PERSPECTIVE • OPEN ACCESS

How far can we stretch the timescale with RETIS?

To cite this article: Titus S. van Erp 2023 *EPL* **143** 30001

A Letters Journal Exploring the Frontiers of Physics

August 2023

EPL, **143** (2023) 30001
doi: 10.1209/0295-5075/ace9f6

www.epljournal.org

Perspective

How far can we stretch the timescale with RETIS?

Titus S. van Erp[(a)]

Norwegian University of Science and Technology (NTNU) - Høgskoleringen 5, 7491 Trondheim, Norway

received 20 April 2023; accepted in final form 24 July 2023
published online 1 August 2023

Abstract – Molecular dynamics (MD) and Monte Carlo (MC) have long coexisted as two main independent branches of molecular simulation. In the late eighties, however, algorithms based on the combination of both were created such as hybrid Monte Carlo which uses large MD steps as MC moves. An entirely different kind of combination emerged a decade later via the transition path sampling (TPS) method in which MD trajectories are not just part of the MC move, but also form the state space being sampled. Algorithms like replica exchange transition interface sampling (RETIS) exploit this idea to compute reaction rates via a series of TPS simulations. RETIS yields results identical to hypothetical long MD runs, but with exponentially reduced computation time. This perspective describes the RETIS method and discusses recent and future advancements that will enable the study of even longer molecular timescales with reasonable computational resources.

Introduction. – Due to the vastly increasing speed of computer hardware and the rapid algorithmic development, molecular modeling, with molecular dynamics (MD) and Monte Carlo (MC) as its main branches, has undeniably established itself as an essential third pillar for addressing future scientific challenges next to theory and experimentation [1]. Enormous progress has been made to attack its three main difficulties: accuracy, system size, and time range. Accuracy has drastically been improved due to relentless efforts to develop force fields and the development of efficient *ab initio* MD [2] methods. System sizes of present biomolecular systems are routinely several hundreds of thousands of molecules as parallel computing algorithms exploit the power of present-day high-performance computing [3]. Except for certain replica-based accelerated MD methods [4], parallel computing has played a smaller role in extending the time range, as time evolution is fundamentally a sequential process and not easily parallelizable. Even if the aforementioned developments allow us to perform more MD steps per wall time, this is often still not enough to study rare molecular transitions such as chemical reactions, protein folding, membrane permeation, and nucleation.

In condensed systems, MD typically outperforms MC due to its ability to move particles collectively and the effect of inertia which results in less diffusive sampling. MC usually has an advantage in porous media where the non-locality of moves, such as those implemented in configurational bias Monte Carlo (CBMC) [5], can be exploited. In an effort to combine the best of both worlds, hybrid MC [6] simulations use (sequences of) large MD steps with a Metropolis acceptance/rejection decision [7]. The aim here is not to accurately describe the system's dynamics, which can become unphysical due to rejections and a time step that would be unstable in standard MD, but rather to optimize the sampling of the configurational space. This approach provides additional flexibility in optimizing sampling by assigning artificial masses to certain degrees of freedom or principal components [8].

In a different approach, transition path sampling (TPS) [9] combines MD and MC with the specific aim to study the dynamics of rare transitions. In this method, generating MD trajectories is not only the way how sampling is done, but also what is being sampled. The notion that MD trajectories can be assigned a statistical weight, similar to the Boltzmann weight in configurational space, allowing them to be sampled akin to states in a Markov chain, was first introduced by Pratt [10]. The advantage compared to normal MD is that trajectories can now be subjected to importance sampling techniques, for instance by requiring that trajectories start in the reactant state (A) and end in the product state (B).

[(a)]E-mail: titus.van.erp@ntnu.no (corresponding author)

The requirements define a statistical path ensemble, reflecting the same statistics as cutting relevant trajectory segments from an infinite MD run. Provided that an initial path can be established that is part of the path ensemble, a trial path is generated by making some random modifications of the previous path. This trial path is accepted or rejected based on detailed balance and is always rejected if it does not meet the path ensemble's requirements.

Quantitative dynamical properties like rates can be computed from a series of path ensembles. TPS's original rate evaluation [9] relied on fixed path length ensembles and umbrella sampling (US) [11]. Transition interface sampling (TIS) [12] presented a more efficient algorithm that utilizes flexible length path ensembles and phase space partitioning via interfaces. TIS introduced a rate expression and various theoretical concepts that have been adopted by other methods like forward flux sampling (FFS) [13]. These include the notion of overall states, which captures the most recent stable state visited by the system (see fig. 1), as well as the overall crossing probability and local history-dependent conditional crossing probabilities, which provide insight into the probability of crossing a certain state given the system's past behavior.

