PNAS RESEARCH ARTICLE | CHEMISTRY
BIOPHYSICS AND COMPUTATIONAL BIOLOGY

Highly parallelizable path sampling with minimal rejections using asynchronous replica exchange and infinite swaps

Daniel T. Zhang[a,1], Lukas Baldauf[a,1], Sander Roet[b], Anders Lervik[a], and Titus S. van Erp[a,2]

Edited by Pablo Debenedetti, Princeton University, Princeton, NJ; received October 26, 2023; accepted January 9, 2024

Capturing rare yet pivotal events poses a significant challenge for molecular simulations. Path sampling provides a unique approach to tackle this issue without altering the potential energy landscape or dynamics, enabling recovery of both thermodynamic and kinetic information. However, despite its exponential acceleration compared to standard molecular dynamics, generating numerous trajectories can still require a long time. By harnessing our recent algorithmic innovations—particularly subtrajectory moves with high acceptance, coupled with asynchronous replica exchange featuring infinite swaps—we establish a highly parallelizable and rapidly converging path sampling protocol, compatible with diverse high-performance computing architectures. We demonstrate our approach on the liquid–vapor phase transition in superheated water, the unfolding of the chignolin protein, and water dissociation. The latter, performed at the ab initio level, achieves comparable statistical accuracy within days, in contrast to a previous study requiring over a year.

rare events | path sampling | asynchronous replica exchange | infinite swapping | Markov-chain Monte Carlo

Significance

Current molecular dynamics simulations have the capability to faithfully describe chemical reactions, nucleation, and protein folding events, providing essential information to drive the progress of technologies like catalysis and drug discovery. Nevertheless, even with the fastest supercomputers, a significant portion of these phenomena remains beyond the reach of standard simulations due to waiting times that can exceed billions of CPU years. While enhanced sampling techniques like path sampling allow for exponentially faster study of these events, obtaining converged results in experimentally relevant systems still typically spans from months to years. Here, we apply four innovative algorithmic enhancements that transform the state-of-the-art path sampling algorithm, replica exchange transition interface sampling (RETIS), into ∞RETIS, allowing convergence in a matter of days.

The capacity to rapidly and accurately simulate molecular transition phenomena holds the potential to significantly enhance chemical discoveries, thereby advancing catalytic processes (1), optimizing drug molecule design (2), and guiding self-assembly for various applications, such as organic photovoltaics (3). However, dynamic processes like chemical reactions, nucleation, or protein (un)folding usually hinge on rare molecular events, rendering direct molecular dynamics (MD) simulations ineffective (4). A way to bridge the time gap is to use rare event sampling techniques like the Markov chain Monte Carlo (MC)-based transition path sampling (TPS) method, which involves the collection of numerous short MD trajectories (5).

Transition interface sampling (TIS) (6) and, even more efficiently, replica exchange TIS (RETIS) (7) build upon this idea to calculate quantitative dynamical properties through a series of path sampling simulations, each targeting a distinct path ensemble reflecting different stages of the transition. Each trajectory evolves on the true potential energy surface, and the sampling of trajectories follows the same distributions as what would result if the relevant trajectories were extracted from a hypothetically long MD run. Yet, the distinctive feature of path sampling simulations lies in their computational emphasis on actual barrier-crossing events, which stands in contrast to plain MD where the computational effort is primarily directed toward explorations within stable states. Despite exponential speedup compared to direct MD, the study of complex systems can still require months of simulation time due to the necessity of generating numerous trajectories for achieving the required statistical precision.

In this paper, we leverage recent algorithmic innovations that achieve such results in a matter of days or weeks. This transformative progress is driven by harnessing four recent algorithmic innovations delineated in refs. 8 and 9. Initially, we improve the core MC path generation move, opting for a sequence of intermediate short subtrajectories, yielding enhanced decorrelation from the preceding trajectory upon acceptance. Subsequently, by slightly adjusting the sampling distribution and compensating through reweighting, we maximize the acceptance. Third, the integration of an asynchronous replica exchange scheme facilitates seamless swapping between path ensemble simulations, tackling the challenge of RETIS parallelization attributed to varying central processing unit (CPU) costs in path-generating MC moves. Last, we amplify the impact of computationally efficient replica exchange moves through the embrace of the infinite swapping limit (10), all while circumventing the need for combinatorially explosive computations.

While the mathematical proofs establishing the exactness of these algorithms were published in refs. 8 and 9, this article demonstrates their first implementation and efficient management of numerous realistic molecular processes in parallel, leveraging

Author contributions: A.L. and T.S.v.E designed research; D.T.Z. and L.B. performed research; D.T.Z., L.B., S.R., A.L., and T.S.v.E. analyzed data; and D.T.Z., L.B., S.R., A.L., and T.S.v.E. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]D.T.Z. and L.B. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: titus.van.erp@ntnu.no.

Published February 5, 2024.

both classical and ab initio dynamics on high-performance computing (HPC) systems. Notably, the article showcases a large-scale simulation utilizing 40 graphics processing units (GPUs) in parallel— overcoming the challenge of effectively utilizing multiple GPUs for parallel molecular simulations.

Results

Path Sampling and RETIS Path Ensembles. MC techniques are valuable in various fields like statistical physics, finance, and artificial intelligence, where the common goal is to sample states following specific probability distributions. Within molecular simulations, MC is generally used to sample configuration space following the Boltzmann distribution. By an original insight from Pratt (11), the concept emerged that MC sampling could be applied to target the sampling of dynamic trajectories. In this framework, a trajectory referred to as X is depicted as a discrete sequence of phase points, called time slices $X = [x_0, x_1, \ldots x_L]$. Each time slice x_i encapsulates the coordinates and velocities of atoms at a specific time $t = i\Delta t$. Here, L denotes the trajectory's length, and Δt represents a small time step. The path probability distribution equals $\rho(X) \propto \rho(x_0)p(x_0 \to x_1)p(x_1 \to x_2)\ldots p(x_{L-1} \to x_L)$, where $\rho(x_0)$ is the equilibrium (Boltzmann) distribution of the first phase point and $p(x \to y)$ is the probability that the system's dynamics produces y after a Δt time step from x.

Applying this within an MC algorithm does not yet yield advantages over MD. However, the approach enables focusing on specific path ensembles defined by initial and terminal conditions, and/or reaction progress through sampling from a truncated distribution $\rho_E(X) = \rho(X) \cdot \mathbf{1}_E(X)$, where $\mathbf{1}_E(X)$ equals 1 if the path adheres to the ensemble E conditions, and 0 otherwise. To obtain dynamical quantitative results such as rates, a series of overlapping path ensemble simulations is needed. The RETIS ensembles possess both initial and final conditions, as well as a minimum progress requirement, that is gauged through a series of nonintersecting interfaces: $\lambda_0, \lambda_1, \ldots, \lambda_n$ (Fig. 1*A*). These interfaces are hypersurfaces within phase space, often characterized by an order parameter $\lambda(x)$ that assigns a progress value to the reaction. The ith interface corresponds to the collection of phase points $\{x|\lambda(x) = \lambda_i\}$. The first and last interfaces define the reactant state A, $\{x|\lambda(x) < \lambda_A = \lambda_0\}$, and the product state B, $\{x|\lambda(x) > \lambda_B = \lambda_n\}$, respectively. The RETIS path ensemble $[i^+]$ encompasses all paths that commence by crossing λ_A toward the barrier region and conclude by either re-entering A or entering B. Moreover, each path within the ensemble is required to cross λ_i. This implies that the value of L is not fixed but varies for each path and the average path length typically increases with i. Alongside the $[i^+]$ ensembles, there exists an additional $[0^-]$ ensemble that explores the internal realm of state A (Fig. 1*B*).

