Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

HAN CHENG LIE AND KONSTANTIN FACKELDEY AND
MACRUS WEBER

A square root approximation of transition rates for a Markov State Model

ZIB-Report 13-43 (August 2013)

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: http://www.zib.de

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

A square root approximation of transition rates for a Markov State Model

Han Cheng Lie and Konstantin Fackeldey*and Marcus Weber†

August 28, 2013

Abstract

Trajectory- or mesh-based methods for analyzing the dynamical behavior of large molecules tend to be impractical due to the curse of dimensionality - their computational cost increases exponentially with the size of the molecule. We propose a method to break the curse by a novel square root approximation of transition rates, Monte Carlo quadrature and a discretization approach based on solving linear programs. With randomly sampled points on the molecular energy landscape and randomly generated discretizations of the molecular configuration space as our initial data, we construct a matrix describing the transition rates between adjacent discretization regions. This transition rate matrix yields a Markov State Model of the molecular dynamics. We use Perron cluster analysis and coarse-graining techniques in order to identify metastable sets in configuration space and approximate the transition rates between the metastable sets. Application of our method to a simple energy landscape on a two-dimensional configuration space provides proof of concept and an example for which we compare the performance of different discretizations. We show that the computational cost of our method grows only polynomially with the size of the molecule. However, finding discretizations of higher-dimensional configuration spaces in which metastable sets can be identified remains a challenge.

*fackeldey@zib.de
†weber@zib.de

Abstract

Trajectory- or mesh-based methods for analyzing the dynamical behavior of large molecules tend to be impractical due to the curse of dimensionality - their computational cost increases exponentially with the size of the molecule. We propose a method to break the curse by a novel square root approximation of transition rates, Monte Carlo quadrature and a discretization approach based on solving linear programs. With randomly sampled points on the molecular energy landscape and randomly generated discretizations of the molecular configuration space as our initial data, we construct a matrix describing the transition rates between adjacent discretization regions. This transition rate matrix yields a Markov State Model of the molecular dynamics. We use Perron cluster analysis and coarse-graining techniques in order to identify metastable sets in configuration space and approximate the transition rates between the metastable sets. Application of our method to a simple energy landscape on a two-dimensional configuration space provides proof of concept and an example for which we compare the performance of different discretizations. We show that the computational cost of our method grows only polynomially with the size of the molecule. However, finding discretizations of higher-dimensional configuration spaces in which metastable sets can be identified remains a challenge.

1 Introduction

In statistical mechanics, a molecule is described by a suitably chosen ensemble of systems having the same equations of motion but different initial conditions. The ensemble is specified by a probability density function which gives the probability that a system lies in a certain subset of the phase space. If the dynamical behaviour of a system exhibits two or more time scales, then we can consider the metastabilities of each system in the following way: in the fine scale (smaller time scale) dynamics, the system fluctuates but relaxes for certain given initial conditions to subsets which are almost stable; in the coarse scale (larger time scale) dynamics, the system appears to be a Markovian process jumping between the almost stable subsets, where the transition probability depends only on the current state. A subset is 'almost stable' if the expected exit time of a system from the subset is large, relative to the fluctuations occurring at the short time scale. One sometimes uses the terms 'metastabilities' or 'conformations' to refer to the almost stable subsets of the phase space. In Figure 1 we see a trajectory of a system exhibiting the kind of multiscale dynamics we have described. The identification of the metastabilities of a system and the study of its coarse scale dynamics are among the primary goals of coarse graining and multiscale methods [3, 13, 16].

Methods based on Markov State Models (MSMs) [15, 24, 26] seek to identify the metastabilities and study the coarse scale dynamics by constructing a Markov chain on a finite state space. These methods work by studying certain eigenfunctions of a semigroup of transfer operators or the associated infinitesimal generator [6, 12, 38]. The Markov chain is obtained by constructing a matrix which contains the transition probabilities between the metastabilities [8, 10, 14, 28]. In order to assemble such a matrix, the operators (or the generator) must be suitably discretized, or the transition probabilities must

somehow be approximated, e.g. via molecular dynamics simulations. For systems of practical interest, the high dimension of the phase space leads to difficulties for either approach: if one wishes to discretize an operator, the number of basis functions required (e.g. for an application of the finite element method) grows exponentially with the dimension; if one wishes to simulate the molecular dynamics, one must apply a grid or mesh to the phase space, and the computational cost of applying such a mesh grows exponentially with the dimension of the phase space. The exponential increase in computational cost is also known as the *curse of dimensionality* [2] and features prominently in computational methods for studying molecular dynamical systems.

In this paper we propose a method to break the curse of dimensionality. We adopt the operator approach, but construct a transition rate matrix instead of a transition probability matrix. Starting from a theorem in [38], we compute the rates in terms of quantities of the discretized energy landscape of the system - the Boltzmann weights of discretization sets, the Boltzmann weights of the surfaces of intersection between adjacent sets, and the instantaneous flux across intersections. We present and explain this theorem in Section 2.1.

In Section 2.2 we assume that the flux between discretization sets is constant over all pairs of discretization sets. Using a linear interpolation or averaging argument, we then express the weights of the surfaces of intersection as square roots of the products of the weights of adjacent discretization sets. This yields a square root approximation of transition rates. We show the sequence of approximations leading from the continuous to the discretized dynamics below and italicize the novel steps in the sequence.

$\{ \mathcal{P}(\tau) \}_{\tau>0}$	semigroup of generalized transfer operators [38]
$\downarrow$	
Q	transition rate matrix from [38]
$\downarrow$	*(Assumption of constant flux)*
Q'	
$\downarrow$	*(Square root approximation)*
Q''	

The novel steps in the approximation - the assumption of constant flux and the square root approximation - are important in breaking the curse of dimensionality. By the assumption of constant flux, we can avoid simulating the dynamics. By the square root approximation, we only need to compute the Boltzmann weights of discretization sets, and we approximate these weights by Monte Carlo quadrature.

It remains to choose a discretization method. We need to assign randomly sampled points on the energy landscape to discretization regions for Monte Carlo quadrature, and we need to determine when two regions are adjacent in order to apply the square root approximation. In this article we will use Voronoi tesselations because they are not mesh-based discretizations. In Section 2.3 we present a formulation of Voronoi tesselations in terms of convex polyhedra. This formulation provides us a discretization method that does not incur the curse of dimensionality.

The method proceeds as follows: discretize the state space by Voronoi tesselations and compute the adjacency relations; randomly sample points on the

energy landscape and compute the Boltzmann weights of the Voronoi regions by Monte Carlo quadrature; use the adjacency relations and the square root approximation to compute the transition rate matrix Q''; apply Perron Cluster Analysis and coarse-graining techniques to Q'' in order to identify the metastabilities and coarse-scale dynamics. We describe the method in greater detail in Section 3.

In Section 4 we present the results of applying our method to a low-dimensional model. In Section 5 we critique the method, analyzing the computational cost in Section 5.1 and the performance of our method on higher-dimensional models in Section 5.2. We present our conclusions in Section 6.

2 Theory

In this section we present the key theoretical ideas behind our method. We discuss the sequence of approximations which leads from transfer operators to the transition rate matrix Q'' of our method. We present a formulation of Voronoi tesselations in terms of convex polyhedra and end with a brief overview of the Perron cluster analysis and coarse-graining methods.

2.1 Transfer operators and transition matrices

Modelling molecular dynamical systems according to Newton's laws of motion gives rise to differential equations. These differential equations in turn lead to operators which describe how systems evolve in some state space. In order to study systems using computational methods, one needs to first choose operator which describes the continuous dynamics. Discretizing this operator yields a matrix that describes the discretized dynamics.

Consider the transfer operator $\mathcal{P}(\tau)$ which describes how the configuration (the vector of spatial coordinates of the atoms of a molecule) evolves over a time lag $\tau > 0$ [34, 38]. For computation purposes, one projects the operator, which acts on a high-dimensional continuous space, to a lower-dimensional space Ω, which one then discretizes using closed discretization sets $\{\Omega_i\}_{i=1,\ldots,n}$ which overlap only at their boundaries. We use Voronoi tesselations as our discretization method because they are a meshless discretization method and because we can compute adjacencies between the resulting discretization sets without incurring the curse of dimensionality, as we will show in Section 2.3. The projection of $\mathcal{P}(\tau)$ to the discretization yields a transition probability matrix $(P(\tau)_{ij})_{i,j\in I}$ associated with the discretization, where $P(\tau)_{ij}$ gives the probability of a transition from Ω_i to Ω_j after τ units of time.