The TIS method has been commonly used to investigate complex two-state reaction kinetics. It allows for the study of these kinetics without introducing any additional approximations, except for those already inherent in the MD model, such as the force field and periodic boundaries. TIS is particularly efficient for transitions that are very rare but occur rapidly once initiated. It differs from other advanced MD methods, such as kinetic Monte Carlo [14], temperature accelerated MD (TAMD) [4], and activation-relaxation technique (ART) [15], which aim to describe processes involving sequences of rare events that can be characterized by a basin hopping mechanism at a coarser-grained level such as the process of crystal formation, where each monomer addition involves a reaction barrier and can be considered a rare event. Similarly, TIS differs from Markov state models [16], which are typically employed when the process being studied is not only rare but also slow which is the case for many biomolecular systems. Markov state models aim to determine a rate network between stable states and relatively long-lived metastable states at the barrier region. However, it is worth noting that both types of methods can be combined with TIS to further enhance their efficiency [17,18]. The previously mentioned FFS method also bears similarities to TIS but requires stochastic dynamics and assumes that the relevant initial distribution for transition trajectories can be generated by plain MD. This makes FFS less suitable when the relevant transitions necessitate rare initial conditions instead of rare fluctuations in the stochastic nature of the dynamics. While FFS has the advantage over TIS that it can be used for driven non-equilibrium systems, there is a high chance of not sampling the correct mechanism [19].

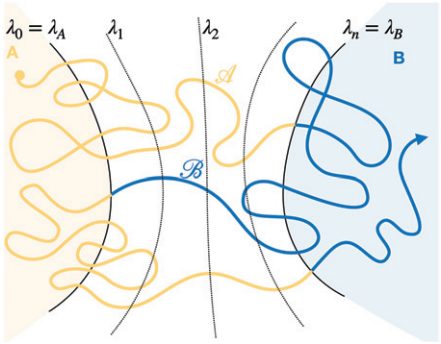

Fig. 1: A hypothetical MD run, with the arrow indicating the direction of time. The two non-adjacent *stable* states for reactant (A) and product (B) are marked in yellow and blue, respectively. The TIS method determines the *overall* state ($\mathcal{A}$ or $\mathcal{B}$) based on the most recent stable state visited by the system. The rate constant k in TIS is computed as the frequency of $\mathcal{A}$ to $\mathcal{B}$ transitions divided by the time spent in $\mathcal{A}$, which is insensitive to the precise stable state definitions and avoids correlated recrossings. The interfaces between λ_0 and λ_n allow for efficient calculation of this rate by sampling path ensembles where the $[i^+]$ ensemble consists of paths crossing at least λ_i.

In 2007, a significant improvement upon TIS was made with Replica Exchange TIS (RETIS) [20,21], which enhances the efficiency by incorporating the *minus* ensemble and replica exchange moves between path ensembles. RETIS has proven effective in modeling complex processes such as gas hydrate formation [22] and *ab initio*-level water dissociation [23]. Simulations like this necessitate thousands of centuries of CPU time if done with traditional MD, but even with RETIS, they still remain expensive, requiring several months of wall time. In this article, I delve into algorithmic advancements of path sampling that have the potential to boost the efficiency of RETIS simulations even further, while oftentimes preserving the exactness of the method.

TIS and RETIS. – The TIS algorithm defines a set of non-intersecting interfaces $\{\lambda_0, \lambda_1, \lambda_2, \ldots, \lambda_n\}$ based on one or a few collective variables (see fig. 1). The rate of the transition from reactant to product is expressed as

$$k = f_A P_A(\lambda_B|\lambda_A) = f_A \prod_{i=0}^{n-1} P_A(\lambda_{i+1}|\lambda_i), \tag{1}$$

where f_A is the flux through $\lambda_A = \lambda_0$ and $P_A(\lambda_B|\lambda_A)$ is the overall crossing probability, the chance that after a crossing with λ_A, $\lambda_B = \lambda_n$ will be crossed without recrossing λ_A. This is often an extremely small number that cannot be directly measured, and is instead expressed using the exact factorization in the second equality. The local history-dependent conditional crossing probability $P_A(\lambda_{i+1}|\lambda_i)$ is a measure of the likelihood that λ_{i+1} will be crossed before λ_A, given that a first crossing with λ_i has occurred since exiting state A. It is crucial to note that the distribution of these *first crossing points* differs from

the equilibrium distribution at λ_i. This fact incorporates history dependence into eq. (1) and makes it exact.