Subtrajectory Moves. The primary MC move for generating paths has been the shooting move (12). It evolves by modifying the velocities of a random time slice of the old path, which is then propagated forward and backward in time using the MD time step integrator. Fine-tuning the shooting move necessitates a delicate balance between maximizing decorrelation and maintaining a satisfactory acceptance. When the adjustment to the shooting point is minimal, the resulting path often closely resembles its predecessor. Although this enhances the chance of the trial path being valid for the considered path ensemble, the substantial correlation among sampled paths necessitates a large number of trajectories to achieve low statistical errors. In subtrajectory moves, the creation of complete trial trajectories involves preceding them with several intermediate short subtrajectories. This ensures that successive accepted full trajectories do not share any configuration points and exhibit a greater degree of distinctiveness compared to shooting. The subtrajectories are only part of the inner workings of the MC move and are not stored or used for statistical analysis.

The initial subtrajectory moves (13) encountered certain implementation challenges which led to the development of the more flexible wire-fencing (WF) (8). The WF move involves a parameter N_s, representing the number of subtrajectories, and potentially a cap interface (Fig. 1*C*). Within the ensemble $[i^+]$, the WF move entails releasing a sequence of N_s short subtrajectories with termination criteria at λ_i or λ_{cap} (or λ_n if no cap is set). The first subtrajectory originates from a time slice of the previous full path with randomized Maxwell–Boltzmann velocities. Subsequent subtrajectories are generated from the previous successful subtrajectory until N_s attempts have been made. A subtrajectory ending at λ_{cap} in both time directions is classified as unsuccessful. After the completion of N_s subtrajectory trials, the last successful subtrajectory is integrated forward and backward in time until reaching either λ_0 or λ_n, resulting in the formation of a new full trajectory. Based on a Metropolis–Hastings acceptance/rejection scheme (14, 15), this new path can be accepted with a probability equal to

$$P_{\text{acc}} = \mathbf{1}_{[i^+]}(X^{(n)}) \times \min\left[1, \frac{M_{\lambda_i}(X^{(o)})}{M_{\lambda_i}(X^{(n)})}\right], \qquad [1]$$

where $X^{(o)}$ and $X^{(n)}$ are, respectively, the old and new paths, and $M_{\lambda_i}(X)$ is the number of possible shooting points for releasing a first subtrajectory from X. It was found that subtrajectory moves (with high acceptance) have the potential to improve RETIS' CPU efficiency by a factor of twelve compared to shooting (13).

High Acceptance. The high acceptance technique represents a significant enhancement for advanced shooting moves by introducing two algorithmic modifications: i) Reverse the time direction of the full trial path if it starts at λ_n and ends at λ_0 (Fig. 1*C*). ii) Modify the targeted sampling distribution from $\rho_{[i^+]}(X)$ to $\tilde{\rho}_{[i^+]}(X) = \rho_{[i^+]}(X) \cdot w_i(X)$, where $w_i(X)$ denotes the high-acceptance weight defined as:

$$w_i(X) = M_{\lambda_i}(X)q(X) \quad \text{with} \qquad [2]$$

$$q(X) = \begin{cases} 1 & \text{if } X \text{ is of the type } \lambda_0 \to \lambda_0 \\ 2 & \text{if } X \text{ is of the type } \lambda_0 \to \lambda_n \end{cases}.$$

These two actions culminate in the outcome that nearly all trajectories become viable for acceptance as P_{acc} of Eq. **1** becomes identical to the indicator function, $P_{\text{acc}} = \mathbf{1}_{[i^+]}(X^{(n)})$, which is always 1 for any trial trajectory within this scheme except for those ending at λ_n in both time directions. That chance is mostly negligible for all $[i^+]$ path ensembles except where λ_i is near the peak of the barrier or beyond. The only other reason for a rejection is when all N_s subtrajectories reach the cap-interface in both temporal directions. In that case, the "last successful" subtrajectory essentially comprises the path segment of the old path from which the initial shooting occurred. Extending this path segment likely generates a highly similar path to the old one, leading us to reject it in our WF implementation instead of investing CPU time in producing a significantly correlated path.