In practice, one computes $P(\tau)_{ij}$ by computing many molecular dynamics simulation (MD) trajectories and counting the proportion of trajectory segments starting in Ω_i that are in Ω_j after τ units of time. As computing trajectories suffers from the curse of dimensionality, we need a different approach to studying the transition behaviour over the discretization sets - one that does not need trajectory segments. The key is to avoid estimating transition probabilities over τ and to consider instantaneous transition rates instead. We will need Theorem 4 from [38], which we abbreviate as:

Theorem 1 *Given a Voronoi tesselation $\Omega = \cup_{i=1}^{n} \Omega_i$ and the associated matrix*

$P(\tau) \in \mathbb{R}^{n\times n}$ of a transfer operator $\mathcal{P}(\tau)$, the matrix $Q := \frac{\partial}{\partial\tau}P(\tau)\big|_{\tau=0}$ satisfies

$$Q_{ij} = \int_{\Omega_i\cap\Omega_j} z(q)\pi_i(q)dS(q)$$

for $i \neq j$, where $dS(q)$ is the surface measure on $\Omega_i \cap \Omega_j$, $\pi_i : \Omega \to [0,\infty)$ is the normalized restriction to Ω_i of the Boltzmann density over Ω, and $z(q)$ denotes the flux of configurations from Ω_i to Ω_j through the point $q \in \Omega_i \cap \Omega_j$.

The quantity Q_{ij} describes the instaneous rate of flux of molecular configurations moving from Ω_i to Ω_j across the surface $\Omega_i \cap \Omega_j$, given that states in Ω_i are distributed according to $\pi_i(\cdot)$. To see how, we follow [38], and rewrite Q_{ij} as $Q_{ij} = s_{ij}\langle z\rangle_{ij}/w_i$, where $w_i = \int_\Omega 1_{\Omega_i}(\tilde{q})\pi_q(\tilde{q})d\tilde{q}$ is the Boltzmann weight of Ω_i, $s_{ij} = \int_{\Omega_i\cap\Omega_j}\pi_q(q)dS(q)$ is the Boltzmann weight of the surface $\Omega_i\cap\Omega_j$, and

$$\langle z\rangle_{ij} = \int_{\Omega_i\cap\Omega_j} z(q)\frac{\pi_q(q)}{\int_{\Omega_i\cap\Omega_j}\pi_q(\hat{q})dS(\hat{q})}\,dS(q)$$

is the flux per unit area from Ω_i to Ω_j across $\Omega_i\cap\Omega_j$. Thus, $s_{ij}\ \langle z\rangle_{ij}$ quantifies the flux of configurations from Ω_i to Ω_j across $\Omega_i\cap\Omega_j$, and we establish that Q_{ij} describes the conditional transition rate of configurations from Ω_i to Ω_j.

2.2 Square root approximation

Now we present the novel idea of our method: the construction of the square root approximation transition rate matrix. We assume that the flux per unit area term $\langle z\rangle_{ij}$ from Theorem 1 is equal to a constant $\hat{z}$ that is independent of i and j. This assumption is equivalent to assuming that the flux per unit area across any surface $\Omega_i\cap\Omega_j$ is the same for all pairs $i \neq j$ for any given discretization of Ω, and yields a matrix Q' with off-diagonal entries given by $Q'_{ij} = \hat{z}\ s_{ij}/w_i$.

Let the Boltzmann density on configuration space $\pi_q : \Omega \to (0,\infty)$ be given by $\pi_q(\cdot) = Z^{-1}\exp\left[-\beta V(\cdot)\right]$, where Z is the partition function or the normalizing constant, $\beta = (k_B T)^{-1}$ is the inverse of the product of temperature T with Boltzmann's constant k_B and $V : \Omega \to \mathbb{R}$ is the potential energy function of the molecule. Recall that a Voronoi tesselation of $\mathbb{R}^d$ generated by n distinct points $\{g_i\}_{i=1}^n$ is a partition of $\mathbb{R}^d$ into n closed Voronoi regions $\{\Omega_i\}_{i=1}^n$, where the Voronoi region Ω_i generated by g_i is the set of points which are closer to g_i than to other g_j in the Euclidean distance metric:

$$\Omega_i := \left\{x \in \mathbb{R}^d\ :\ \|x-g_i\| \leq \|x-g_j\|,\ i \neq j\right\}. \tag{1}$$

In particular, $\Omega_i\cap\Omega_j = \partial\Omega_i\cap\partial\Omega_j$ for $i \neq j$. Now consider a simple one-dimensional configuration space and a Voronoi tesselation of that space, as in Fig. 2. The intersection $q = \Omega_i\cap\Omega_j$ of adjacent Voronoi regions Ω_i and Ω_j is the midpoint between their respective generators g_i and g_j. If we approximate the potential energy function between g_i and g_j by a linear interpolant between g_i and g_j, we have that $V(q) \approx 0.5(V(g_i)+V(g_j))$ and obtain

$$s_{ij} = \frac{1}{Z}\exp\left[-\beta V(q)\right] \approx \frac{1}{Z}\sqrt{\exp(-\beta V(g_i))\exp(-\beta V(g_j))} \tag{2}$$

What if we have the energy values of scattered points in Ω_i and Ω_j, but do not wish to create a Voronoi region for each point? One possible solution would be to let

$$s_{ij} \approx \sqrt{\bar{w}_i \bar{w}_j} N_{ij} \tag{3}$$

where $N_{ij} = 1$ if Ω_i and Ω_j are adjacent (i.e. if $\dim(\Omega_i \cap \Omega_j) = \dim(\Omega_i) - 1$) and $N_{ij} = 0$ otherwise, and $\bar{w}_i$ denotes the arithmetic mean of Boltzmann density values of points in each Voronoi region,

$$\bar{w}_i := M_i^{-1} \left(\sum_{k=1}^{M} 1_{\Omega_i}(q_k) \pi_q(q_k) \right) := \left(\sum_{k=1}^{M} 1_{\Omega_i}(q_k) \right)^{-1} \left(\sum_{k=1}^{M} 1_{\Omega_i}(q_k) \pi_q(q_k) \right). \tag{4}$$

In Eq. (4), the $\{q_k\}_{k=1}^{M}$ are the scattered data points in the configuration space Ω, M is their total number, M_i denotes the number of data points in Ω_i, and $1_{\Omega_i}(x) = 1$ if $x \in \Omega_i$ and 0 otherwise. We will refer to the symmetric matrix $(N_{ij})_{i,j=1}^{n}$ as the 'matrix of adjacency relations'.

The rationale for the approximation in Eq. (3) is as follows. Let Ω_i and Ω_j be adjacent and recall the definition of s_{ij} given in Section 2.1. In the limit of arbitrarily small Voronoi regions, the Boltzmann density values of points on $\partial\Omega_i \cap \partial\Omega_j$ will be arbitrarily close to the Boltzmann density values of the generators g_i and g_j, by continuity of the energy function V; the Boltzmann density values of g_i and g_j will also be arbitrarily close to $\bar{w}_i$ and $\bar{w}_j$ for the same reason.

Now we apply the square root approximation. Replace s_{ij} and w_i in the expression $Q'_{ij} = \hat{z}\ s_{ij}/w_i$ from Section 2.1 with $\sqrt{\bar{w}_i \bar{w}_j} N_{ij}$ and $\bar{w}_i$ respectively. The latter replacement holds since $\bar{w}_i$ converges in probability to the true Boltzmann weight w_i as $M_i \to \infty$ by the Law of Large Numbers. We obtain

$$Q''_{ij} = \hat{z} \sqrt{\bar{w}_i \bar{w}_j} N_{ij} / \bar{w}_i = \hat{z} \sqrt{\bar{w}_j / \bar{w}_i} N_{ij}, \tag{5}$$

which is our *square root approximation* for the transition rate of configurations from Ω_i to Ω_j for $i \neq j$. An immediate advantage of (5) is that by taking the ratio $\bar{w}_j/\bar{w}_i$ we no longer need to compute the partition function Z in Eq. (4). It also follows from (5) that detailed balance between the transition rate matrix Q'' and the measure $\bar{w} = (\bar{w}_1, \ldots, \bar{w}_n)$ holds:

$$\bar{w}_i Q''_{ij} = \bar{w}_j Q''_{ji}. \tag{6}$$

The equation above implies time-reversibility of the associated Markov chain [27].

Interesting consequences follow from the interpretation of the off-diagonal entries of Q'' as transition rates. Since the total concentration of states leaving Ω_i for adjacent regions must be deducted from the concentration of states in Ω_i, we must have

$$Q''_{ii} = -\sum_{j \neq i}^{n} Q''_{ij}. \tag{7}$$

Furthermore, Q'' must also describe the evolution of distributions $x(t) \in [0,1]^n$ over the $\{\Omega_i\}_{i=1}^{n}$ by $\frac{d}{dt} x^\top(t) = x^\top Q''$. Since the Boltzmann distribution is an equilibrium distribution of the continuous dynamics, we expect that its discretization w is an equilibrium distribution of the discretized dynamics: $0 = w^\top Q$.

