NUMBER PARTITIONING LANDSCAPE

The number partitioning problem (NPP) is a decision problem in the theory of computation and computational complexity [32–34]. It asks if a given set A of N real non-negative numbers can be partitioned into two subsets B, C such that numbers in B have the same sum as those in C . In an equivalent formulation, we label the numbers in A as a_1, \dots, a_N and use spin variables x_1, \dots, x_N to encode if a_i is in subset B ($x_i = +1$) or in subset C ($x_i = -1$). This system has the set of micro-states $X = \{-1, +1\}^N$. We define the energy of state $x \in X$ as

$$E(x) = \left| \sum_{i=1}^N x_i a_i \right|. \quad (5)$$

Then the NPP amounts to the question if the ground state energy of this system is zero.

The number partitioning *landscape* is obtained by using the hypercube as the neighborhood structure. For each $x \in X$ we have

$$M(x) = \{y \in X \mid d(x, y) = 1\} \quad (6)$$

as the set of neighbors. The usual Hamming distance d is used, so $d(x, y)$ is the number of entries i such that $x_i \neq y_i$. A local move on the landscape means flipping one of the N spin variables x_i .

Random instances are typically generated by drawing the a_i as statistically independent random variables uniformly distributed in the unit interval. Then the expected number of local minima grows exponentially with N , more precisely $\langle |B| \rangle \sim 2^N N^{-3/2}$ [35]. Here we use special instances of the NPP where

$$a_i = (i - 1)^{-\alpha} \quad (7)$$

with $\alpha = 0.55$. For these instances, we have found the number of local minima to grow exponentially with N for $N \leq 40$. However, the growth is much slower than for randomly generated instances. At $N = 40$, the instance of Equation (7) has 318 local minima, to be compared with an expected number of $\approx 10^{15}$ local minima for randomly generated instances. Each landscape with $10 \leq N \leq 40$ has at least one basin with an energy barrier ≥ 0.1 . We explore the landscape at temperature $1/\beta = 0.1$.

Figure 1 shows the convergence of the probability estimates. For each system size N , the sampling error decreases inversely proportional to the number of sampling steps performed per basin. Larger systems need more computational effort to reach a certain precision. The inset of Fig. 1 indicates that the total computational effort required for the error to fall below a given value grows sub-exponentially with N , to be compared with a number of micro-states increasing as 2^N . Thus under growing N , sampling a strongly decreasing fraction of micro-states is sufficient in order to reach a given precision.

VI. FOLDING LANDSCAPE OF AN RNA SWITCH

As a real-world example of folding landscapes of biopolymers we consider RNA molecules [37]. The primary structure of an RNA molecule is a finite sequence (a string) over the alphabet of the four nuclear bases $\{\mathbf{A}, \mathbf{C}, \mathbf{G}, \mathbf{U}\}$. An RNA secondary structure is a list of pairs (i, j) of positions in the primary structure such that the following conditions hold. (1) Base combinations at pairing positions must be $\mathbf{A-U}$ or $\mathbf{G-C}$ (Watson-Crick pairs) or $\mathbf{G-U}$ (wobble pair); (2) each position i can pair with at most one other position j ; (3) there are no two pairs (i, j) and (k, l) with $i < k < j < l$. The latter condition forbids so-called pseudoknots and makes the graph representation of a secondary structure outer-planar (see Fig. 2).




FIG. 1: (Color online) The deviation of estimated transition probabilities from the exact values is inversely proportional to the number of sampling steps (main panel). Shown are the analyzed special instances of number partitioning landscapes (see Eq. 7) for various sizes N . Error bars ($N = 15$ and $N = 40$) indicate the standard deviation between errors for different basins. The inset shows the N -dependence of the total number of sampling steps required for reaching a given precision, *i.e.* lowering the error below r . Given a macro-state a , we employ the Kullback-Leibler (KL) divergence $D(\cdot||\cdot)$ [36] to define the error as $\epsilon(s, a) := D(q'(s, a)||q'(2s, a))$, making a comparison of the estimate of the outgoing transition probability vector $q'(s, a) = (q'_{a \rightarrow b})_{b \in B}(s)$ after s sampling steps with its estimate after $2s$ sampling steps. The plotted values are the equally weighted average of the errors $\epsilon(s, a)$ over all macro-states $a \in B$.