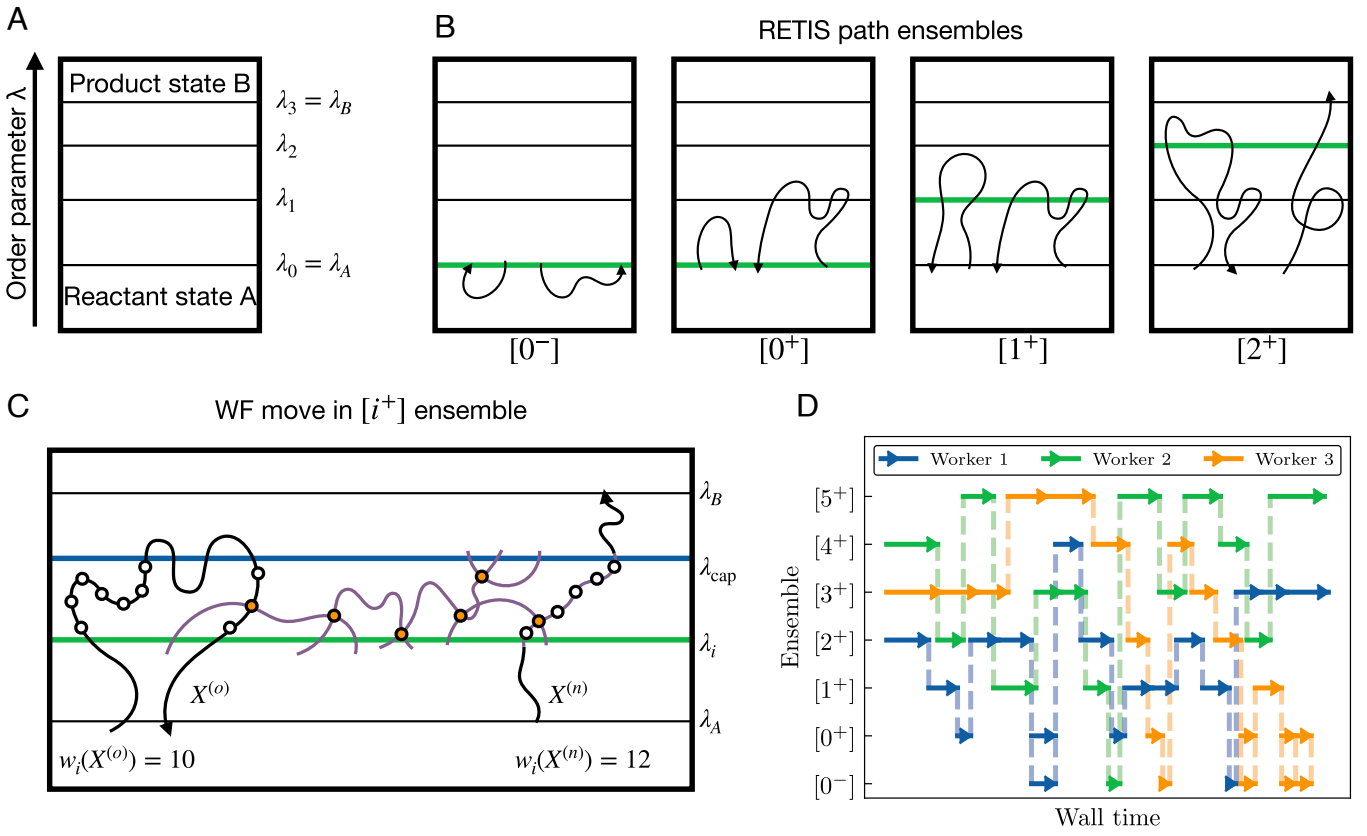

Fig. 1. RETIS path ensembles and path sampling methodology. (*A*) Concept of interfaces and states based on an order parameter λ (reaction coordinate). The horizontal axis represents an arbitrary additional order parameter. In this example, four interfaces are defined: λ_0, λ_1, λ_2, and λ_3. The first and last interfaces define the reactant and product states, respectively. (*B*) RETIS path ensembles. The minimal progress interface is highlighted in green. Two representative trajectories are shown for each ensemble. Two trajectories in $[0^+]$ and $[1^+]$ are identical, illustrating that ensembles overlap and paths may be sampled in several ensembles via a swapping move. (*C*) Demonstration of the WF move with $N_s = 6$. The fourth subtrajectory is unsuccessful and subsequently dismissed. Shooting points are indicated as orange circles and additional potential shooting points on both the old (X^o) and new (X^n) paths are represented by white circles. Shooting from a $\lambda_{\text{cap}} \rightarrow \lambda_{\text{cap}}$ segment is disallowed. With the extension of the last subtrajectory, a time direction is randomly chosen, which is flipped in the high-acceptance scheme if the resulting path is of the type $\lambda_B \rightarrow \lambda_A$. The resulting high-acceptance weights for the old and new paths are, respectively, 10 and 12, based on $M_{\lambda_i}(X^{(0)}) = 10$, $q(X^{(0)}) = 1$, $M_{\lambda_i}(X^{(n)}) = 6$, and $q(X^{(n)}) = 2$. (*D*) Time spent per ensemble per worker in an actual asynchronous replica exchange simulation double-well system (9). Arrows denote the moments when a worker completes a path to initiate the exchange of replicas between free ensembles. Minimal computational time is consumed during this process and concludes when the worker is randomly reassigned to another (or the same) free ensemble for a new path generation move. When both $[0^+]$ and $[0^-]$ ensembles are free, a point-exchange move is also incorporated into the random reassignment. In this move, the worker creates two new paths: one in $[0^+]$ by extending the endpoint of the $[0^-]$-path forward in time and another in $[0^-]$ by extending the starting point of the $[0^+]$-path backward in time (7).

The post-simulation analysis counteracts the impact of the distorted distribution by employing weighted averages for the sampled paths, assigning each path X a weight proportional to $1/w_i(X)$. In a simple one-dimensional double-well potential, the acceptance rate of the WF move stood at 100% for the path ensembles $[i^+]$ when λ_i was near state A and only slightly decreased to 99.2% in the ensemble closest to state B (8).

Asynchronous Replica Exchange. Despite TIS being significantly less efficient than RETIS, it has the advantage that its separate path ensembles can be simulated entirely autonomously, allowing parallel execution without communication overhead. In RETIS, however, path-generating MC moves within a single ensemble are alternatively succeeded by replica exchange moves between ensemble simulations. The irregular CPU costs of path-generating MC moves, stemming from the diverse path lengths, introduce synchronization challenges within a parallel RETIS simulation setup.

Assigning individual path ensembles with their own hardware setup leads to instances where hardware managing faster ensembles frequently remain idle, awaiting the completion of MC moves by their slower counterparts. For this reason, open-source path-sampling codes (16, 17) have implemented RETIS as a fully sequential algorithm. However, this design choice limits its potential to run simulations in a massively parallel manner.

In ref. 9, this challenge was addressed through an asynchronous replica exchange approach, where the number of path ensembles is set to be approximately double the number of hardware groups (referred to as "workers") that are assigned to execute path generation moves. This design ensures that, at any given moment, about half of the ensembles are "busily" engaged in path creation, while the other half is labeled as "free." Following the completion of an MC move by a worker, the ensemble it was assigned to and the newly formed path change status to free. Before the worker is randomly reassigned to one of the free ensembles for performing a new path generation move, a series of swapping moves take place between randomly selected pairs of free ensembles in which they attempt to exchange their current paths (Fig. 1*D*). For a selected ensemble pair, $[i^+]$ and $[j^+]$ with $j > i$ and, respectively, current paths X_i and X_j are swapped with an acceptance probability:

$$P_{\text{acc}} = \mathbf{1}_{[j^+]}(X_i) \times \min\left[1, \frac{w_i(X_j)w_j(X_i)}{w_i(X_i)w_j(X_j)}\right], \qquad [3]$$

where term $\mathbf{1}_{[i^+]}(X_j)$ is omitted as $X_j \in [j^+]$ is always a valid trajectory for ensemble $[i^+]$ if $\lambda_j \geq \lambda_i$.

Ref. 9 demonstrated that asynchronous replica exchange significantly enhances wall time efficiency while minimally impacting CPU efficiency. Surprisingly, it even led to occasional improvements in CPU efficiency due to a more efficient distribution of CPU resources among different path ensembles. The algorithm tends to generate more trajectories in path ensembles with shorter average path lengths, which contributes positively to the overall efficiency.

Infinite Swapping. Generating a complete path may span minutes or hours, while evaluating Eq. 3 takes sub-seconds. This allows for numerous swap moves, but when does it become excessive? Plattner et al. (10) showed the feasibility of replicating the impact of executing an infinite number of swaps within finite CPU time, potentially maximizing the benefit of each swapping opportunity. To determine the frequency of sampling a specific state (path) in a particular ensemble after an infinite number of swaps, one only needs to sum over the probabilities of all permutations in which the considered state and ensemble are linked. While this method is efficient for a modest number of participating ensembles ($\lesssim$10), computational costs increase dramatically, transitioning from approximately a single second of wall time to millions of years as the number of ensembles grows from 7 to 20 (9).

Remarkably, this factorial scaling obstacle can be addressed by employing an expression based on weight matrices' permanents, which is equivalent to the summation of permutations (9). Despite being similar to the determinant, commonly taught in high school mathematics textbooks, the permanent is relatively unfamiliar among scientific researchers, potentially contributing to the lack of prior discovery of this relationship.

Like determinants, a matrix's permanent is recursively defined as the sum of permanents of reduced matrices with a row and column removed, but unlike determinants, it lacks alternating plus and minus signs. Recursive relations also involve factorial scaling, but faster methods exist for large matrices, such as Gaussian elimination for determinants, leading to third-order scaling.

Unfortunately, this technique does not extend to the computation of permanents, for which more complex approaches are necessary, characterized by steeper scaling (18, 19). Nonetheless, these methods are still considerably faster than factorial computations. Furthermore, since many elements of the weight matrices are zero, permanents only need to be computed for a limited number of low-dimensional sub-blocks of the weight matrix. As demonstrated in this paper, this enables us to conduct infinite swapping replica exchanges involving 80 ensembles and 40 workers, with the infinite swaps constituting only a minor portion of the CPU cost compared to that of path creation.

Application I: (Superheated) Water Boiling. We employ RETIS with the aforementioned algorithmic advancements, hereafter referred to as ∞RETIS (9), to study the liquid–vapor phase transition. Boiling phenomena have previously been explored using TPS (20, 21). However, the ∞RETIS approach offers a notable advantage in quantification, enabling the calculation of rates– a feat not easily achieved with the previous TPS method, even for the more common occurrence of surface boiling. Hence, the previous TPS studies were more qualitative than our current investigation and did not provide information on transition rates.

In the first boiling study, we aimed to compute the boiling rate in superheated water at 573.15 K. Superheated liquid is produced by gently heating a liquid beyond its boiling point (22). While establishing superheated water at this temperature is difficult experimentally, in our nano-sized simulation system devoid of nucleation sites like walls or impurities, the metastable liquid has a long lifetime. Consequently, this transition to vapor serves as a unique test for assessing our path sampling protocol's hardware scaling capabilities.

