



HAL
open science

Self-optimized construction of transition rate matrices from accelerated atomistic simulations with Bayesian uncertainty quantification

T D Swinburne, D Perez

► To cite this version:

T D Swinburne, D Perez. Self-optimized construction of transition rate matrices from accelerated atomistic simulations with Bayesian uncertainty quantification. *Physical Review Materials*, 2018, 2 (5), pp.053802. <10.1103/physrevmaterials.2.053802>. <hal-01962002>

HAL Id: hal-01962002

<https://hal.science/hal-01962002v1>

Submitted on 11 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Self-optimized construction of transition rate matrices from accelerated atomistic simulations with Bayesian uncertainty quantification

Thomas D. Swinburne and Danny Perez

Theoretical Division T-1, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA



(Received 7 March 2018; published 21 May 2018)

A massively parallel method to build large transition rate matrices from temperature-accelerated molecular dynamics trajectories is presented. Bayesian Markov model analysis is used to estimate the expected residence time in the known state space, providing crucial uncertainty quantification for higher-scale simulation schemes such as kinetic Monte Carlo or cluster dynamics. The estimators are additionally used to optimize where exploration is performed and the degree of temperature acceleration on the fly, giving an autonomous, optimal procedure to explore the state space of complex systems. The method is tested against exactly solvable models and used to explore the dynamics of C15 interstitial defects in iron. Our uncertainty quantification scheme allows for accurate modeling of the evolution of these defects over timescales of several seconds.

DOI: [10.1103/PhysRevMaterials.2.053802](https://doi.org/10.1103/PhysRevMaterials.2.053802)

I. INTRODUCTION

The vast size and complexity of the potential energy landscape of materials make the investigation of their long-time dynamical evolution extremely difficult, as significant free-energy barriers between different regions of configuration space prohibit the use of direct simulation methods. Indeed, molecular dynamics (MD) simulations of materials are typically restricted to sub-microsecond timescales, a time that is often much too short for a trajectory to cross the barriers that determine the long-time behavior. This makes extrapolation of long-time behavior based on short simulations fraught with danger.

Overcoming the extremely restrictive timescale limitation of MD is a long-standing challenge and numerous solution strategies have been proposed over the years. In *open-ended* situations where the goal is to generate dynamically correct evolution from a given initial condition without regards to possible end states, these methods often adopt one of two philosophies. First, trajectory-based methods such as accelerated molecular dynamics (AMD) [1–4] and adaptive kinetic Monte Carlo [5,6] generate individual trajectories that span long timescales without having to extensively explore configuration space. They do so by breaking up the problem of generating a long trajectory into that of generating a proper sequence of state-to-state transitions, which can be effectively carried out using specifically crafted MD simulations. The second class of techniques, including methods such as discrete path sampling [7] or Markov state models [8], instead begin by thoroughly exploring the energy landscape, thereby producing a kinetic model that can then be postprocessed to infer long-time behaviors.

While the local nature of the exploration required by the first class of approaches typically leads to more accurate and affordable results, it produces only one (or a few) of an astronomically large number of possible trajectories; the representativity of the results it generates can therefore be difficult to assess. On the other hand, the second approach

produces a comprehensive *global* model of the dynamics that can account for the contribution of large ensembles of trajectories, but the accuracy of its prediction requires that the underlying model be complete (or at least “sufficiently” so), an assumption that can be hard to assess, as fully sampling configuration space is typically impossible for nontrivial systems. Quantifying the completeness of models of the potential energy landscape has therefore recently emerged as a critical issue [9–11]. It is important to note that this same challenge also affects trajectory-based methods that rely on having a complete *local* description of the landscape (e.g., as in adaptive kinetic Monte Carlo [5,6]). A further challenge that has received comparatively less attention is that generating a sufficiently complete model that is accurate enough to make long-time predictions is likely to be an extremely computationally costly endeavor. Finding optimal strategies to allocate computational resources, in particular on massively parallel architectures, can therefore be expected to be paramount in making such approaches practical and scalable.