The computation of f_A in TIS involves a straightforward MD run. The system is initialized in state A, and the number of positive crossings with λ_A is counted and divided by the simulation time. Since λ_A is near stable state A, a relatively high number of crossings is expected to occur within reasonable simulation time. It is not anticipated that the system will transition to overall state $\mathcal{B}$, as this would be a rare event. However, if such an event were to occur, the system would need to be reinitialized into state A before continuing the MD run. RETIS on the other hand relies solely on path simulations, achieved by introducing the *minus* path ensemble $[0^-]$ which has all of its time frames at the left side of λ_0 except for the endpoints. In conjunction with $[0^+]$, the average path length τ in both ensembles is used to determine the flux [20],

$$f_A = \left(\tau_{[0^-]} + \tau_{[0^+]}\right)^{-1}. \tag{2}$$

The primary advantage of replacing the traditional MD run with an additional path ensemble in RETIS is that it creates a more uniform set of simulations. This enables effective implementation of replica exchange moves across all path ensembles, resulting in a more comprehensive exploration of the path space and better sampling.

The crossing probabilities $P_A(\lambda_{i+1}|\lambda_i)$ are computed by generating path ensembles $[i^+]$, where each path starts with a positive crossing with λ_0 and ends by crossing either λ_n or λ_0 again. For a path to be included in the ensemble $[i^+]$, it must have at least one frame to the right of λ_i. TIS and RETIS estimate $P_A(\lambda_{i+1}|\lambda_i)$ by counting the fraction of paths in $[i^+]$ reaching λ_{i+1} in addition to λ_i after generating a substantial number of paths.

The primary MC move for generating paths in TPS, TIS and RETIS has been the shooting move [24]. In this move, a random frame (also called a time slice) from the previous path is randomly perturbed, usually affecting only velocities, to create a new phase point. The new point is then propagated forwards and backwards in time until reaching either state A or B, resulting in a trial path. To be accepted, the trial path must satisfy the ensemble criteria, such as starting at λ_0 (state A) and crossing λ_i at least once for ensemble $[i^+]$. Additional acceptance/rejection rules may be necessary to obey detailed balance, depending on the perturbation and the range of shooting point selection. If the perturbation is small, the trial path will resemble the old path, which increases its likelihood of being valid for the ensemble, but slows down decorrelation. Usually, a 30–50% acceptance probability is aimed for, which is often achieved by uniformly selecting shooting points from any frame except endpoints and regenerating velocities from a new Maxwell-Boltzmann distribution, which yields the following acceptance probability [21]:

$$P_{\text{acc}} = \mathbf{1}^{(n)}_{[i^+]} \times \min\left[1, \frac{L^{(o)}}{L^{(n)}}\right]. \tag{3}$$

The first term is the indicator function, which is 1 if the new path meets the criteria of $[i^+]$, and 0 otherwise. The last term involves the ratio of the path lengths, expressed as the number of frames in the old and new paths. To reduce computational cost, an early rejection scheme can be used to replace the probabilistic decision step (eq. (3)) by a maximum allowed path length for the new path as $L^{(o)}/\alpha$ where α is a random number between 0 and 1. If the trial path exceeds the maximum length, the move is rejected without completing the path. By first integration backward, the path can also be early rejected if the backward trajectory ends at state B since the ensembles criteria include starting at A.

Uniform shooting allows trajectories to bypass barriers that are perpendicular to the reaction coordinate λ [25], but replica exchange moves between path ensembles can enhance this. When a path in the $[i^+]$ ensemble crosses λ_{i+1}, it is also valid in the $[(i+1)^+]$ ensemble, while the path in $[(i+1)^+]$ is always valid for $[i^+]$. This means that the paths can be swapped successfully whenever the $[i^+]$ path crosses λ_{i+1}, without doing costly MD steps. Swapping the last two frames of $[0^-]$ and the first two frames of $[0^+]$ allows for successful swaps between $[0^-]$ and $[0^+]$. The pair of frames are integrated at the other side of λ_0 forward and backward in time to establish new paths in both ensembles. Although this swap requires MD integration, it is a very effective way to decorrelate the sampling. By employing these swapping moves, RETIS typically enhances the efficiency of TIS by more than an order of magnitude [20].

Optimizing performance and initialization. – Setting the interfaces is a crucial aspect of TIS and RETIS. When two interfaces λ_i and λ_{i+1} are placed close together, $P_A(\lambda_{i+1}|\lambda_i)$ increases, leading to faster convergence of its computation. However, if all interfaces are closely placed, many path ensemble simulations are needed to cover the space between λ_A and λ_B, resulting in an overall statistical error that is still large due to error propagation of many small statistical errors. Solvable one-dimensional models suggest that the ideal interface positioning for TIS is achieved when $P_A(\lambda_{i+1}|\lambda_i) \approx 0.2$ for all i [21], indicating that one of the five paths reaches the next interface. For RETIS, the optimal value is presumably higher as this increases the number of successful swapping moves. The first interface λ_0 is typically placed slightly uphill along the free energy barrier. The last interface, λ_B, is usually placed at the *point of no return*, where the system is unlikely to return to state A. At this point, the overall crossing probability shows a horizontal plateau.