In the folding landscape of an RNA sequence, the set of micro-states X contains the valid secondary structures. The energy $E(x)$ of a secondary structure $x \in X$ is a sum over binding energies of stacks (contiguous regions of binding) and entropic contributions from open (unbound) sections of the RNA chain. For details of energy calculations, we refer to the literature [39, 41, 42]. Micro-states $x, y \in X$ are adjacent, *i.e.* $y \in M(x)$ and $x \in M(y)$, if y can be generated from x by adding or removing a single base pair. Shift moves [37] are not considered in this contribution. When the lowest energy neighbor of a structure is not unique the degeneracy is resolved by the lexicographic ordering on string representations of the structures [6, 7, 9, 37].

Multistable RNAs, so called RNA-switches, are essential for the regulation of cellular processes. Thus, an understanding of the folding kinetics of such molecules is of high importance. For a detailed overview see [40]. Specifically, we work with the bistable RNA d33 sequence




FIG. 2: (Color online) RNA secondary structures (left/right) with energies $-14.4 \frac{\text{kcal}}{\text{mol}}$ and $-14.3 \frac{\text{kcal}}{\text{mol}}$ of the tested bistable RNA d33 and their outer-planar Feynman diagrams (middle) (drawn using jViz.Rna v1.77 [38]). Energy evaluation and sequence design is based on Vienna RNA package v1.8.2 [39] and the method from [40].




FIG. 3: (Color online) Sampling precision in terms of the predicted average time $\tau(b)$ to reach the ground state from a macro-state b for RNA d33. For each b , the corresponding data point gives the ratio between τ_{sampling} based on the sampled transitions q' and the value τ_{exact} from the exact ones q versus τ_{exact} itself. Symbols indicate number of sampling steps per macro-state as 10^3 (squares), 10^4 (diamonds), and 10^5 (crosses). The target set contains both ground states.

shown in Fig. 2. It has 29,759,371 micro-states, allowing for full enumeration and thus for a comparative analysis with our method. Out of the 3,223 local minima, the two lowest are the secondary structures given in Fig. 2. These two ground states have practically the same energy. A walk between the ground states involves breaking all base pairs, resulting in an energy barrier of height $\Delta E = 1.18 \times 10^{-19}$ J. The temperature for both sampling and energy calculation is $T = (273.15 + 37.00)$ K. Therefore $1/\beta = k_B T = 4.28 \times 10^{-21}$ J is more than one order of magnitude below the barrier height ΔE of the RNA switch.

A comparison between exact and sampled transition probabilities is made in terms of the average time $\tau(b)$ from macro-state b to one of the ground states. For a biopolymer as considered here, $\tau(b)$ is the folding time when starting in an initial state b such as the open chain.

Given a set of target states $A \subset B$, the time to target is




FIG. 4: (Color online) Time evolution of the sampling error (KL divergence Eq. 9) for the folding landscape of the RNA switch molecule d33. Mean, median, and maximum are for the distribution of KL values over the $|B| = 3223$ macro-states (local minima).

$\tau(a) = 0$ when starting in one of the target states $a \in A$ (boundary condition). For a starting state $b \in B \setminus A$, the average time $\tau(b)$ until first reaching one of the target states obeys the recursion

$$\tau(b) = 1 + \sum_{c \in B} q_{b \rightarrow c} \tau(c). \quad (8)$$

The average time to target from b is one time step plus the time to target from the state c following b . The distribution of c is given by the transition probability $q_{b \rightarrow c}$. Time to target is also called exit time [43].

Figure 3 shows that $\tau(b)$ -values based on the sampled transition probabilities have small relative error for all starting macro-states $b \in B$. With 10^4 sampling steps per basin, the ratios between exact and approximate times τ are in the range [0.75; 1.15]. They fall into [0.96; 1.07] when using 10^5 steps per basin.