In this paper, we introduce a self-optimizing scheme called TAMMBER (temperature-accelerated Markov models with Bayesian estimation of rates) that comprehensively addresses these challenges. As illustrated in Fig. 1, TAMMBER relies on an AMD method, namely temperature-accelerated dynamics (TAD) [3,12], as an efficient local exploration tool. The local completeness of the TAD exploration is assessed using a Bayesian framework. TAMMBER then invokes the mathematics of absorbing continuous-time Markov chains (CTMCs) [13–18] to provide a global exploration completeness metric, the expected *residence time* in the known configuration space. This completeness metric is then systematically optimized using a parallel resource allocation protocol.

To put the central concepts of this paper in a concrete setting, consider a system with a total discrete state space \mathcal{S} . States are here defined as basins of attraction under energy minimization, as is customary for hard materials. After a given period of exploration with TAD, we will have discovered a subset $\mathcal{K} \subset \mathcal{S}$ of the total state space, the known states. While




FIG. 1. TAMMBER workflow. TAD MD produces interstate transition trajectories which are analyzed by Bayesian rate estimators and static calculation. An absorbing Markov chain then gives the expected residence time and optimally allocates resources and the degree of temperature acceleration. The cycle is then repeated until the target residence time is achieved.

an observed system state $i \in \mathcal{K}$ will be connected to a subset of states $\mathcal{S}_i \subset \mathcal{S}$, in general we will have observed only a subset of connections $\mathcal{K}_i \subset \mathcal{K}$ in the explored state space [19]. Defining the transition rate from a state i to a state j at a temperature $T = 1/(k_B\beta)$ as $k_{ij}(\beta)$, the *total* escape rate for a state i reads

$$k_i^{\text{tot}}(\beta) \equiv \sum_{j \in \mathcal{S}_i} k_{ij}(\beta). \quad (1)$$

As discussed above, due to incomplete exploration we will only have access to the *observed* escape rate

$$k_i^{\text{obs}}(\beta) \equiv \sum_{j \in \mathcal{K}_i} k_{ij}(\beta), \quad (2)$$

which immediately defines the statewise *unknown* escape rate

$$k_i^{\text{un}}(\beta) \equiv k_i^{\text{tot}}(\beta) - k_i^{\text{obs}}(\beta) = \sum_{j \in \mathcal{S}_i \setminus \mathcal{K}_i} k_{ij}(\beta), \quad (3)$$

where $\mathcal{S}_i \setminus \mathcal{K}_i \equiv \{x : x \in \mathcal{S}_i, x \notin \mathcal{K}_i\}$ is the set difference between \mathcal{S}_i and \mathcal{K}_i . In an absorbing CTMC, the unknown rates k_i^{un} are encoded as transition rates to single or multiple absorbing states (sinks) that represents the entire unexplored space and unobserved connections within the explored state space. Standard results [20] can be used to obtain the *residence time* of the model, which quantifies the expected amount of time before an unknown transition should statistically occur. The residence time can be interpreted as a typical duration over which model trajectories are a valid representation of the true system trajectories, providing an important uncertainty quantification metric when using the calculated rate matrices in coarse-grained methods such as kinetic Monte Carlo or cluster dynamics. The direct optimization of this metric with respect to additional computational work then provides an optimal allocation strategy to maximally improve the quality of the model at the smallest possible computational cost. Upon completion of a batch of TAD simulations, the model is updated and the cycle repeats.

The mathematics of absorbing CTMCs have previously been used to accelerate kinetic Monte Carlo simulations of superbasin escape [13] and highly heterogeneous glassy systems [14,15] though in both of these cases the chains were fully specified and this partitioning into two groups was made for computational convenience. Estimation of the unknown rate for each state has previously been investigated in molecular dynamics simulations of biological systems [17,18], while high-temperature dynamics has also been used to estimate the degree of sampling completeness in individual states [16] which is closely related to estimation of the unknown rate. The central development of this work is both the robust form of our estimators for the unknown escape rate from each state and an expression for the expected decrease in the unknown rate with additional computational work. Using these expressions we are able to determine both the optimal degree of temperature acceleration for each state on the fly and the response of the residence time to additional computational effort applied to a given *distribution* of states, an essential feature for application to massively parallel computers. Importantly, by optimizing the distribution of computational resources to grow the residence time as fast as possible, we optimize a *global* metric of sampling completeness, a point we return to below.