Lemma 1 *Let $v \in \mathbb{R}^n$ be a strictly positive vector and let $R = R(v) \in \mathbb{R}^{n\times n}$ be the matrix whose entries are given by*

$$R_{ij} = \begin{cases} \sqrt{\frac{v_j}{v_i}} & i \neq j \\ -\sum_{k\neq i} R_{ik} & i = j \end{cases}.$$

Then the vector v satisfies $0 = v^\top R$.

Proof: Fix $j \in \{1, \ldots, n\}$.

$$(v^\top R)_j = \sum_{i\neq j} v_i R_{ij} + v_j \left(-\sum_{k\neq j} R_{jk} \right) = \sum_{i\neq j} v_i \sqrt{\frac{v_j}{v_i}} + v_j \left(-\sum_{k\neq j} \sqrt{\frac{v_k}{v_j}} \right) = 0. \; \square$$

Note that Lemma 1 also follows from detailed balance of v and R [27].

Corollary 1 *The sample mean approximation $\bar{w}$ of the discrete Boltzmann measure given in Eq. (4) is the invariant measure of the transition rate matrix Q'' given by Eqs. (5,7): $0 = \bar{w}^\top Q''$.*

2.3 Convex polyhedra formulation of Voronoi tesselations

In this section we present the second important idea of our method: a formulation of Voronoi tesselations in terms of convex polyhedra, which we use to avoid the curse of dimensionality when we discretize the configuration space Ω and compute the matrix N of adjacency relations. In Eq. (1) we defined a Voronoi region Ω_i associated with a generator g_i as the set of points which were closer to g_i than any other generator. We will present an alternative definition and show that these two definitions agree.

For the set of n distinct generating points $\{g_i\}_{i=1}^n \subset \mathbb{R}^d$, let $g_{i,j}$ denote the j-th coordinate of g_i, construct $b \in \mathbb{R}^n$ and $A \in \mathbb{R}^{n\times(d+1)}$ according to

$$b_i := \|g_i\|^2 \quad \text{and} \quad A_{ij} := \begin{cases} 2g_{i,j} & 1 \leq j \leq d \\ -1 & j = d+1, \end{cases} \tag{8}$$

and use these to construct the $(d+1)$-dimensional convex polyhedron $P \subset \mathbb{R}^{d+1}$,

$$P := \left\{ \tilde{x} = (x, x_{d+1}) \in \mathbb{R}^d \times \mathbb{R} \; : \; b - A\tilde{x} \geq 0 \right\}. \tag{9}$$

For $i = 1, \ldots, n$, define $f^i : \mathbb{R}^{d+1} \to \mathbb{R}$ by $f^i(\tilde{x}) := b_i - \sum_{\ell=1}^{d+1} A_{i\ell}\tilde{x}_\ell$. The supporting hyperplane h_i of P and the facet F_i associated with g_i are given by

$$h_i := \left\{ \tilde{x} \in \mathbb{R}^{d+1} \; : \; f^i(\tilde{x}) = 0 \right\} \quad \text{and} \quad F_i := P \cap h_i.$$

(Recall that a facet of a convex k-polyhedron has dimension $k-1$, so that F_i has dimension d.) Let Π_d be the operator which projects subsets of $\mathbb{R}^{d+1}$ to the first d coordinates. Consider the d-dimensional set Ω_i',

$$\Omega_i' := \Pi_d F_i. \tag{10}$$

Lemma 2 *Let Ω_i be as in Eq. (1) and Ω_i' be as in Eq. (10). Then $\Omega_i' = \Omega_i$.*

Results similar to Lemma 2 appear in [1, 23]. The proof below is our own.
Proof: By definition of h_i, b and A, every $\tilde{x} = (x, x_{d+1}) \in h_i$ satisfies $x_{d+1} = 2\langle g_i, x\rangle - \|g_i\|^2$. Substitute $\tilde{x} = (x, 2\langle g_i, x\rangle - \|g_i\|^2)$ into $b_j - \sum_k A_{jk}\tilde{x}_k \geq 0$ for $j \neq i$:

$$\|g_j\|^2 - 2\langle g_j, x\rangle + 2\langle g_i, x\rangle - \|g_i\|^2 \geq 0.$$

Since $F_i = P \cap h_i$, we have

$$\begin{aligned}
\Omega_i' &:= \Pi_d \left\{(x, 2\langle g_i, x\rangle - \|g_i\|^2) \in \mathbb{R}^d \times \mathbb{R} \ : \right. \\
&\qquad \left. \|g_j\|^2 - 2\langle g_j, x\rangle + 2\langle g_i, x\rangle - \|g_i\|^2 \geq 0 \text{ for } j \neq i\right\} \\
&= \left\{x \in \mathbb{R}^d \ : \ \|g_j\|^2 - 2\langle g_j, x\rangle + 2\langle p_i, x\rangle - \|g_i\|^2 \geq 0 \text{ for } j \neq i\right\} \\
&= \left\{x \in \mathbb{R}^d \ : \ \|g_j\|^2 - 2\langle g_j, x\rangle + \|x\|^2 - \|x\|^2 + 2\langle p_i, x\rangle - \|g_i\|^2 \geq 0 \text{ for } j \neq i\right\} \\
&= \left\{x \in \mathbb{R}^d \ : \ \|x - g_j\|^2 \geq \|x - g_i\|^2 \text{ for } j \neq i\right\} = \Omega_i. \ \square
\end{aligned}$$

Recall from Section 2.2 that two Voronoi regions were defined to be adjacent if $\dim(\Omega_i \cap \Omega_j) = \dim(\Omega_i) - 1$. Now consider

Definition 1 *Two d-dimensional Voronoi regions $\Omega_i = \Pi_d F_i$ and $\Omega_j = \Pi_d F_j$ are* adjacent *if their liftings F_i and F_j have the property that $dim(F_i \cap F_j) = d-1$. Equivalently, Ω_i and Ω_j are adjacent if $F_i \cap F_j$ is a facet of F_i.*

Together, Lemma 2 and Definition 1 lead to a method for computing the matrix N of adjacency relations via linear programming. The linear programming approach allows us to discretize configuration spaces while avoiding the curse of dimensionality. We first found the idea of using linear programs to compute N in [9].

Fix $i \in \{1, \ldots, n\}$ and $j \neq i$. Let $\{e_k\}_{k=1}^n \subset \mathbb{R}^n$ be the set of canonical orthonormal basis vectors and define the linear program

$$LP(i,j): \quad \begin{cases} \text{minimize} & f^i(\tilde{x}) \\ \text{subject to} & (b + e_i) - A\tilde{x} \geq 0 \\ & f^j(\tilde{x}) = 0. \end{cases} \tag{11}$$

For $v \geq 0$, consider the closed half-spaces

$$H_{i,\geq -v} := \left\{\tilde{x} \in \mathbb{R}^{d+1} \ : \ f^i(\tilde{x}) \geq -v\right\} \text{ and } H_{i,\leq -v} := \left\{\tilde{x} \in \mathbb{R}^{d+1} \ : \ f^i(\tilde{x}) \leq -v\right\}.$$

The half-spaces are important because they define the polyhedron P via $P = \cap_k H_{k,\geq 0}$. The idea from [9] behind the linear program LP(i,j) is that only the Voronoi regions which are adjacent to Ω_i will change when Ω_i changes. We change Ω_i by translating the supporting hyperplane h_i down along the x_{d+1}-axis by unity to obtain the parallel hyperplane h_i'. Translating h_i to h_i' perturbs $H_{i,\geq 0}$ to $H_{i,\geq -1}$ and P to $P' := (\cap_{k \neq i} H_{k,\geq 0}) \cap H_{i,\geq -1}$. The facets of P' are given by $F_j' := P' \cap h_j$ for $j \neq i$ and $F_i' := P' \cap h_i'$. In LP(i,j) above, we minimize over F_j'; adjacency of F_j to F_i implies that the image F_j' of F_j under the perturbation will properly contain F_j, in which case $f^i(\cdot)$ will attain strictly negative values. On the other hand, if $f^i(\cdot)$ attains strictly negative values, then we can also show that F_j must be adjacent to F_i. We now make these arguments rigorous.