To investigate the sampling error we compare, separately for each macro-state b , the exact with the estimated transition probability vectors for leaving b . We quantify the discrepancy between the two vectors by the Kullback-Leibler divergence [36]

$$D(q'_b || q_b) = \sum_{c \in B} q'_{b \rightarrow c} \ln \frac{q'_{b \rightarrow c}}{q_{b \rightarrow c}}. \quad (9)$$

Figure 4 summarizes the evolution of the sampling error for increasing sampling steps per basin. As in the number partitioning landscape, Kullback-Leibler divergence (KL) decreases inversely proportional to the number of sampling steps. The mean of the distribution of KL values across basins is larger than the median by a factor of 3 indicating a broad distribution. This is due to a broad distribution of macro-state sizes. Probability vectors for macro-states comprising one or a few micro-states reach a low error after fewer sampling steps than




FIG. 5: (Color online) Time evolution of macro-state concentrations for RNA d33 from the exact transition probabilities q (curves) and from the estimates q' via sampling for 10^4 (\diamond) and 10^5 steps per macro-state ($+$). The concentrations both for the exact and the sampled transition rates are at the stationary values for $t \geq 10^{15}$. See main text for discussion of the discrepancy in these stationary concentrations.

those for large macro-states. Still also the maximum of the error across all basins decreases proportionally to the average. One of the extensions of the method outlined in Sec. VII chooses a number of sampling steps individually for each macro-state based on its estimated partition function.

In Figure 5, we compare the kinetics of the molecule for the approximated transition probabilities via sampling and the exact ones obtained by enumeration of all micro-states of the landscape. As an initial condition we choose the whole ensemble to be in the macro-state of the open chain (structure without base pairs). As a qualitative description, the ensemble first populates the first and, somewhat later, the second ground state. On an intermediate time scale (10^6), an almost constant concentration vector is reached with the second ground state dominating the first. However, this plateau concentration vector is transient. Probability mass flows from the second to the first ground state on a slow time scale (10^{15}) to reach the stationary concentrations.

With transition probabilities obtained by sampling for 10^4 steps per macro-state, the kinetics is reproduced with high precision both in the timing as well as the absolute concentration in the plateau where the relative error is below 10^{-2} . Since we hash the probabilities for already

visited structures, the computational effort per basin is dominated by the number of visited states instead of overall sampling steps. Thus, small basins are sampled faster than larger ones. When sampling 10^5 steps per macro-state, which renders the kinetics with even higher accuracy, computation time is still reduced by a factor of ≈ 9 compared with full enumeration.

The stationary concentrations found at time $t \geq 10^{15}$, however, do not agree with the exact solution. A much larger number of sampling steps is required to match these. Further tests with other RNA switch molecules yield the same qualitative result for moderate number of sampling steps per basin. Both the concentration levels and the time scales are faithfully reproduced by the transition rates from sampling, except for the stationary concentrations.

A closer look at the particular structure of the landscape of RNA switches hints at an explanation for the discrepancy. Both ground states have large and deep basins. In RNA d33, barriers to neighboring basins are all at least 8.6×10^{-20} J above ground state energy, which is more than $20 k_B T$. Exits from one ground state basin towards the other lead through a small number of micro-states. When sampling the large ground state basins, these few salient micro-states are likely to be missed. The subtle balance between incoming and outgoing probability flow in the equilibrium is distorted. Due to the symmetry of the move set, however, those missed non-zero transition probabilities can be identified to some extent. When the forward rate is non-zero then the backward rate must be non-zero as well. For dynamics with detailed balance, as considered here, even quantitative correction of missed or undersampled rates is possible. This is suggested as one of the extensions in the following section.

VII. EXTENSIONS AND MODIFICATIONS OF THE METHOD

We outline ideas for varying the method to potentially increase efficiency and applicability in various settings. These are not used in the applications in Sec. V and VI.

A. Guided sampling

The stopping criterion (iv) of the outer loop for full transition matrix estimation (Sec. IV) may be modified if we do not aim to explore the whole landscape but only a subset of the set of macro-states [18]. Then Q may be handled as a priority queue. For instance, we may be interested only in transitions between macro-states below a certain energy threshold or those involved in typical trajectories. In the latter case, the next macro-state to be explored is the one that is reached from already explored macro-states with the largest probability.