The paper is organized as follows. In Sec. II we recall the temperature-accelerated dynamics method [3] and detail how the method may be extended to allow for a variable high temperature. In Sec. III we derive a Bayesian estimator for the $k_{ij}(\beta)$ of observed transitions ($j \in \mathcal{K}_i$) at any desired temperature and the *unknown* escape rate $k_i^{\text{un}}(\beta)$ from each state. In Sec. IV we derive an analytical expression to determine the statewise optimum temperature to reduce the unknown rate for each state and use these results to derive the residence time and optimal control protocol using an absorbing CTMC in Sec. V. Details of the numerical implementation are described in Sec. VI, along with a test against known rate matrices (using kinetic Monte Carlo to generate trajectories) and a demonstrative study of C15 interstitial defects in iron.

II. TEMPERATURE-ACCELERATED DYNAMICS

The temperature-accelerated dynamics (TAD) method [3,12] is an AMD technique that exploits the Poisson distribution of rare-event escape times [21] and the approximations of harmonic transition state theory (HTST) [22] to generate statistically correct low-temperature trajectories from high-temperature MD data alone. When the transition barriers are sufficiently large, TAD can provide a very significant acceleration of the state-to-state dynamics as compared to MD, because the first event to occur at low temperature will typically occur after only a much shorter time at a higher temperature. TAD provides a statistically sound way of assessing when the said first event has indeed been observed at high temperature, and hence of selecting a proper low-temperature transition.

We recall that when the free-energy barrier ΔF_{ij} for some state transition $i \rightarrow j$ is much larger than the thermal energy β^{-1} , the transition rate $k_{ij}(\beta)$ is well approximated by the Arrhenius expression [22]

$$k_{ij}(\beta) = \omega_{ij} \exp[-\beta \Delta F(\beta)] \simeq \nu_{ij} \exp[-\beta \Delta E_{ij}]. \quad (4)$$

The second equality in Eq. (4) constitutes the HTST approximation, where the entropic contribution to the barrier ΔS_{ij} is assumed to be constant, leading to a constant prefactor $v_{ij} = \omega_{ij} \exp(\Delta S_{ij}/k_B)$ and a potential-energy barrier ΔE_{ij} . The extension of the approach developed here to incorporated anharmonic entropic effects [23] will be the topic of a future publication. HTST (4) can be exploited in the present context by noting that the event times for a Poisson process of rate $k(\beta)$ are distributed as

$$\tau_{ij}(\beta) \sim -\log[\mathcal{U}(0,1)]/k_{ij}(\beta), \quad (5)$$

where $\mathcal{U}(0,1)$ is the uniform distribution on the unit interval; from this functional form it is clear that a valid event time $\tau_{ij}(\beta')$ at a different temperature can be obtained from a sample $\tau(\beta)$ through

$$\tau_{ij}(\beta') = \tau_{ij}(\beta) \frac{k_{ij}(\beta)}{k_{ij}(\beta')} \simeq \tau_{ij}(\beta) \exp[(\beta' - \beta)\Delta E_{ij}], \quad (6)$$

where the HTST approximation was used to obtain the final relation. As ΔE_{ij} is readily calculated using minimum-energy-path algorithms such as the NEB method [24], after a process has been observed for the first time at high temperature, we can thus generate a corresponding first-passage time at other temperatures.

In TAD, this remapping of first-passage times is exploited as follows. Consider a state i that has dynamically accessible pathways to a set of connected states $j \in \mathcal{S}_i$, with escape rates $k_{ij}(\beta) = v_{ij} \exp[-\beta \Delta E_{ij}]$. TAD uses high-temperature MD to produce high-temperature escape times $\{\tau_{ij}(\beta_H)\}$ to a subset of connected states $\mathcal{K}_i \subset \mathcal{S}_i$. Once an escape is detected, the system is put back into state i , accumulating a total effective state time $\tau_i(\beta)$. The escape times along each pathways can then be rescaled to yield a set of low-temperature first-passage times $\{\tau_{ij}(\beta_L)\}$, which will in general have a different ordering given the nonlinear character of (6). In conventional TAD, the goal is to identify the transition that should have occurred first, i.e., the transition which corresponds to the minimum value of $\tau_{ij}(\beta_L)$. The central difficulty is the observed escape times are only to a subset of all possible final states \mathcal{K}_i . It is therefore important to avoid prematurely choosing a low-temperature transition from the set transitions so far observed at high temperature. TAD achieves this through a Poisson uncertainty bound; defining a minimum prefactor $v_{\min} \simeq 0.1$ THz, high-temperature MD is carried out until the probability that the proper first escape pathway at low temperature has yet to be observed at high temperature is less than $\delta \sim 0.05$. The worst possible case in this setting is that of a low-barrier and low-prefactor process with rate $v_{\min} \exp(-\beta_H E_i^{\min})$, where E_i^{\min} is the smallest barrier that could potentially remain unobserved after running dynamics at high temperature for a time $\tau_i(\beta_H)$. It is simple to show that [12]

