MSM/RD: Coupling Markov state models of molecular kinetics with reaction-diffusion simulations

Manuel Dibak[1,†], Mauricio J. del Razo[1,†], David De Sancho[2], Christof Schütte[1] and Frank Noé[1,a)]

Abstract: Molecular dynamics (MD) simulations can model the interactions between macromolecules with high spatiotemporal resolution but at a high computational cost. By combining high-throughput MD with Markov state models (MSMs), it is now possible to obtain long-timescale behavior of small to intermediate biomolecules and complexes. To model the interactions of many molecules at large lengthscales, particle-based reaction-diffusion (RD) simulations are more suitable but lack molecular detail. Thus, coupling MSMs and RD simulations (MSM/RD) would be highly desirable, as they could efficiently produce simulations at large time- and lengthscales, while still conserving the characteristic features of the interactions observed at atomic detail. While such a coupling seems straightforward, fundamental questions are still open: Which definition of MSM states is suitable? Which protocol to merge and split RD particles in an association/dissociation reaction will conserve the correct bimolecular kinetics and thermodynamics? In this paper, we make the first step towards MSM/RD by laying out a general theory of coupling and proposing a first implementation for association/dissociation of a protein with a small ligand ($A + B \rightleftharpoons C$). Applications on a toy model and CO diffusion into the heme cavity of myoglobin are reported.

†Equal contribution
[1]Freie Universität Berlin, Department of Mathematics and Computer Science, Arnimallee 6, 14195 Berlin, Germany
[2]Kimika Fakultatea, Euskal Herriko Unibertsitatea (UP-V/EHU), and Donostia International Physics Center (DIPC), P.K. 1072, 20080 Donostia, Euskadi, Spain
a)Corresponding author. Electronic mail: frank.noe@fu-berlin.de.

I. INTRODUCTION

Life processes such as cellular signaling, control and regulation arise from complex interactions and reactions between biomolecules. A fundamental challenge of understanding and controlling life processes is that they are inherently multiscale – cellular signaling alone involves 6 orders of magnitude in lengthscales (0.1 nanometers to 100 micrometers) and 18 orders of magnitude in timescales (femtoseconds to hours). Unfortunately, these scales are tightly coupled – a single-point mutation in a protein can disturb the biochemical interactions such that this results in disease or death of the organism. No single experimental or simulation technique can probe all time- and lengthscales at a resolution required to understand such a process comprehensively.

In computer simulations, this dilemma can be mitigated by multiscale techniques – different parts of the system are described by a high-resolution and a low-resolution model, and these parts are coupled to give rise to a hybrid simulation. A famous example of such a multiscale model in biophysical chemistry is the coupling of quantum mechanics and molecular mechanics (QM/MM) [1]. Here we lay the foundations for a hybrid simulation technique that couples two scales that are particularly useful to model intracellular dynamics: a Markov state model (MSM) of the molecular dynamics (MD) scale that describes structural changes of biomolecules and their complexes, and the reaction-diffusion scale that describes diffusion, association and dissociation on the lengthscale of a cell. We call this approach MSM/RD, due to the combination of the simulation models chosen at these scales:

1. MSMs of the molecular scale: MD simulation allows us to probe molecular processes at atomic detail, but its usefulness has long been limited by the sampling problem. Recently, the combination of hard- and software for high-throughput MD simulations [2–5] with MSMs [6–8] has enabled the extensive statistical description of protein folding and conformation changes [9–12], as well as the association of proteins with ligands [13–17] and even other proteins [18]. Using multi-ensemble Markov models (MEMMs) [19–22], MSMs can be derived that even capture the kinetics of ultra-rare events beyond the seconds timescale at atomistic resolution [23, 24]. MSM approaches can thus model the long-lived states and transition rates of molecular detail interactions, but the cost of atomistic MD sampling limits them to relatively small biomolecules and complexes.

2. Reaction-diffusion (RD) scale: While atomic detail is relevant for some processes that affect the cellular scale, it is neither efficient nor insightful to maintain atomic resolution at all times for cellular processes. We choose particle-based reaction-diffusion (PBRD) dynamics kinetics as a reference model for the cellular scale. PBRD simulates particles, representing individual copies of proteins, ligands or other metabolites. Particles move in space via diffusion and reactive species will react with a probability according to their reaction rate when being close by. Here, a reaction may represent molecular processes such as binding, dissociation, conformational change, or actual enzymatic reactions. PBRD acknowledges that chemical reactions are in-

herently discrete and stochastic in nature [25], and that diffusion in cells is often not fast enough to justify well-stirred reaction kinetics [26–28]. A large number of recent software packages and codes implement some form of PBRD [29–36], see also the reviews [37, 38]. Hydrodynamic interactions at this scale could be incorporated by particle-based coupling terms [39, 40]. The effect of crowders and complicated boundaries such as membranes on the particle diffusion can be represented by including interaction forces on the RD scale [34].

In the limit that the conformational transitions of all molecules are fast, the MSM dynamics of each molecule effectively averages, and the interaction between the molecules (e.g. association) occurs with suitably averaged rates, reducing the problem to PBRD. However, when the lifetimes of some conformations are long compared to the typical time between two molecular interactions, or even the time between successive rebinding events of two molecules, the conformation dynamics of molecules described by the MSM part couples with the RD dynamics. MSM/RD opens up the possibility to simulate and analyze such effects quantitatively. For example, bimolecular binding rates from MD-derived MSMs can be inaccurate due to periodic boundary effects and a short-lived dissociated state in comparison to the MSM lag-time [18]. MSM/RD can overcome these issues by extending the diffusion domain available lessening the periodic boundary effects and increasing the lifetime of the dissociated state.

The ultimate aim of MSM/RD is to produce an efficient multiscale simulation that reproduces the essential statistical behavior of a practically unaffordable large-scale MD simulation by employing only statistics obtained from simulations of the constituent biomolecules in small solvent boxes. As developing a full theory involving rotational diffusion, three- or more-body interactions, hydrodynamics will be highly complex, we here aim to make a first step towards this goal by coupling MSM and RD scales for bimolecular systems without large-scale hydrodynamic interactions.

We first derive a theory of MSM/RD for bimolecular systems, as depicted in Fig. 1. When the two molecules are far from each other, they both undergo a diffusion process. When they come close to each other, molecular interactions, modeled with MD-derived MSMs, need to be taken into account. We further develop an algorithm to couple the MSM and RD scales for the special case of a protein interacting with a ligand, which is one of the main advances in this paper. This is not a trivial undertaking since one needs to solve two problems: to couple the MSM and RD part in such a way that the correct macroscopic rates and equilibrium probabilities are recovered, and to develop a suitable MSM discretization such that this coupling can be made. We demonstrate the validity of our theory and algorithms on a toy model of protein-ligand interaction and on binding of carbon monoxide to myoglobin.
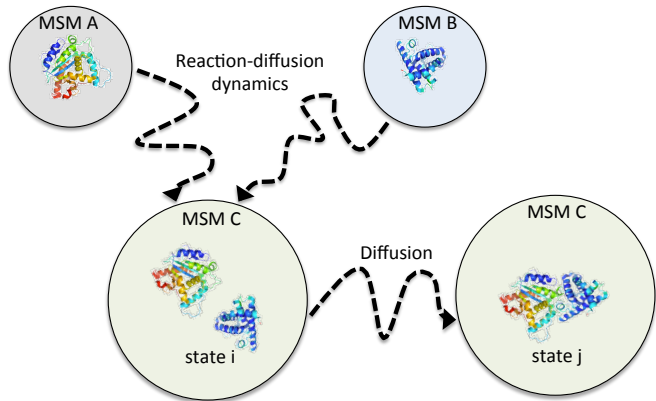
FIG. 1: Sketch of the MSM/RD scheme. When molecules A and B are not in close proximity, they diffuse freely. When A and B are close, they merge into a complex particle C which itself diffuses and whose internal dynamics are encoded by coupled MSM state transitions. When the molecules transition into a dissociated state, they are again separated into two separately diffusing particles A and B with initial positions depending on the last MSM state. Note that in the dissociated state, molecules A and B could also potentially undergo conformational changes encoded in independent MSM state transitions.

In related work, [41, 42] have coupled MD with a diffusion scheme. The work [43] further incorporates milestoning theory [44] to compute the local kinetic information in terms of transitions between milestones via short MD runs. In contrast with their work, we do not employ direct MD simulations at the "small" scale, but represent the small scale by an MSM as this allows us to operate on roughly the same timesteps for the small and the large scales. Other works have proposed alternative schemes to couple random walks (MSMs) with Brownian diffusion schemes, some examples can be seen in [45–47]. However, these works focus on specific contexts that are not directly applicable for coupling MD-derived MSMs with reaction-diffusion schemes.

II. MSM/RD: COUPLING MARKOV STATE MODELS AND REACTION-DIFFUSION

We develop a theoretical description for MSM/RD. The relevant scenarios for MSM/RD can be classified by the number of interacting particles, or the related reaction order:

1. First-order reactions: isolated diffusing particles can be modeled by an MSM obtained from MD simulations in a solvent box. The MSM directly translates into a set of unimolecular reactions that can be implemented in standard PBRD software. As long as the particles don't interact, the only effect of different states on the dynamics are changes between different diffusion constants/tensors.

2. Second-order reactions: interactions between two molecules that can be modeled as bimolecular re-

actions including protein-ligand or protein-protein association ($A + B \rightarrow C$). As soon as the complex C has been formed, its dynamics may be described by state transitions of an MSM of the complex.

3. Higher-order reactions: simultaneous interactions between more than two molecules.

In this work, we will focus on second-order reactions. First-order reactions are trivial state changes of a particle that are occurring as part of the MSM dynamics. Consistent with current conventions in PBRD frameworks, we follow the convention of breaking down higher-order reactions to second-order reactions, although in Sec. V we suggest possible extensions to treat these explicitly.

In order to derive the theory for second-order reactions, we concentrate on the dynamics of two molecules, A and B. For the sake of simplicity, we assume the two molecules do not have conformational changes of their own, so they can only diffuse and interact with each other. However, it is straightforward to extend MSM/RD to include conformational changes (first-order reactions) coupled with second-order reactions.