Lemma 3 *If $F_i \cap F_j$ is not a facet of F_j, then $\cap_{k\neq i}\left(H_{k,\geq 0}\cap h_j\right) \subseteq H_{i,\geq 0}\cap h_j$.*

Proof: We prove the contrapositive. Suppose $\cap_{k\neq i}\left(H_{k,\geq 0}\cap h_j\right) \supset H_{i,\geq 0}\cap h_j$. Intersecting both sides with $H_{i,\geq 0}\cap h_j$ yields $F_j = H_{i,\geq 0}\cap h_j$, and intersecting both sides of $F_j = H_{i,\geq 0}\cap h_j$ with h_i yields $F_j \cap h_i = h_i \cap h_j$. But $F_j \cap h_i = F_j \cap F_i$ and $\dim(h_i \cap h_j) = d-1$ since h_i and h_j are hyperplanes in $\mathbb{R}^{d+1}$. $\square$

Lemma 4 *It holds that $F_j \subseteq F_j'$. If $\cap_{k\neq i}\left(H_{k,\geq 0}\cap h_j\right) \subseteq H_{i,\geq 0}\cap h_j$ then $F_j = F_j'$.*

Proof: Since $H_{i,\geq 0} \subset H_{i,\geq -1}$, for $j \neq i$ we have

$$\left[\cap_{k\neq i}\left(H_{k,\geq 0}\cap h_j\right)\right]\cap\left(H_{i,\geq 0}\cap h_j\right) \subseteq \left[\cap_{k\neq i}\left(H_{k,\geq 0}\cap h_j\right)\right]\cap\left(H_{i,\geq -1}\cap h_j\right). \quad (12)$$

The second conclusion follows immediately from $H_{i,\geq 0}\cap h_j \subseteq H_{i,\geq -1}\cap h_j$. $\square$

Lemma 5 *The optimal value of $f^i(\cdot)$ in $LP(i,j)$ is strictly negative if and only if $F_j \subset F_j'$.*

Proof: '$\Leftarrow$': Since $F_j' \subset H_{i,\geq -1}$ and $F_j \subset H_{i,\geq 0}$, it holds that $F_j \subset F_j'$ implies $\emptyset \neq F_j' \setminus F_j \subset H_{i,\geq -1} \setminus H_{i,\geq 0}$. Thus there exists a $\tilde{x}$ in F_j' for which $f^i(\tilde{x}) < 0$.

'$\Rightarrow$': Suppose that $F_j' = F_j$. Then LP(i,j) is a minimization problem over $F_j \subset H_{i,\geq 0}$. Over $H_{i,\geq 0}$, $f^i(\cdot)$ has only nonnegative values. $\square$

We now state and prove the main result of this section.

Theorem 2 *Let $\{\Omega_i\}_{i=1}^n \subset \mathbb{R}^d$ be a Voronoi tesselation of $\mathbb{R}^d$ generated by $\{g_i\}_{i=1}^n$. Fix $i \neq j \in \{1,\ldots,n\}$ and define $LP(i,j)$ as in (11). Then the optimal value of $f^i(\cdot)$ is strictly negative if and only if Ω_i and Ω_j are adjacent.*

Proof of Theorem 2: Since the optimization problem in LP(i,j) is over the facet $F_j' = P' \cap h_j \neq \emptyset$ and since $P' \subseteq H_{i,\geq -1}$ it holds that the objective function $f^i(\cdot)$ is bounded from below by -1, so the minimization problem is well-defined.

'$\Leftarrow$': Suppose Ω_i and Ω_j are adjacent. Then $F_j \cap F_i = F_j \cap h_i$ is a facet of F_j, by Definition 1. Since $F_j \cap h_i \neq F_j' \cap h_i'$, it follows that a facet of F_j changes under the perturbation which sends P to P'. Changing a facet of a convex polyhedron changes the polyhedron itself and F_j is a convex polyhedron, so $F_j' \neq F_j$. Applying the first part of Lemma 4 and then Lemma 5 yields the desired conclusion.

'$\Rightarrow$': Suppose Ω_i and Ω_j are not adjacent. Then $F_j \cap h_i$ is not a facet of F_j, so we can use Lemma 3, the second part of Lemma 4 and Lemma 5 (in that order). $\square$

Theorem 2 appears without proof in [9]. The proof we have shown above, (including Lemmas 3-5) is our own. Connections between linear programs and Voronoi tesselations have been studied for other problems in [11, 33, 36].

In this section, we computed adjacency relations by solving a linear program that was expressed in terms of (b, A). We can also use (b, A) to assign a randomly sampled configuration $q \in \Omega$ to a Voronoi region. Let $\tilde{q} := (q, \|q\|) \in \mathbb{R}^d \times \mathbb{R}$. The index $i = \text{argmin}_k(b - A\tilde{q})_k$ gives the index of the Voronoi region

which contains q. In Section 4, we compute Centroidal Voronoi tesselations [30] using uniformly distributed generators, because such tesselations are simple to generate via an iterative method and yield discretization regions which are close to uniform. We leave the question of how to best choose the generators for future work. In Section 5.1 we compare the computational cost of different methods of discretizing via Voronoi tesselations, and show that the linear programming formulation we presented above enables us to break the curse of dimensionality.

2.4 Identifying conformations

In this section we review some results from [37] on identifying conformations from transition matrices. In the previous section we used non-overlapping, closed Voronoi regions to discretize the configuration space,

$$\Omega = \cup_{k=1}^{n} \Omega_k, \quad \text{int}(\Omega_i \cap \Omega_j) = \text{int}(\partial\Omega_i \cap \partial\Omega_j) = \emptyset \text{ for } i \neq j, \tag{13}$$

and we used indicator functions to assign configurations to Voronoi regions: $1_{\Omega_i}(q) = 1$ if $q \in \Omega_i$ and $1_{\Omega_i}(q) = 0$ otherwise. We wish to find a partitioning of configuration space into n_c overlapping conformations, $\{C_j\}_{j=1}^{n_c}$

$$\Omega = \cup_{\ell=1}^{n_c} C_\ell, \quad \text{int}(C_i \cap C_j) \neq \emptyset \text{ for adjacent } C_i \text{ and } C_j, \ i \neq j, \tag{14}$$

where the regions of overlap are the *transition regions*.

To describe the conformations, we use membership functions $\chi_i : \Omega \to [0, 1]$, which assign a degree of membership to each conformation C_i. The χ_i generalize indicator functions by permitting configurations to belong to more than one conformation, and form a partition of unity: $\forall q \in \Omega$, $\sum_{j=1}^{n_c} \chi_j(q) = 1$. Discretizing the set of membership functions yields a matrix $\chi' \in [0, 1]^{n \times n_c}$ where the i-th column is the discretization of the i-th membership function and χ'_{ij} gives the degree of membership of the jth discretization region to the i-th conformation. It follows from the nonnegativity and partition of unity properties of the membership functions $\{\chi_i\}_{i=1}^{n_c}$ that χ' is non-negative and row-stochastic, which is equivalent to the following property:

(P0): Each row $(\chi'^{\top})_i$ as a point in $\mathbb{R}^{n_c}$ must lie on or in the (n_c)-standard simplex, i.e. the simplex in $\mathbb{R}^{n_c}$ spanned by the canonical orthonormal basis vectors.

We want to compute χ' from the transition rate matrix Q'' that we constructed in Section 2.2, which is an approximation of $Q = \frac{d}{d\tau} P(\tau)|_{\tau=0}$ from Theorem 1. The matrix $P(\tau)$ is a discretization of the generalized transfer operator $\mathcal{P}(\tau)$ which describes the transfer of configurations q over a lag time τ. We must have $\mathcal{P}(\tau)\chi_i(q) \approx \chi_i(q)$ because configurations in C_i tend to stay in C_i over τ units of time, i.e. since C_i is a metastable set. If $\{\mathcal{P}(\tau)\}_{\tau>0}$ is a semigroup with an infinitesimal generator $\mathcal{Q}$, then

$$\mathcal{Q}\chi_i(q) = \lim_{\tau \searrow 0} \frac{\mathcal{P}(\tau)\chi_i(q) - \chi_i(q)}{\tau} \approx \lim_{\tau \searrow 0} \frac{\chi_i(q) - \chi_i(q)}{\tau} = 0. \tag{15}$$

Thus, the membership functions are close to eigenfunctions of $\mathcal{Q}$ associated with eigenvalues near zero, and (15) implies the following property:

(P1): the columns of χ' are close to right eigenvectors of Q'' corresponding to eigenvalues $\{\theta_i\}_{i=1}^{n_c}$ near zero:

$$Q''\chi' \approx \chi'\Theta, \quad \Theta := \mathrm{diag}(\theta_1, \ldots, \theta_{n_c}),\ \theta_i \approx 0,\ i = 1, \ldots, n_c. \tag{16}$$