$$E_i^{\min} = \beta_H^{-1} \log \left[\frac{v_{\min} \tau_i(\beta_H)}{\log(1/\delta)} \right], \quad (7)$$

which produces a low-temperature effective state time

$$\tau_i(\beta_L) = \tau_i(\beta_H) \exp[(\beta_L - \beta_H)E_i^{\min}], \quad (8)$$

after which we have a confidence $1 - \delta$ to have seen all relevant first passages up to this time.




FIG. 2. Left: Illustration of the TAD method developed here. Low-temperature first-passage times become valid as they are swept past, while the high temperature can be changed to accommodate trajectory data at a new temperature. Left: Qualitatively representative posterior for the total escape rate from a state. The unknown rate is the difference between the mean total rate and the observed rate.

In the original TAD method, the goal is to follow the first valid escape process; i.e., state time is accumulated until $\tau_i(\beta_L)$ is greater than the smallest rescaled first-passage time. In the present case we continue accumulating state time, producing an ever greater catalog of valid low-temperature escape times [i.e., all of those whose rescaled event times are smaller than $\tau_i(\beta_L)$], for use in our rate estimators detailed in the next section. As the total state time τ_i and first-passage times τ_{ij} are defined at any temperature, we can incorporate multitemperature data by using (8). An illustration of this procedure is detailed in Fig. 2.

III. DETERMINATION OF THE KNOWN AND UNKNOWN ESCAPE RATES FROM A STATE

In order to apply the absorbing CTMC analysis which is central to our approach, we need to produce an estimate for the individual rates $k_{ij}(\beta)$ between known states at any given temperature and for the *unknown* escape rate $k_i^{\text{un}}(\beta)$ from each known state. In the following we derive Bayesian likelihood estimators for the individual and total escape rates from a given state using the first-passage trajectories $\tau_{ij}(\beta)$ and state time $\tau_i(\beta)$.

A. Estimation of individual escape rates

Once an individual escape process from a state i to a state j has been observed, the NEB method can be used to obtain the minimum-energy pathway and hence the energy barriers ΔE_{ij} and ΔE_{ji} . To calculate the individual escape rates k_{ij} and k_{ji} we therefore only require calculation of the rate prefactors v_{ij} and v_{ji} .

It is possible to directly calculate an estimate for the rate prefactors using harmonic transition state theory [22]. A key advantage is that the HTST approximation to v is often accurate and produces a rate matrix which satisfies detailed balance, but calculation requires computationally expensive diagonalization of the Hessian matrix at the end points and saddle point of the transition pathway [25]. An alternative approach is to directly estimate the rate prefactor from the transitions observed during MD simulation. A disadvantage of this approach is that this requires multiple observed transitions to give reliable results and that the resultant prefactors have no

guarantee of satisfying detailed balance. Nevertheless, when transitions are sufficiently rapid (which can be expected when using accelerated approaches such as TAD) sufficient data can often be obtained to produce accurate estimates.