We conduct two simulations using different numbers of identical, GPU-equipped nodes. In the first simulation i), 20 interfaces and 10 workers run on one node, utilizing the NVIDIA Multi-Process Service (MPS) feature, which enables multiple compute unified device architecture (CUDA) processes or applications to share and utilize a single GPU. The other simulation ii) involves 40 interfaces and 20 workers, each operating on their exclusive nodes without MPS.

Fig. 2*A* illustrates the crossing probabilities, $P_A(\lambda|\lambda_A)$, of these simulations. This is the probability that the system's order parameter function $\lambda(x)$ reaches the order parameter value

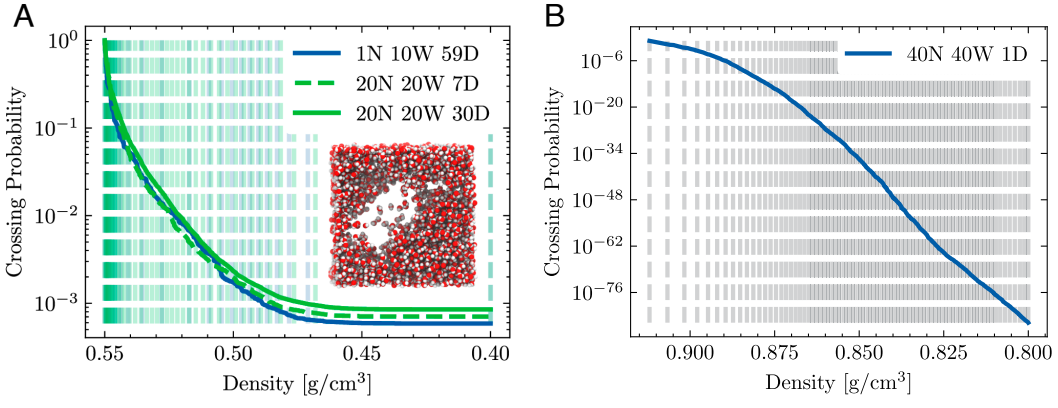

Fig. 2. The crossing probability for density reduction, with the snapshot in (*A*) describing the system at 0.40 g/cm^3. The legend acronyms represent (N)odes, (W)orkers, and (D)ays. The vertical dashed lines indicate the interface placements. (*A*) Metastable liquid water at 573.15 K. Two ∞RETIS simulations are performed on identical, GPU-equipped nodes, where the single node simulation utilizes MPS (*Optimizing GPU Utilization* and Fig. 5) while the multi-node simulation does not. The 30-d plot is the continuation of the 7-d plot. (*B*) The results of a 1-d ∞RETIS simulation of liquid water at the SPC/E boiling temperature equal to 396.0 K based on 40 workers each utilizing their own individual nodes. Based on equilibrium MD runs, an average density of 0.57 g/cm^3 is obtained for *A* and 0.92 g/cm^3 for (*B*).

$\lambda > \lambda_A$ after crossing λ_A in the positive direction without recrossing λ_A (6). In our case, the order parameter was defined as minus the density of the system (the minus signs are omitted in Fig. 2). Hence, the computed probability is reflecting the likelihood of a small density fluctuation below the metastable density of 0.55 g/cm^3 causing the density to continue decreasing until reaching the point of no return. At this point, the system transitions to the gas phase with minimal probability of returning to the metastable liquid state, which occurs when the density falls below 0.45 g/cm^3. This trend is evident as the crossing probability converges to a consistent horizontal plateau beyond this density threshold.

Simulation (ii) requires just 7 d to generate the same total MD steps as simulation (i) would in 59 d. Additionally, from the graph, it is evident that simulation (ii) achieves excellent convergence, with minimal differences in the crossing probability observed after continuing the 7-d simulation up to 30 d. However, simulation (i) benefits from MPS utility, enabling it to run with 10 workers in parallel on a single GPU, maximizing output per node.

To explore the transformative potential of our methodology, we also assessed ∞RETIS's capability in investigating the liquid-to-vapor transition at the actual SPC/E boiling temperature of 396.0 K (23). At the phase transition temperature, the critical nucleus size for the vapor bubble diverges, leading to a vanishingly small rate in the thermodynamic limit. We therefore examined exceedingly rare density fluctuations that do not yet indicate an irreversible transition to the vapor phase, akin to a point of no return. Although these fluctuations are likely to be dependent on system size, their occurrence rate presents an exceptional computational challenge. In this work, we have used this as a litmus test for ∞RETIS, probing its ability to converge calculations of exceedingly small probabilities within a short wall time period when operating on a massively parallel GPU computer.

By employing 40 workers on 40 individual nodes and 80 interfaces for 1 d, ∞RETIS manages to compute the crossing probability and the corresponding rate for the scenario in which random fluctuations lead the system to reach a density below 80% of the stable liquid phase, see Fig. 2*B* and Table 1. To the best of our knowledge, the final value of the crossing probability, astonishingly low in the order of 10^{-86}, represents a world record for the lowest computed crossing probability in a realistic molecular system.

Application II: Chignolin Unfolding. The CLN025 mutant of chignolin is a popular test system for rare event methods, and we examine the unfolding of this mini protein with two different order parameters. The first-order parameter is the RMSD of the protein backbone from the folded state. During this simulation, the system quickly began exploring a set of long-lived misfolded states.

Paths going through these misfolded states were characterized by low weights in their corresponding path ensembles due to their length. In addition, an actual experiment will hardly be able to discriminate between the misfolded and native states, and grouping them into an ensemble of folded structures is more meaningful (28). Based on these arguments, we include the misfolded structures in the folded state definitions and perform an additional simulation with a second-order parameter; a neural network trained on a diffusion map created from a couple of reactive trajectories from previous simulation data, an approach we denote Deep-DM. In Fig. 3*A*, we illustrate the conditional free energy from the first simulation mapped onto the two deep-DM coordinates, in which the misfolded states are apparent.

Fig. 3*A* sheds light on the vast conformational landscape that even a mini protein like chignolin covers during its transition to the unfolded state and illustrates the extensive sampling enabled by our rare event protocol. Notably, we observe a total of 1,000 and 1,600 reactive trajectories during the RMSD and deep-DM simulations, respectively, which allows us to sample a wide range of transition paths. In comparison, brute force Anton simulations of almost twice the length observed around 30 to 40 transitions (24), and we observed 14 transitions in the course of an 80 μs equilibrium simulation.

We also see that instead of a single unfolding route, the protein explores a variety of configurations during the unfolding process, which is not characterized by a single well-defined free energy barrier. This gives rise to the complex kinetic behavior reported previously in long unbiased simulations (24). Such processes can be challenging to model with other methods that rely on assumptions regarding the transition state and the reaction coordinate, which in our protocol can be obtained post-simulation (29).