TIS allows for a hierarchical approach by running a $[0^+]$ ensemble simulation without defining λ_1. The paths are continued until reaching stable states, and after generating a sufficient number of paths, the crossing probability can be analyzed to determine the value of λ_1 at which $P_A(\lambda_1|\lambda_0) = 0.2$. A path from the $[0^+]$ ensemble that reaches this level can then be used to start MC sampling

in $[1^+]$. This process is repeated to fix λ_2 and so on. In practice, however, it is desirable to run the ensembles simultaneously in parallel and the hierarchical approach is also not feasible for RETIS because of the swapping moves between ensembles. Instead, initialization is typically done in an *ad hoc* fashion with interface repositioning in a preliminary phase. During a production run, interface changes are typically avoided unless local crossing probabilities deviate from an acceptable range of (0.05, 0.6). To bootstrap the MC sampling in each path ensemble in the absence of a hierarchical approach, an initial path must be provided, which can be generated using various techniques. This initial path need not be highly probable and may even be unphysical like non-connected frames as the MC sampling should automatically drift towards the more likely path space. Yet, higher quality initial paths are likely to converge faster [26].

Advanced shooting. – Despite the significant algorithmic advancements in path sampling, the main path generating MC move *shooting* [24] has remained largely unchanged [24] apart from variations in velocity generation and shooting point selection. Many auxiliary MC moves have been developed, such as the swapping moves in RETIS, or the mirror move and target-swap moves for permeation studies [27]. While they provide speed-ups when combined with shooting, they cannot provide ergodic sampling on their own. Clearly, performing many replica exchange moves (swaps) eventually leads to resampling of the same paths repeatedly.

Still, the efficiency of shooting can be hampered by rejections and the degree of similarity between consecutive accepted paths. The recently developed *subtrajectory moves with high acceptance* address both issues [28,29]. Three subtrajectory moves have been devised: web throwing (WT) [28], stone skipping (SS) [28], and wire fencing (WF) [29]. The implementation of the former two can be difficult, particularly when the MD portion of the algorithm is outsourced to dedicated MD programs. To reduce communication overhead with these external MD engines, the time step in TIS or RETIS is usually comprised of 5 to 1000 actual MD steps, but WT and SS ideally require single MD step resolution. To address these implementation challenges, the WF move was developed. While the WF move may be less thrifty with MD steps compared to the other two methods, it is still substantially more efficient than standard shooting without requiring significant modifications to external molecular dynamics software or the need to increase the communication overhead.

As shown in fig. 2, a new path in WF is generated via a series of shorter subpaths. While the WF move itself is more costly than a shooting move, the resulting new path is more decorrelated from its source such that fewer trajectories are needed to achieve a desired statistical error. Additionally, the high-acceptance technique increases the acceptance typically above 95%. In high-acceptance advanced shooting, any new trial path that starts at λ_n and ends at λ_0 is not rejected but time-reversed. In addition, the sampling distribution is adjusted so that the statistical weight of a path is multiplied by a biasing factor that aims to maximize the acceptance. In the analysis, the effect of this artificial bias is negated by weighting each sampled path with the inverse of this biasing factor. Mesoscopic DNA simulations indicate that advanced shooting moves can lead to a 12-fold increase in efficiency [28].

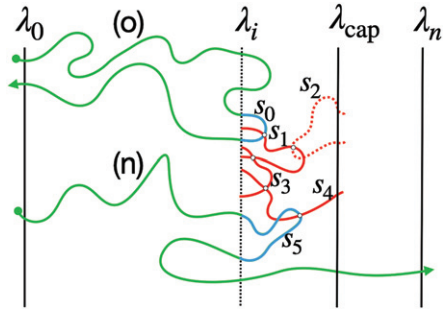

Fig. 2: The sampling process in the $[i^+]$ ensemble using the WF move: the process begins by selecting a segment, subpath s_0 (in blue), randomly from the old path (in green) that connects λ_i to λ_n or λ_i to λ_i. Multiple subpaths (red) are then generated by consecutive shooting moves within the restricted interval $(\lambda_i, \lambda_{\text{cap}})$. If a subpath ends at λ_{cap} in both time directions (like s_2), it is rejected, and the next subpath is released from the shooting point of the previous subpath (s_1). Once the desired number of subpaths has been generated (5 in this case), the last accepted one, s_5 (blue), is extended in both time directions (green), until reaching λ_0 or λ_n.