A. The *ground truth* model with full dynamics

Ground truth is a term often used in machine learning that refers to a reference model with respect to which modeling errors are measured. In the present context, the *ground truth* model contains the two (or more) solute molecules whose interactions will be later approximated by an MSM in a *large-scale* simulation, i.e. a simulation box that is not truncated after a small solvent boundary as customary for MD simulation. Importantly, there is no universally correct ground truth, but this model employs the MD simulation setup and dynamical model chosen by the user for the modeling task at hand. This choice includes the MD force field, solvation conditions and ion concentration, the protonation state at the pH of interest or even constant-pH simulations [48], the treatment of electrostatics, the thermostat, the integrator and time step, etc.

If such a large-scale model were simulated for a long time or with many trajectories, it would give rise to statistical properties of the solute molecules that we want to reproduce, such as their equilibrium constants and association rates. However, such a simulation is in general inefficient or infeasible, and our aim is that to reproduce its statistical properties using an MSM/RD model that is parametrized only using small MD simulations of the constituent solute molecules and complexes.

For simplicity, we derive the MSM/RD theory using all-atom explicit solvent MD simulations with a Langevin thermostat as the ground truth, as this setup is frequently used for MD simulations. However, the MSM/RD results apply more generally, e.g. to different choices of thermostats or integrators, as the MSM limit

for long-time description of the dynamics and the over-damped limit for long-time and large-scale description of the solute transport are achieved from a large family of ground truth models.

Langevin dynamics evolve as:

$$m_k \frac{d^2}{dt^2} x_k(t) = -\nabla_k U(\mathbf{x}_t) - \gamma_k \frac{d}{dt} x_k + \sqrt{2k_B T \gamma_k} \boldsymbol{\xi}_k(t),$$
(1)

where x_k represents the three-dimensional position of the k^{th} atom in the system (including the solvent), $\mathbf{x}_t = [x_1(t), \dots, x_k(t), \dots, x_N(t)]$, N the total number of atoms, U is the potential energy and $-\nabla_k U$ is the force acting on the k^{th} particle, m_k is the k^{th} particle mass, γ_k is the k^{th} damping coefficient, and $\boldsymbol{\xi}_k(t)$ is a Gaussian random force such that the expectations of its components satisfy $E[\xi_{k,i}(t)] = 0$ (zero mean) and $E[\xi_{k,i}(t)\xi_{k,j}(s)] = \delta_{ij}\delta(t-s)$ (white noise) with $k_B T$ being the thermal energy. In simulations, we use finite-time-step approximations of (1) and use it to generate stochastic trajectories. For the theoretical analysis, it is more useful to look at the ensemble dynamics, i.e., the propagation of probability densities in time. For this, we can ask: If we start the dynamical system in phase space point $\mathbf{y}$ and let it run, with which probability will we find it in a point $\mathbf{x}$ a time τ later? We call this probability the transfer probability $p(\mathbf{y} \rightarrow \mathbf{x}; \tau)$, and we will use it to describe the action of the ground truth dynamics [49]. The transfer probability $p(\mathbf{y} \rightarrow \mathbf{x}; \tau)$ subsumes the full complexity of the MD model, including interaction energies of the molecules with each other and external fields, and it can be constructed regardless of which thermostat or integrator is used. The propagation of probability densities $\varrho(\mathbf{x}; t)$ in time is formally described by the propagator $\mathcal{P}_\tau$:

$$\varrho(\mathbf{x}; t+\tau) = \mathcal{P}_\tau \varrho(\mathbf{x}; t)$$
$$= \int p(\mathbf{y} \rightarrow \mathbf{x}; \tau)\varrho(\mathbf{y}; t) \, d\mathbf{y} \qquad (2)$$

We want to find an efficient algorithm to approximate these dynamics. More specifically we want to approximate certain aspects of these dynamics, such as the long-time behavior.

It is often useful to consider densities relative to the stationary density $\pi(\mathbf{x})$ given by

$$u(\mathbf{x}; t) = \frac{\varrho(\mathbf{x}; t)}{\pi(\mathbf{x})},$$

which defines the propagator relative to the stationary density, or transfer operator [49]:

$$u(\mathbf{x}; t+\tau) = \mathcal{T}_\tau u(\mathbf{x}; t)$$
$$= \int \frac{\pi(\mathbf{y})}{\pi(\mathbf{x})} p(\mathbf{y} \rightarrow \mathbf{x}; \tau)u(\mathbf{y}; t) \, d\mathbf{y}$$
$$= \int p(\mathbf{x} \rightarrow \mathbf{y}; \tau)u(\mathbf{y}; t) \, d\mathbf{y} \qquad (3)$$

The third row follows from detailed balance. For reversible systems, where detailed balance is fulfilled, $\mathcal{T}_\tau$ is often called backward propagator, as it appears to evolve densities backward in time.

We will now introduce a scale separation by treating molecules A and B different when they are close (interacting) and far apart (non-interacting). More specifically these scales are defined by the distance between the centers of mass of A and B, r_{AB}:

1. MSM domain: molecules are in the *interaction* region $I = \{\mathbf{x} \mid r_{AB}(\mathbf{x}) < R\}$.

2. RD domain: molecules are in the *outside* region $O = \{\mathbf{x} \mid r_{AB}(\mathbf{x}) \geq R\}$.

The definition of R will be investigated later. Next, we take a closer look at the dynamics valid in these respective domains.

B. Markov state models for interacting molecules

We consider molecules that are closer than R to be interacting, hence we call the corresponding subset of state space I. The kinetics in I are fully described by Eq. (3), which can be approximated by an MSM derived from a MD simulation that fully includes I (usually plus some extra space, because MD simulations typically employ periodic rather than spherical boundary conditions). We implicitly assume that the interaction forces between proteins or protein-ligand pairs have decayed to zero at distances R or greater. Note that this assumption requires that the MD simulation conducted to parametrize an MSM/RD model has a sufficiently large simulation box, suitable electrostatics treatment and solvation conditions (ions etc.) such that in the dissociated state the solutes can be in any orientation without significantly interacting with each other or with their periodic images.

The interaction region I will here be approximated by an MSM. We perform a spectral decomposition of (3), assuming that there exists a unique stationary density π and the dynamics obey detailed balance. Furthermore, we truncate the spectral decomposition after a finite number of k terms:

$$\varrho(\mathbf{x};\, t+\tau) \approx \pi(\mathbf{x}) \sum_{i=1}^{k} \lambda_i(\tau) \langle \psi_i(\mathbf{x}), \varrho(\mathbf{x};\, t) \rangle \psi_i(\mathbf{x}) \quad (4)$$

Here, $\langle , \rangle$ denotes the scalar product with respect to $\pi(\mathbf{x})$, ψ_i are the eigenfunctions of $\mathcal{T}_\tau$ and its leading eigenvalues have the form

$$\lambda_i(\tau) = \mathrm{e}^{-\tau/t_i},$$

where t_i is a characteristic relaxation timescale. The truncation after k terms in Eq. (4) assumes that τ is long compared to t_{k+1}, where $3t_k < \tau$ is sufficient for practical purposes. Most of the solvent dynamics correspond to the fast coordinates that are averaged out (large-scale hydrodynamics are not part of the MSM term, while solvent molecules with long-lived interactions with the solute molecule can be considered to be part of that solute molecule). Now we can perform a Galerkin projection of the transfer operator by discretizing the phase space using basis functions $\chi_i(\mathbf{x})$, $i = 1, ..., n$. In MSMs, these are characteristic functions

$$\chi_i(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in S_i \\ 0 & \mathbf{x} \notin S_i \end{cases}$$

where the S_i form a complete partition of phase space, i.e. $\Omega = \{S_1 \cup S_2 \cup \cdots \cup S_n\}$. The phase space has now been discretized into a finite state space. The local densities become vectors simply given by

$$\pi_j = \int_{x \in S_j} \pi(\mathbf{x}) d\mathbf{x}, \qquad \varrho_j(t) = \int_{x \in S_j} \varrho(\mathbf{x};\, t) d\mathbf{x}.$$

Furthermore, we want the transfer operator to be approximated by a matrix. We can obtain this matrix by noting that the eigenfunctions of the transfer operator also become vectors in state space

$$\psi_i^j = \frac{1}{\int_\Omega \chi_j(\mathbf{x}) d\mathbf{x}} \int_{x \in S_j} \psi_i(\mathbf{x}) d\mathbf{x}.$$

Inserting these equations into (4), and rewriting it in matrix form, we obtain the Chapman-Kolmogorow equation

$$\boldsymbol{\varrho}(t + \tau) = \mathbf{T}^\top(\tau)\, \boldsymbol{\varrho}(t), \quad (5)$$

with $\lambda_i(\tau)$ and $\boldsymbol{\psi_i} = [\psi_i^1, \ldots, \psi_i^n]$ the i^{th} eigenvalue and eigenvector of the transition probability matrix $\mathbf{T}(\tau)$, respectively, and with $\boldsymbol{\varrho}(t) = [\varrho_1(t), \ldots, \varrho_n(t)]$ the probability mass function [6].

Estimating a high-quality MSM from MD simulation data can be quite complex. It typically involves (i) mapping the MD coordinates to a set of features, such as residue distances, contact maps or torsion angles, (ii) reducing the dimension to slow collective variables (CVs), often based on the variational approach or conformation dynamics [50, 51] or its special case time-lagged independent component analysis (TICA) [52, 53] – see [54, 55] for an overview, (iii) optionally, embedding the resulting coordinates in a metric space whose distances correspond to some form of dynamical distance [56, 57], (iv) discretizing the result space using data-based clustering [7, 58, 59], typically resulting in 100-1000 discrete states, and (v) estimating the transition matrix $\mathbf{T}(\tau)$ or a transition rate matrix $\mathbf{K}$ with $\mathbf{T}(\tau) = \exp{(\tau\mathbf{K})}$ at some lag time τ, and validating it [6, 10, 60, 61]. Finally, the MSM may be coarse-grained to few metastable states [62–65]. The MSM software packages PyEMMA [58] and MSM-builder [59] can greatly help to simplify this process and make it reproducible.