It is interesting to note the resemblance of the D state to an α-helical structure even though chignolin is a β-hairpin in its native state. Experimental evidence suggests that β-hairpin formation may occur competitively with α-helical formation (30). Fig. 3*C* presents the running averages of the rates for both

Table 1. The results and setup for all the simulations ran in this paper

Simulation	Rate	Flux	$P_A(\lambda_B\|\lambda_A)$	Nodes	Workers	Days
$\text{Boiling}_{573.15K}$	4.73e+06 s^{-1} ± 102%	8.05e+09 s^{-1} ± 71%	5.88e−04 ± 45%	1	10	59
$\text{Boiling}_{573.15K}$	4.76e+06 s^{-1} ± 54%	5.59e+09 s^{-1} ± 52%	8.51e−04 ± 14%	20	20	30
$\text{Boiling}_{396.0K}$	4.03e−75 s^{-1} ± 61%	5.33e+10 s^{-1} ± 13%	7.56e−86 ± 64%	40	40	1
Chignolin_{RMSD}	0.17 μs^{-1} ± 59%	14,000 μs^{-1} ± 19%	1.2e−05 ± 50%	1	16	19
$\text{Chignolin}_{Deep\text{-}DM}$	0.17 μs^{-1} ± 29%	17,000 μs^{-1} ± 6 %	9.9e−06 ± 32%	1	10	15
Eq. sim.	0.18 μs^{-1}					
Anton (24)	0.45 μs^{-1}					
Deep-TICA (27)	0.31 μs^{-1}					
HLDA (27)	0.16 μs^{-1}					
Dissociation*	1.17e−01 s^{-1} ± 92%	2.22e+12 s^{-1} ± 13%	5.28e−14 ± 90%	20	40	10
Ref. (25)*	1.29e−01 s^{-1}	2.92e+12 s^{-1}	4.40e−14			

The number of ensembles for each reported simulation is twice the number of workers. *The reported rate, flux, and crossing probability for the dissociation simulation and ref. 25 in this table is for $\lambda_B = 5.0$, which is different from the results reported in ref. 25. Additionally, due to a difference in subcycles (26), the flux and crossing probability between the simulation data and ref. are not directly comparable, but the rate is. Reported errors are estimated based on block-averaging procedures using single SDs.

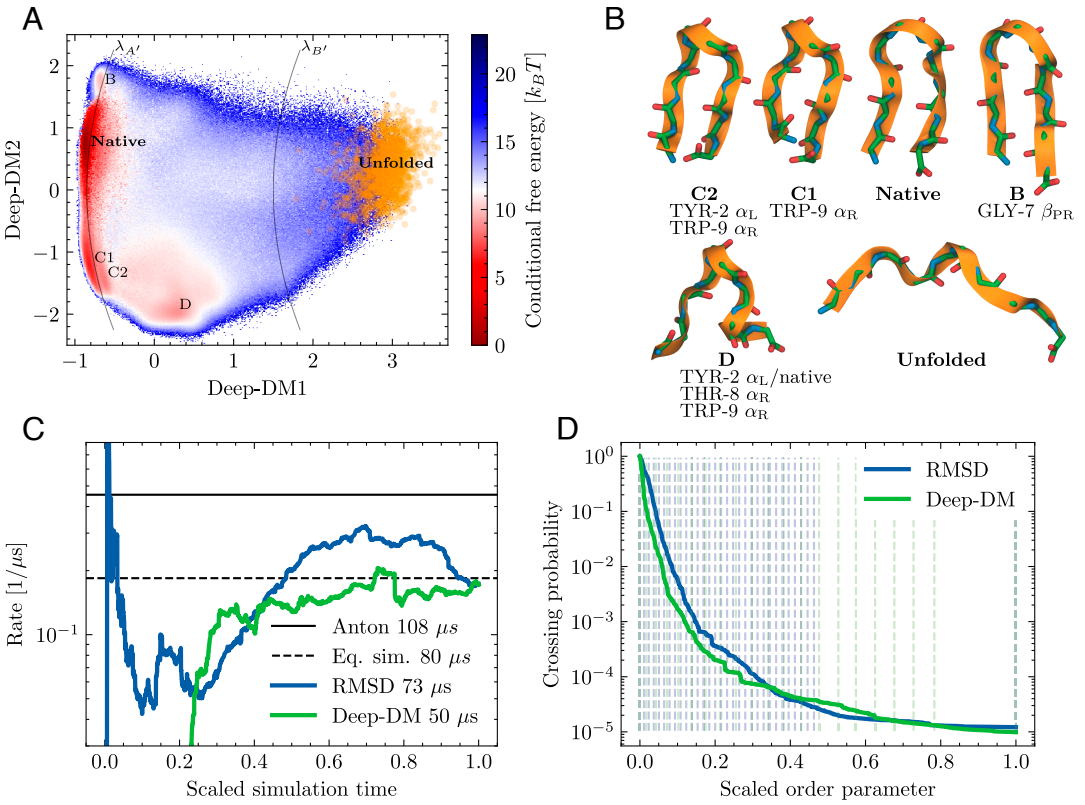

Fig. 3. (*A*) The conditional free energy (i.e., based on phasepoints lying on paths coming from the folded state) mapped onto the deep-DM coordinates, where a set of metastable states is apparent. These results correspond to the simulation with the RMSD order parameter. We do not see a minimum in the path free energy around the unfolded region because of the decreasing probability of reaching such high-order parameter values, given that the path starts in A. We also plot the final phasepoint of each reactive path in orange, which corresponds to a backbone RMSD ≥ 6.0 Å from the folded structure. The curved black lines represent the interface positions of a second set of path simulations with another order parameter, which is a combination of the two deep-DM coordinates. In this second simulation, $\lambda_{A'}$ refers to the folded state interface, which is now a collection of three structures, and $\lambda_{B'}$ refers to the unfolded state. (*B*) Chignolin conformations illustrating the native state, a set of misfolded states (B, C1, C2), a metastable state (D), and a representative sample of the unfolded state that was observed during our simulations. For each metastable state, we also annotate the difference in amino acid conformations compared to the native state. (*C*) The running average of the unfolding rates using the two order parameters, and comparisons with rates obtained from unbiased simulations with Anton (24) and our own equilibrium simulation (Eq. sim.). The legend gives the total simulation time used in the running averages. (*D*) The crossing probabilities from the two simulations and the corresponding interface locations. The order parameter is scaled for comparability.

simulations, while Fig. 3*D* displays the crossing probabilities and interface locations. The calculated rates are in good agreement with those obtained from extensive unbiased simulations and enhanced sampling simulations, even when one of the order parameters incorporates two misfolded states, underscoring the robustness of our approach. A summary of the results can be found in Table 1.

Application III: Ab Initio Water Dissociation. We replicate the RETIS study (25) of calculating the water dissociation reaction rate at 300 K using ab initio MD with the CP2K (31) engine. Satisfactory agreement with the RETIS simulation is obtained from a 10-d ∞RETIS simulation using 40 workers and 80 ensembles, as shown in Fig. 4. A subtle qualitative difference becomes evident within the 3.0 to 4.0 Å shoulder region. The presence of intermediate horizontal plateaus can be attributed to the Grotthuss mechanism, involving a simultaneous double proton transfer (25). This leads to the excess proton residing at an oxygen atom not in direct proximity to the hydroxyl group. Notably, while the original RETIS results suggest that this subprocess consistently reaches completion once initiated, the new ∞RETIS findings paint a more nuanced picture. They reveal a slightly shorter plateau, implying that some double proton transfers may fail and reverse, despite being nearly completed. In terms of wall time, over 1 y was spent running the RETIS simulation, so a rough estimate of the increased wall time efficiency when using ∞RETIS would be 365/10 = 36.5. A considerable contributor to this difference lies in the sequentiality of the RETIS algorithm. As the average path length generally increases with the ensembles number [i^+], the wall time required to generate new paths increases as well. For this system, trajectories generated by ensembles in the gradual 3.0 to 5.0 Å range can be up to 100 times longer than those from ensembles in the steeper 1.0 to 1.5 Å region (Fig. 4). Consequently, in the context of RETIS, even though the lower ensembles hold the potential for rapid sampling due to their shorter average path lengths, the sequential sampling of higher ensembles leads to extended periods between each ensemble update, as these higher ensembles demand significantly more wall time to generate new trajectories. This sequential challenge is circumvented in the ∞RETIS simulation, where any of the participating workers can initiate MC moves in the lower ensembles whenever they are available. In addition to these algorithmic enhancements, other factors may have contributed to the performance increase, including hardware and software advancements since the original study in 2018.