∞RETIS. – TIS improves both efficiency and accuracy compared to the original TPS rate evaluation. Although TPS has been advocated as an exact method, in practice, the fixed path ensembles employed in TPS introduce a cut-off such that a small but significant part of the tail in the path distribution cannot be sampled. The flexible path lengths in TIS, on the other hand, are on average lower, but occasionally higher than the TPS cut-off. RETIS, besides being substantially more CPU efficient than TIS, also improves accuracy for moderately rare events due to the slightly more accurate flux evaluation that does not require reinitialization when a spontaneous event occurs. Hence, it seems logical that RETIS should always be preferred over TIS and the TPS rate method. However, TIS still has the advantage that it is easier to parallelize since the path ensembles can be run in an embarrassingly parallel fashion. This means that while RETIS significantly outperforms TIS in terms of CPU efficiency, TIS could still win the battle in terms of wall time if a very powerful parallel computer cluster is available.

The main challenge in parallelizing the RETIS algorithm is that the standard MC moves do not require the same amount of CPU time. Specifically, creating a new path via a shooting or a WF move has a varying CPU cost due to the variable path length and the ensemble's average path lengths differ as well. If each path ensemble

in RETIS is assigned to a specific processor or a group of processors (referred to as a *worker*), the workers in the fast ensembles must wait for the slow ones to finish before they can swap. Therefore, programs like OpenPathSampling [30] and PyRETIS [31] implement RETIS as a fully sequential algorithm. Alternatively, serial replicas [32] can be used in single replica TIS [33].

To effectively compute ensembles in parallel without idle time, we recently developed an asynchronous replica exchange method [34] that executes multiple swaps upon completion of each path by a worker. Previous asynchronous replica exchange methods for parallel tempering or Hamiltonian exchange [35–37] differ in that they maintain a fast state generation move, typically comprising a single MD step, while aiming to solve the issue of hardware diversity (slow *vs.* fast workers), rather than addressing inherent CPU cost imbalances of the states (long *vs.* short paths). Workers can therefore run strides of many MD steps in isolation and occasionally switch to an interactive state to allow swaps with neighboring ensembles.

It is non-trivial to mathematically demonstrate the unbiased sampling property of an asynchronous replica exchange method when the computational cost of a MC move depends on the characteristics of the generated states. The conventional proof for replica exchange assumes that the M replicas can be treated as a single superstate residing in a dimension that is M times higher than the dimension of a normal state. This conceptual view allows replica exchange to be treated as a regular Markov chain operating in this higher-dimensional space. However, when updates within ensembles occur in a non-cohort manner, this interpretation breaks down.

The algorithm of ref. [34] can be proven in a fundamental different way based on a *twisted balance* relation in which the superstate view is replaced by an *ensemble plus environment* based view. A crucial aspect of the method is that, like ref. [37], the availability of states that can be swapped is ensured by having more ensembles than workers. After a worker completes its standard move within an ensemble, the ensemble and its state are freed. At this point, a series of swaps occur between the free ensembles to exchange their states before the worker is randomly reassigned to another available ensemble (see fig. 3).

Furthermore, it is possible to perfectly replicate the hypothetical outcome of performing an infinite number of swapping moves by calculating the sum of all probabilities for every potential permutation between free ensembles and free states [34]. The idea of doing an infinite number of swaps was proposed before [38], but the steep factorial increase in the number of permutations drastically restricts the number of replicas that can be involved. Addressing this problem, ref. [34] formulates the summations as permanents of weight matrices. This approach provides a versatile solution that outperforms the factorial method [39] and enables even faster computation by exploiting specific characteristics of the weight matrix. The asynchronous infinite swapping replica exchange method

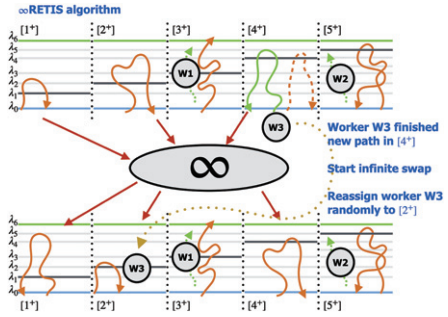

Fig. 3: Visualisation of the ∞RETIS algorithm, showcasing how three parallel workers navigate seven path ensembles ($[0^+]$ and $[0^+]$ not shown). Top: worker W3 just generated a new path (in green) in the $[4^+]$ ensemble, which was subsequently accepted (in orange below). Middle: the next step involves swapping the paths of all free ensembles in every possible way, mimicking an infinite number of swaps. Bottom: W3 is randomly reassigned to the $[3^+]$ ensemble to initiate a new advanced shooting move.

can be applied to various MD or MC methods, but it is particularly beneficial for improving efficiency and scalability in RETIS, leading to the ∞RETIS method.

