

Computing transition path theory quantities with trajectory stratification

Bodhi P. Vani,¹ Jonathan Weare,² and Aaron R. Dinner^{*1}

¹*Department of Chemistry and James Franck Institute, University of Chicago, Chicago IL 60637*

²*Courant Institute of Mathematical Sciences, New York University, New York 10012*

(*Electronic mail: dinner@uchicago.edu)

(Dated: 23 August 2023)

Transition path theory computes statistics from ensembles of reactive trajectories. A common strategy for sampling reactive trajectories is to control the branching and pruning of trajectories so as to enhance the sampling of low probability segments. However, it can be challenging to apply transition path theory to data from such methods because determining whether configurations and trajectory segments are part of reactive trajectories requires looking backward and forward in time. Here, we show how this issue can be overcome efficiently by introducing simple data structures. We illustrate the approach in the context of nonequilibrium umbrella sampling (NEUS), but the strategy is general and can be used to obtain transition path theory statistics from other methods that sample segments of unbiased trajectories.

I. INTRODUCTION

One of the main purposes of simulations is to learn how processes occur. However, most processes of interest occur on timescales that are orders of magnitude longer than the numerical integration time step. In the case of molecular systems, on which we focus here, the time step for all-atom models is in the femtosecond range, while conformational transitions of functional significance typically take place in the microsecond to seconds range. As a result, at best, only a small number of events can be observed by direct simulation. Moreover, even if one obtains examples of events of interest, the number of dynamical variables is generally sufficiently large that it is not obvious which are the key steps¹, leave alone their likelihoods and the distributions of times that they occur.

Transition Path Theory (TPT) provides a means of computing these statistics for transitions between two metastable states from the ensemble of reactive trajectories (i.e., those that contain an event)²⁻⁴. A key idea in TPT is that these statistics can be expressed as products over the steady-state distribution and the probabilities of coming from or going to the metastable states (known as commitment probabilities or committers). This factorization allows estimating these probabilities separately, which opens the door to computing statistics of reactive trajectories from a broader range of data.

Markov State Models (MSMs) are a popular means of computing TPT statistics^{5,6}. In MSMs, dynamics are coarse-grained to transitions between a set of discrete states, the probabilities of which are assumed to be independent of the sequence of states visited. The advantage relative to direct simulation comes from the fact

that there is considerable freedom in how one estimates the state-to-state transition probabilities^{7,8}. Dynamical Galerkin Approximation (DGA)^{9,10} is a generalization of MSMs for TPT statistics that accounts for the boundary conditions of the statistics¹¹ and represents them through expansions over sets of basis functions (which can take a form beyond indicator functions on discrete states). Recently, TPT statistics have also been computed with milestoning¹², which instead assumes that the system reaches an equilibrium within selected parts of the configuration space (the milestones)^{13,14}.

Although often reasonable, the above approximation schemes can break down, and, even if they do not, they ultimately limit the resolution with which statistics can be computed. Methods that sample the ensemble of trajectories that connect the metastable states without such assumptions can in principle provide results with arbitrary resolution. One such method is Transition Path Sampling (TPS), a Monte Carlo procedure that accepts and rejects whole trajectories¹⁵⁻¹⁷. However, generally it is more efficient to divide the configuration space and sample segments of trajectories that transition between regions. This idea is at the foundation of the Weighted Ensemble (WE)¹⁸⁻²⁰, Transition Interface Sampling (TIS)^{21,22}, Forward Flux Sampling (FFS)^{23,24}, Nonequilibrium Umbrella Sampling (NEUS)²⁵⁻²⁹, and Exact Milestoning (EM)³⁰ methods. With the exception of TIS, these methods allow treatment of microscopically irreversible dynamics, which is also not straightforward in TPS.

The differences between these methods are in the details of the trajectory segments they select and how they track their probabilities to enable reconstruction of the overall statistics. Generally the interfaces are speci-

fied through collective variables that combine information from multiple coordinates used for the underlying numerical integration. TIS and FFS require non-intersecting interfaces because they are formulated in terms of the probabilities of going from one to another in order. WE, NEUS, and EM do not have this restriction and thus readily allow control of sampling in more than one collective variable. In its traditional form, WE tracked the weight of each trajectory segment (equivalent to a member of the ensemble) individually. NEUS instead tracked the weight of a subpopulation within a region (termed a stratum in statistics) and redistributed the weight to satisfy a global flux balance condition. Later implementations of WE that incorporate this idea³¹ are very similar to NEUS, as is EM, which relaxes the distributions on the milestones. Several of these algorithms can thus be described by a common framework, trajectory stratification³².

Because these methods sample trajectory segments rather than whole trajectories, care must be used to weight trajectory segments appropriately when computing statistics. Vanden-Eijnden and Venturoli²⁷ showed that the rate can be obtained exactly from NEUS by augmenting the dynamics with information about the last metastable state visited, and this scheme was further refined and applied in ref. 26. However, the rate requires only tracking transitions into the metastable states, making it easier to compute than statistics that require information about intervening states. Vanden-Eijnden and Venturoli²⁷ provided a formula for computing the (backward) committor from the region weights, but, to the best of our knowledge, statistics such as committors and probability currents of reactive trajectories (reactive currents for short), which reveal microscopic likelihoods of reaction and reaction pathways, respectively, have not been computed from such methods. The challenge presented by these quantities is that each trajectory segment can contribute to an infinite number of complete trajectories, and determining the ones that contribute to a given statistic requires looking both forward and backward in time.