B. Partition function estimation

In addition to transition probabilities, the canonical partition function

$$Z_b = \sum_{x \in F^{-1}(b)} \exp(-\beta E(x)) \quad (10)$$

of the macro-state b may be estimated at any time during the sampling. Consider a subset $X \subseteq F^{-1}(b)$ of the micro-states of macro-state b . When sampling in the macro-state b , the fraction of time the Markov chain spends in X is

$$r = \frac{1}{Z_b} \sum_{x \in X} \exp(-\beta E(x)). \quad (11)$$

Therefore Z_b can be calculated when knowing r and the energies of all states in X . Now X can be taken as the set of states visited in the first t^* steps of the Markov chain, $X = \{x_t \mid t = 1, \dots, t^*\}$. An estimate r' of r is obtained by counting how often the Markov chain visits states in X during a sufficiently long time interval $[t_{\text{start}}, t_{\text{stop}}[$,

$$r' = \frac{|\{t \mid t_{\text{start}} \leq t < t_{\text{stop}} \wedge x_t \in X\}|}{t_{\text{stop}} - t_{\text{start}}}. \quad (12)$$

In order to obtain an unbiased estimate of r , this time interval must not overlap with the time steps during which X is recorded, thus $t_{\text{start}} > t^*$. By solving Eq. (11) for Z_b and replacing r with the estimate r' , we obtain

$$Z'_b = \frac{1}{r'} \sum_{x \in X} \exp(-\beta E(x)) \quad (13)$$

as an unbiased estimate of the partition function Z_b .

C. Sampling time adjustment

The estimate $r(t)$ may also be used for adapting the length of the Markov chain exploring macro-state b to the size of b . The sampling will be run until the fraction of covered probability mass exceeds a certain threshold, *e.g.* stopping as soon as $r(t) > 0.5$.

D. Landscape coarse-graining

The state space can be coarse-grained beyond the initially chosen macro-state state partitioning by dynamically merging macro-states [44]. For merging macro-state b into macro-state a , the affected entries in the matrix q are replaced by

$$\hat{q}_{a \rightarrow c} = \frac{Z_a q_{a \rightarrow c} + Z_b q_{b \rightarrow c}}{Z_a + Z_b} \quad (14)$$

$$\hat{q}_{c \rightarrow a} = q_{c \rightarrow a} + q_{c \rightarrow b} \quad (15)$$

for all macro-states $c \notin \{a, b\}$. The new diagonal element $\hat{q}_{a \rightarrow a}$ is obtained by normalization of probability. The row and column of macro-state b are set zero (or deleted). In a separate index, the mapping of macro-state b to macro-state a is stored.

Strategies for the choice of macro-states to be merged need to be explored yet. A reasonable starting point is to choose pairs of macro-states with high overlap in successor states, *e.g.* choosing a and b such that

$$\sum_{c \in B} q_{a \rightarrow c} q_{b \rightarrow c} \quad (16)$$

is maximal.

E. Balancing the transition probability matrix

If the micro-state dynamics in terms of the transition probabilities $p_{x \rightarrow y}$ fulfills detailed balance, then so does the macro-state dynamics with transition probabilities $q_{b \rightarrow c}$. Detailed balance means

$$Z_l q_{l \rightarrow s} = Z_s q_{s \rightarrow l} . \quad (17)$$

for all macro-state pairs (l, s) . The transition probabilities $q'_{l \rightarrow s}$ obtained by the sampling, however, need not fulfill the same condition. By the transformation

$$q_{l \rightarrow s}^* = \frac{1}{2} q'_{l \rightarrow s} + \frac{1}{2} q'_{s \rightarrow l} \frac{Z_s}{Z_l} \quad (18)$$

transition probabilities q^* with detailed balance are obtained. The transformation also serves to impose a known stationary distribution of concentrations on the transition probability matrix.

VIII. CONCLUSION AND DISCUSSION

When coarse-graining the state space of an energy landscape into macro-states, transition probabilities be-

tween macro-states have to be obtained in order to capture the coarse-grained stochastic dynamics. Here we have introduced a sampling method that allows for a fast yet accurate estimation of these transition probabilities. We have demonstrated the scalability of the approach with system size for special instances of the number partitioning problem. As a real-world application, we have analyzed the folding landscape of the secondary structure of an RNA switch as an example of a biopolymer. Its rich dynamic behavior on separate fast and slow timescales is accurately rendered by transition probabilities obtained with low computational cost.