Reducing memory: PPTIS and REPPTIS. – As discussed above, TPS, TIS, and RETIS can become less efficient when the transition is not only rare but also slow. In such cases, it may be necessary to sacrifice the exactness of the methodology in favor of computational efficiency. By assuming that memory is lost after traveling a certain distance along the λ parameter, path ensembles can be redefined, terminating trajectories before they reach stable states. This approach gives rise to the partial path variant of TIS (PPTIS) [40], which shares similarities with the milestoning method [41].

It is interesting to note that while the idea of performing replica exchange moves between path ensembles was initially proposed for PPTIS before RETIS [42], it was only recently that a replica exchange variant of PPTIS (REPPTIS) was realized and demonstrated [43]. Although PPTIS and REPPTIS are no longer exact, they still retain some memory, unlike milestoning or Markov state models. Moreover, as replica exchange moves in REPPTIS require extending the paths, similar to the $[0^-] \leftrightarrow [0^+]$ swap in RETIS, efforts are being made to explore whether this information can be used to increase the amount of memory further and enhance the accuracy without additional computational cost.

Path sampling and machine learning. – Recent advancements in machine learning (ML) are expected to significantly enhance path sampling in terms of both sampling efficiency and path data analysis. A systematic mathematical approach for gathering mechanistic information about transitions can be based on committor analysis, employing techniques such as neural networks [44], autoencoder models [45], or symbolic regression [46]. A related approach is to find so-called reaction triggers

based on the RETIS output via the predictive power method [47], which has been combined with decision trees [23]. An interesting recent development is the use of conditional normalizing flows to generate statistically independent paths from a given distribution [48]. However, the efficacy of this approach in tackling high-dimensional complex systems is still an unresolved inquiry.

Although the development of ML force fields has become well-established [49–51], accurately creating reactive force fields remains a challenge. Molecular interactions during bond breaking and formation require a more complex mathematical description due to the subtle dependence of charge redistribution on the atomic coordinates of large collective groups of molecules. Incorporating this level of complexity into an ML force field can be difficult, and developing and validating such a force field requires a large dataset of chemical reactions that can be used to train the model.

RETIS can offer a double benefit when used alongside ML in the development of force fields. At the front end, *ab initio* MD-based RETIS simulations can generate trajectories that offer the most relevant set of configurations for the reaction under study, providing valuable training data for the creation of reactive force fields. At the back end, RETIS can also be used with classical MD-based simulations and the developed force field when conventional MD simulations are still too slow, despite electronic structure calculations not being used.

Ab initio-based RETIS simulations have identified collective phenomena in water structures, such as precise [52] or distorted [23] tetrahedral orientation around solute molecules, that create specific local electric field conditions acting as catalysts for chemical reactions, and it is uncertain whether the ML force field can accurately replicate such phenomena. When fitting classical reactive force fields like ReaxFF [53], we found that mutual comparative testing methods are crucial to ensure trustworthiness [54]. Relying solely on fitting to forces and energies from an *ab initio* data set may result in the force field producing unphysical complexes or reactions that were not included in the data set. Further research is therefore needed to determine whether present training routines for generating ML force fields are capable of capturing both the behavior of molecular systems in solution and avoiding the generation of artifacts, such as reactions or reaction mechanisms that are not present in an *ab initio*-based RETIS simulation.

One interesting approach is not to eliminate *ab initio* computations entirely, but to use them when classical force fields are potentially problematic, such as when chemical bonds are formed or broken. This forms the basis of the QuanTIS algorithm [55], which can be viewed as a QM/MM method in which a quantum level and a classical level description are not glued together in space but in time. Specifically, QuanTIS simulates the $[0^-]$ path ensemble at the lower level of theory as these typically reflect rearrangements of reactants, solvent structure, and catalysts without any real chemistry occurring. All the $[i^+]$ ensembles that describe the barrier crossing event and involve the breaking and making of bonds in case of a reaction are described at the higher level of theory. In this algorithm, the $[0^-] \leftrightarrow [0^+]$ swap has to go through a Metropolis acceptance/rejection step as the exchanged configuration points move to another potential energy surface yielding a double energy difference. The acceptance is the highest if the sum of the two energy differences is low and, therefore, this acceptance ratio could potentially be used to optimize both the efficiency of the method and simultaneously the force field. An additional advantage of this approach is that the exchange automatically includes mutual comparative testing of the *ab initio* and ML potential energy surfaces. With the current ease of creating ML force fields on the fly, the QuanTIS method can become a powerful tool for studying highly complex industry- or biology-relevant chemical processes.