Here, we introduce a simple bookkeeping procedure and practical formulas for computing committors and reactive currents. While we present our methods in terms of the NEUS algorithm, they can be adapted to the other algorithms for sampling trajectory segments described above. The paper is organized as follows. In Section II, we define these quantities in the TPT framework, and in an associated appendix we provide an explicit justification for an estimator of the current. In Section III we provide a succinct description of the NEUS algorithm and show how it must be modified. We demonstrate the algorithm for isomerization of a peptide in Section IV.

II. TRANSITION PATH THEORY

We are interested in learning the statistics of reactive trajectories. In TPT, these are defined to be trajectories that start in a state A and end in a non-overlapping state B without returning to A . To delimit their ends, we define

$$t_+(t) = \min(t' > t, X_{t'} \in A \cup B) \quad (1)$$

$$t_-(t) = \max(t' < t, X_{t'} \in A \cup B), \quad (2)$$

where $X_{t'}$ is the configuration at time t' , and $X_t \in (A \cup B)^c$. The forward committor $q_+(x)$ is the probability that the system goes to B before A starting from state x :

$$q_+(x) = \mathbf{P}[X(t_+(0)) \in B | X(0) = x]. \quad (3)$$

Similarly, the backward committor $q_-(x)$ is the probability that a system at x came from A after B :

$$q_-(x) = \mathbf{P}[X(t_-(0)) \in A | X(0) = x]. \quad (4)$$

In the present study, for visualization, we project the committors onto the space defined by a vector of collective variables, $\theta(x)$, weighted by $\pi(x)$:

$$q_\pm^\theta(\Theta) = \int q_\pm(x) \pi(x) \delta(\theta(x) - \Theta) dx, \quad (5)$$

A key element of TPT is that the probability that a trajectory is reactive can be expressed in terms of quantities that are local in space. Specifically, the probability that a trajectory that passes through x is reactive is proportional to $\pi(x)q_-(x)q_+(x)$, where $\pi(x)$ is the steady-state distribution. The fact that TPT is based on local statistics aids in interpretation, but it also can make obtaining certain statistics challenging. We discuss how TPT of an augmented process can be used to obtain additional statistics that require knowledge of sequences of states visited in ref. 33. Of course, the analysis we propose here can be complemented by direct analysis of the reactive trajectories as well.

A. Reactive current

The reactive current is a central quantity in TPT^{2,3,34}. It can be used to define reactive pathways, and averages over reactive trajectories—in particular the rate—can be computed from it. The reactive current is defined by considering a surface S that divides the space into two regions: C containing A and its complement C^c containing B . Then, the reactive current, $\mathbf{I}_{AB}(x)$, is the vector field

whose integral is the reactive flux across S :

$$\int_S \mathbf{l}_{AB} \cdot \hat{n}_S d\sigma_S = \lim_{\tau \rightarrow 0} \lim_{T \rightarrow \infty} \frac{1}{2T\tau} \int_{-T}^T [\mathbf{1}_C(X(t)) \mathbf{1}_{C^c}(X(t+\tau)) - \mathbf{1}_{C^c}(X(t)) \mathbf{1}_C(X(t+\tau))] \mathbf{1}_A(X(t-\tau)) \mathbf{1}_B(X(t+\tau)) dt \quad (6)$$

where $\hat{n}_S$ is a unit vector normal to S (pointing from C to C^c), $d\sigma_S$ is a surface element, τ is a lag time, and

$$\mathbf{1}_D(x) = \begin{cases} 1 & \text{if } x \in D \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

In general, the full state space of the system has many dimensions. $\mathbf{l}_{AB}(x)$ is a vector field in the full space of dynamical variables, of which there are many in general. However, most of these are irrelevant for determining whether the reaction proceeds, and we seek a small number of key dimensions. Because these are often combinations of the original dynamical variables, we term them collective variables. We show in the Supplemental Information that

$$\begin{aligned} \mathbf{l}_{AB} \cdot \nabla \theta(x) &= \pi(x) q_-(x) q_+(x) \\ &\times \lim_{\tau \rightarrow 0} \frac{1}{2\tau} \mathbf{E}[(\theta(X(\tau)) - \theta(X(-\tau))) \\ &X(t_-(0)) \in A, X(t_+(0)) \in B, X(0) = x]. \end{aligned} \quad (8)$$

We use $\mathbf{E}$ to denote expectations over the steady-state trajectory ensemble. This formula is simply the product of the probability of being on a reactive trajectory and the average increment per unit time of θ along a reactive trajectory. We project this current onto the space of collective variables:

$$\begin{aligned} \mathbf{l}_{AB}^\theta(\Theta) &= \int \mathbf{l}_{AB}(x) \cdot \nabla \theta(x) \delta(\theta(x) - \Theta) dx. \\ &= \lim_{|d\Theta| \rightarrow 0} \frac{1}{|d\Theta|} \int_{\{\theta(x) \in d\Theta\}} \mathbf{l}_{AB}(x) \cdot \nabla \theta(x) dx. \end{aligned} \quad (9)$$

We recently showed that this projected reactive current can also be used to compute the reactive flux through an integral similar in form to (6) but in the θ space; the projected current $\mathbf{l}_{AB}^\theta(\Theta)$ can in turn be used to compute the rate¹⁰. The expressions in (8) and (9) allow us to obtain the reactive current in the space of collective variables from reactive trajectories.

In the next section, we discuss how reactive trajectories can be sampled efficiently. Although we treat time as continuous above, with a view toward developing a practical algorithm, we henceforth assume a dynamics with discrete time steps; the lag time τ is thus an integer multiple of the time step in practice.