In this section we derive a simple Bayesian estimator for the rate prefactor which incorporates prior knowledge of the prefactor and dynamical information from an ensemble of escape-replace trajectory data. The prior estimate for the prefactor can either be set to a typical value of $v_0 = 1$ THz or a static HTST calculation. In a Bayesian setting, this knowledge can be encoded in an unnormalized prior distribution

$$\pi_0(v_{ij}) = \exp[-\alpha(v_{ij}/v_0 - 1)^2/2], \quad (9)$$

where v_0 is the prior estimate and α will turn out to control the number of data points that are needed to override the influence of the prior. As a result, if a full HTST calculation is undertaken, α should be large as we are confident that our prior is accurate. In practice, as a full prefactor calculation is computationally intensive, we only undertake such calculations when we expect dynamical data to be rare, i.e., when ΔE_{ij} is large, though many strategies can be envisaged, for example performing an approximate calculation with the degree of approximation reflected in the prior distribution.

We represent escape-replace trajectory data as $\{\beta_i, \tau_i, N_{ij}\}$, where β_i is the inverse temperature, τ_i is the total effective state time at that temperature, and N_{ij} is the total number of $i \rightarrow j$ transitions observed [26]. For clarity of presentation we also define the dimensionless, Boltzmann-scaled trajectory times

$$\tilde{\tau}_{i,j} = \tau_i v_0 \exp[-\beta_i \Delta E_{ij}], \quad (10)$$

where the notation distinguishes $\tilde{\tau}_{i,j}$ from the first-passage times τ_{ij} . Using the Poisson likelihood for N events in a time τ , $(k\tau)^N \exp(-k\tau)/N!$, the HTST relation (4), and the prior distribution (9), the unnormalized posterior for the rate prefactor reads

$$\pi(v_{ij} | \tilde{\tau}_{i,j}, N_{ij}) = \pi_0(v)(v_{ij} \tilde{\tau}_{i,j})^{N_{ij}} \exp(-v_{ij} \tilde{\tau}_{i,j}/v_0). \quad (11)$$

While the posterior distribution is quite cumbersome, we can produce an estimator for v_{ij} using the maximum log likelihood (MLL) technique, where the logarithm of the unnormalized posterior (11) is maximized with respect to v_{ij} , a well known procedure in parameter estimation [27]. Through elementary operations one obtains from $\partial_v \log \pi = 0$ a quadratic equation for v_{ij} which has the unique positive solution

$$v_{ij} = \frac{v_0}{2} \left[1 - \frac{\tilde{\tau}_{i,j}}{\alpha} + \sqrt{\left(1 - \frac{\tilde{\tau}_{i,j}}{\alpha}\right)^2 + 4 \frac{N_{ij}}{\alpha}} \right]. \quad (12)$$

In the small time and data limit $\tilde{\tau}_{i,j} \ll \alpha, N_{ij} \ll \alpha$, we find $v_{ij} = v_0$, as one would expect, while at long times $\tilde{\tau}_{i,j} \gg \alpha$ we recover $v_{ij} = N_{ij} \exp[\beta_i \Delta E_{ij}]/\tau_i$, which is the minimum-variance estimator for this Poisson process [28].

We have found $\alpha \simeq 10$ to give robust sampling behavior using a standard initial prefactor $v_0 = 0.1$ THz. A key advantage of the Bayesian approach is that if a more detailed HTST prefactor calculation is undertaken to give a more reliable prior estimate, we make the prior distribution sharper by increasing the α parameter. As a result, a much larger number of dynamical data are required to significantly change the posterior

prediction of the prefactor, thus naturally incorporating the two estimation methods.

B. Estimation of the unknown escape rate from a state

With calculated prefactors and energy barriers $\{v_{ij}, \Delta E_{ij}\}$ for each observed escape process, we can readily calculate the corresponding escape rates $k_{ij}(\beta)$. Furthermore, using the procedure described above, we can also obtain an effective state time $\tau_i(\beta)$ at any given temperature. In this section, we show how this information, taken together with the sampled first-passage time $\tau_{ij}(\beta)$ obtained with TAD, can be used to produce a Bayesian estimator for the *unknown* escape rate from the generated first-passage-time trajectory, again at any temperature.

In anticipation of the results below, we time-order the individual escapes labels such that $\tau_{i(j-1)}(\beta) < \tau_{ij} < \tau_{i(j+1)}$ and then define the running total rate

$$k_{i,j}^{\text{obs}}(\beta) \equiv \sum_{\tau_{ik}(\beta) \leq \tau_{ij}(\beta)} k_{ik}(\beta); \quad (13)$$

i.e., the running total rate $k_{i,j}^{\text{obs}}(\beta)$ includes all events that occurred at times $\tau_{ij}(\beta)$ that are lower or equal to the effective residence time at β , $\tau_i(\beta)$. As all rates are evaluated at a constant temperature for the entirety of this section, we now omit β for clarity of presentation.