The general method introduced here may serve as a flexible framework for stochastic exploration of energy landscapes. As laid out in the Sec. VII, several extensions and modifications may be made to obtain increased performance and wider applicability. In particular, the high variation of macro-state sizes may be exploited in a scheme for an automatic choice of sampling effort. Furthermore, the merging of small macro-states with larger neighbors during the sampling may lead to more manageable and potentially more meaningful partitions of the landscape akin to metabasins [8].

In ongoing and future work, the method shall be applied to other energy landscapes including those of state-discrete protein folding dynamics [45–47]. Such landscapes have been shown to be amenable to sampling approaches [48]. Another field of application of our method is the clarification of concepts for dynamics on energy surfaces, such as the notion of a folding funnel [11, 49, 50].

Acknowledgments

KK gratefully acknowledges funding from Volkswagen-Stiftung.

-
- [1] C. M. Reidys and P. F. Stadler, *SIAM Rev.* **44**, 3 (2002), ISSN 0036-1445.
 - [2] D. Wales, *Energy Landscapes* (Cambridge University Press, 2003).
 - [3] Q. Zhou and W. H. Wong, *Phys. Rev. E* **79**, 051117 (2009).
 - [4] C. B. Anfinsen, *Science* **181**, 223 (1973), ISSN 0036-8075.
 - [5] D. Ming, M. Anghel, and M. E. Wall, *Phys. Rev. E* **77**, 021902 (2008).
 - [6] C. Flamm, I. L. Hofacker, P. F. Stadler, and M. T. Wolfinger, *Z.Phys.Chem* **216**, 155 (2002).
 - [7] M. T. Wolfinger, W. A. Svrcek-Seiler, C. Flamm, I. L. Hofacker, and P. F. Stadler, *J. Phys. A: Math. Gen.* **37**, 4731 (2004).
 - [8] A. Heuer, *Journal of Physics: Condensed Matter* **20**, 373101 (56pp) (2008).
 - [9] M. T. Wolfinger, S. Will, I. L. Hofacker, R. Backofen, and P. F. Stadler, *Europhys. Lett.* **74**, 726 (2006).
 - [10] M. Baiesi, L. Bongini, L. Casetti, and L. Tattini, *Phys. Rev. E* **80**, 011905 (2009).
 - [11] P. Garstecki, T. X. Hoang, and M. Cieplak, *Phys. Rev. E* **60**, 3219 (1999).
 - [12] W. Fontana, P. F. Stadler, E. G. Bornberg-Bauer, T. Griesmacher, I. L. Hofacker, M. Tacker, P. Tarazona, E. D. Weinberger, and P. Schuster, *Phys. Rev. E* **47**, 2083 (1993).