III. TRAJECTORY STRATIFICATION

The idea of stratification in general is to obtain better overall statistics by controlling the sampling of subpopulations (strata). The most familiar form of stratification in molecular simulations is umbrella sampling^{29,35–38}, which is frequently used to compute free energies. In this case, the strata are defined by regions in a space of collective variables. Each of the strata is sampled independently by a copy of the system, and the information is then combined to obtain unbiased averages. Trajectory stratification extends this strategy from states to trajectories and thus enables treating microscopically irreversible systems and estimating dynamical statistics (for microscopically reversible or irreversible systems)^{25–29}.

Here, we focus on the case of steady-state, time-independent quantities. To formulate the algorithm mathematically, we associate with $X(t)$ a variable $J(t)$ that reports the index of the stratum at time t . For example, a trajectory that spent one time step in stratum 1, two time steps in stratum 2, went back to stratum 1 for one time step, and then spent three time steps in stratum 5 would have $J(t=1)=1, J(t=2)=2, J(t=3)=2, J(t=4)=1, J(t=5)=5, J(t=6)=5$, and $J(t=7)=5$. Then, denoting the time that a trajectory first exits a stratum as $s = \min\{t : J(t) \neq J(0)\}$, we can write expectations as

$$\mathbf{E}[f] = \frac{1}{\Omega} \sum_i z_i f_i, \quad (11)$$

where Ω is a normalization factor and

$$f_i = \int_x \sum_t f(X(t)) \mathbf{P}_{x,i}[X(t), t < s] \pi_i(dx). \quad (12)$$

$\mathbf{P}_{x,i}$ indicates that we are conditioning the probability on starting at configuration x in stratum i ; $\pi_i(dx) \equiv \pi_i(x)dx$ is the probability of being in the differential volume element dx conditioned on entering stratum i at that state, and z_i is the normalization factor for $\pi_i(dx)$. Operationally, (12) corresponds to initiating trajectories from steady-state entry points to stratum i , terminating them when they leave the stratum, and averaging over the sampled points. We show in ref. 32 that, for ergodic averages, the vector z can be obtained by solving an eigenequation:

$$z^T G = z^T. \quad (13)$$

The matrix G tracks the number of transitions between strata; we define it precisely below in (17). Physically, (13) can be viewed as a statement of global flux balance.

The key idea in the present work is that we define the index process so as to track the last metastable state visited, in the spirit of refs. 26 and 27. In other words, if

there are K strata with domains D_j defined in terms of a set of order parameters, then $J(t)$ runs from 1 to $2K$:

$$J(t) = K\mathbf{1}[X(t_-(t)) \in B] + \sum_{j=1}^K j\mathbf{1}[X(t) \in D_j]. \quad (14)$$

Thus strata with indices from 1 to K capture data for trajectories that last came from state A , and strata with indices $K+1$ to $2K$ represent trajectories that last came from state B .

A. A self-consistent iteration

Given the index process, we can now define z_i , π_i , and G_{ij} more precisely (we note that these quantities correspond to the steady-state averages of $\bar{z}_i$, $\bar{\pi}_i$, and $\bar{G}_{ij}$ defined in (24), (31), and (34), respectively in ref. 32; we suppress the overbars here for simplicity of notation):

$$\pi_i(dx) = \lim_{T \rightarrow \infty} \frac{1}{T} \frac{\sum_{t=0}^T \mathbf{P}[J(t) = i, J(t-1) \neq i, X(t) \in dx]}{z_i} \quad (15)$$

and

$$z_i = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbf{P}[J(t) = i, J(t-1) \neq i]. \quad (16)$$

To obtain G_{ij} , we draw states from $\pi_i(dx)$ and evolve them until they leave stratum i :

$$G_{ij} = \lim_{T \rightarrow \infty} \frac{1}{T} \frac{\sum_{t=0}^T \mathbf{P}[J(t) = j, J(t-1) = i]}{z_i} \quad (17)$$

$$= \int \mathbf{P}_{x,i}[J(s) = j] \pi_i(dx). \quad (18)$$

Furthermore, we add to the distribution of entry points to stratum j that come from stratum i :

$$\gamma_j(dx) = \lim_{T \rightarrow \infty} \frac{1}{T} \frac{\sum_{t=0}^T \mathbf{P}[J(t) = j, J(t-1) = i, X(t) \in dx]}{G_{ij} z_i} \quad (19)$$

$$= \int \mathbf{P}_{y,i}[J(s) = j, X(s) \in dx] \pi_i(dy). \quad (20)$$

Thus, given $\pi_i(dx)$, we can compute G_{ij} and $\gamma_j(dx)$, and in turn z_i from (13). Because the entries of G depend on the dynamics within the strata, and those are initialized using G , the algorithm is iterative in nature. To complete the iteration, we start from (15) and write $\pi_j(dx)$ in

terms of the other quantities:

$$\begin{aligned} \pi_j(dx) &= \lim_{T \rightarrow \infty} \frac{1}{T} \frac{\sum_{t=0}^T \mathbf{P}[J(t) = j, J(t-1) \neq j, X(t) \in dx]}{z_j} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \frac{\sum_{i \neq j} \sum_{t=0}^T \mathbf{P}[J(t) = j, J(t-1) = i, X(t) \in dx]}{z_j} \\ &= \frac{1}{z_j} \sum_{i \neq j} z_i G_{ij} \gamma_j(dx). \end{aligned} \quad (21)$$

B. Algorithm

If one had a long trajectory that passed between the metastable states many times, one could compute (8) by identifying the reactive trajectories from A to B and accumulating the increment per unit time in θ as a function of θ . However, in the stratification scheme above, we generally do not know whether a copy of the system (random walker, or walker for short) is part of a reactive trajectory as we are evolving it.