Summary. – TPS has been renowned for its conceptual elegance in studying rare events, but accurately computing dynamical properties such as rates and permeation coefficients can be computationally demanding. However, recent algorithmic innovations such as TIS, RETIS, and ∞RETIS, as well as new types of path-generating MC moves, have increased the efficiency of quantitative path sampling by orders of magnitude. Remarkably, these developments have not sacrificed accuracy, but have instead improved it. However, if metastable states or long-lived quasi-stable regions exist, it can be wise to make a compromise between accuracy and efficiency by introducing a soft Markovian assumption of memory loss as in the PPTIS and REPPTIS methods. Additional speed-ups can be expected from synergies between ML and RETIS methods, specifically for training and using reactive force fields. Furthermore, the use of ML methods in path analysis is becoming a dominant instrument.

How far we can stretch the time scale depends on factors like system size, level of theory, desired statistical accuracy, hardware and patience. Nonetheless, RETIS has already produced accurate rate estimates for challenging transitions. *Ab initio* RETIS simulations [23] revealed the process of water dissociation, which occurs once per 11 hours for a water molecule. A classical MD-based RETIS study focused on the formation of a critical nucleus in hydrates [22], which physically takes 3 years (real time) given the size of the simulation box. Although such simulations are still expensive requiring months of computations, the recent algorithmic advancements and untapped potential of ML tools may transform quantitative path sampling into a standard technique for unraveling the complexities of many dynamical molecular processes in the near future.

* * *

I thank Daniel T. Zhang and Anders Lervik for fruitful feedback and comments.

Data availability statement: No new data were created or analysed in this study.