To build a posterior distribution for the unknown rate, consider the likelihood of observing a first passage $i \rightarrow j$ after waiting a time $\tau_{ij} - \tau_{i(j-1)}$ since the last event. For a postulated total rate k , the remaining rate in this time interval is simply $k - k_{i,j-1}^{\text{obs}}$, giving a likelihood for τ_{ij} of

$$\begin{aligned} \pi(\tau_{ij} | k_i^{\text{tot}} = k) &= [k - k_{i,j-1}^{\text{obs}}](\tau_{ij} - \tau_{i(j-1)}) \\ &\times \exp\left\{-[k - k_{i,j-1}^{\text{obs}}](\tau_{ij} - \tau_{i(j-1)})\right\}. \end{aligned} \quad (14)$$

We note the use of the remaining total rate in the interval $[\tau_{i(j-1)}, \tau_{ij}]$ is essential to give the correct likelihood. In addition, as we know that the total rate satisfies $k_i^{\text{tot}} = k_i^{\text{obs}} + k_i^{\text{un}}$, we can write the same likelihood for τ_{ij} for a postulated unknown rate k as

$$\begin{aligned} \pi(\tau_{ij} | k_i^{\text{un}} = k) &= [k + k_i^{\text{obs}} - k_{i,j-1}^{\text{obs}}](\tau_{ij} - \tau_{i(j-1)}) \\ &\times \exp\left\{-[k + k_i^{\text{obs}} - k_{i,j-1}^{\text{obs}}](\tau_{ij} - \tau_{i(j-1)})\right\}, \end{aligned} \quad (15)$$

where k_i^{obs} [as defined in (2)] is the sum of the escape rates for events at any temperature, independently of their first-passage times at β . A total likelihood for all the observed event times for a postulated unknown rate is simply the product of (15) for each event satisfying $\tau_{ij}(\beta) \leq \tau_i(\beta)$, multiplied by the likelihood of not seeing any other events over a time $\delta\tau_i^W = \tau_i - \max_{\tau_{ij} < \tau_i}(\tau_{ij})$ to give

$$\pi(\{\tau_{ij}\} | \tau_i, k_i^{\text{un}} = k) = \exp(-k\delta\tau_i^W) \prod_{\tau_{ij} \leq \tau_i} \pi(\tau_{ij} | k_i^{\text{un}} = k). \quad (16)$$

We can now use Bayes' formula to construct an unnormalized posterior for the unknown rate, using the Jeffreys prior [29] $\pi_0(k) = 1/k$ for the initial likelihood function $k \exp(-kt)$. Removing all multiplicative factors independently of the postulated unknown rate, as these will disappear under renormalization, we obtain a central result of this paper, an unnormalized posterior distribution for the unknown rate

$$\pi(k_i^{\text{un}}|\tau_i, \{\tau_{ij}\}) = \frac{\exp(-k_i^{\text{un}}\tau_i)}{k_i^{\text{un}} + k_i^{\text{obs}}} \times \prod_{\tau_{ij} < \tau_i} (k_i^{\text{un}} + k_i^{\text{obs}} - k_{i;j-1}^{\text{obs}}). \quad (17)$$

We emphasize that although all trajectory information is used in the individual rate calculations, we only use first-passage information in the Bayesian posterior (17). It can in fact be shown that subsequent passages in fact do not contribute additional information, as we assume that the rate for a given process can be calculated once it has been observed. This is ideal for implementation in a TAD setting, as multitemperature MD data can be incorporated to produce an effective first-passage trajectory at a wide range of desired temperatures.