The key to addressing this issue is based on the fact that, as described previously^{25,26,32,39}, in practice, $\gamma_j(dx)$ is a list of configurations that are saved when walkers enter stratum j from stratum i . Consequently, for each walker, we can save the θ points visited, the increments in θ at those points, and a pointer to the configuration in $\gamma_j(dx)$ from which the walker originated; then, when a walker enters A or B , we reconstruct the full sequence of walkers connecting the metastable states and in turn the sequence of θ values visited and the associated increments, and we add these data to the appropriate sums. The algorithm is as follows.

1. Initialize simulation quantities as follows.

- (a) Define the $2K$ strata as in (14).
- (b) Use preliminary trajectories (obtained from either unbiased simulation or an enhanced sampling method such as STePS⁴⁰) to populate $\gamma_j(dx)$. In practice, $\gamma_j(dx)$ is a list of configurations used to represent the distribution of entry points to stratum j from stratum i , and each element contains a full configuration and a pointer (set to null for the configurations from the preliminary trajectories). In the present study, we allow the list to grow throughout the simulation (i.e., we do not overwrite configurations, as doing so may bias the statistics as discussed in ref. 26).
- (c) Estimate G_{ij} as the number of transitions from stratum i to stratum j in the prelimi-

- nary trajectories, normalized by the number of transitions out of stratum i .
- (d) Initialize z by solving (13).
- (e) Create a grid over the collective variables and denote the volume of each grid element by $\Delta\Theta$. Denoting the grid points by Θ , initialize all elements of the following arrays to zero:
- $p_A^+(\Theta)$ (resp. $p_B^+(\Theta)$), which accumulates the weight of trajectories entering state A (resp. B);
 - $p_A^-(\Theta)$ (resp. $p_B^-(\Theta)$), which accumulates the weight of trajectories exiting state A (resp. B);
 - $v_{AB}(\Theta)$ (resp. $v_{BA}(\Theta)$), which accumulates the weighted CV increments of trajectories exiting state A and entering state B (resp. exiting state B and entering state A).
2. Run NEUS essentially as described in Section 3.3 of ref. 32. Namely, at each iteration l :
- Initialize all elements of a matrix $g^{(l)}$ with the same dimensions as G (i.e., $2K \times 2K$) to zero.
 - For each stratum j , draw N walkers from π_j . To be precise, for each walker, select an element of $\gamma_j(dx)$ with probability $z_j G_{ij}$ and use the saved configuration to initialize the walker; this procedure represents (21). Associate with the walker a pointer to the element of $\gamma_j(dx)$.
 - Evolve each walker by unbiased dynamics until it exits its initial stratum (j in step 2b).
 - When a walker exits from stratum j to stratum k , save its final configuration and associated pointer to $\gamma_{jk}(dx)$ and add 1 to $g_{jk}^{(l)}$. Also save the sequence of θ values visited by the walker trajectory.
 - Compute G_{ij} as
$$G_{ij} = \frac{\sum_{l'=L}^l g_{ij}^{(l')}}{(l-L+1)N},$$
where $L = \max[1, \min(l-L', L')]$, and L' is chosen to minimize bias from the preliminary trajectories (i.e., to provide a burn-in period for the simulation).
 - Update z by solving (13).
 - Check convergence criteria and go to step 2a if not satisfied. We examine both the change in the z vector and the number of complete reactive trajectories.
3. When a walker enters $A \cup B$, reconstruct the sequence of walkers leading to that event from the last exit of $A \cup B$ and in turn the sequence of θ values and strata visited. For each stratum i and each point $\theta(t)$ in the sequence, determine the nearest grid point Θ and
- increment $p_A^+(\Theta)$ (resp. $p_B^+(\Theta)$) by z_i if the sequence terminated in A (resp. B);
 - increment $p_A^-(\Theta)$ (resp. $p_B^-(\Theta)$) by z_i if the sequence originated in A (resp. B);
 - increment $v_{AB}(\Theta)$ (resp. $v_{BA}(\Theta)$) by $z_i[\theta(t+\tau) - \theta(t-\tau)]$ if the sequence originated in A and terminated in B (resp. originated in B and terminated in A).
- Above, z_i accounts for the weight of each trajectory segment, which is necessary to correct for the enhanced sampling.
4. Construct the TPT quantities for reactive trajectories from A to B as follows (exchange A and B below for reactive trajectories from B to A).
- Compute the forward committor to B as
$$q_+^\theta(\Theta) = \frac{p_B^+(\Theta)}{p_A^+(\Theta) + p_B^+(\Theta)}.$$
 - Compute the backward committor from A as
$$q_-^\theta(\Theta) = \frac{p_A^-(\Theta)}{p_A^-(\Theta) + p_B^-(\Theta)}.$$

Note that $p_A^- + p_B^- = p_A^+ + p_B^+$ is the total weight associated with all trajectories that start and end in $A \cup B$.
 - Compute the current from A to B as
$$j_{AB}^\theta(\Theta) = \frac{v_{AB}(\Theta)}{2\tau\Delta\Theta \sum_{\Theta'} [p_A^+(\Theta') + p_B^+(\Theta')]}.$$
- Figure 1 shows how the algorithm samples the ensemble of reactive trajectories in a piecewise fashion and how they are reconstructed.