In the case where there are well-defined meta-stable regions in phase space, we can greatly reduce the number of states in the MSM. One way to simplify the MSM construction process above and to directly end up with a few-state MSM is to employ VAMPnets, where the complex MSM construction pipeline is replaced by a neural network that is trained using the variational approach for Markov processes [66]. Alternatively, one can replace the discretization step (iv) above by employing a core set approach that was derived in [60] and further analyzed in [67]. The essential idea is to define the states as cores around the metastable regions. Due to the metastability, the probability of finding the system outside of the metastable regions is very small, so to a good approximation the kinetics can be described as a core-to-core jump process [67]. This approach will be employed throughout this paper and and explained in more detail in Sec. III.

C. Reaction-diffusion dynamics for noninteracting molecules

When molecules are far apart, and thus in the RD domain defined by $r_{AB}(\mathbf{x}) \geq R$, they are not directly interacting. As the dynamics of the two molecules are independent, it is convenient to only track the net diffusion of the centers of mass, $\boldsymbol{r}_A$ and $\boldsymbol{r}_B$. Furthermore, we assume that the dynamics in the RD domain can be tracked by coarse timesteps of at least Δt which exceeds the typical velocity autocorrelation time (picoseconds). At such timescales, the fast dynamics corresponding to the solvent are averaged out. It is possible that even longer timesteps are made using an event-based integration scheme such as first-passage kinetic Monte Carlo (FPKMC) algorithm, Green's function reaction dynamics (GFRD) or MD-GFRD [28, 31, 35, 36, 42]. At such timesteps, the Langevin equation (1) becomes an overdamped Langevin equation for the centers of mass of the two molecules, i.e. the motion is governed by pure diffusion:

$$\frac{d\boldsymbol{r}_A(t)}{dt} = \sqrt{2D_A}\boldsymbol{\xi}_A(t), \quad \frac{d\boldsymbol{r}_B(t)}{dt} = \sqrt{2D_B}\boldsymbol{\xi}_B(t), \quad (6)$$

where $\boldsymbol{\xi}_A(t)$ and $\boldsymbol{\xi}_B(t)$ are independent white noise vectors with each of their components satisfying $E[\xi_{K,i}(t)] = 0$ and $E[\xi_{K,i}(t)\xi_{K,j}(s)] = \delta_{ij}\delta(t-s)$. D_A and D_B are the net diffusion coefficients for the centers of mass, which can be obtained from MD simulations. In general, as we are tracking the center of mass, we also need to track the rotational diffusion of the molecules. However, as rotational diffusion is not relevant for the examples discussed in this manuscript, we refer to [68, 69].

In the present case, we can simply fix the frame of reference in $\boldsymbol{r}_A(t)$, assume the rotation of A is slower than the diffusion of B, which is true for protein-ligand systems, and fix the orientation of the axis to that of molecule A. We further assume that B is a small molecule such that

its orientation is not very relevant, as it will be the case in our implementation of the scheme. This simplifies Eqs. 6 into a simple diffusion in $\boldsymbol{r}_B$ only

$$\frac{d\boldsymbol{r}_B(t)}{dt} = \sqrt{2(D_B + D_A)}\boldsymbol{\xi}(t), \quad (7)$$

with the components of $\boldsymbol{\xi}(t)$ satisfying $E[\xi_i(t)] = 0$ and $E[\xi_i(t)\xi_j(s)] = \delta_{ij}\delta(t-s)$.

D. MSM/RD coupled dynamics

The present coupled model only considers interactions between up to two molecules. This is a frequent assumption in PBRD [30, 31, 34, 36] but may be restrictive from a molecular standpoint. We assume that simultaneous reactions between three or more molecules such as $A + B + C \rightarrow D$ can always be broken down into $A + B \rightarrow AB$; $AB + C \rightarrow D$ or other bimolecular pathways, and therefore focus on MSM/RD involving two molecules. In order to do the coupling, as the dynamics in the I and O region are given in terms of states and coordinates respectively, we need to recognize that $\mathbf{x}$ and $\mathbf{y}$ in the transfer density $p(\mathbf{y} \rightarrow \mathbf{x}; \tau)$ can be either coordinates $\mathbf{c}$ (center-of-mass position and perhaps orientation of the molecule) or states s (metastable regions in the coordinate space). In order to implement the coupling, we suggest defining two quantities:

- $p_{\text{entry}}[\mathbf{c}_t \rightarrow \mathbf{x}_{t+\Delta}; \Delta]$, transfer probability of starting in coordinates $\mathbf{c}_t$ just inside the MSM domain ($r_{AB}(\mathbf{c}_t) < R$) conditioned on hitting only one state $\mathbf{x}_{t+\Delta} = s_{t+\Delta}$ in the MSM domain (transition event) OR on exiting once the MSM domain $\mathbf{x}_{t+\Delta} = \mathbf{c}_{t+\Delta}$ (return event)

- $p_{\text{exit}}[s_t \rightarrow \mathbf{x}_{t+\Delta}; \Delta]$, transfer probability of starting in state s_t conditioned on exiting once the MSM domain $\mathbf{x}_{t+\Delta} = \mathbf{c}_{t+\Delta}$ (exit event) OR hitting once any other state ($\mathbf{x}_{t+\Delta} = s_{t+\Delta}$).

Once we know these transfer probabilities, we can introduce the basic algorithm where τ_{RD} and τ_{MSM} correspond to the diffusion and MSM time-step, respectively:

Input: Initial mode (RD or MSM), initial condition (coordinates $\mathbf{c}_0$ or state s_0, respectively) and $t = 0$:

While $t \leq t_{\text{final}}$:

1. If in RD mode:

   (a) Propagate $\mathbf{c}_t \rightarrow \mathbf{c}_{t+\tau_{\text{RD}}}$ by diffusion
   (b) Update time $t \mathrel{+}= \tau_{\text{RD}}$
   (c) If $r_{AB}(\mathbf{c}_t) < R$ (enter MSM domain):
      - Sample next event $(\mathbf{x}_{t+\Delta}, \Delta)$ from $p_{\text{entry}}[\mathbf{c}_t \rightarrow \mathbf{x}_{t+\Delta}; \Delta]$.

- If transition event:
  Map to state $s_{t+\Delta} = \mathbf{x}_{t+\Delta}$
  Update time $t \mathrel{+}= \Delta$
  Switch to MSM mode
- Else (return event):
  Map to coordinates $\mathbf{c}_{t+\Delta} = \mathbf{x}_{t+\Delta}$
  Update time $t \mathrel{+}= \Delta$

2. Else (MSM mode):

   (a) If $s_t \neq s_{t-\tau_{MSM}}$ or previous mode $\neq$ MSM mode:

   - Sample next event $(\mathbf{x}_{t+\Delta}, \Delta)$ from $p_{\text{exit}}[s_t \to \mathbf{x}_{t+\Delta}; \Delta]$.
   - If exit event:
     Map to coordinates $\mathbf{c}_{t+\Delta} = \mathbf{x}_{t+\Delta}$
     Update time $t \mathrel{+}= \Delta$
     Switch to RD mode and break current loop iteration

   (b) Propagate $s_t \to s_{t+\tau_{MSM}}$ using the MSM

   (c) Update time $t \mathrel{+}= \tau_{MSM}$

There are additional issues in specific scheme implementations, such as estimating the unknown conditional transfer probabilities, and choosing the MSM discretization and R such that the overall discretization error is small, among others. These issues are non-trivial and could potentially be tackled with different approaches. In order to quantify the accuracy of a given approach, we quantify how well is our scheme approximating the ground truth by comparing relevant macroscopic observables. We present one possible implementation of the scheme in Sec. III.

III. AN MSM/RD IMPLEMENTATION FOR PROTEIN-LIGAND SYSTEMS

Now we develop an implementation of the MSM/RD scheme for a special class of systems: the binding of a small ligand to a protein – a case that is relevant in the study of protein-drug binding kinetics [70]. While the theory described before is more general, implementations to more challenging systems such as protein-protein interaction will be treated in future contributions. We begin by considering the macromolecule A fixed at the origin with fixed orientation and the ligand B freely diffusing around it with an overall diffusion constant $D = D_A + D_B$. The macromolecule has several possible binding sites given by some interaction potential. In order to present the MSM/RD scheme in detail, we distinguish three different simulations:

1. **Reference simulation** (ground truth, if available): MD simulation of B and its interaction with A in a large spherical domain with radius R_s. Unfortunately, reference simulations of realistic systems are in general not computationally feasible

due to the time and lengthscales of the simulation. Nonetheless, reference simulations of simple systems are used to verify the MSM/RD scheme and validate its use in more complex systems.
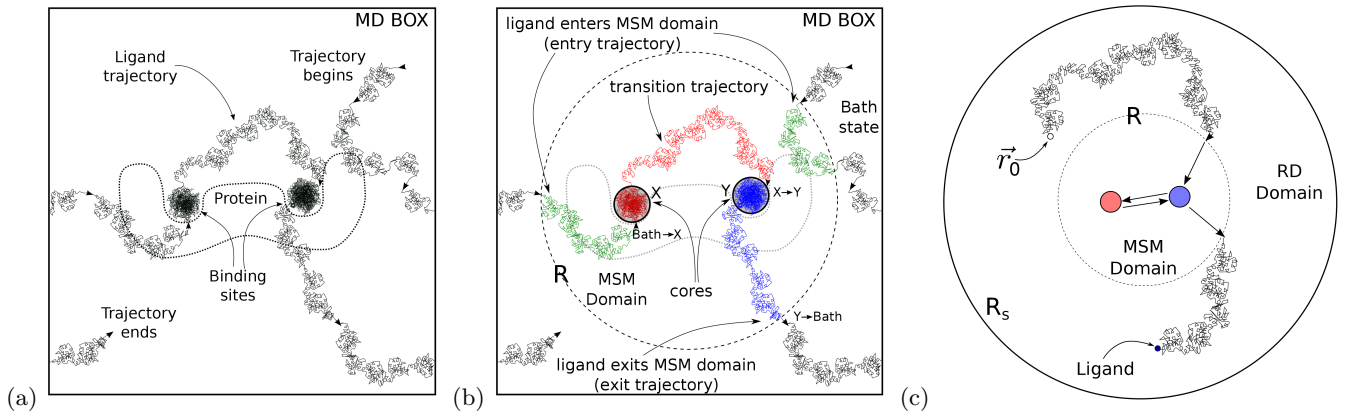
2. **Small-scale simulation** (MD simulation): analogous to the reference simulation with the difference that B is constrained to a small box with periodic boundary conditions, see Fig. 2a. As the potential is negligible outside this box, the main interaction dynamics are extracted from this simulation's data into an MSM. This simulation is used to parametrize the MSM/RD model.