If $\Theta = \theta\ \mathrm{Id} \in \mathbb{R}^{n_c \times n_c}$ for some θ near 0, then given an eigenvector matrix X satisfying $Q''X = \theta X$, any transformation $T \in GL(n_c)$ yields another eigenvector matrix:

$$Q''(XT) = (Q''X)T = (\theta X)T = \theta(XT).$$

We know by (P1) that $\Theta \approx 0\ \mathrm{Id}$. Thus, for any matrix X computed from standard eigenvector algorithms satisfying $Q''X = X\Theta$, we have $\chi' \approx XT$. Combining this fact with (P0) implies that the rows of X are close to a simplex. In order to obtain a membership vector matrix χ' from an eigenvector matrix X, it suffices to map the vertices of the simplex formed by the rows of X to the canonical orthonormal basis vectors in $\mathbb{R}^{n_c}$. The idea behind the Inner Simplex Algorithm is to construct T from

$$T^{-1} = \begin{pmatrix} X(\rho_1, 1) & \ldots & X(\rho_1, n_c) \\ \vdots & \ddots & \vdots \\ X(\rho_{n_c}, 1) & \ldots & X(\rho_{n_c}, n_c) \end{pmatrix}, \tag{17}$$

where the $\{\rho_i\}_{i=1}^{n_c}$ are indices of the rows of the eigenvector matrix X that are furthest away from each other (i.e. indices of the rows which form the vertices of a simplex) [37]. If the resulting matrix $\chi' := XT$ has negative entries, one may use Robust Perron Cluster Analysis (PCCA+) to turn χ' into a non-negative matrix [7, 20].

In the discussion above, we assumed that we knew the number of conformations n_c. If one does not know the value of n_c, one identifies a cluster of eigenvalues of Q'' which are close together near zero (of smallest magnitude) and separated from the rest of the spectrum of Q'' by a spectral gap, and sets n_c equal to the number of eigenvalues in the cluster. The smallest eigenvalues are related to the time scales of the slow dynamics [24], and in order to obtain a spectral gap, one must identify a configuration space in which one can distinguish between fine scale dynamics and coarse scale dynamics [15]. One calls the eigenvalues in the cluster 'Perron eigenvalues' and the corresponding right eigenvectors 'Perron eigenvectors'.

2.5 Conformation dynamics

So far, we have seen that with Q'' we can describe the evolution of distributions of configurations over the Voronoi regions and identify the conformations of the system. In this section, we identify the equilibrium distribution on the conformations and compute conformational transition rates by coarse-graining [21, 22].

We obtain distributions x_c over conformations from distributions x over discretization sets by projection. Let $R : \mathbb{R}^n \to \mathbb{R}^{n_c}$ be given by $Rx := {\chi'}^{\top} x \in [0, 1]^{n_c}$ and set $x_c := Rx$; it holds that x_c is non-negative and its elements sum to one. We obtain the equilibrium distribution over conformations

$$\bar{w}_c := R\bar{w}. \tag{18}$$

Define $D := \mathrm{diag}(\bar{w})$, $D_c := \mathrm{diag}(\bar{w}_c)$, the interpolation operator $I : \mathbb{R}^{n_c} \to \mathbb{R}^n$ by $Iy := D\chi' D_c^{-1} y$ and the reduction Q_c'' of Q'' by $Q_c'' := \left(R(Q'')^\top I(RI)^{-1}\right)^\top$. Using the definitions of R and I and using $\chi' = XT$ one can show that

$$Q_c'' = T^{-\top}\Theta T^\top. \tag{19}$$

Note that the entries of Q_c'' are not true transition rates, because the conformations overlap. We can mitigate, but not fully eliminate, this problem by reducing the overlap between conformations [18]. However, the matrix Q_c'' that we obtain from the coarse-graining procedure is 'correct' in the sense that one obtains the same coarse-grained distributions, regardless of whether one first evolves $x(t)$ forward in time using Q'' and then coarse-grains using R, or first coarse-grains $x(t)$ and then evolves the resulting $x_c(t)$ forward in time using Q_c'' (see Theorem IV.2 of [22]). Eq. (19) is useful if one wishes to compute $P_c(\tau) = \exp(\tau Q_c'')$ - the matrix of 'transition probabilities' between conformations over a lag time τ - since then one only needs to compute $\exp(\tau\Theta)$.

Note also that since we obtain the rate matrix Q'' after the assumption of constant flux (see Section 2.2) Q_c'' is a scalar multiple of the true coarse-grained propagator. Thus, one can only use the entries of Q_c'' to compute ratios between transition rates. For example, the ratio $Q_{c,ik} : Q_{c,ij}$ - the ratio of the k-th entry of the i-th row of Q_c to the j-th entry in the i-th row of Q_c - is the ratio of the absolute rate of transition from the i-th to the k-th metastable set, relative to the absolute rate of transition from the i-th to the j-th metastable set. The diagonal entries are related to the average or expected time that a molecule stays in the conformation [22].

The idea of coarse-graining - that some Markov chains can be simplified by partitioning their state spaces into groups - has been studied before under different names, e.g. 'lumping' [4] or 'aggregation' [35], and chains which can be simplified this way are sometimes called 'nearly completely decomposable'.

3 Method

We now provide a more detailed description of our method to supplement that given in Section 1. We assume that one has a potential energy function $V : \Omega \to \mathbb{R}$ described over a configuration space Ω and has fixed the value of the inverse temperature β. We ignore the computation of the partition function Z in steps 1 and 3 below.

1. Generate a collection $\{q_j\}_{j=1,\ldots,M} \subset \Omega$ of M configurations. Compute the Boltzmann densities $\{p_j := \exp(-\beta V(q_j))\}_{j=1,\ldots,M}$.

2. Using n random generators drawn from the uniform distribution on a subset of the configuration space, compute a centroidal Voronoi tesselation [30] such that every Voronoi region contains at least one configuration.

3. Compute the $\{\bar{w}_i\}_{i=1,\ldots,n}$ as given in Eq. (4). Normalize the resulting vector to get a discretized distribution.

4. Compute the matrix of adjacency relations $N \in \{0,1\}^{n\times n}$ using Theorem 2.

5. Construct the matrix Q'' as described in Eqs. (5) and (7).

6. Given n_c, compute the Perron eigenvector matrix $X \in \mathbb{R}^{n \times n_c}$ for Q''. If n_c is not known, identify a cluster of Perron eigenvalues in the spectrum of Q'' near zero and set n_c to the number of elements in this cluster.

7. Transform the Perron eigenvector matrix X into a matrix χ' of membership vectors, using the Inner Simplex Method or Robust Perron Cluster Analysis.

8. Compute the vector w_c of weights of conformations using (18) and the matrix Q_c'' of relative conformational transition rates using (19).

4 Examples

In this section we investigate the performance of the method on some examples.

4.1 A simple model of an energy landscape

We applied the method to a model of a molecule with the potential energy function

$$\begin{aligned} V(q) =& 3\exp\left(-q_1^2 - (q_2 - 1/3)^2\right) - 3\exp\left(-q_1^2 - (q_2 - 5/3)^2\right) \\ & - 5\exp\left(-(q_1 - 1)^2 - q_2^2\right) - 5\exp\left(-(q_1 + 1)^2 - q_2^2\right) \\ & + 0.2q_1^4 + 0.2\,(q_2 - 1/3)^4\,. \end{aligned} \tag{20}$$

This energy landscape has been studied before, e.g. in [29, 32]. This system has three metastable sets: C_2 and C_3, centered at $(-1, 0)$ and $(1, 0)$ respectively; and C_1 centered at $(0, 5/3)$. The basins of low energy at C_2 and C_3 are deeper than the basin at C_1. A small bump centered at $(0, 1/3)$ gives an energy barrier between the basins at C_2 and C_3. The configuration space is $\Omega = \mathbb{R}^2$.

We applied the method to the above model. We used the value $\beta = 3.34$ for the inverse temperature, and we used $n = 100$ generators and $M = 2000$ configurations drawn from the uniform distribution on $[-2, 2] \times [-1, 2]$. We show some of our results in Figures 3 to 8.

In Figure 3, a Voronoi region has a darker shade if its Boltzmann measure is higher. Two clusters of dark regions correspond to the basins at C_2 and C_3. The regions clustered around $(0, 5/3)$ are not as dark because the basin at C_1 is shallower than the basins at C_2 and C_3, so the Voronoi regions in C_1 have smaller Boltzmann measure. In the plot of the smallest nine eigenvalues of Q'' (Figure 4) a spectral gap separates the three smallest eigenvalues from the rest of the spectrum, indicating that this application of our method has resolved three conformations, i.e. that $n_c = 3$. In Figure 5, the rows of the reduced Perron eigenvector matrix $X \in \mathbb{R}^{100 \times 2}$ almost perfectly span a simplex in $\mathbb{R}^2$ (by the 'reduced Perron eigenvector matrix', we refer to the Perron eigenvector matrix without the constant right eigenvector corresponding to the zero eigenvalue of Q''). We map this simplex to the standard simplex in $\mathbb{R}^2$ to obtain the membership vectors for each conformation. We show the membership vectors in Figures 6, 7, and 8, where a Voronoi region is darker if its membership to the conformation is higher (closer to unity).