Optimizing GPU Utilization. Asynchronous replica exchange effectively harnesses the benefits of current and future developments of HPC by allowing the initialization of high worker

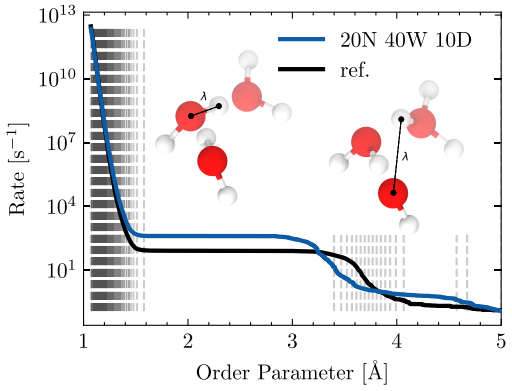

Fig. 4. The rate of water dissociation (crossing probability multiplied by the flux) is compared between the previously reported result, ref. 25, and the result generated from a 10 d ∞RETIS simulation using 20 (N)odes, 40 (W)orkers and 80 ensembles, where each node employs two workers each. The *x*-axis is the composite order parameter described in ref. 25 and in *Materials and Methods*. While the MD time step is identical for both simulations, there is a slight variation in the time between frames, Δt, caused by differences in the frame-saving rate. Although this discrepancy may influence factors like flux and crossing probability, the product remains independent of it (26), allowing for a direct comparison.

numbers to the high amounts of compute hardware (CPU, GPU, and nodes) available on HPCs. However, an additional benefit is also the effective utilization of NVIDIA MPS (32) when running GPU-accelerated MD, as MPS allows multiple independent processes to concurrently run on the same GPU. For our GPU acceleratable application examples, we observe a 2.4-fold increase in the effective throughput (total ns/day) when running a 12,165 particle boiling system on a node with a 12-core Intel Xeon E5-2690 v3 CPU and an NVIDIA Tesla P100 16 GB GPU, as seen in Fig. 5. With even better scaling, we observe a 6.0-fold increase for a 5,889 particle Chignolin system on a node with a 16-core Intel Xeon E5-2687W CPU and NVIDIA GeForce RTX 3090 GPU. Therefore, large worker numbers (i.e., 10 to 16 in our case) can be readily initialized without necessarily

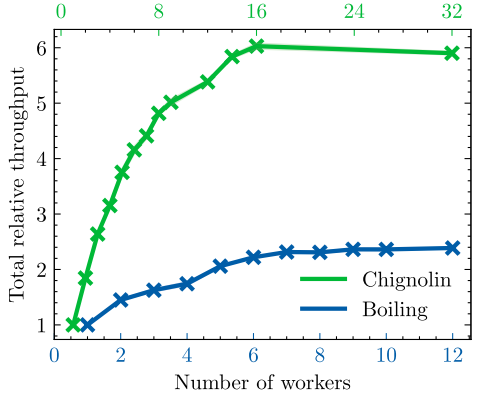

Fig. 5. Total relative throughput (ns/day) for the studied water boiling and chignolin systems running on one GPU-equipped node, with a varying number of concurrently running simulations. At zero parallelization, i.e., when one simulation employs all of the hardware resources of one node, MD throughput averages at 75.0 and 646.2 ns/day. Optimally, with the use of NVIDIA's MPS service, throughputs of 178.8 ns/day (14.9 ns/day × 12 parallels) and 3896.0 ns/day (243.5 ns/day × 16 parallels) are achieved. The irregular spacing between the data points is due to the sharing of CPU cores (constant) between the number of workers (variable). The hardware specification for each system is detailed in *Materials and Methods*. The data points are averages based on 10 repeated trials per data point, with a low to insignificant SD.

requiring multi-node hardware. MPS on multiple parallel nodes would be even more powerful but was not feasible on the available computing resources.

Discussion

Utilizing the power of recent path sampling innovations, we have developed an efficient path sampling protocol referred to as ∞RETIS. In challenging realistic applications, our protocol demonstrated outstanding scalability across diverse GPU and CPU computing platforms using both classical and Ab Initio dynamics. Its remarkable sampling efficiency enabled swift convergence of transition rates within high-dimensional systems that previously would require months to years for convergence. With the ∞RETIS algorithm deployed on potent HPC systems, they now succumb within mere days or weeks. This is a significant advancement, as path sampling offers a distinct advantage over other rare event techniques, such as metadynamics (33) and steered MD (34), by enabling the study of completely unbiased dynamics. However, its computational costs have slowed down the widespread adoption of quantitative path sampling simulations in large molecular systems. The algorithmic innovations detailed in this paper are poised to revolutionize this landscape, making previously unattainable systems accessible and potentially guiding experimental discoveries.

Materials and Methods

Rate Calculation. Rates were computed from the RETIS ensembles by writing $k_{AB} = f_A P_A(\lambda_B|\lambda_A)$ where f_A is the frequency for the system to exit state *A*, and $P_A(\lambda_B|\lambda_A)$ is the crossing probability, the very small chance that after an exit, the system manages to reach state *B* without revisiting state *A*. In RETIS, the flux is determined from the average path lengths in the $[0^-]$ and $[0^+]$ path ensembles. The total crossing probability is obtained from the product of local crossing probabilities, $P_A(\lambda_B|\lambda_A) = \prod_{i=0}^{n-1} P_A(\lambda_{i+1}|\lambda_i)$ where $P_A(\lambda_{i+1}|\lambda_i)$ is estimated from the fraction of sampled paths in the $[i^+]$ ensemble that cross the next interface λ_{i+1}. Further improvement in the statistical analysis (29, 35) has been obtained using the weighted histogram analysis method (WHAM) (36). All estimated errors on computed properties have been based on a block-averaging procedure. Additional properties like the conditional free energy (Fig. 3*A*) were obtained using a WHAM (36) reweighting procedure on the collective phase points of the trajectories of all RETIS ensembles (29, 35).

Initialization. Like in standard RETIS and TIS, the interface positions in ∞RETIS are initially configured so that $P_A(\lambda_{i+1}|\lambda_i)$ is approximately the same across different *i* values, a tuning process conducted during preliminary initialization runs. However, ∞RETIS does not aim for a specific target value like the rule of thumb value of 0.2 (37), which fixes the number of required interfaces. Instead, the number of interfaces $(n + 1)$ is based on the available hardware, i.e., the number of workers that can be launched. To ensure plenty availability of free ensembles with sufficient overlap at each infinite swapping step, *n* is set to be twice the number of workers. Once *n* is fixed, we aim to place the interfaces such that $P_A(\lambda_{i+1}|\lambda_i) \approx P_A(\lambda_B|\lambda_A)^{(1/n)}$ for all *i* using an estimation for $P_A(\lambda_B|\lambda_A)$ from the short initialization run.

Sampling. The WF move was employed in all $[i^+]$ ensemble simulations for $0 < i < n$. We determined the parameters N_s and λ_{cap} without conducting an extensive optimization analysis; instead, we chose values that appeared reasonable. This approach led us to use the same N_s value for all ensembles, rather than aiming for ensemble-specific values based on the ratio of each ensemble's average path length to the average path length of the subtrajectories (8). While significantly enhancing the efficiency compared to standard shooting, more efficient parameter sets likely exist. We plan to explore automatic parameter adjustment and initialization in the future. The $[0^-]$ and $[0^+]$ ensembles

employed normal shooting without high acceptance. In these ensembles, where the path length of subtrajectories matches that of full paths, the WF move has a reduced impact. Furthermore, the absence of high acceptance implies that the MD-intensive point exchange move $[0^-] \leftrightarrow [0^+]$ can always be accepted. Instead of high acceptance, these ensembles use an early rejection scheme (6) that allows for the interruption and rejection of the generation of excessively long paths, which would have been rejected anyway in the Metropolis–Hastings step.