3. **MSM/RD simulation** (hybrid model): couples the MSM for short-range interactions derived from the MD simulation 2. with a diffusion scheme for the long-range, see Fig. 2c. The goal of the scheme is to approximate the ground truth dynamics given by the reference simulation 1.

A. Estimation

In order to parametrize the MSM/RD scheme, we need to estimate quantities from the small-scale simulation 2. that characterize the state-to-state dynamics and the coupling between the MSM and RD domain. The state to state dynamics are estimated using an MSM, and the coupling is given in terms of entry and exit events from the MSM domain. These might happen on different timescales, so we would like to be free from the fixed time-step the MSM requires to be well equilibrated. Therefore, we use trajectory statistics for entry and exit events.

1. MSM

As a first step for the construction of the Markov model and MSM/RD parametrization, we need to find a discrete representation of the underlying data. In this work we use the core MSM approach [67], which requires the definition of cores as metastable regions of phase space. Cores are given by spherical domains around the metastable regions in the MD simulation and can be found using a clustering algorithm. In the core MSM approach a discrete trajectory is constructed by assigning the last visited state-index to each point in the trajectory. Note the trajectory may leave the core of a given state and re-enter multiple times without transitioning to other states. Using this discretization technique we truncate the discrete trajectory into three types of trajectories as shown in Fig. 2b: *i*) entry trajectories that start just inside the MSM domain and either leave the domain next or hit a core inside the domain. *ii*) transition trajectories that start in a state and hit another state as next event and *iii*) exit trajectories that start in

FIG. 2: Illustrations of a MD trajectory, its classification to extract the relevant dynamics and the MSM/RD scheme. **a)** Illustration of a trajectory of a ligand in a MD simulation within a box with periodic boundaries (small-scale simulation). Note that there are two metastable regions, e.g. binding sites on a protein, where the ligand stays for a longer time. **b)** Illustration of truncation and classification of the trajectory. The MSM domain is chosen so the interaction potential is effectively zero outside this region (bath state); the cores X and Y are chosen to represent the metastable regions in phase space. The truncated trajectories are classified into entry trajectories (green), transition trajectories (red) and exit trajectories (blue), which are used for the coupling in the MSM/RD scheme. In order to obtain the MSM for the MSM/RD scheme, the system is also classified into three states, the bath state and the two cores X and Y; it is also shown when the transition between these states occur along a trajectory. **c)** Representation of the MSM/RD scheme. The full trajectories from the MD-simulation are used to derive an MSM to model the dynamics in the MSM domain. The entry and exit trajectories from the MD-simulation are used to couple the Brownian dynamics in the diffusion domain with the dynamics in the MSM domain.

a state and leave the MSM domain as next event. These trajectories are used to estimate the transfer densities and to parametrize the MSM/RD simulation.

The MSM for the short-range interactions is built using the full discrete trajectories and the exit trajectories (Fig. 2a). We follow the methods from [6] to estimate a transition matrix $\mathbf{T}(\tau)$, where the entries are the transition probabilities $T_{i,j}$ from state i to j. Using the discrete trajectories, we create count matrices $C_{i,j}^{\text{full}}(\tau)$ from the complete data set and $C_{i,j}^{\text{exit}}(\tau)$ from the exit trajectories, which count all the transitions from state i to j at a lag time τ observed in the respective datasets. As the coupling between the MSM and RD domain is handled separately, the MSM dynamics only accounts for transitions amongst the cores and therefore the counts arising from exit trajectories have to be subtracted

$$C_{i,j}(\tau) = C_{i,j}^{\text{full}}(\tau) - \delta_{i,j} C_{i,j}^{\text{exit}}(\tau), \qquad (8)$$

where $\delta_{i,j}$ denotes the Kronecker delta. We then use a maximum likelihood estimator to obtain a transition matrix from the given counts C_{ij}. Note that here we have chosen an irreversible estimator, as we can no longer assume that detailed balance holds for this count matrix.

2. Entering the MSM domain

The protocol to enter the MSM domain from the RD domain is constructed with the entry trajectories as defined above. It consists of generating a list $L_{\text{entry}} = \{\mathbf{c}_{\text{entry}}, \mathbf{x}_{\text{end}}, \Delta\}$ of all start coordinates $\mathbf{c}_{\text{entry}}$ (just inside the MSM domain) and endpoints $\mathbf{x}_{\text{end}}$ of entry trajectories and their corresponding times Δ. The endpoints

may be either MSM states or coordinates in the RD domain, see Fig. 2b. The ensemble of trajectories in this list estimates the conditional transfer probability $p_{\text{entry}}[\mathbf{c}_t \rightarrow s_{t+\Delta}; \Delta]$ (Sec. II D) for several times Δ. In the MSM/RD simulation samples are drawn from this list of entry points.

3. Exiting the MSM domain

For each state $\mathbf{s}$ of the MSM, we collect all exit and transition trajectories and save their end coordinate or state along with their respective exit time in the lists $L_{\text{exit},s} = \{\mathbf{c}_{\text{exit}}, \Delta\}$ and $L_{\text{trans},s} = \{s_{\text{trans}}, \Delta\}$. The ensemble of trajectories in these list estimates the conditional transfer probability $p_{\text{exit}}[s_t \rightarrow \mathbf{c}_{t+\Delta}; \Delta]$ (Sec. II D) for several times Δ. The probability of an exit event $P_{\text{exit},s}$ is simply estimated as the ratio of exiting trajectories over the total numbers of trajectories,

$$P_{\text{exit},s} = \frac{\# \text{ of trajectories in } L_{\text{exit},s}}{\# \text{ of trajectories in } L_{\text{exit},s} \text{ and } L_{\text{trans},s}}. \qquad (9)$$

The MSM/RD scheme for this implementation is based on the scheme from Sec. II D and is shown in the Appendix.

B. Verification of the MSM/RD scheme

In order to verify the MSM/RD scheme, we use systems where a reference simulation is available. We verify the internal dynamics by comparing the first passage

times (FPTs) distributions and mean first passage times (MFPTs) for each pair of metastable states within region I between the MSM/RD and reference simulations. We estimate the ground truth MFPTs by computing the FPTs $t_{i,j}^{\text{ref}}$, where the initial conditions are chosen as the minima $\boldsymbol{\mu}_i$ and the system is propagated following the reference simulation until hitting state j (conditioned on not leaving the MSM domain). For the MSM/RD scheme, we compute the FPTs $t_{i,j}^{\text{MSM}}$ by placing the particle in state i and propagating the system following the MSM/RD scheme until state j is hit. If the particle exits the MSM domain before reaching state j, the trajectory is not taken into account. When a sufficiently large sample is generated, we can estimate the distributions of FPTs by histograms. The MFPTs are estimated as $\tau_{ij}^{\text{ref}} = \overline{t_{ij}^{\text{ref}}}$ and $\tau_{ij}^{\text{MSM}} = \overline{t_{ij}^{\text{MSM}}}$, respectively. The MFPT relative error between the MSM/RD and the reference simulations is estimated as

$$(E_{\text{rel}})_{\text{ij}} = \frac{\tau_{ij}^{\text{ref}} - \tau_{ij}^{\text{MSM}}}{\tau_{ij}^{\text{ref}}}. \tag{10}$$

In order to verify the coupling between the RD and MSM domain, we also estimate and compare the unbinding rate, binding rate and equilibrium constant. The two latter are calculated for different particle concentrations c by fixing the radius R_s of the simulation domain such that $c = 1/V_{\text{RD}}$, with V_{RD} the volume of the RD domain.

IV. RESULTS

In this section, we implement the MSM/RD scheme from Sec. III in two systems. The first is a simple model of a ligand diffusing in a potential landscape, which is used to verify that the MSM/RD scheme reproduces the correct dynamics. The second corresponds to a more realistic MD system, where we study the binding of carbon monoxide to myoglobin.

A. Ligand diffusion in potential landscape

We implement the MSM/RD scheme in a simple model, where the reference simulation is available. The model consists of a ligand B under over-damped Langevin dynamics in a three-dimensional potential landscape

$$\frac{d\mathbf{x}(t)}{dt} = -\frac{1}{\gamma}\nabla U(\mathbf{x}) + \sqrt{2D}\boldsymbol{\xi}(t), \tag{11}$$

with U the interaction potential with some macromolecule A fixed at the origin, γ the damping, and each component of the noise satisfies $E[\xi_i(t)] = 0$ and $E[\xi_i(t)\xi_j(s)] = \delta_{ij}\delta(t-s)$ with $D = k_B T/\gamma$ the diffusion coefficient. A trajectory density plot of the potential

landscape chosen is shown in Fig. 3a, and it consists of nine Gaussians with different depths and widths

$$U(\mathbf{r}) = -\sum_{i=1}^{9} s_i \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i), \tag{12}$$

where $\mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ denotes a Gaussian centered at minimum $\boldsymbol{\mu}_i$ with covariance matrix $\boldsymbol{\Sigma}_i$, s_i denotes a scale factor. The small-scale simulation consists of Euler-Maruyama numerical realizations of Eq. (11) under this potential constrained to a box with an edge length of 6 units with periodic boundary conditions. The reference simulation is analogous to the small-scale simulation with the difference that it uses a larger spherical domain with reflective boundary conditions at a range of radii corresponding to simulations at different ligand concentrations.

1. Parametrization of the MSM/RD scheme

We use a radius of $R = 2.5$ nm for the MSM domain (I region) since outside this domain the potential (Eq. 12) is essentially zero. We generate 120 small-scale simulation trajectories, each with a length of 10^7 steps, a time-step of $\Delta t = 10^{-4}$ ns, and sampled every tenth step. This results in a total simulation time of $t = 1.2 \cdot 10^5$ ns.