We evaluated the method by comparing the statistical weights and ratios of transition rates to the reference values obtained by Gauss-Legendre quadrature on a regular 9×9 grid as given in [38]. We used a regular 201×151 mesh of data points in $[-2, 2] \times [-1, 2]$ as our set of data on the energy landscape. On this data set we applied four different discretizations: 9×9 and 18×18 mesh-based discretizations, and random CVTs consisting of 81 and 324 centers computed after 500 iterations.

In Table 1 we present the results from mesh-based discretizations and from CVTs. Refining the discretization led to better approximations of weights and transition rates for both the mesh-based and centroidal Voronoi tesselation discretizations. However, mesh-based discretizations yielded better approximations of statistical weights, while CVTs yielded better approximations of ratios of transition rates.

The spectral gap $|\theta_4| - |\theta_3|$ of the transition rate matrix Q'' decreased when the discretization was refined for both discretization methods, from 1.5249 to 0.6984 for mesh-based discretizations and from 1.7700 to 0.6107 for CVTs. This observation suggests that refining discretizations uniformly need not lead to a better description of the slow dynamics.

4.2 Entropic effects

To illustrate that our method can also resolve the entropic effects as mentioned in [29], we applied a CVT to $[-2, 2] \times [-1, 2]$, while using the same mesh-based data sites as in Table 1. We used two different values of the inverse temperature, $\beta = 1.67$ and $\beta = 6.67$. We summarize our results in Table 2.

In Table 2 the ratio $2 \to 1 : 2 \to 3$ is $1.4969 : 1$ when $\beta = 1.67$ and is $3.1027 : 1$ when $\beta = 6.67$. That is, at both temperatures, more molecules in C_2 transition to C_1 than to C_3, because they need less energy to move from C_2 and C_1. Decreasing the temperature (changing from $\beta = 1.67$ to $\beta = 6.67$) reduced the transition rate from C_2 to C_3, because at lower temperatures the molecules have less energy on average to overcome the energy barrier between C_2 and C_3.

We observed that increasing the temperature reduced the spectral gap $|\theta_4| - |\theta_3|$ from 1.0330 at $\beta = 6.67$ to 0.3354 at $\beta = 1.67$. That is, increasing the temperature reduces the separation between the short and large time scales. This agrees with the physical intuition that at higher temperatures, molecules have more energy on average and thus are able to leave the metastable sets sooner.

5 Critique

In this section we analyze our method. We discuss the computational cost of the method, and study the performance of the method when the dimension of the problem increases.

5.1 Computational cost

The bottleneck in our method lies in the computation of the matrix N of adjacency relations given the generators of a tesselation. A straightforward approach would be to compute the full Voronoi tesselation in order to find all the vertices

of the Voronoi regions and compare vertices of each Voronoi region. However, algorithms for computing full Voronoi tesselatons suffer from restrictions due to dimensionality. For example, the Qhull algorithm does not a priori suffer from the curse of dimensionality but works only for spaces of dimension $d \leq 9$ [5]. An incremental algorithm for computing a full Voronoi tesselation of $\mathbb{R}^d$ into n regions runs in $O(n^{\lceil d/2 \rceil})$ [17]. Methods for finding all the vertices in Voronoi tesselations of spaces of any dimension seem to lead to the curse of dimensionality, given the tight upper bound for the number of vertices of $O(n^{\lfloor (d+1)/2 \rfloor})$ [31].

The discretization method we presented in Section 2.3 has an important advantage in that we need not find any vertices at all in order to find the adjacency relations. The cost of the linear program used to solve $LP(i,j)$ in Theorem 2 depends on the choice of linear program solver. A straightforward, practical choice would be the simplex algorithm, which is known to have at worst exponential computational cost but on average polynomial computational cost [25]. Furthermore, it was shown in [19] that linear programs can be solved in polynomial time. We conclude that the discretization method shown in Section 2.3 does not lead to the curse of dimensionality.

The most expensive part of computing the transition rate matrix Q'' after computing N is the computation of $\bar{w}$. For a Voronoi region Ω_i, the Monte Carlo quadrature approximation $\bar{w}_i$ converges in probability with order $O(1/\sqrt{M_i})$ to the true Boltzmann weight w_i as the number M_i of configurations belonging to Ω_i goes to infinity. Note that approximating the potential energy $V(q_j)$ of a configuration q_j may be very expensive for large molecules, given that such potential energies are often computed using pairwise atomic interactions, and given that the number of pairwise interactions increases exponentially with the number of atoms.

In Section 2 we showed how we obtained Q'' from the rate matrix Q defined in Theorem 1. In [38], the complexity of this approach is $O(n(f_n+1))$ (see Table 4.1 of [38] under 'Theorem of Gauss'). The quantity f_n is the average number of nonzero entries in $N \in \{0,1\}^{n \times n}$, i.e. the average number of neighbours of a Voronoi region. Compared to the complexity of $O(1/(1-\lambda_2))$ of thermostated molecular dynamics simulations from the same table (where $\lambda_2 \in (0,1)$ is the eigenvalue of the transfer operator $\mathcal{P}(\tau)$ closest to unity), the approach we use is more favourable in terms of complexity.

We conclude our critique by observing that our method is robust with respect to the dimension of the configuration space of the molecule, since the computational cost is effectively polynomial in the dimension d and the number of Voronoi regions n. Furthermore, the computational complexity of our method is smaller than using thermostated molecular dynamics simulations.

5.2 Higher-dimensional models

Now we investigate how our method performs when we consider higher-dimensional models. Rewrite the potential energy function in (20) as the sum of Gaussian and polynomial terms,

$$V(q) = \sum_{i=1}^{s} a_i \exp\left[-\|q - w_i\|^2\right] + \sum_{j=1}^{d} b_j (q_j - c_j)^4 \tag{21}$$

where $d = 2$ is the dimension of the configuration space and $s = 4$ is the number of Gaussian terms. The coefficient vectors are $a = (3, -3, -5, -5)^\top$, $b = (0.2, 0.2)^\top$, and the centering terms are $c = (0, 1/3)^\top$ and $W \in \mathbb{R}^{d\times s}$, where

$$W = \begin{pmatrix} 0 & 0 & 1 & -1 \\ 1/3 & 5/3 & 0 & 0 \end{pmatrix}.$$

We call c and W the 'centering terms' because they provide the centers for the polynomial and Gaussian terms.

We investigated the effect of increasing the dimension of the simple model given by (20) by appropriately augmenting the centering terms c and W with zero rows. For example, if we consider the model into $\mathbb{R}^4$ by adding two extra dimensions, we obtain the parameter set

$$a = \begin{pmatrix} 3 \\ -3 \\ -5 \\ -5 \end{pmatrix}, \quad b = \begin{pmatrix} 0.2 \\ 0.2 \\ 0.2 \\ 0.2 \end{pmatrix}, \quad c = \begin{pmatrix} 0 \\ 1/3 \\ 0 \\ 0 \end{pmatrix}, \quad W = \begin{pmatrix} 0 & 0 & 1 & -1 \\ 1/3 & 5/3 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

For $d = 2, \ldots, 7$ we applied our method by using a CVT with 324 regions and 3000 data sites drawn from the uniform distribution on $[-2, 2] \times [-1, 2] \times \prod^{d-2}[-0.005, 0.005]$, and specifying $\beta = 3.34$. In Figures 9-14 we plot the rows of the corresponding Perron eigenvector matrices. We observe that the simplex structure persists for $d = 2$ and $d = 3$, but for $4 \leq d \leq 7$, there is very little simplex structure to the rows. Recall that the simplex structure of the Perron eigenvector matrix is important in identifying the matrix χ' of membership vectors (see Section 2.4). When the eigenvector matrix lacks a clear simplex structure, it becomes difficult to identify the conformations.

The examples above emphasize that working with high-dimensional energy landscapes is difficult, because it is difficult to find a discretization in which conformations can be identified when the dimension of the landscape is high.

5.3 Voronoi tesselations with other metrics

An important aspect of
Voronoi tesselations which we have not investigated in this article is the application of Voronoi tesselations to spaces other than $\mathbb{R}^d$. Given that some molecules may have configuration spaces with torsion angle coordinates, the choice of metric for computing distances to generators may lead to difficulties. For example, if the configuration space of a molecule is given in terms of two torsion angles, then this configuration space corresponds to the 2-torus, $\mathbb{R}^2/\mathbb{Z}^2$, and the metric on the 2-torus is different from the Euclidean metric we have used in this article.