A prediction for the unknown rate $\langle k_i^{\text{un}}(\beta) \rangle$ and total rate $\langle k_i^{\text{un}}(\beta) \rangle$ at an inverse temperature β can now be produced by evaluating (17), yielding moments

$$[\langle k_i^{\text{un}}(\beta) \rangle^n] = \frac{\int_0^\infty k^n \pi(k|\beta, \tau_i, \{\tau_{ij}\}) dk}{\int_0^\infty \pi(k|\beta, \tau_i, \{\tau_{ij}\}) dk}, \quad (18)$$

where we have reintroduced the temperature dependence explicitly. In Appendix A we show that these integrals can be expressed analytically by exploiting properties of exponential integrals and a recursive scheme to expand the product, avoiding numerical quadrature issues.

This is the first important result of this paper: the first moment, namely the mean, will be used as an estimator of the unknown rate out of a given state given an observed sequence of first-passage times generated with TAD. This provides a crucial local completeness metric. The higher moments also prove critical to solve the important question of the choice of the optimal high temperature at which the TAD procedure should be carried out in order to maximize computational efficiency, a problem which we discuss next.

IV. OPTIMAL TAD TEMPERATURE

The TAD method uses an elevated temperature $T_H = 1/(k_B\beta_H)$ to reduce the computational effort required to produce a valid set of first-passage times and pathways at some lower temperature $T_L = 1/(k_B\beta_L)$. When all barriers are sufficiently large compared to $k_B T_L$, the efficacy of the TAD method initially increases with increasing T_H away from T_L . However, if T_H becomes too high, transitions with very large energy barriers will become more frequent. As characterizing these transitions incurs a cost but contributes very little to the low-temperature total rate, the computational efficiency of the procedure should ultimately decrease with increasing T_H . In addition, known events will reoccur more frequently at higher T_H , increasing the frequency at which the system must be re-prepared in the initial state in order to accumulate additional effective state time.

These arguments indicate that there will in general exist an optimum high temperature T_H , the precise value of which depends on the desired outcome. Recent work [30] has investigated finding the optimal T_H in TAD to produce a single valid escape event from a given state, i.e., a single rescaled first-passage time less than the effective state time $[\tau_{ij}(\beta_L) < \tau_i(\beta_L)]$. In this section, we instead ask for the temperature which maximizes the decrease of the expected low-temperature unknown rate $\langle k_i^{\text{un}}(\beta_L) \rangle$ with respect to additional computational effort $c_i(\beta_H)$ that consists of carrying out the TAD procedure at temperature β_H , namely

$$\beta_i^{\text{TAD}} = \arg \max_{\beta_H} \left[-\frac{d\langle k_i^{\text{un}}(\beta_L) \rangle}{dc_i(\beta_H)} \right]. \quad (19)$$

Given that the simulation cost is dominated by force calculation (also known as force calls), the total computational effort per unit high-temperature MD time can be written in units of force calls as

$$\frac{dc_i(\beta_H)}{d\tau_i(\beta_H)} = \dot{c}_{\text{MD}} + c_{\text{ST}}k_i^{\text{obs}}(\beta_H) + c_{\text{NEB}}k_i^{\text{un}}(\beta_H), \quad (20)$$

where \dot{c}_{MD} is the number of force calls per unit MD time in frequency units, c_{ST} is the cost of state identification and preparation in force calls, and c_{NEB} is the cost of an NEB calculation in force calls. In a typical example, where transition rates are quoted in THz and the MD time step is a femtosecond, we have $\dot{c}_{\text{MD}} = 1000$, $c_{\text{ST}} \simeq 1000$, and $c_{\text{NEB}} \simeq 10\,000$. By the chain rule we make the useful expansion

$$\frac{d\langle k_i^{\text{un}}(\beta_L) \rangle}{dc_i(\beta_H)} = \frac{d\langle k_i^{\text{un}}(\beta_L) \rangle}{d\tau_i(\beta_H)} \left(\frac{dc_i(\beta_H)}{d\tau_i(\beta_H)} \right)^{-1}. \quad (21)$$

To evaluate the first term in (21) we first consider the expected change in the low-temperature unknown rate from a small interval $\delta\tau_i(\beta_H)$ of high-temperature MD when \mathcal{E} , a new transition, is observed, or $!\mathcal{E}$, when no new transition occurs. The corresponding change in the low-temperature state time, $\delta\tau_i(\beta_L)$, is readily evaluated through the use of (7) as

$$\delta\tau_i(\beta_L) = \delta\tau_i(\beta_H) \frac{\beta_L}{\beta_H} \left(\frac{\log(1/\delta)}{v_{\text{min}}\tau_i(\beta_H)} \right)^{\beta_L/\beta_H - 1}. \quad (22)$$