The cores are defined as spheres with radius 0.2 nm around the minima $\boldsymbol{\mu}_i$, and the count matrix of transition between cores is generated from the trajectories following Eq. (8). A maximum likelihood estimator (implemented in PyEMMA [58]) is then applied to the count matrix to yield the MSM. From the trajectories, we also generate the lists L_{entry}, $L_{\text{exit},s}$, $L_{\text{trans},s}$ and $P_{\text{exit},s}$, introduced in Secs. III A 2 and III A 3. We then estimate the timescales of the eigenmodes for different MSM lag times to test how well the underlying process is estimated by the MSM. The timescales have small variations for different lag times (Fig. 3d), which means the system can be considered Markovian for all lag times. However, we have to be careful not to choose the lag time too large, such that relevant fast timescales are neglected resulting in significant errors. For all further analyses, we consider a lag time of $\tau_{MSM} = 500\Delta t = 0.05$ ns to be an optimal compromise.

2. Comparison of dynamic properties

In order to compute the binding rate, we calculate the first passage time from a uniformly sampled location close to the boundary $r = R_S - \delta$ to any MSM state. We choose $\delta = 0.05$ nm and use 10^4 simulations to average and estimate the MFPT$_{\text{on}}$, from which we calculate the binding rate as $k_{\text{on}}^* = 1/\text{MFPT}_{\text{on}}$. This procedure is performed for both the MSM/RD and the reference simulation, and we observe excellent agreement between the two (Fig. 3e).
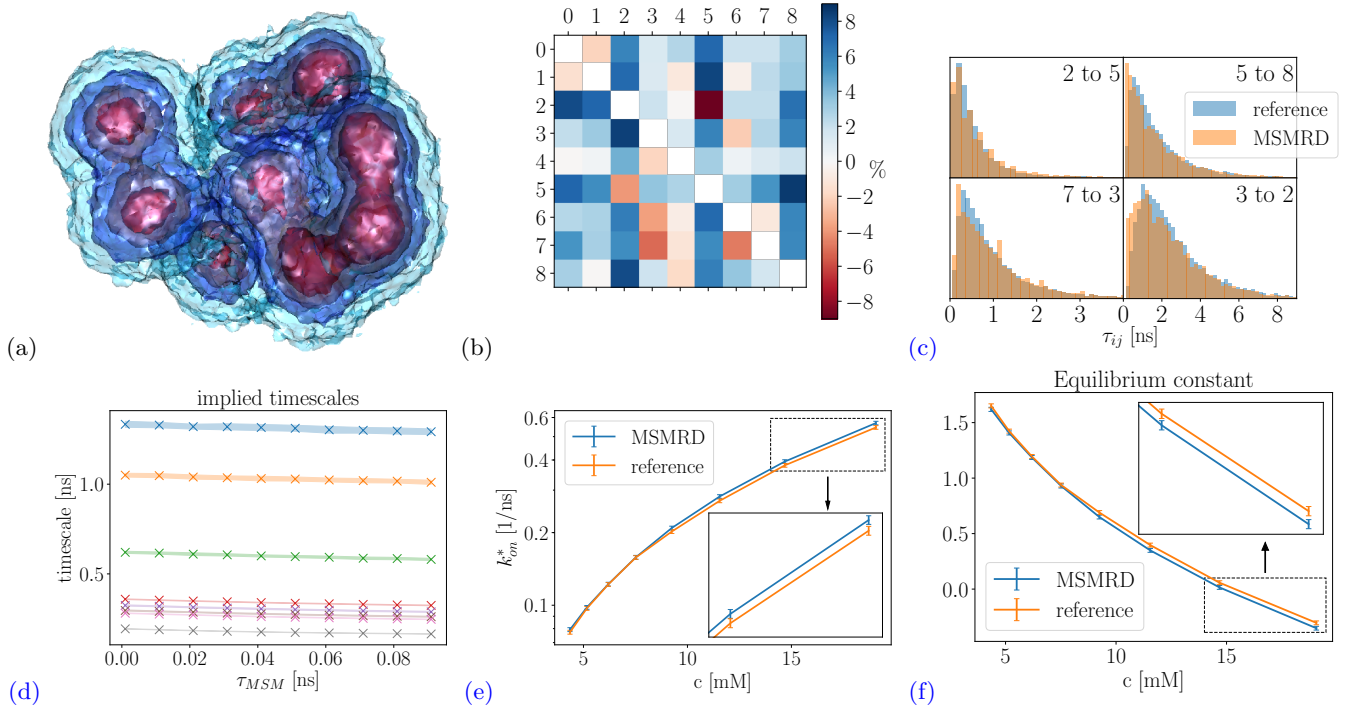
FIG. 3: Visualization and verification results for the simple model of ligand diffusing in a potential landscape. **(a)** Density plot of the position of the ligand in the three dimensional potential. Red indicates regions of higher density while blue indicates regions of lower density. **(b)** Relative error of MFPTs conditioned on not leaving the MSM domain between the MSM/RD and the reference simulation. **(c)** Comparison of first passage times distribution histograms for the transitions with the highest error in (b). The left pane corresponds to transitions with negative relative error, and the right pane to transitions with positive relative error. **(d)** Implied timescales of the MSM. The shaded area represents the standard deviation of the bootstrapping sample. We observe well converged timescales for all considered lag times. **(e)** The rate k_{on} as function of the concentration of the system for the MSM/RD and reference simulations. **(f)** Same plot as e) but for the logarithm of the equilibrium constant $\log(K_{eq})$. The error bars in (e) and (f) represent the 95% confidence interval using a bootstrapping approach.

For the unbinding rate, we consider the inverse process by starting in an MSM state and propagating the dynamics until crossing a boundary defined by a sphere with radius $2.7\,\mathrm{nm} > R$. We obtain a reference value of $0.402^{0.404}_{0.400}\,\mathrm{ns}^{-1}$ (Sub- and superscript indicate lower and upper bound of the 95% precentile) and an MSM/RD simulation value of $k_{\mathrm{off}} = 0.400^{0.402}_{0.398}\,\mathrm{ns}^{-1}$. We further compute the logarithm of the equilibrium constant $\log(K_{\mathrm{eq}}) = \log(k_{\mathrm{off}}/k^*_{\mathrm{on}})$ for both models and for the chosen values of concentrations, resulting in accurate reproduction of the reference values by the MSM/RD scheme(Fig. 3f). Thus we verify that the coupling between the MSM domain and the RD domain works consistently in the MSM/RD simulation scheme.

Next, we want to ensure that also the dynamics between the states inside the MSM are reproduced to a high accuracy. We compare MFPTs between all pairs of states conditioned on not leaving the MSM domain. In the reference simulation this is done by placing the particle at position $\boldsymbol{\mu}_i$ and propagating the system until state j is reached. If the particle leaves the MSM domain before reaching state j, this trajectory is discarded. For the MSM/RD simulation, we simply start in state i and

propagate until state j is hit, while discarding trajectories that leave the MSM domain. This procedure is repeated until 10^4 successful trajectories are found for both simulations, which are averaged to obtain the MFPTs. The relative errors are calculated with Eq. (10); all relative errors are below 9% (Fig. 3b). We further observe that negative errors arise for state pairs that are close together and thus have short passage times. For these transitions, we tend to overestimate the MFPT in the MSM/RD simulation as short processes are truncated in the MSM estimation. Moreover, we observe that the highest positive errors arise for transitions which are far apart. These are the hardest to sample since for these transitions there are a very high number of possible long and non-direct transition trajectories, which are less likely to be observed . We chose the four transitions with the highest relative error and compared their FPTs distribution histograms (Fig. 3c). Even though these transitions have the highest errors, we observe the distributions match well. Therefore, we verify MSM/RD scheme also describes the internal dynamics accurately.

B. Binding of CO to myoglobin

As an application of the MSM/RD scheme, we study the binding of carbon monoxide (CO) to myoglobin. Myoglobin is a globular protein which is responsible for the transport of oxygen in muscle tissue. The binding process of CO to myoglobin has recently been studied by de Sancho et al. [16], whose data we use to parametrize the MSM/RD scheme. The dataset consist of MD trajectories of 20 CO molecules and one myoglobin protein for a total simulation time of 500 ns. The MD simulation is confined to a periodic box with edge length of 5 nm. Despite the fact that only one CO molecule can reside in the binding pocket, the error of treating 20 CO molecules as being statistically independent is small within statistical uncertainty (see [16] for details). We therefore extract 20 independent CO trajectories, effectively increasing the total simulation time to 10 μs.

1. Parametrization of MSM/RD scheme

In order to parametrize the scheme, all frames are first aligned using the C_α atoms of the myoglobin as reference. On the aligned data, we run the density-based spatial clustering of applications with noise algorithm (DBSCAN) [71], which finds a total of 16 metastable regions/cores. The positions and size of the cores are shown in Fig. 4a, where it can be observed that the algorithm correctly identifies regions of high ligand density, including the myoglobin bound state indicated in red. The radius of the spherical cores is the radius at which 80 % of the datapoints that were assigned to the respective state are inside the core. Four states are discarded as they are not part of the largest connected set. As the simulation box had been set up to just contain the protein and a 1 nm solvent layer, we choose the largest MSM domain that still fits inside the box ($R = 2.5$ nm). Analogous to the previous example, we follow Sec. III to estimate an MSM for the close-range dynamics and generate L_{entry}, $L_{\text{exit},s}$, $L_{\text{trans},s}$ and $P_{\text{exit},s}$ to couple the dynamics in the two domains.

We compute the implied timescales for the MSM and choose a lag time of 150 ps where timescales are sufficiently converged (Fig. 4b). The diffusion constant is computed using the mean squared displacement (MSD) of the parts of the CO trajectories that are far from the protein, with $D = \Delta\text{MSD}(t)/6\Delta t$. We find a diffusion constant of $D_{\text{CO}} = 2.5\,\text{nm}^2\text{ns}^{-1}$, which is comparable to the experimental value which is in the range of $D_{\text{CO}} = 2.03\,\text{nm}^2\text{ns}^{-1}$ (at 20 $C°$) to $D_{\text{CO}} = 2.43\,\text{nm}^2\text{ns}^{-1}$ (at 30 $C°$) [72].