FIG. 1. Schematic of algorithm. (upper left) Define the metastable states and the strata (black grid) and initialize the data structures (step 1). In the example shown, trajectories are initialized from the upper left metastable state boundary, so the last metastable state visited is known. (upper right) In each stratum, draw configurations from $\gamma_j(dx)$, save pointers to the associated $\gamma_j(dx)$ elements, and then run unbiased dynamics from these configurations until the trajectories exit the stratum (steps 2b and 2c). A single example stratum is shown. (lower right) As the simulation progresses, each point in $\gamma_j(dx)$ may give rise to segments of reactive or unreactive trajectories. Save the collective variable values associated with these trajectory segments (step 2d). (lower left) When a reactive or unreactive trajectory is realized (i.e., a trajectory segment enters a metastable state), trace back from pointer to pointer to determine the sequences of strata and collective variable grid sites (gray grid) visited (step 3). Use these to increment unnormalized statistics with appropriate weights (steps 3a-c). Note that the collective variable grid sites and the strata need not coincide.

IV. NUMERICAL RESULTS

We demonstrate our method by analyzing the $C_{7ax} \rightarrow C_{7eq}$ transition of the alanine dipeptide (*N*-acetyl-alanyl-*N'*-methylamide) in vacuum. The peptide is represented by the CHARMM36m force field^{41–43}. The simulations are performed at 300 K with the SD (stochastic dynamics) integrator⁴⁴ in GROMACS 5.1.4⁴⁵; LINCS was used to constrain the lengths of bonds to hydrogen atoms; the step size was 1 fs, and the friction coefficient was 10 ps^{-1} . We used PLUMED 2.3^{46–48} to extract collective variables; simulations were terminated when walkers left their strata using the PLUMED COMMITTOR function.

The metastable states and strata are defined in terms of the ϕ and ψ dihedral angles. The metastable states

are taken to be circles of radius 20° around $(\phi, \psi) = (-83^\circ, 75^\circ)$ and $(70^\circ, -70^\circ)$, which correspond to the minima of the potential of mean force determined below. The system was initialized from preliminary trajectories that were generated by using Steered Transition Path Sampling (STePS)⁴⁰ to drive the system from each metastable state to the other. STePS builds up trajectories by repeatedly shooting bursts of short segments and then preferentially selecting those that make forward progress for continuation. Trajectories were initialized from points in the equilibrium ensemble within each metastable state, and progress was measured by the distance to the center of the other metastable state. The segment length was 1 ps, and the bias threshold was 0.75 (i.e., forward trajectories are selected with a probability of 0.75 or their probability estimated from the burst, whichever is higher). Initially, 10 trajectories were released in each burst; subsequently, we varied the number with the progress distance such as to expect two forward segments in each burst based on statistics of the previous run. We do not use the statistics to correct for the STePS bias—i.e., each reactive trajectory is weighted equally.

The strata form a partition of unity; their centers are at $(-7 + 2n)30^\circ$ for integer $n \in [1, 6]$, and their widths are 100° , so they overlap 40° both ways, and the strata wrap around to enforce periodicity. We release 10 walkers in every stratum in each iteration and run each iteration until all 10 walkers exit their stratum. The collective variables ϕ and ψ and their increments are saved every 100 fs. The lag time used to compute the increments (τ in (8)) was 2 ps. The simulation is run until we see at least 1000 crossings in each direction between the metastable states, i.e., walkers that cross from $j \leq K$ to $j > K$ and the other way. We used a burn-in period of $L' = 100$ iterations. The full simulation was approximately 140 million steps. Though we stratify in ϕ and ψ , we can project the results onto other variables as well, and we also consider ω (Figure 2). Statistics are accumulated for 50 values of each dihedral angle.

Potentials of mean force (PMF) computed by NEUS are shown in Figure 3. In Figure 4 we show the forward committor and the $A \rightarrow B$ and $B \rightarrow A$ currents in the $\phi\psi$ -plane. We can see the clear existence of two main pathways, one of which cuts diagonally across the middle of the figure and one of which cuts across the upper left and lower right corners. Three representative trajectories consistent with the pathways are shown in Figure 5. In Figure 6, we show the committor and currents in the $\phi\omega$ -plane. We observe a diagonal structure in these plots that is consistent with coupling between distortion of the peptide plane and ϕ noted previously⁴⁹.

To assess the accuracy of our committor estimates, we compute the committor as a function of all three dihedral angles (not shown), select configurations predicted




FIG. 2. Alanine dipeptide with the three dihedral angles labeled: ϕ (C5-N7-C9(α)-C15), ψ (N7-C9(α)-C15-N17), and (C1-C5-N7-C9(α)). Colors are cyan for carbon, blue for nitrogen, white for hydrogen, and red for oxygen.

to have $0.49 < q_+^\theta < 0.51$, and evaluate their committors by shooting 20 independent simulations from each configuration. The resulting histogram of values is peaked at $q_+ \approx 0.5$, as desired. To characterize the convergence of the committor estimates, we compute the backwards committor and plot $q_+^\theta + q_-^\theta - 1$, which should be close to zero for this system (Figures S1 and S2); the deviation provides an indication of the numerical error. We see that after 250 crossings between the metastable states the error overall is small, with the largest deviations in regions with high free energy close to where $q_+^\theta = 0.5$. These deviations decrease with additional crossings.