REFERENCES

[1] Schlick T. and Portillo-Ledesma S., *Nat. Comput. Sci.*, **1** (2021) 321.
[2] Iftimie R., Minary P. and Tuckerman M. E., *Proc. Natl. Acad. Sci. U.S.A.*, **102** (2005) 6654.
[3] Shaw David E. *et al.*, *Proceedings of SC21: International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE, St. Louis, Mo., USA) 2021.
[4] Zamora R. J., Perez D., Martinez E., Uberuaga B. P. and Voter A. F., *Handbook of Materials Modeling: Methods: Theory and Modeling*, edited by Andreoni W. and Yip S. (Springer Nature, Switzerland) 2020.
[5] Siepmann J. I. and Frenkel D., *Mol. Phys.*, **75** (1992) 59.
[6] Duane S., Kennedy A. D., Pendleton B. J. and Roweth D., *Phys. Lett.*, **195** (1987) 216.
[7] Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H. and Teller E., *J. Chem. Phys.*, **21** (1953) 1087.
[8] Michielssens S., van Erp T. S., Kutzner C., Ceulemans A. and de Groot B. L., *J. Phys. Chem. B*, **116** (2012) 8350.
[9] Dellago C., Bolhuis P. G., Csajka F. S. and Chandler D., *J. Chem. Phys.*, **108** (1998) 1964.
[10] Pratt L. R., *J. Chem. Phys.*, **85** (1986) 5045.
[11] Torrie G. M. and Valleau J. P., *J. Comput. Phys.*, **23** (1977) 187.
[12] van Erp T. S., Moroni D. and Bolhuis P. G., *J. Chem. Phys.*, **118** (2003) 7762.
[13] Allen R. J., Warren P. B. and ten Wolde P. R., *Phys. Rev. Lett.*, **94** (2005) 018104.
[14] Gillespie D. T., *J. Comput. Phys.*, **22** (1976) 403.
[15] Barkema G. T. and Mousseau N., *Comput. Mater. Sci.*, **20** (2001) 285.
[16] Chodera J. D. and Noé F., *Curr. Opin. Struct. Biol.*, **25** (2014) 135.
[17] Eidelson N. and Peters B., *J. Chem. Phys.*, **137** (2012) 094106.
[18] Rogal J. and Bolhuis P. G., *J. Chem. Phys.*, **129** (2008) 224107.
[19] van Erp T., *Adv. Chem. Phys.*, **151** (2012) 27
[20] van Erp T. S., *Phys. Rev. Lett.*, **98** (2007) 268301.
[21] Cabriolu R., Refsnes-Skjelbred K. M., Bolhuis P. G. and van Erp T. S., *J. Chem. Phys.*, **147** (2017) 152722.
[22] Arjun A. and Bolhuis P. G., *J. Phys. Chem. B*, **124** (2020) 8099.
[23] Moqadam M., Lervik A., Riccardi E., Venkatraman V., Alsberg B. K. and van Erp T. S., *Proc. Natl. Acad. Sci. U.S.A.*, **115** (2018) E4569.
[24] Dellago C., Bolhuis P. G. and D. Chandler D., *J. Chem. Phys.*, **108** (1998) 9236.
[25] Geissler P. L., Dellago C. and Chandler D., *Phys. Chem. Chem. Phys.*, **1** (1999) 1317.
[26] Hall S. W., Leines G. D., Sarupria S. and Rogal J., *J. Chem. Phys.*, **156** (2022) 200901.
[27] Ghysels A., Roet S., Davoudi S. and van Erp T. S., *Phys. Rev. Res.*, **3** (2021) 033068.
[28] Riccardi E., Dahlen O. and van Erp T. S., *Phys. Chem. Lett.*, **8** (2017) 4456.
[29] Zhang D. T., Riccardi E. and van Erp T. S., *J. Chem. Phys.*, **158** (2023) 024113.
[30] Swenson D. W. H., Prinz J.-H., Noe F., Choder J. D. and Bolhuis P. G., *J. Chem. Theory Comput.*, **15** (2019) 837.
[31] Riccardi E., Lervik A., Roet S., Aaroen O. and van Erp T. S., *J. Comput. Chem.*, **41** (2020) 370.
[32] Hagen M., Kim B., Liu P., Friesner R. A. and Berne B. J., *J. Phys. Chem. B*, **111** (2007) 1416.
[33] Du W. and Bolhuis P. G., *J. Chem. Phys.*, **139** (2013) 044105.
[34] Roet S., Zhang D. T. and van Erp T. S., *J. Phys. Chem. A*, **126** (2022) 8878.
[35] Gallicchio E., Levy R. M. and Parashar M., *J. Comput. Chem.*, **29** (2008) 788.
[36] Bussi G., *Nuovo Cimento C*, **32** (2009) 283.
[37] Gallicchio E., Xia J., Flynn W. F., Zhang B., Samlalsingh S., Mentes A. and Levy R. M., *Comput. Phys. Commun.*, **196** (2015) 236.
[38] Plattner N., Doll J. D., Dupuis P., Wang H., Liu Y. and Gubernatis J. E., *J. Chem. Phys.*, **135** (2011) 134111.
[39] Glynn D. G., *Eur. J. Comb.*, **31** (2010) 1887.
[40] Moroni D., Bolhuis P. G. and van Erp T. S., *J. Chem. Phys.*, **120** (2004) 4055.
[41] Faradjian A. K. and Elber R., *J. Chem. Phys.*, **120** (2004) 10880.
[42] van Erp T. S. and Bolhuis P., *J. Comput. Phys.*, **205** (2005) 157.
[43] Vervust W., Zhang D. T., van Erp T. S. and Ghysels A., *Biophys. J.*, **122** (2023) 1.
[44] Ma A. and Dinner A. R., *J. Phys. Chem. B*, **109** (2005) 6769.
[45] Frassek M., Arjun A. and Bolhuis P. G., *J. Chem. Phys.*, **155** (2021) 064103.
[46] Jung H., Covino R., Arjun A., Leitold C., Dellago C., Bolhuis P. G. and Hummer G., *Nat. Comput. Sci.*, **3** (2023) 334.
[47] van Erp T. S., Moqadam M., Riccardi E. and Lervik A., *J. Chem. Theory Comput.*, **12** (2016) 5398.
[48] Falkner S., Coretti A., Romano S., Geissler P. and Dellago C., arXiv:2207.14530 (2023).
[49] Behler J. and Parrinello M., *Phys. Rev. Lett.*, **98** (2007) 146401.
[50] Bartók A. P., Payne M. C., Kondor R. and Csányi G., *Phys. Rev. Lett.*, **104** (2010) 136403.
[51] Batzner S., Musaelian A., Sun L., Geiger M., Mailoa J. P., Kornbluth M., Molinari N., Smidt T. E. and Kozinsky B., *Nat. Commun.*, **13** (2022) 2453.
[52] Van Erp, T. S. and Meijer, E. J., *Angew. Chem. Int. Ed.*, **43** (2004) 1660.
[53] Senftle T. P., Hong S., Islam M. M., Kylasa S. B., Zheng Y. X., Shin Y. K., Junkermeier C., Engel-Herbert R., Janik M. J., Aktulga H. M., Verstraelen T., Grama A. and van Duin A. C. T., *npj Comput. Mater.*, **2** (2016) 15011.
[54] Moqadam M., Riccardi E., Trinh T. T., Astrand P.-O. and van Erp T. S., *J. Chem. Phys.*, **143** (2015) 184113.
[55] Lervik A. and van Erp T. S., *J. Chem. Theory Comput.*, **11** (2015) 2440.