2. Comparison of dynamic properties

As in the previous example, we compute the binding rate by sampling positions sampled uniformly in the

(a)

(b)

(c)
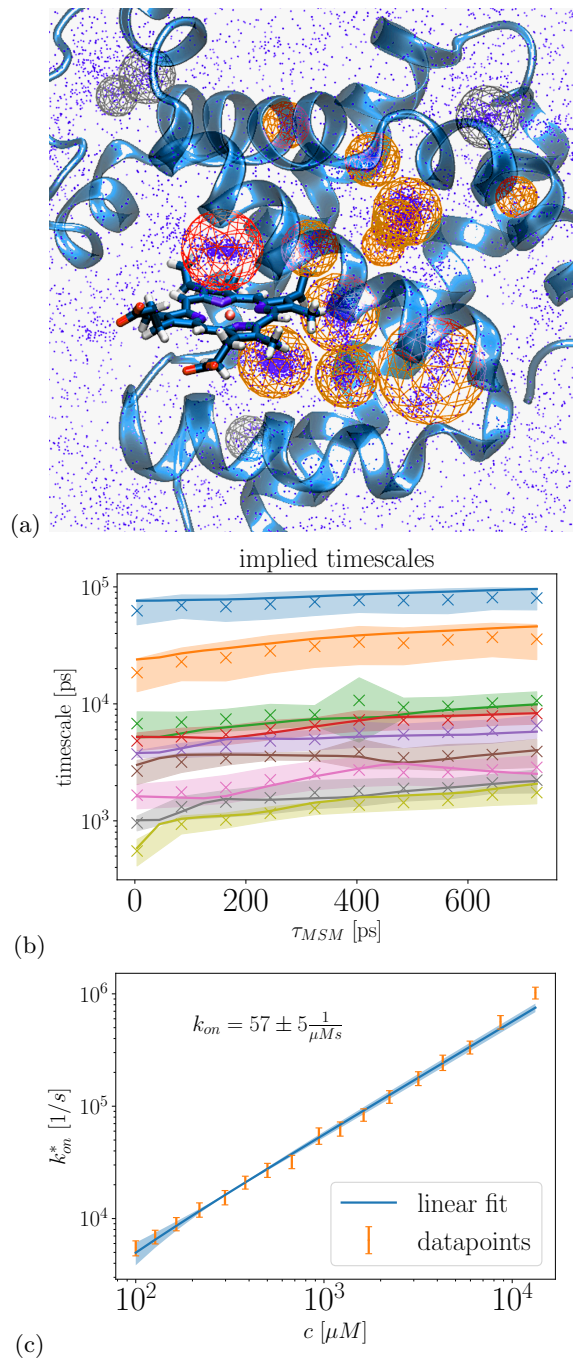
FIG. 4: Discretization and results of the CO-myoglobin system. **(a)** Definition of the cores (wire frame spheres) within the myoglobin. The red sphere indicates the bound state. The gray spheres correspond to the states that were not in the connected set and therefore discarded. The blue dots are positions of the CO molecules for every 50th frame in the vicinity of the protein. **b)** Implied timescales of the dynamics of the CO myoglobin system. The datapoints and shaded area denote the sample mean and standard deviation of the bootstrapping sample over the trajectories: from the 20 given trajectories we resample 20 with replacement. Over this sample we run our discretization process which returns a sample of timescales. The trajectory-samples which are not ergodic or do not lead to a connected count matrix are considered invalid and discarded. Solid lines are found using the full dataset. **c)** Reaction rate as estimated from multiple simulations at different concentrations.

RD domain and simulating the MSM/RD model until it reaches the bound state. For each concentration, 200 trajectories are run to estimate the binding rate k_{on}^*. These rates are plotted against the concentration and shown in Fig. 4c. The reaction rate $k_{on} = 57_{52}^{62} \, \mu M^{-1} s^{-1}$ is obtained as the slope of the linear fit. For the unbinding rate, we start simulations in the bound state and collect MFPTs for leaving the MSM domain; we find a rate of $k_{off} = 19.0_{18.8}^{19.2} \, \mu s^{-1}$. The resulting equilibrium constant $K_{eq} = k_{on}/k_{off} = 3.0_{2.7}^{3.3} \, M^{-1}$ is similar to $3.6 \, M^{-1}$ found by de Sancho et al. [16], both of which are close to the experimental value of $2.2 \, M^{-1}$ [73] (see Tab. I for comparison). The binding rate and unbinding rate found by de Sancho et al. [16], although yielding a similar equilibrium constant, are both nearly an order of magnitude faster than the ones obtained with MSM/RD (Tab. I). The first indication that the present rates are an improved estimate is the fact that the kinetics (both the MSM relaxation timescales and k_{on}) are independent of the lag time (Fig. 4b, c).

To validate that the MSM/RD estimates of k_{off} and k_{on} have been estimated without significant bias, it must be shown that they are statistically consistent with the ground truth (in this case a sufficiently large and sufficiently long MD simulation). Here, k_{off} can be estimated directly by counting the frequency of ligand dissociation events from the binding pocket in the underlying MD simulations. Since there are not sufficient full dissociation pathways from the bound to the dissociated states in the MD data in order to make a statistically relevant comparison, we obtain a more precise estimate by computing the MFPT using an MSM directly constructed from the original MD data with the same discretization as used in the MSM/RD model. This resulted in a reference estimate of $23.4_{11.6}^{46.6} \, \mu s^{-1}$ (95% percentile computed with 1000 bootstrap samples), which is consistent with the MSM/RD estimate (Tab. I).

Unfortunately, this method is not as accurate for the binding rate k_{on}, which is notoriously difficult to estimate from small MD simulation boxes, where the length of trajectory segments in which the ligand stays in the dissociated state without touching the protein or crossing the periodic boundary are short compared to lagtimes τ used in an MSM approach, resulting in biased estimates [18]. Therefore, we performed another Myoglobin MD simulation in an eightfold larger periodic box (edge length 10 nm) with the same CO concentration as in the small MD simulation (resulting in 160 CO molecules) for a total simulation time of 405 ns. For this data, a direct MSM estimate of the binding rate yields $74.7_{29.9}^{130.9} \, \mu M^{-1} s^{-1}$ (95% percentile computed with 1000 bootstrap samples). As a result, the MSM/RD binding and dissociation rates are consistent with standard estimates computed directly from MD simulation, and the MSM/RD modeling error can be concluded to be statistically insignificant.

Given the consistency of the model, we also compare the results to experimental measurements, which is essentially a test of the MD model (e.g. force field, thermostat,

integrator). These are yet a factor 4-5 slower than our estimates ($k_{on} = 12 \, \mu M^{-1} s^{-1}$ and $k_{off} = 5.3 \, \mu s^{-1}$ found in [73]), confirming that the major part of the difference between the estimates in [16] and the experimental values could be removed by the fact that MSM/RD is a significantly more accurate model of the binding kinetics.

V. CONCLUSION

We introduced and developed the MSM/RD scheme, which couples MD-derived MSMs with RD simulations. We showed an implementation for protein-ligand systems and applied it to two simple systems. The main advantage of the algorithm is that it can simulate large time- and lengthscales while conserving molecular resolution and computational efficiency. This is achieved by extracting the characteristic features of the dynamics from several short MD simulations into an MSM, which can produce new data with great accuracy and at a much faster rate than the original MD simulations. This is a clear advantage in comparison to previous works, like [42, 69], since it does not require running MD simulations every time two particles are close to each other. It can further yield more accurate binding rates than traditional MSM methods by extending the diffusion domain available, lessening the periodic boundary effects and increasing the lifetime of the dissociated state. The scheme can be, in principle, coupled to any RD scheme, like over-damped Langevin dynamics, Langevin dynamics, GFRD [35, 36] and FPKMC algorithm [31], which could yield additional efficiency and accuracy or even incorporate long-range hydrodynamic interactions.

We first implemented the MSM/RD scheme for a simple ligand diffusion model (Sec. IV A), which served to verify the scheme. It reproduced the expected dynamics and binding/unbinding rates of the reference simulation. It was also able to generate an accurate MSM for the internal dynamics with a relatively small amount of data, which hints that it is feasible to extract the characteristic dynamics of a computationally feasible amount of MD simulations. Moreover, we implemented the MSM/RD scheme for the binding of CO to myoglobin system. After successfully extracting a self-consistent MSM and a coupling scheme, we found that the equilibrium constant is consistent with previous experimental and computational results [16, 73]. We also showed that the MSM/RD estimates are consistent with the underlying MD simulations – in particular our estimated association rate is consistent with the association rate estimated from a reference MD simulation conducted in a large simulation box that was not used to parametrize the MSM/RD model. This is a significant improvement over Ref. [16], where tenfold higher rates were estimated.

The MSM/RD theory we introduced provides the framework upon which schemes for more complex systems can be constructed. In particular, the next steps are to include association of two macromolecules, which may

	MSM/RD	Reference (approx. ground truth)	MSM in [16]	Experiment [73]	Unit
k_{on}	$57.0^{62.0}_{52.0}$	$74.7^{130.9}_{27.9}$	647	12	$M^{-1}\mu s^{-1}$
k_{off}	$19.0^{19.2}_{18.8}$	$23.4^{46.6}_{11.6}$	179	5.3	μs^{-1}
K_{eq}	$3.0^{3.3}_{2.7}$	$3.19^{3.8}_{2.6}$	3.6	2.2	M^{-1}

TABLE I: Rates and equilibrium constants for Myglobin-CO estimated from different methods. The reference values approximate the ground truth by conducting a standard MSM-based MFTP estimate from the MD simulation (for k_{on} a larger simulation box was used to allow for a generous definition of the dissociated state).

require to account for rototranslational diffusion, and the coupling between protein-ligand association and conformational changes. With the addition of these features, biologically relevant scenarios can be simulated. For example, if conformational changes of the protein are rare events and have different ligand association / dissociation rates, then the conformational dynamics and the ligand binding dynamics are nontrivially coupled at high ligand concentrations – see [15] for the example of Trypsin and Benzamidine. A biological relevant example is the activation of the Calcium sensor Synaptotagmin in neuronal synapses [74]. Here, a locally very high Calcium concentration is created by the opening of voltage-gated Calcium channels as a response to an electric signal. Synaptotagmin then binds up to five Calcium ions while going through different conformations, while the local Calcium concentration is reduced by diffusion. If Synaptotagmin successfully binds enough Calcium ions and transitions into an active conformation, it can catalyze the fission of neuronal vesicles, which transduces the signal to the postsynaptic side. Such scenarios can be simulated with MSM/RD simulations, in which the channels, the Synaptotagmin proteins and the ions are resolved as individual particles, and the binding/dissociation kinetics and conformational changes of Synaptotagmin is encoded in an MSM.