In quantitative terms, the acceptance of the WF in the boiling simulation at 396.0 K reached 100% due to the absence of a stable *B* state attainable from *A*. Likewise, in the water dissociation study, the WF move demonstrated a similarly remarkable acceptance rate of 98.9%. However, in boiling simulations conducted at the higher temperature of 573.15 K, the WF move exhibited lower acceptances of 74.1% and 84.9% for the simulations utilizing 10 workers (21 interfaces) and 20 workers (41 interfaces) respectively. Rejections predominantly occurred within the last set of easily converging path ensembles, where the λ_i interface required for crossing is already proximate to state *B*. Focusing on the challenging part wherein the system ascends in free energy, disregarding the latter path ensembles where paths have over a 50% likelihood of reaching state *B*, the WF acceptance escalates to, respectively, 89.3% and 94.2%. This hints at the potential for even greater efficiency by adhering to a slightly modified protocol than the one described in the previous section, aligning the $(n-1)$th interface such that $P_A(\lambda_n|\lambda_{n-1}) < 0.5$.

The protein unfolding study exhibited a similar trend in the acceptance rate of the WF move, demonstrating nearly 100% acceptance in the initial path ensembles before decreasing for interfaces closer to state *B*. Across the simulations shown in Fig. 3, the overall WF acceptance rates were 73.8% (RMSD) and 45.3% (Deep-DM). The relatively lower acceptance observed in the latter case is attributed not only to the suboptimal positioning of λ_{n-1}, but also arises from the asymmetric shape of the free energy landscape, requiring λ_{cap} to be placed farther from state *B*, closer to the peak of the barrier. With the current interfaces, the generation of a single $A \rightarrow B$ trajectory tends to provide shooting points predominantly within the basin of attraction of state *B*, which can lead to dramatically low shooting acceptance (38). Consequently, there is a high likelihood that all N_s subtrajectories become unsuccessful. This shows that the very high acceptance and optimal efficiency is achieved with fine tuning of the method's parameters, but even with suboptimal WF parameters, both acceptance and decorrelation are still superior to those achieved with standard shooting especially for the asymmetric barrier case (38).

Code Implementation and Availability. We run an in-house ∞RETIS Python code which mainly consists of PyRETIS (17) function imports together with the use of the Dask (39) package which handles scheduling worker tasks. To start a simulation, the user determines the number of workers to be employed based on the hardware available and the type of system to be simulated. An additional user variable is subcycles, which controls the number of frames to be saved between the number of generated MD steps. For instance, if a trajectory comprises 200,000 MD steps, the trajectory in the path ensemble is delineated by 200 time slices, each corresponding to every 1,000th MD frame when subcycles is set to 1,000. Once the setup is completed, the Python code schedules available workers to perform MD-based MC moves. The running of MD, engine input/output, and data storage are mainly handled by the PyRETIS functions that externally start and stop GROMACS/CP2K simulations. The code used to generate the paper data is available at https://doi.org/10.5281/zenodo.8380343, but an updated code that is under development is accessible via GitHub https://github.com/infretis/.

Simulation Details on Superheated Water Boiling. Superheated liquid water in the form of 4055 H2O water molecules is simulated with periodic boundary conditions and a timestep of 0.5 fs in the NPT ensemble at 1 bar and the two temperatures 396.0 and 573.15 K using Gromacs 2021.5 (40). The temperature and pressure are kept constant by applying a V-rescale thermostat (41) of 2.5 ps relaxation time and a C-rescale barostat (42) with a relaxation time of 10 ps. As with the previous TPS studies (20, 21), the SPC/E water model (43) is also used. The order parameter is simply the water density, and the initial reactive trajectories are obtained by quickly heating an equilibrated system. The ∞RETIS simulation ran with a subcycle of 1,000 and the number of WF subtrajectories equals 4. The simulations were run on HPC nodes consisting of 12 core Intel Xeon E5-2690 v3 CPU and NVIDIA Tesla P100 16 GB GPU.

Simulation Details on Chignolin Unfolding. The simulations of the CLN025 mutant of chignolin are performed with the setup described by Bonatti (44), using their provided input files available online. The mini protein is modeled with the CHARMM22* force field and solvated with TIP3P water molecules at 340 K. The terminal amino acids and the ASP and GLU amino acids are modeled in their charged states, and the system is neutralized by adding two sodium ions. The equations of motion are integrated with a timestep of 2 fs using the velocity Verlet scheme, and canonical sampling is achieved with the V-rescale thermostat (41).

The first path sampling simulation is performed with an order parameter defined as the RMSD between the protein backbone and the folded structure (the average structure from a long simulation, not the crystal structure). The folded and unfolded interfaces were given by an RMSD of 0.6 and 6.0 Å, respectively, and we used an interface cap at 4.0 Å. Multiple misfolded states were observed during this simulation. We train a neural network that includes the B and C1 states as part of the folded ensemble in the following manner; we first construct a diffusion map with the approach outlined in ref. 45 using chignolin configurations representing the native folded state, the misfolded states observed in the first path simulations, unfolded configurations, and a set of configurations from the transition path ensemble. The protein backbone RMSD is used as a distance metric to construct the diffusion map. We then train the neural network directly on the two leading eigenvectors of the diffusion map. We use the 741 interatomic distances between backbone atoms as input features for the network, with architecture 741-50(ReLu)-25(ReLu)-12(ReLu)-2. This is motivated and similar in spirit to the approach described in ref. 46, except that we do not fit the network to pre-assigned positions, but rather on the output of the diffusion map. Using this approach, we can discriminate between the folded, misfolded, and unfolded states as well as some other metastable states. The A and B interfaces were given by an order parameter of −0.85 and 1.5, respectively, and we used an interface cap of 1.2.

For all systems, we use 16 Intel Xeon E5-2687W CPUs and partition 1 NVIDIA447 GeForce RTX 3090 GPU among the workers using MPS. For the RMSD, we use 16 workers, and for the deep-DM system, we use 10 workers. The ∞RETIS timestep was 4 ps (2000 MD integration steps with a 2 fs timestep), and the number of WF subtrajectories was 3.

Simulation Details on Ab Initio Water Dissociation. We replicate the previous RETIS study (25), i.e., 32 water molecules are simulated in a periodic 9.85 Å cubic box with ab initio DFT MD using CP2K 9.1 (31). The DFT calculations use Becke–Lee–Yang–Parr functional (47, 48) with a DZVP-MOLOPT (49) basis set and a plane-wave cutoff of 280 Ry. The MD simulations were run with a timestep of 0.5 fs, the number of subcycles is set to 5 for the ∞RETIS simulation and the number of WF subtrajectories is set to 2. The new velocities generated by the shooting/WF move are drawn from a Maxwell–Boltzmann distribution corresponding to an average temperature of 300 K. The order parameter is the longest O-H distance in the case where no dissociated species exist in the system. When $OH^- + H_3O^+$ pair(s) are detected, the order parameter becomes the shortest distance between the oxygen in OH^- and hydrogens in H_3O^+. The ∞RETIS simulation was run on the Sigma2 HPC system Saga with 20 nodes equipped with Intel Xeon-Gold 6138 CPUs, where two workers were run on each node.

Data, Materials, and Software Availability. All study data are included in the main text.

ACKNOWLEDGMENTS. We thank Zhiliang Zhang for the helpful comments on improving our paper. We acknowledge funding support from the Research Council of Norway (Grant No. 275506). We also thank Sigma2, the National Infrastructure for high-performance computing (HPC) and Data Storage in Norway and the HPC infrastructure IDUN at the NTNU.