6 Conclusion

In this paper we presented a method for constructing a Markov State Model of a molecular system. The novel ideas in our method are the assumption of constant flux and the square root approximation of transition rates between two adjacent discretization regions. These ideas enabled us to avoid trajectories and reduce the computation of transition rates between discretization sets

to the computation of Boltzmann weights of discretization sets and the computation of adjacency relations. We approximate the Boltzmann weights by Monte Carlo quadrature, which gives dimension-independent but only probabilistic convergence. We obtain the adjacency relations using a linear programming formulation of Voronoi tesselations, which allows us to break the curse of dimensionality when discretizing the configuration space. Finding discretizations of high-dimensional configuration spaces in which conformations can be identified remains a challenge.

The key idea in our method is that the equilibrium dynamics of a molecular system are determined completely by its energy landscape. In particular, if one only wishes to approximate the conformation statistical weights and relative conformation transition rates, then it is not necessary to simulate trajectories; simulations are only necessary if one wishes to obtain dynamic quantities. It suffices to randomly sample the energy landscape at points and to construct a Markov State Model using these point data on the energy landscape. By avoiding trajectory-based methods (and more generally, mesh-based methods) one can break the curse of dimensionality and in principle study large molecules which are too large to be simulated.

Acknowledgments

The authors thank the referees for their careful reading of the manuscript and helpful comments. The first author would like to thank the Berlin Mathematical School and the second author the DFG - Research Center MATHEON.

References

[1] A. Okabe, B. Boots, K. Sugihara and S.N. Chiu, *Spatial tessellations: Concepts and applications of Voronoi diagrams*, 2nd ed., Wiley Series in Probability and Statistics, Wiley, 2000.

[2] R. Bellman, *Dynamic programming*, Princeton University Press, Princeton, New Jersey, 1972.

[3] A. Bovier, *Metastability: a potential theoretic approach.*, Marta Sanz-Solé (ed.) et al., Proceedings of the International Congress of Mathematicians (ICM), Madrid, Spain, 2006. Volume III: Invited lectures. Zürich: European Mathematical Society (EMS), 499-518.

[4] P. Buchholz, *Exact and Ordinary Lumpability in Finite Markov Chains*, J. Appl. Prob. **31** (1994), 59–75.

[5] C. B. Barber, D.P. Dobkin and H. Huhdanpaa, *The Quickhull Algorithm for Convex Hulls*, ACM Trans. Math. Softw. **22** (1996), 469–483.

[6] C. Schütte, W. Huisinga, and P. Deuflhard, *Transfer operator approach to conformational dynamics in biomolecular systems*, Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems (B. Fiedler, ed.), Springer., 2001, pp. 191–223.

[7] P. Deuflhard and M. Weber, *Robust Perron cluster analysis in conformation dynamics*, 2005, Technical Report ZIB 03-19, Zuse Institut Berlin.

[8] F. Rao and A. Caflisch, *The protein folding network.*, J. Mol. Biol. **342** (2004), 299–306.

[9] K. Fukuda, *Frequently Asked Questions in Polyhedral Computation*, http://www.ifor.math.ethz.ch/ fukuda/polyfaq.html.

[10] G.R. Bowman and V.S. Pande, *Protein folded states are kinetic hubs.*, Proc Natl Acad Sci U S A. **107** (2010), 16749.

[11] D. Hartvigsen, *Recognizing Voronoi Diagrams with Linear Programming*, INFORMS Journal on Computing **4** (1992), 369–374.

[12] W. Huisinga, *Metastability of Markovian systems: A transfer operator based approach in application to molecular dynamics*, 2001, Doctoral Thesis, Department of Mathematics and Computer Science, Freie Universität Berlin.

[13] W. Huisinga and B. Schmidt, *Metastability and Dominant Eigenvalues of Transfer Operators*, New Algorithms for Macromolecular Simulation (Benedict Leimkuhler et al., ed.), Lecture Notes in Computational Science and Engineering, vol. 49, Springer-Verlag, Berlin/Heidelberg, 2006, pp. 167–182.

[14] J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill., and W. C. Swope, *Automatic discovery of metastable states for the construction of Markov models of macromolecular conformational dynamics*, J. Chem. Phys. **126** (2007), 155101.

[15] J. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. Chodera, C. Schütte, F. Noe, *Markov models of molecular kinetics: Generation and validation*, J. Chem. Phys. **134** (2011), 174105.

[16] A. Jain and G. Stock, *Identifying metastable states of folding proteins*, Journal of Chemical Theory and Computation **8** (2012), no. 10, 3810–3819.

[17] J.E. Goodman and J. O'Rourke (eds.), *Handbook of Discrete and Computational Geometry*, CRC Press, 1997.

[18] K. Fackeldey, S. Roeblitz, O. Scharkoi, M. Weber, *Soft Versus Hard Metastable Conformations in Molecular Simulations*, 2011, ZIB Report 02-12, Zuse Institut Berlin.

[19] L.G. Khachiyan, *A polynomial algorithm in linear programming*, Soviet Mathematics Doklady **20** (1979), 191–194.

[20] S. Kube and P. Deuflhard, *Errata on Robust Perron Cluster Analysis in Conformation Dynamics*, 2006.

[21] S. Kube and M. Weber, *Coarse grained molecular kinetics*, 2006, ZIB Report 06-35, Zuse Institut Berlin.

[22] ———, *A coarse graining method for the identification of transition rates between molecular conformations*, J. Chem. Phys. **126** (2007), 024103.

[23] M. de Berg, O. Cheong, M. van Kreveld and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed., Springer-Verlag TELOS, 2008.

[24] M. Sarich, F. Noe, C. Schütte, *On the approximation quality of Markov state models*, Mult. Mod. Sim. **8** (2010), 1154–1177.

[25] J. Matousek and B. Gärtner, *Understanding and Using Linear Programming*, Springer, 2007.

[26] N. Djurdjevac, M. Sarich, C. Schütte, *Estimating the eigenvalue error of Markov state models*, Mult. Mod. Sim. **10** (2012), 61–81.

[27] J.R. Norris, *Markov chains*, Cambridge University Press, Cambridge, 1997.

[28] P. Deuflhard, W. Huisinga, A. Fischer, and C. Schütte, *Identification of almost invariant aggregates in reversible nearly uncoupled Markov chains*, Lin. Alg. Appl. **315** (2000), 39–59.

[29] P. Metzner, C. Schütte and E. Vanden-Eijnden, *Illustration of transition path theory on a collection of simple examples.*, J. Chem. Phys. **125** (2006), 084110.

[30] Q. Du, V. Faber and M. Gunzburger, *Centroidal Voronoi Tesselations: Applications and Algorithms*, SIAM Review **41** (1999), 637–676.

[31] R. Seidel, *Exact upper bounds for the number of faces in d-dimensional Voronoi diagrams*, Applied Geometry and Discrete Mathematics - The Victor Klee Festschrift (P. Gritzmann and B. Sturmfels, ed.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 1991, pp. 517–529.

[32] S. Park, M.K. Sener, D. Lu and K. Schulten, *Reaction paths based on mean first-passage times.*, J. Chem. Phys. **119** (2003), 1313–1319.

[33] S. Rakovic, P. Grieder, and C. Jones, *Computation of Voronoi Diagrams and Delaunay Triangulation via Parametric Linear Programming*, Tech. report, 2004, http://infoscience.epfl.ch/record/169767/files/cued_control_779.pdf.

[34] C. Schütte, *Conformational Dynamics: Modelling, Theory, Algorithm, and Application to Biomolecules.*, 1999, Habilitation Thesis, Department of Mathematics and Computer Science, Freie Universität Berlin.

[35] H.A. Simon and A. Ando, *Aggregation of variables in dynamic systems*, Econometrica **29** (1961), 111–138.

[36] T. Hiroshima, Y. Miyamoto and K. Sugihara, *Another Proof of Polynomial-Time Recognizability of Delaunay Graphs*, IEICE Trans. Fund. **E83-A** (2000), 627–638.

[37] M. Weber, *Meshless methods in conformation dynamics*, 2006, Doctoral Thesis, Department of Mathematics and Computer Science, Freie Universität Berlin.

[38] _____, *A Subspace Approach to Molecular Markov State Models via an Infinitesimal Generator*, 2010, Habilitation Thesis, Department of Mathematics and Computer Science, Freie Universität Berlin.