To assess the reactive current estimates, we compute the rate as a function of the number of crossings two ways. First we directly sum the flux into $B^{26,27}$:

$$R_{\text{NEUS}} = \frac{1}{\sum_{j=1}^{2K} z_j} \sum_{j \leq K} z_j \frac{n_j^B}{T_j}, \quad (22)$$

where n_j^B is the number of reactive trajectories that enter B from stratum j , T_j is the total time simulated in that stratum, and the sum over $j \leq K$ selects for contributions from trajectories that were last in A . Second, we sum the reactive currents that cross a dividing surface S :

$$R_{\text{TPT}} = \sum_S \int_{AB}^\theta(\Theta) \cdot \hat{n}_S \Delta_S, \quad (23)$$

where $\hat{n}_S$ is the unit vector normal to S , and Δ_S is a differential element with the same dimension as S . Since the collective variable space is periodic, there are two pathways, and we average over a set of surfaces for each. Specifically, we use vertical lines in the $\phi\psi$ plots: (i)




FIG. 3. Potentials of mean force computed from the NEUS simulations: projections on (top) ϕ and ψ and (bottom) ϕ and ω . The stable states are marked with white circles on the top plot.

$\phi = (-57.6 + 7.2n)^\circ$ for integer n in $[0, 15]$, which cut through the middle of Figure 4, and (ii) $\phi = (-172.8 + 7.2n)^\circ$ for integer n in $[0, 9]$ and $[37, 49]$, which line the left and right sides of Figure 4.

The two estimates are about a factor of two different from each other (Figure 8); we also obtain a direct estimate of $1.4 \times 10^{-6} \text{ ps}^{-1}$ from an unbiased trajectory of length $2.5 \mu\text{s}$. We consider the three estimates to be good agreement given that the rate can be a very challenging statistic to converge to even its order of magnitude. We expect (22) to be most precise (as evidenced by the scale of fluctuations in Figure 8) and recommend its use; we consider (23) only as a means of validating the reactive currents. That said, we can use the currents to obtain




FIG. 4. Committer (top) and forward (middle) and backward (bottom) reactive currents for the C7ax $\rightarrow$ C7eq transition of the alanine dipeptide, plotted on the $\phi\psi$ -plane.






FIG. 5. Example trajectories plotted on the $\phi\psi$ -plane. Dots mark the centers of the metastable states.

the fluxes associated with each of the pathways. We find them to be $R_{\text{TPT}} = 4.6 \times 10^{-7} \text{ ps}^{-1}$ for the pathway that crosses $\phi = 0^\circ$ and $R_{\text{TPT}} = 8.1 \times 10^{-7} \text{ ps}^{-1}$ for the pathway that crosses $\phi = 180^\circ$.

To further assess the reactive currents, we also compute $J_{AB}^\theta(\Theta) + J_{BA}^\theta(\Theta)$ (Figure S3). We expect this sum to be close to but not exactly equal to zero given that the integrator is underdamped. We see that this is the case, with the only significant deviations close to the metastable states. Put together, these results validate the method.

V. CONCLUSIONS

In this paper, we investigated the computation of committers and reactive currents with trajectory stratification. The fundamental challenge is that these quantities require knowledge of the metastable states in which trajectories begin and end, but this information cannot be obtained directly from most walkers within the simulation. To address this issue, we store the collective variable values visited by each walker and the associated collective variable increments sampled, as well as a pointer to the previous walker. This enables us to reconstruct the sequence of walkers that gave rise to a trajectory once it reaches a metastable state, despite the fact that the number of trajectories to which a single walker can contribute grows exponentially with the number of walkers that follow it. The current work builds on previous trajectory stratification studies that showed that computation of the rate requires separating the ensemble of trajectories based on the metastable state in which they originate^{26,27}.



Histogram of committors computed by shooting for 137 structures predicted to have $0.49 < q_+^\theta < 0.51$ for reaction from A to B . 20 independent simulations were used for each configuration.







FIG. 8. Rates computed directly from the flux into B (top) and by integrating the reactive currents (bottom). An equal number of crossings between metastable states each way is used for each point on the horizontal axis.

FIG. 6. Committor (top) and forward (middle) and backward (bottom) reactive currents for the $C7_{ax} \rightarrow C7_{eq}$ transition of the alanine dipeptide, plotted on the $\phi\omega$ -plane. For the reactive currents, only grid points with data from at least 5 trajectories are shown.

As in previous trajectory stratification studies, the results are exact in the limit of infinite sampling, in contrast to MSMs and DGA.

Although we demonstrated the approach with NEUS,

it is general and can be applied to other algorithms that sample trajectory segments^{18–24,30}. By the same token, we focused on a simple numerical example that permits validation of the results. Extending the approach to more complex systems should be straightforward, so long as collective variables that enable an efficient stratification and the accumulation of statistics can be identified. This challenge is system specific, though methods for facilitating the selection of collective variables based on simulation data exist (e.g., refs. 1 and 10). In the case of NEUS, the most significant challenge may be that the lists representing $\gamma_{ij}(dx)$ can become large (up to 5500 elements in the present example). *Although this issue has not been limiting for complex systems treated to date*⁵⁰, it is an important practical consideration that we leave for future work.

VI. SUPPLEMENTARY MATERIAL

Justification of (8), committor convergence plots, sums of forward and backward reactive currents.

VII. ACKNOWLEDGMENTS

The authors wish to thank Chatipat Lorpai boon and John Strahan for useful discussions and critical readings of the manuscript. This work was supported by National Institutes of Health award R35 GM136381 and National Science Foundation award DMS-2054306.