MSM/RD could be extended to deal with higher-order reactions. The most direct approach is to treat interactions of order 2, 3, etc., by different MSMs which are then coupled in a regular MSM/RD framework. The question then is how the higher-order MSMs are obtained. The brute-force approach would be to simulate the dynamics between three or more molecules with MD – e.g. with the help of enhanced sampling methods – and to extract corresponding higher-order MSMs. A cheaper, but approximate approach would be to ignore coupling between different states and assume that multiple ligands can bind and transition between binding sites independently, perhaps except for multiple occupation of the same binding site. Based on such an assumption, higher-order MSMs could be constructed by tensor products of MSMs with one protein and one ligand. In practice, conducting *some* but not all higher-order simulations and combining them to a generative model via machine learning methods may present a feasible pathway.

Finally, when considering protein interactions at high concentrations, the diffusion dynamics and long-range interactions of proteins are expected to be more compli-

cated and involve hydrodynamic effects and anomalous diffusion. To include such effects, appropriate dynamical schemes should be included in the RD part.

In future developments, we will extend the MSM/RD scheme to address these issues; however, it should be acknowledged that some of these extensions come with their own set of challenges that are not trivial to address.

APPENDIX: MSM/RD SCHEME FOR SEC. III

Based on the estimated quantities defined in the Sec. III, we introduce an implementation of the MSM/RD algorithm from Sec. II D.

Input: Initial mode (RD or MSM), initial condition (coordinates $\mathbf{c}_0$ or state s_0, respectively) and $t = 0$:

While $t \leq t_{\text{final}}$:

1. If in RD mode:

    (a) Propagate $\mathbf{c}_t \to \mathbf{c}_{t+\tau_{\text{RD}}}$ by diffusion
    (b) Update time $t \mathrel{+}= \tau_{\text{RD}}$
    (c) If $r_{AB}(\mathbf{c}_t) < R$ (enter MSM domain):
      - Select trajectory from $L_{\text{entry}} = \{\mathbf{c}_{\text{entry}}, \mathbf{x}_{\text{end}}, \Delta\}$ with $\mathbf{c}_{\text{entry}}$ closest to $\mathbf{c}_t$
      - If $\mathbf{x}_{\text{end}}$ is a state:
        Map to state $s_{t+\Delta} = \mathbf{x}_{\text{end}}$
        Update time $t \mathrel{+}= \Delta$
        Switch to MSM mode
      - If $\mathbf{x}_{\text{end}}$ are coordinates:
        Map to coordinates $\mathbf{c}_{t+\Delta} = \mathbf{x}_{\text{end}}$
        Update time $t \mathrel{+}= \Delta$

2. `Else (MSM mode):`

    (a) `If` $s_t \neq s_{t-\tau_{MSM}}$ `or previous mode` $\neq$ `MSM mode:`

- `Sample exit event with` $P_{\text{exit},s}$
- `If exit event:`
  `Uniformly select trajectory from` $L_{\text{exit},s} = \{\mathbf{c}_{\text{exit}}, \Delta\}$
  `Map to coordinates` $\mathbf{c}_{t+\Delta} = \mathbf{c}_{\text{exit}}$
  `Update time` $t \mathrel{+}= \Delta$
  `Switch to RD mode and break current loop iteration`

(b) `Propagate` $s_t \to s_{t+\tau_{\text{MSM}}}$ `using` $\mathbf{T}(\tau_{\text{MSM}})$

(c) `Update time` $t \mathrel{+}= \tau_{\text{MSM}}$.

The diffusion in the RD domain is done using a Euler-Maruyama discretization of Eq. (7) [75]. Note the diffusion step can be simulated more efficiently with event-based algorithms, like FPKMC or eGFRD [28, 31, 35, 42] for systems with low particle concentrations. In order to optimize the efficiency of the algorithm, the entry points of entry trajectories are classified into equal area bins on the sphere. This allows the algorithm to find the closest trajectory to a given entry point more efficiently. The partition of the sphere was done following [76].

[1] A. Warshel and M. Levitt. Theoretical studies of enzymic reactions: Dielectric, electrostatic and steric stabilization of the carbonium ion in the reaction of lysozyme. *J. Mol. Biol.*, 103:227–249, 1976.

[2] M. Shirts and V. S. Pande. Screen savers of the world unite! *Science*, 290:1903–1904, 2000.

[3] I. Buch, M. J. Harvey, T. Giorgino, D. P. Anderson, and G. De Fabritiis. High-throughput all-atom molecular dynamics simulations using distributed computing. *J. Chem. Inf. Model.*, 50:397–403, 2010.

[4] D. E. Shaw, P. Maragakis, K. Lindorff-Larsen, S. Piana, R.O. Dror, M.P. Eastwood, J.A. Bank, J.M. Jumper, J.K. Salmon, Y. Shan, and W. Wriggers. Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science*, 330:341–346, 2010. ISSN 1095-9203. doi: 10.1126/science.1187409. URL http://dx.doi.org/10.1126/science.1187409.

[5] S. Doerr, M. J. Harvey, F. Noé, and G. De Fabritiis. HTMD: High-Throughput Molecular Dynamics for Molecular Discovery. *J. Chem. Theory Comput.*, 12:1845–1852, 2016.

[6] J.-H. Prinz, H. Wu, M. Sarich, B. G. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé. Markov models of molecular kinetics: Generation and validation. *J. Chem. Phys.*, 134:174105, 2011.

[7] G. R. Bowman, V. S. Pande, and F. Noé, editors. *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation.*, volume 797 of *Advances in Experimental Medicine and Biology.* Springer Heidelberg, 2014.

[8] M. Sarich and C. Schütte. *Metastability and Markov State Models in Molecular Dynamics*. Courant Lecture Notes. American Mathematical Society, 2013.

[9] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl. Constructing the full ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. USA*, 106:19011–19016, 2009.

[10] G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande. Progress and challenges in the automated construction of Markov state models for full protein systems. *J. Chem. Phys.*, 131:124101, 2009.

[11] K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw. How fast-folding proteins fold. *Science*, 334:517–520, 2011.

[12] K. J. Kohlhoff, D. Shukla, M. Lawrenz, G. R. Bowman, D. E. Konerding, D. Belov, R. B. Altman, and V. S. Pande. Cloud-based simulations on google exacycle reveal ligand modulation of gpcr activation pathways. *Nat. Chem.*, 6:15–21, 2014.

[13] I. Buch, T. Giorgino, and G. De Fabritiis. Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations. *Proc. Natl. Acad. Sci. USA*, 108:10184–10189, 2011. ISSN 1091-6490. doi: 10.1073/pnas.1103547108.

[14] D.-A. Silva, G. R. Bowman, A. Sosa-Peinado, and X. Huang. A role for both conformational selection and induced fit in ligand binding by the lao protein. *PLoS Comput. Biol.*, 7:e1002054, 2011.

[15] N. Plattner and F. Noé. Protein conformational plasticity and complex ligand binding kinetics explored by atomistic simulations and markov models. *Nat. Commun.*, 6:7653, 2015.

[16] D. De Sancho, A. Kubas, P. Wang, J. Blumberger, and R. B. Best. Identification of mutational hot spots for substrate diffusion: Application to myoglobin. *J. Chem. Theory. Comput.*, 11(4):1919–1927, 2015.

[17] A. Kubas, C. Orain, D. De Sancho, L. Saujet, M. Sensi, C. Gauquelin, I. Meynial-Salles, P. Soucaille, H. Bottin, C. Baffert, et al. Mechanism of O2 diffusion and reduction in FeFe hydrogenases. *Nat. Chem.*, 2017.

[18] N. Plattner, S. Doerr, G. De Fabritiis, and F. Noé. Complete protein–protein association kinetics in atomic detail revealed by molecular dynamics simulations and markov modelling. *Nat. Chem.*, 9:1005–1011, 2017.

[19] H. Wu, A. S. J. S. Mey, E. Rosta, and F. Noé. Statistically optimal analysis of state-discretized trajectory data from multiple thermodynamic states. *J. Chem. Phys.*, 141:214106, 2014.

[20] E. Rosta and G. Hummer. Free energies from dynamic weighted histogram analysis using unbiased markov state model. *J. Chem. Theory Comput.*, 11:276–285, 2015.

[21] H. Wu, F. Paul, C. Wehmeyer, and F. Noé. Multiensemble markov models of molecular thermodynamics and kinetics. *Proc. Natl. Acad. Sci. USA*, 113(23):E3221–E3230, 2016. doi:10.1073/pnas.1525092113.

[22] A. S. J. S. Mey, H. Wu, and F. Noé. xTRAM: Estimating equilibrium expectations from time-correlated simulation data at multiple thermodynamic states. *Phys. Rev. X*, 4:041018, 2014.

[23] F. Paul, C. Wehmeyer, E. T. Abualrous, H. Wu, M. D.

Crabtree, J. Schöneberg, J. Clarke, C. Freund, T. R. Weikl, and F. Noé. Protein-peptide association kinetics beyond the seconds timescale from atomistic simulations. *Nat. Commun.*, 8:1095, 2017.

[24] R. Casasnovas, V. Limongelli, P. Tiwary, P. Carloni, and M. Parrinello. Unbinding kinetics of a p38 map kinase type ii inhibitor from metadynamics simulations. *J. Am. Chem. Soc.*, 139:4780–4788, 2017.

[25] H. Qian. Cellular biology in terms of stochastic nonlinear biochemical dynamics: Emergent properties, isogenetic variations and chemical system inheritability. *J. Stat. Phys*, 141(6):990–1013, 2010.

[26] R. Erban and S. J. Chapman. Stochastic modelling of reaction–diffusion processes: algorithms for bimolecular reactions. *Phys. Biol.*, 6(4):046001, 2009.

[27] D. Fange, O. Berg, P. Sjöberg, and J. Elf. Stochastic reaction-diffusion kinetics in the microscopic limit. *Proc. Natl. Acad. Sci. USA*, 107(46):19820–19825, 2010.