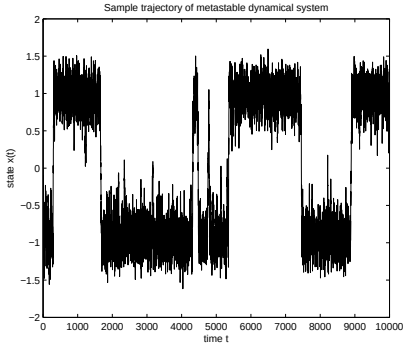

Figure 1: Sample trajectory of a metastable dynamical system. Metastable sets of the dynamical system are the intervals $(-1.7, 0.2)$ and $(0.2, 1.5)$. In each metastable set, the dynamical system evolves at a smaller time scale, i.e. the dynamical system exhibits very rapid fluctuations. On the other hand, the dynamics between the metastable sets occurs at a larger time scale - the system appears to be a Markovian process which jumps from one metastable set to another at a slower rate.

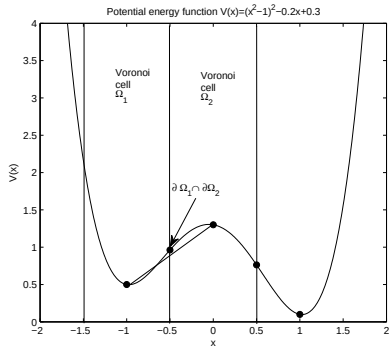

Figure 2: Potential energy function, with Voronoi regions $\Omega_1 = [-1.5, -0.5]$ and $\Omega_2 = [-0.5, 0.5]$ generated by $g_1 = -1$ and $g_2 = 0$, and the surface of intersection $\partial\Omega_1 \cap \partial\Omega_2 = \{-0.5\}$. The potential energy function between $g_1 = -1$ and $g_2 = 0$ is approximated by the average of $V(-1)$ and $V(0)$.

	Gauss-Legendre	9×9 Mesh	81 centers CVT	18×18 Mesh	324 centers CVT
$w_c(1)$	0.0024	0.0023	0.0028	0.0025	0.0025
$w_c(2)$	0.4987	0.4990	0.5360	0.4988	0.5157
$w_c(3)$	0.4987	0.4987	0.4612	0.4987	0.4819
$1 \to 2 : 1 \to 3$	1:1	0.8728:1	0.5647:1	1.1130:1	0.9059:1
$2 \to 1 : 2 \to 3$	1.7647:1	0.4211:1	0.4118:1	1.0028:1	1.6754:1
$3 \to 1 : 3 \to 2$	1.7647:1	0.2857:1	0.75:1	0.9146:1	1.3967:1

Table 1: Comparison for discretizations. Data set: points on 201×151 mesh. Inverse temperature $\beta = 3.34$.

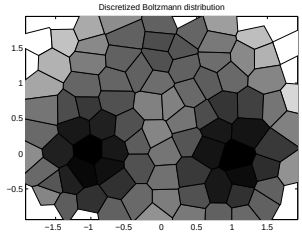

Figure 3: Discretized Boltzmann distribution corresponding to Eq. (20).

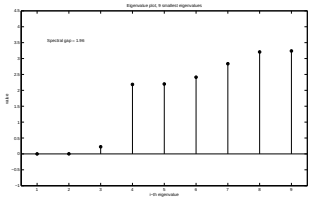

Figure 4: Absolute value, first 9 eigenvalues of Q''. Spectral gap $|\theta_4| - |\theta_3| \approx 1.96$.

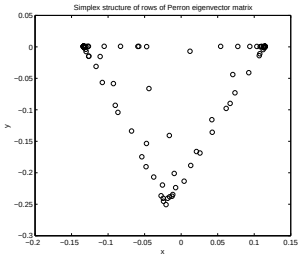

Figure 5: Rows of Perron eigenvector matrix X are close to a simplex.

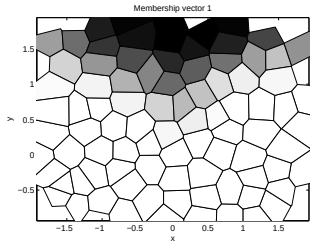

Figure 6: Conformation C_1 at $(0, 5/3)$. Dark regions belong strongly to C_1.

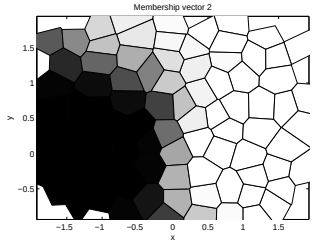

Figure 7: Conformation C_2 at $(-1, 0)$. Dark regions belong strongly to C_2.

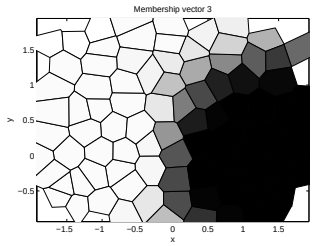

Figure 8: Conformation C_3 at $(1, 0)$. Dark regions belong strongly to C_3.

	$\beta = 1.67$	$\beta = 6.67$
$w_c(1)$	0.0500	0.0000
$w_c(2)$	0.4894	0.4826
$w_c(3)$	0.4606	0.5174
$1 \to 2 : 1 \to 3$	0.8790:1	0.9674:1
$2 \to 1 : 2 \to 3$	1.4969:1	3.1027:1
$3 \to 1 : 3 \to 2$	1.0757:1	2.6857 :1

Table 2: Entropic effects due to different temperatures $\beta = 1.67$ and $\beta = 6.67$, based on results using 324-center CVT and 202×151 mesh-based data set as in Table 1. Decreasing the temperature (changing from $\beta = 1.67$ to $\beta = 6.67$) leads to increases in the ratios in the last two rows because molecules have on average less energy at lower temperatures, so fewer molecules can overcome the high energy barrier between C_2 and C_3 and as a result prefer to transition into C_1.

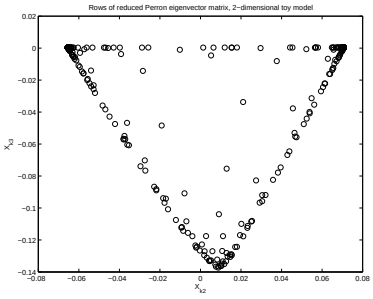

Figure 9: Dimension 2 model, rows of Perron eigenvector matrix. Clear simplex structure.

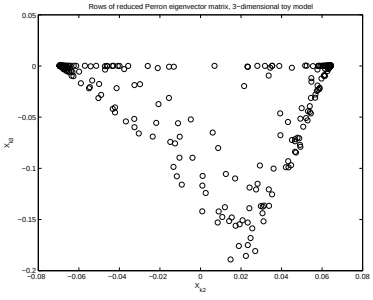

Figure 10: Dimension 3 model, rows of reduced Perron eigenvector matrix. Observable simplex structure.

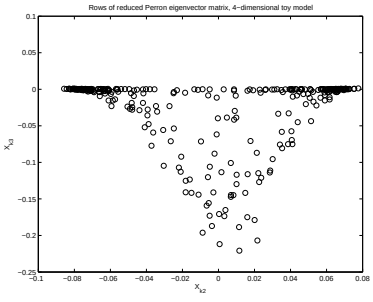

Figure 11: Dimension 4 model, rows of reduced Perron eigenvector matrix. Simplex structure difficult to distinguish.

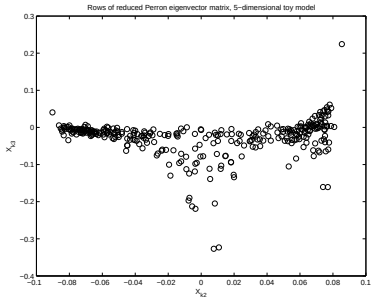

Figure 12: Dimension 5 model, rows of reduced Perron eigenvector matrix. No clear simplex structure.

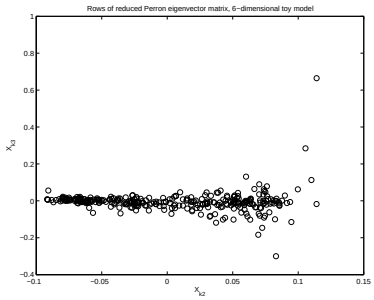

Figure 13: Dimension 6 model, rows of reduced Perron eigenvector matrix. No clear simplex structure.

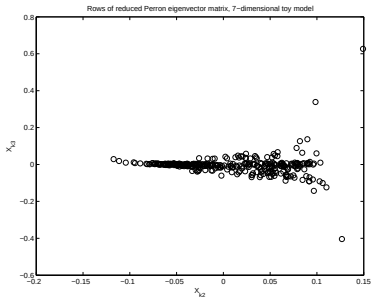

Figure 14: Dimension 7 model, rows of reduced Perron eigenvector matrix. No clear simplex structure.