- ¹A. Ma and A. R. Dinner, “Automatic method for identifying reaction coordinates in complex systems,” *Journal of Physical Chemistry B* **109**, 6769–6779 (2005).
- ²E. Vanden-Eijnden, “Transition path theory,” in *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, edited by M. Ferrario, G. Ciccotti, and K. Binder (Springer Berlin Heidelberg, Berlin, Heidelberg, 2006) pp. 453–493.
- ³W. E. and E. Vanden-Eijnden, “Towards a theory of transition paths,” *Journal of Statistical Physics* **123**, 503 (2006).
- ⁴P. Metzner, C. Schütte, and E. Vanden-Eijnden, “Illustration of transition path theory on a collection of simple examples,” *Journal of Chemical Physics* **125**, 084110 (2006).
- ⁵F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, “Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations,” *Proceedings of the National Academy of Sciences* **106**, 19011–19016 (2009).
- ⁶G. R. Bowman, V. S. Pande, and F. Noé, *An Introduction to Markov State Models and their Application to Long Timescale Molecular Simulation*, Vol. 797 (Springer Science & Business Media, 2013).
- ⁷H. Wu, F. Nüske, F. Paul, S. Klus, P. Koltai, and F. Noé, “Variational Koopman models: Slow collective variables and molecular kinetics from short off-equilibrium simulations,” *Journal of Chemical Physics* **146**, 154104 (2017).
- ⁸N.-V. Buchete and G. Hummer, “Coarse master equations for peptide folding dynamics,” *Journal of Physical Chemistry B* **112**, 6057–6069 (2008).

- ⁹E. H. Thiede, D. Giannakis, A. R. Dinner, and J. Weare, “Galerkin approximation of dynamical quantities using trajectory data,” *Journal of Chemical Physics* **150**, 244111 (2019).
- ¹⁰J. Strahan, A. Antoszewski, C. Lorpai boon, B. P. Vani, J. Weare, and A. R. Dinner, “Long-time-scale predictions from short-trajectory data: A benchmark analysis of the trp-cage miniprotein,” *Journal of chemical theory and computation* **17**, 2948–2963 (2021).
- ¹¹J. D. Russo, J. Copperman, D. Aristoff, G. Simpson, and D. M. Zuckerman, “Unbiased estimation of equilibrium, rates, and committors from Markov state model analysis,” arXiv:2105.13402 (2021).
- ¹²R. Elber, J. M. Bello-Rivas, P. Ma, A. E. Cardenas, and A. Fathizadeh, “Calculating iso-committor surfaces as optimal reaction coordinates with milestoning,” *Entropy* **19**, 219 (2017).
- ¹³A. K. Faradjian and R. Elber, “Computing time scales from reaction coordinates by milestoning,” *Journal of Chemical Physics* **120**, 10880–10889 (2004).
- ¹⁴R. Elber, “Milestoning: An efficient approach for atomically detailed simulations of kinetics in biophysics,” *Annual Review of Biophysics* **49**, 69–85 (2020).
- ¹⁵C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, “Transition path sampling and the calculation of rate constants,” *Journal of Chemical Physics* **108**, 1964–1977 (1998).
- ¹⁶C. Dellago, P. G. Bolhuis, and D. Chandler, “Efficient transition path sampling: Application to Lennard-Jones cluster rearrangements,” *Journal of Chemical Physics* **108**, 9236–9245 (1998).
- ¹⁷C. Dellago, P. Bolhuis, P. L. Geissler, *et al.*, “Transition path sampling,” *Advances in Chemical Physics* **123** (2002).
- ¹⁸G. A. Huber and S. Kim, “Weighted-ensemble Brownian dynamics simulations for protein association reactions,” *Biophysical Journal* **70**, 97–110 (1996).
- ¹⁹B. W. Zhang, D. Jasnow, and D. M. Zuckerman, “Efficient and verified simulation of a path ensemble for conformational change in a united-residue model of calmodulin,” *Proceedings of the National Academy of Sciences* **104**, 18043–18048 (2007).
- ²⁰D. M. Zuckerman and L. T. Chong, “Weighted ensemble simulation: review of methodology, applications, and software,” *Annual Review of Biophysics* **46**, 43–57 (2017).
- ²¹P. G. Bolhuis, “Transition-path sampling of β -hairpin folding,” *Proceedings of the National Academy of Sciences* **100**, 12129–12134 (2003).
- ²²T. S. van Erp, D. Moroni, and P. G. Bolhuis, “A novel path sampling method for the calculation of rate constants,” *Journal of Chemical Physics* **118**, 7762–7774 (2003).
- ²³R. J. Allen, P. B. Warren, and P. R. Ten Wolde, “Sampling rare switching events in biochemical networks,” *Physical Review Letters* **94**, 018104 (2005).
- ²⁴R. J. Allen, C. Valeriani, and P. R. ten Wolde, “Forward flux sampling for rare event simulations,” *Journal of Physics: Condensed Matter* **21**, 463102 (2009).
- ²⁵A. Warmflash, P. Bhimalapuram, and A. R. Dinner, “Umbrella sampling for nonequilibrium processes,” *Journal of Chemical Physics* **127**, 114109 (2007).
- ²⁶A. Dickson, A. Warmflash, and A. R. Dinner, “Separating forward and backward pathways in nonequilibrium umbrella sampling,” *Journal of Chemical Physics* **131**, 154104 (2009).
- ²⁷E. Vanden-Eijnden and M. Venturoli, “Exact rate calculations by trajectory parallelization and tilting,” *Journal of Chemical Physics* **131**, 044120 (2009).
- ²⁸A. Dickson and A. R. Dinner, “Enhanced sampling of nonequilibrium steady states,” *Annual Review of Physical Chemistry* **61**, 441–459 (2010).
- ²⁹A. R. Dinner, E. H. Thiede, B. Van Koten, and J. Weare, “Stratification as a general variance reduction method for Markov chain Monte Carlo,” *SIAM/ASA Journal on Uncertainty Quantification* **8**, 1139–1188 (2020).