- [13] J. Manuch, C. Thachuk, L. Stacho, and A. Condon, in *Proc. of DNA Computing and Molecular Programming* (Springer, 2009), vol. 5877 of *LNCS*, pp. 106–115.
- [14] C. Thachuk, J. Manuch, L.-A. M. Arash Rafiey, L. Stacho, and A. Condon, in *Proc. of Pacific Symposium on Biocomputing* (2010), vol. 15, pp. 108–119.
- [15] H. S. Chan and K. A. Dill, *The Journal of Chemical Physics* **100**, 9238 (1994).
- [16] S. Wuchty, W. Fontana, I. L. Hofacker, and P. Schuster, *Biopolymers* **49**, 145 (1999).
- [17] P. Sibani, R. van der Pas, and J. C. Schüi, *Computer Physics Communications* **116**, 17 (1999).
- [18] Z. Burda, A. Krzywicki, O. C. Martin, and Z. Tabor, *Phys. Rev. E* **73**, 036110 (2006).
- [19] P. Sibani, J. C. Schön, P. Salamon, and J.-O. Andersson, *EPL (Europhysics Letters)* **22**, 479 (1993).
- [20] P. Sibani and P. Schriver, *Phys. Rev. B* **49**, 6667 (1994).
- [21] J. C. Schön and P. Sibani, *J. Physics A: Mathematical and General* **31**, 8165 (1998).
- [22] A. Xayaphoummine, T. Bucher, F. Thalmann, and H. Isambert, *Proc. Natl. Acad. Sci. USA* **100**, 15310 (2003).
- [23] X. Tang, B. Kirkpatrick, S. Thomas, G. Song, and N. M. Amato, *J. Comp. Biol.* **12**, 862 (2005).
- [24] S. Thomas, G. Song, and N. M. Amato, *Physical Biology* **2**, S148 (2005).
- [25] S. Thomas, X. Tang, L. Tapia, and N. M. Amato, *J. Comp. Biol.* **14**, 839 (2007).
- [26] M. Geis, C. Flamm, M. T. Wolfinger, A. Tanzer, I. L. Hofacker, M. Middendorf, C. Mandl, P. F. Stadler, and C. Thurner, *J. Mol. Biol.* **379**, 160 (2008).
- [27] D. Gfeller, P. De Los Rios, A. Cafilisch, and F. Rao, *Proceedings of the National Academy of Sciences* **104**, 1817 (2007).
- [28] D. Prada-Gracia, J. Gómez-Gardeñes, P. Echenique, and F. Falo, *PLoS Comput Biol* **5**, e1000415 (2009).
- [29] H. A. Kramers, *Physica* **7**, 284 (1940).
- [30] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, and A. H. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [31] M. Mann, S. Will, and R. Backofen, in *Proc. of BIRD'07 (OCG, 2007)*, vol. 217, pp. 83–86.
- [32] M. R. Garey and D. S. Johnson, *Computers and intractability* (Freeman, 1979).
- [33] S. Mertens, *Theor. Comp. Sci.* **265**, 79 (2001).
- [34] P. F. Stadler, W. Hordijk, and J. F. Fontanari, *Phys. Rev. E* **67**, 056701 (2003).
- [35] F. F. Ferreira and J. F. Fontanari, *Journal of Physics A: Mathematical and General* **31**, 3417 (1998).
- [36] S. Kullback, *The American Statistician* **4**, 340 (1987).
- [37] C. Flamm, W. Fontana, I. Hofacker, and P. Schuster, *RNA* **6**, 325 (2000).
- [38] K. C. Wiese and E. Glen, *IEEE Symposium on Computer-Based Medical Systems* **0**, 659 (2006), ISSN 1063-7125.
- [39] I. L. Hofacker, W. Fontana, P. F. Stadler, L. S. Bonhoeffer, M. Tacker, and P. Schuster, *Chemical Monthly* **125**, 167 (1994).
- [40] C. Flamm, I. L. Hofacker, S. Maurer-Stroh, P. F. Stadler, and M. Zehl, *RNA* **7**, 254 (2000).
- [41] I. Tinoco, O. C. Uhlenbeck, and M. D. Levine, *Nature* **230**, 362 (1971).
- [42] S. M. Freier, R. Kierzek, J. A. Jaeger, N. Sugimoto, M. H. Caruthers, T. Neilson, and D. H. Turner, *Proceedings of the National Academy of Sciences of the United States of America* **83**, 9373 (1986).
- [43] D. R. S. Geoffrey R. Grimmett, *Probability and Random Processes* (Oxford Science Publications, 1982).
- [44] L. Bongini, L. Casetti, R. Livi, A. Politi, and A. Torcini, *Phys. Rev. E* **79**, 061925 (2009).
- [45] K. A. Dill, *Biochemistry* **24**, 1501 (1985).
- [46] M. Mann, S. Will, and R. Backofen, *BMC Bioinformatics* **9**, 230 (2008).
- [47] M. Mann, D. Maticzka, R. Saunders, and R. Backofen, *HFSP Journal* **2**, 396 (2008).
- [48] T. Wüst and D. Landau, *Computer Physics Communications* **179**, 124 (2008), ISSN 0010-4655, special issue based on the Conference on Computational Physics 2007 - CCP 2007.
- [49] P. E. Leopold, M. Montal, and J. N. Onuchic, *Proc. Natl. Acad. Sci. USA* **89**, 8721 (1992).
- [50] K. Klemm, C. Flamm, and P. F. Stadler, *The European Physical Journal B* **63**, 387 (2008).