Author affiliations: [a]Department of Chemistry, Norwegian University of Science and Technology, Trondheim N-7491, Norway; and [b]Department of Chemistry, Utrecht University, Utrecht 3584 CH, Netherlands

1. L. Grajciar *et al*., Towards operando computational modeling in heterogeneous catalysis. *Chem. Soc. Rev.* **47**, 8307–8348 (2018).
2. D. W. Borhani, D. E. Shaw, The future of molecular dynamics simulations in drug discovery. *J. Comput. Aided Mol. Des.* **26**, 15–26 (2012).
3. K. Shmilovich *et al*., Discovery of self-assembling π-conjugated peptides by active learning-directed coarse-grained molecular simulation. *J. Phys. Chem. B* **124**, 3873–3891 (2020).
4. B. Peters, *Reaction Rate Theory and Rare Events* (Elsevier, 2017).
5. P. G. Bolhuis, D. Chandler, C. Dellago, P. L. Geissler, Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Ann. Rev. Phys. Chem.* **53**, 291–318 (2002).
6. T. S. van Erp, D. Moroni, P. G. Bolhuis, A novel path sampling method for the calculation of rate constants. *J. Chem. Phys.* **118**, 7762–7774 (2003).
7. T. van Erp, Reaction rate calculation by parallel path swapping. *Phys. Rev. Lett.* **98**, 268301 (2007).
8. D. T. Zhang, E. Riccardi, T. S. van Erp, Enhanced path sampling using subtrajectory Monte Carlo moves. *J. Chem. Phys.* **158**, 024113 (2023).
9. S. Roet, D. T. Zhang, T. S. van Erp, Exchanging replicas with unequal cost, infinitely and permanently. *J. Phys. Chem. A* **126**, 8878–8886 (2022).
10. N. Plattner *et al*., An infinite swapping approach to the rare-event sampling problem. *J. Chem. Phys.* **135**, 134111 (2011).
11. L. R. Pratt, A statistical method for identifying transition states in high dimensional problems. *J. Chem. Phys.* **85**, 5045 (1986).
12. C. Dellago, P. G. Bolhuis, F. S. Csajka, D. Chandler, Transition path sampling and the calculation of rate constants. *J. Chem. Phys.* **108**, 1964 (1998).
13. E. Riccardi, O. Dahlen, T. S. van Erp, Fast decorrelating Monte Carlo moves for efficient path sampling. *J. Phys. Chem. Lett.* **8**, 4456–4460 (2017).
14. N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, E. Teller, Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).
15. W. Hastings, Monte-Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109 (1970).
16. D. W. H. Swenson *et al*., A Python framework for path sampling simulations. 2. Building and customizing path ensembles and sample schemes. *J. Chem. Theory Comput.* **15**, 837–856 (2019).
17. E. Riccardi, A. Lervik, S. Roet, O. Aarøen, T. S. van Erp, PyRETIS 2: An improbability drive for rare events. *J. Comput. Chem.* **41**, 370–377 (2020).
18. D. G. Glynn, The permanent of a square matrix. *Eur. J. Comb.* **31**, 1887–1891 (2010).
19. P. Lundow, K. Markström, Efficient computation of permanents, with applications to boson sampling and random matrices. *J. Comput. Phys.* **455**, 110990 (2022).
20. D. Zahn, How does water boil? *Phys. Rev. Lett.* **93**, 227801 (2004).
21. K. Karalis, D. Zahn, N. I. Prasianakis, B. Niceno, S. V. Churakov, Deciphering the molecular mechanism of water boiling at heterogeneous interfaces. *Sci. Rep.* **11**, 19858 (2021).
22. R. E. Apfel, Water superheated to 279.5 °C at atmospheric pressure. *Nat. Phys. Sci.* **238**, 63–64 (1972).
23. M. Fugel, V. C. Weiss, A corresponding-states analysis of the liquid-vapor equilibrium properties of common water models. *J. Chem. Phys.* **146**, 064505 (2017).
24. K. Lindorff-Larsen, S. Piana, R. O. Dror, D. E. Shaw, How fast-folding proteins fold. *Science* **334**, 517–520 (2011).
25. M. Moqadam *et al*., Local initiation conditions for water autoionization. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E4569–E4576 (2018).
26. P. G. Bolhuis, D. W. H. Swenson, Transition path sampling as Markov Chain Monte Carlo of trajectories: Recent algorithms, software, applications, and future outlook. *Adv. Theory Simul.* **4**, 2000237 (2021).
27. D. Ray, N. Ansari, V. Rizzi, M. Invernizzi, M. Parrinello, Rare event kinetics from adaptive bias enhanced sampling. *J. Chem. Theory Comput.* **18**, 6500–6509 (2022).
28. P. Kührová, A. De Simone, M. Otyepka, R. B. Best, Force-field dependence of chignolin folding and misfolding: Comparison with experiment and redesign. *Biophys. J.* **102**, 1897–1906 (2012).
29. T. S. van Erp, M. Moqadam, E. Riccardi, A. Lervik, Analyzing complex reaction mechanisms using path sampling. *J. Chem. Theory Comput.* **12**, 5398–5410 (2016).
30. C. M. Davis, S. Xiao, D. P. Raleigh, R. B. Dyer, Raising the speed limit for β-hairpin formation. *J. Am. Chem. Soc.* **134**, 14476–14482 (2012).
31. T. D. Kühne *et al*., CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **152**, 194103 (2020).
32. Maximizing GROMACS Throughput with Multiple Simulations per GPU Using MPS and MIG (2021). Accessed 25 September 2023.
33. A. Laio, M. Parrinello, Escaping free-energy minima. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 12562 (2002).
34. S. Izrailev *et al*., "Steered molecular dynamics" in *Computational Molecular Dynamics: Challenges, Methods, Ideas*, P. Deuflhard, Ed. (Springer, Heidelberg, 1999), pp. 39–65.
35. J. Rogal, W. Lechner, J. Juraszek, B. Ensing, P. G. Bolhuis, The reweighted path ensemble. *JCP* **133**, 174109 (2010).
36. A. Ferrenberg, R. Swendsen, Optimized Monte-Carlo data-analysis. *Phys. Rev. Lett.* **63**, 1195–1198 (1989).
37. R. Cabriolu, K. M. S. Refsnes, P. G. Bolhuis, T. S. van Erp, Foundations and latest advances in replica exchange transition interface sampling. *J. Chem. Phys.* **147**, 152722 (2017).
38. Z. F. Brotzakis, P. G. Bolhuis, A one-way shooting algorithm for transition path sampling of asymmetric barriers. *J. Chem. Phys.* **145**, 164112 (2016).
39. Dask Development Team, Dask: Library for dynamic task scheduling (2016).
40. M. J. Abraham *et al*., GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **1–2**, 19–25 (2015).
41. G. Bussi, D. Donadio, M. Parrinello, Canonical sampling through velocity rescaling. *J. Chem. Phys.* **126**, 014101 (2007).
42. M. Bernetti, G. Bussi, Pressure control using stochastic cell rescaling. *J. Chem. Phys.* **153** (2020).
43. H. J. C. Berendsen, J. R. Grigera, T. P. Straatsma, The missing term in effective pair potentials. *J. Phys. Chem.* **91**, 6269–6271 (1987).
44. L. Bonati, G. Piccini, M. Parrinello, Deep learning the slow modes for rare events sampling. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2113533118 (2021).
45. S. B. Kim, C. J. Dsilva, I. G. Kevrekidis, P. G. Debenedetti, Systematic characterization of protein folding pathways using diffusion maps: Application to Trp-cage miniprotein. *J. Chem. Phys.* **142** (2015).
46. D. Ray, E. Trizio, M. Parrinello, Deep learning collective variables from transition path ensemble. *J. Chem. Phys.* **158**, 204102 (2023).
47. A. D. Becke, Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **38**, 3098–3100 (1988).
48. C. Lee, W. Yang, R. G. Parr, Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B* **37**, 785–789 (1988).
49. J. VandeVondele, J. Hutter, Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *J. Chem. Phys.* **127**, 114105 (2007).