- ³⁰J. M. Bello-Rivas and R. Elber, "Exact milestoning," *Journal of Chemical Physics* **142**, 03B602_1 (2015).
- ³¹D. Bhatt, B. W. Zhang, and D. M. Zuckerman, "Steady-state simulations using weighted ensemble path sampling," *Journal of Chemical Physics* **133**, 014110 (2010).
- ³²A. R. Dinner, J. C. Mattingly, J. O. Tempkin, B. Van Koten, and J. Weare, "Trajectory stratification of stochastic dynamics," *SIAM Review* **60**, 909–938 (2018).
- ³³C. Lorpaiboon, J. Weare, and A. R. Dinner, "Augmented transition path theory for sequences of events," in preparation (2022).
- ³⁴P. Metzner, C. Schütte, and E. Vanden-Eijnden, "Transition path theory for Markov jump processes," *Multiscale Modeling & Simulation* **7**, 1192–1219 (2009).
- ³⁵G. M. Torrie and J. P. Valleau, "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling," *Journal of Computational Physics* **23**, 187–199 (1977).
- ³⁶C. Pangali, M. Rao, and B. Berne, "A Monte Carlo simulation of the hydrophobic interaction," *Journal of Chemical Physics* **71**, 2975–2981 (1979).
- ³⁷D. Frenkel and B. Smit, *Understanding Molecular Simulation: from Algorithms to Applications* (Elsevier, 2001).
- ³⁸E. H. Thiede, B. Van Koten, J. Weare, and A. R. Dinner, "Eigenvector method for umbrella sampling enables error analysis," *Journal of Chemical Physics* **145**, 084115 (2016).
- ³⁹A. Dickson, A. Warmflash, and A. R. Dinner, "Nonequilibrium umbrella sampling in spaces of many order parameters," *Journal of Chemical Physics* **130**, 02B605 (2009).
- ⁴⁰N. Guttenberg, A. R. Dinner, and J. Weare, "Steered transition path sampling," *Journal of Chemical Physics* **136**, 234103 (2012).
- ⁴¹A. D. MacKerell, D. Bashford, M. Bellott, R. L. Dunbrack, J. D. Evanseck, M. J. Field, S. Fischer, J. Gao, H. Guo, S. Ha, D. Joseph-McCarthy, L. Kuchnir, K. Kuczera, F. T. K. Lau, C. Mattos, S. Michnick, T. Ngo, D. T. Nguyen, B. Prodhom, W. E. Reiher, B. Roux, M. Schlenkrich, J. C. Smith, R. Stote, J. Straub, M. Watanabe, J. Wiórkiewicz-Kuczera, D. Yin, and M. Karplus, "All-atom empirical potential for molecular modeling and dynamics studies of proteins," *Journal of Physical Chemistry B* **102**, 3586–3616 (1998).
- ⁴²R. B. Best, X. Zhu, J. Shim, P. E. M. Lopes, J. Mittal, M. Feig, and A. D. MacKerell, "Optimization of the additive CHARMM all-atom protein force field targeting improved sampling of the backbone ϕ , ψ and side-chain χ_1 and χ_2 dihedral angles," *Journal of Chemical Theory and Computation* **8**, 3257–3273 (2012).
- ⁴³J. Huang, S. Rauscher, G. Nawrocki, T. Ran, M. Feig, B. L. de Groot, H. Grubmüller, and A. D. M. Jr, "CHARMM36m: An improved force field for folded and intrinsically disordered proteins," *Nature Methods* **14**, 71–73 (2017).
- ⁴⁴N. Goga, A. Rzepiela, A. De Vries, S. Marrink, and H. Berendsen, "Efficient algorithms for Langevin and DPD dynamics," *Journal of Chemical Theory and Computation* **8**, 3637–3649 (2012).
- ⁴⁵M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindah, "GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers," *SoftwareX* **1-2**, 19–25 (2015).
- ⁴⁶M. Bonomi, D. Branduardi, G. Bussi, C. Camilloni, D. Provasi, P. Raiteri, D. Donadio, F. Marinelli, F. Pietrucci, R. A. Broglia, and M. Parrinello, "PLUMED: A portable plugin for free-energy calculations with molecular dynamics," *Computer Physics Communications* **180**, 1961–1972 (2009).
- ⁴⁷G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, "PLUMED 2: New feathers for an old bird," *Computer Physics Communications* **185**, 604–613 (2014).
- ⁴⁸M. Bonomi, G. Bussi, C. Camilloni, G. A. Tribello, P. Banáš, A. Barducci, M. Bernetti, P. G. Bolhuis, S. Bottaro, D. Branduardi, R. Capelli, and P. Carloni, "Promoting transparency and reproducibility in enhanced molecular simulations," *Nature Methods* **16**, 670–673 (2019).
- ⁴⁹P. G. Bolhuis, C. Dellago, and D. Chandler, "Reaction coordinates of biomolecular isomerization," *Proceedings of the National Academy of Sciences* **97**, 5877–5882 (2000).
- ⁵⁰A. Dickson, M. Maienschein-Cline, A. Tovo-Dwyer, J. R. Hammond, and A. R. Dinner, "Flow-dependent unfolding and refolding of an RNA by nonequilibrium umbrella sampling," *Journal of Chemical Theory and Computation* **7**, 2710–2720 (2011).