[28] K. Takahashi, S. Tănase-Nicola, and P. R. ten Wolde. Spatio-temporal correlations can drastically change the response of a mapk pathway. *Proc. Natl. Acad. Sci. USA*, 107(6):2473–2478, 2010.

[29] S. S. Andrews and D. Bray. Stochastic simulation of chemical reactions with spatial resolution and single molecule detail. *Phys. Biol.*, 1(3):137, 2004.

[30] J. Biedermann, A. Ullrich, J. Schöneberg, and F. Noé. Readdymm: Fast interacting particle reaction-diffusion simulations using graphical processing units. *Biophys. J.*, 108:457–461, 2015.

[31] A. Donev, V. V. Bulatov, T. Oppelstrup, G. H. Gilmer, B. Sadigh, and M. H. Kalos. A first-passage kinetic monte carlo algorithm for complex diffusion–reaction systems. *J. Comput. Phys.*, 229(9):3214–3236, 2010.

[32] Aleksandar Donev, Chiao-Yu Yang, and Changho Kim. Efficient reactive brownian dynamics. *The Journal of Chemical Physics*, 148(3):034103, 2018.

[33] J. Hattne, D. Fange, and J. Elf. Stochastic reaction-diffusion simulation with mesord. *Bioinformatics*, 21(12):2923–2924, 2005.

[34] J. Schöneberg and F. Noé. Readdy - a software for particle based reaction diffusion dynamics in crowded cellular environments. *PLoS ONE*, 8(e74261), 2013.

[35] J. S. van Zon and P. R. Ten Wolde. Green's-function reaction dynamics: A particle-based approach for simulating biochemical networks in time and space. *J. Chem. Phys.*, 123(23):4910, 2005.

[36] J. S. Van Zon and P. R. ten Wolde. Simulating biochemical networks at the particle level in time and space: Green's function reaction dynamics. *Phys. Rev. Lett.*, 94:128103, 2005.

[37] P. Mereghetti, D. Kokh, J. A. McCammon, and R. Wade. Diffusion and association processes in biological systems: theory, computation and experiment. *BMC biophysics*, 4(1):2, 2011.

[38] J. Schöneberg, A. Ullrich, and F. Noé. Simulation tools for particle-based reaction-diffusion dynamics in continuous space. *BMC Biophysics*, 7:11, 2014.

[39] D. L. Ermak and J. A. McCammon. Brownian dynamics with hydrodynamic interactions. *J. Chem. Phys.*, 69(4):1352–1360, 1978.

[40] T. Geyer and U. Winter. An o (n 2) approximation for hydrodynamic interactions in brownian dynamics simulations. *J. Chem. Phys.*, 130(11):114905, 2009.

[41] Luigi Sbailò and Frank Noé. An efficient multi-scale

[42] green's function reaction dynamics scheme. *J. of Chem. Phys.*, 147(18):184106, 2017.

[42] A. Vijaykumar, P. G. Bolhuis, and P. R. ten Wolde. Combining molecular dynamics with mesoscopic green's function reaction dynamics simulations. *J. Chem. Phys.*, 143(21):214102, 2015.

[43] L. W. Votapka, B. R. Jagger, A. L. Heyneman, and R. E. Amaro. SEEKR: Simulation enabled estimation of kinetic rates, a computational tool to estimate molecular kinetics and its application to trypsin–benzamidine binding. *J. Phys. Chem. B*, 121:3597–3606, 2017.

[44] Anton K Faradjian and Ron Elber. Computing time scales from reaction coordinates by milestoning. *J. Chem. Phys.*, 120(23):10880–10889, 2004.

[45] M. J. del Razo and H. Qian. A discrete stochastic formulation for reversible bimolecular reactions via diffusion encounter. *Comm. Math. Sci.*, 14(6):1741–1772, 2016.

[46] M. B. Flegg, S. J. Chapman, and R. Erban. The two-regime method for optimizing stochastic reaction–diffusion simulations. *J. Royal. Soc. Interface*, 9(70):859–868, 2012.

[47] M. B. Flegg, S. Hellander, and Radek R. Erban. Convergence of methods for coupling of microscopic and mesoscopic reaction–diffusion simulations. *J. Comput. Phys.*, 289:1–17, 2015.

[48] S. Donnini, F. Tegeler, G. Groenhof, and H. Grubmüller. Constant ph molecular dynamics in explicit solvent with lambda-dynamics. *J. Chem. Theory. Comput.*, 7:1962–1978, 2011.

[49] C. Schütte, A. Fischer, W. Huisinga, and P. Deuflhard. A Direct Approach to Conformational Dynamics based on Hybrid Monte Carlo. *J. Comput. Phys.*, 151:146–168, 1999.

[50] F. Noé and F. Nüske. A variational approach to modeling slow processes in stochastic dynamical systems. *Multiscale Model. Simul.*, 11:635–655, 2013.

[51] F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé. Variational approach to molecular kinetics. *J. Chem. Theory Comput.*, 10:1739–1752, 2014.

[52] G. Perez-Hernandez, F. Paul, T. Giorgino, G. D Fabritiis, and Frank Noé. Identification of slow molecular order parameters for markov model construction. *J. Chem. Phys.*, 139:015102, 2013.

[53] C. R. Schwantes and V. S. Pande. Improvements in markov state model construction reveal many non-native interactions in the folding of ntl9. *J. Chem. Theory Comput.*, 9:2000–2009, 2013.

[54] F. Noé and C. Clementi. Collective variables for the study of long-time kinetics from molecular trajectories: theory and methods. *Curr. Opin. Struc. Biol.*, 43:141–147, 2017.

[55] S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé. Data-driven model reduction and transfer operator approximation. *J. Nonlinear Sci.*, pages 1–26, 2017.

[56] F. Noé and C. Clementi. Kinetic distance and kinetic maps from molecular dynamics simulation. *J. Chem. Theory Comput.*, 11:5002–5011, 2015.

[57] F. Noé, R. Banisch, and C. Clementi. Commute maps: separating slowly-mixing molecular configurations for kinetic modeling. *J. Chem. Theory Comput.*, 12:5620–5630, 2016.

[58] M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Perez-Hernandez, M. Hoffmann, N. Plattner, J.-H. Prinz, and F. Noé. PyEMMA 2: A software package for

estimation, validation and analysis of Markov models. *J. Chem. Theory Comput.*, 11:5525–5542, 2015.

[59] M. P. Harrigan, M. M. Sultan, C. X. Hernández, B. E. Husic, P. Eastman, C. R. Schwantes, K. A. Beauchamp, R. T. McGibbon, and V. S. Pande. Msmbuilder: Statistical models for biomolecular dynamics. *Biophys J.*, 112: 10–15, 2017.

[60] N. V. Buchete and G. Hummer. Coarse Master Equations for Peptide Folding Dynamics. *J. Phys. Chem. B*, 112: 6057–6069, 2008.

[61] B. Trendelkamp-Schroer, H. Wu, F. Paul, and F. Noé. Estimation and uncertainty of reversible markov models. *J. Chem. Phys.*, 143:174101, 2015.

[62] S. Kube and M. Weber. A coarse graining method for the identification of transition rates between molecular conformations. *J. Chem. Phys.*, 126:024103, 2007. doi: 10.1063/1.2404953. URL `http://dx.doi.org/10.1063/1.2404953`.

[63] G. Hummer and A. Szabo. Optimal dimensionality reduction of multistate kinetic and markov-state models. *J. Phys. Chem. B*, 119:9029–9037, 2015.

[64] S. Orioli and P. Faccioli. Dimensional reduction of markov state models from renormalization group theory. *J. Chem. Phys.*, 145:124120, 2016.

[65] F. Noé, H. Wu, J.-H. Prinz, and N. Plattner. Projected and hidden markov models for calculating kinetics and metastable states of complex molecules. *J. Chem. Phys.*, 139:184114, 2013.

[66] A. Mardt, L. Pasquali, H. Wu, and F. Noé. VAMPnets for deep learning of molecular kinetics. *Nat. Commun.*, page 5, 2018.

[67] C. Schütte, F. Noé, J. Lu, M. Sarich, and E. Vanden-Eijnden. Markov state models based on milestoning. *J. Chem. Phys.*, 134(20):05B609, 2011.

[68] J. Schluttig, D. Alamanova, V. Helms, and U. S. Schwarz. Dynamics of protein-protein encounter: A langevin equation approach with reaction patches. *J. Chem. Phys.*, 129:155106, 2008.

[69] A. Vijaykumar, T. E. Ouldridge, P. R. ten Wolde, and P. G. Bolhuis. Multiscale simulations of anisotropic particles combining brownian dynamics and green's function reaction dynamics. *J. Chem. Phys.*, 146(11):114106, 2017.

[70] Doris A Schuetz, Wilhelmus Egbertus Arnout de Witte, Yin Cheong Wong, Bernhard Knasmueller, Lars Richter, Daria B Kokh, S Kashif Sadiq, Reggie Bosma, Indira Nederpelt, Laura H Heitman, et al. Kinetics for drug discovery: an industry-driven effort to target drug residence time. *Drug discovery today*, 22(6):896–911, 2017.

[71] M. Ester, H. P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996.

[72] D. L. Wise and G. Houghton. Diffusion coefficients of neon, krypton, xenon, carbon monoxide and nitric oxide in water at 10–60 °C. *Chem. Eng. Sci.*, 23(10):1211–1216, 1968.

[73] T. E. Carver, R. J. Rohlfs, J. S. Olson, Q. H. Gibson, R. S. Blackmore, B. A. Springer, and S. G. Sligar. Analysis of the kinetic barriers for ligand binding to sperm whale myoglobin using site-directed mutagenesis and laser photolysis techniques. *J. Biol. Chem.*, 265 (32):20007–20020, 1990.

[74] Thomas C. Südhof. Neurotransmitter release: The last millisecond in the life of a synaptic vesicle. *Neuron*, 80: 675–690, 2013.

[75] D. J. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM review*, 43(3):525–546, 2001.

[76] P. Leopardi. A partition of the unit sphere into regions of equal area and small diameter. *Electron. Trans. Numer. Anal.*, 25(12):309–327, 2006.