We evaluate changes in $k_i^{\text{un}}(\beta_L)$ through perturbation theory applied to expectation values over the low-temperature posterior for the total rate, $\pi(k_i^{\text{tot}}|\beta_L, \tau_i)$. If no event is seen in high-temperature MD, the new posterior is given by

$$\pi(k_i^{\text{un}}|\beta_L, \tau_i + \delta\tau_i) = \pi(k_i^{\text{un}}|\beta_L, \tau_i) \exp(-[k_i^{\text{un}}]\delta\tau_i). \quad (23)$$

To leading order in $\delta\tau_i(\beta_L)$ the expected change in the unknown rate takes the simple form

$$\langle \delta k_i^{\text{un}}(\beta_L) | !\mathcal{E} \rangle = -\delta\tau_i(\beta_L) [\langle [k_i^{\text{un}}(\beta_L)]^2 \rangle - \langle k_i^{\text{un}}(\beta_L) \rangle^2]. \quad (24)$$

If an event \mathcal{E} is seen in high-temperature MD, to a state p with a rescaled low-temperature rate $k_{\text{new}} = k_{ip}(\beta_L)$, the new posterior distribution is given by

$$\pi(k_i^{\text{un}}|\beta_L, \tau_i + \delta\tau_i) = \pi(k_i^{\text{un}} + k_{\text{new}}|\beta_L, \tau_i) \times (k_i^{\text{un}} + k_i^{\text{obs}} - \max k_{i;j}^{\text{obs}}). \quad (25)$$

While we can progress without any assumptions, to simplify the expectation value over this new distribution we take the

mild assumption that $\max k_{i,j}^{\text{obs}} \simeq k_i^{\text{obs}}$, i.e., that the majority of the rate has been seen at the temperature of interest. We have found this to hold in practice, and can be expected from the form of the rescaled state time τ_i . Under this approximation, the expected change in the unknown rate reads

$$\langle \delta k_i^{\text{un}}(\beta_L) | \mathcal{E} \rangle = -k_{\text{new}} + \frac{\langle [k_i^{\text{un}}(\beta_L)]^2 \rangle - \langle k_i^{\text{un}}(\beta_L) \rangle^2}{\langle k_i^{\text{un}}(\beta_L) \rangle}. \quad (26)$$

To complete this expression we require an estimate for the new low-temperature rate $k_{\text{new}} = k_{i_p}(\beta_L)$, ideally without making any additional assumptions on the spectrum of escape rates. We base our assumption on the expected first-passage-time relation $\langle \tau_{i_p} \rangle = 1/k_{\text{new}}$. New events are therefore expected to be first observed in order of descending rate. If the barrier spectrum is dense, then a reasonable estimate for the next new event rate is simply the minimum of all the observed rates so far, $\min\{k_{ij}(\beta_L)\}$. However, if the spectrum has a large spectral gap, we would expect long periods without any new events, meaning the minimum of the seen rates could significantly overestimate the next event rate. In this long-waiting-time limit, it can be shown that the Bayesian estimator gives a max log likelihood unknown rate of $\langle k_i^{\text{un}}(\beta_L) \rangle \sim 1/\tau_i(\beta_L)$. As the new rate is expected to occur at a time $\tau_i(\beta_L)$, we see that the unknown rate estimate is expected to be a slight overestimate, i.e., our estimates tend to be conservative. Combining these two cases, our estimate for the next observed rate is therefore

$$\langle k_{\text{new}} \rangle \simeq \min \left[\langle k_i^{\text{un}}(\beta_L) \rangle \cup \{k_{ij}(\beta_L)\} \right]. \quad (27)$$

Given that the expected probability of seeing a new event in high-temperature MD is simply $P(\mathcal{E}) = \delta\tau_i(\beta_H)k_i^{\text{un}}(\beta_H)$ in the limit of small $\delta\tau_i(\beta_H)$, with $P(!\mathcal{E}) = 1 - P(\mathcal{E})$, we can write the expected change in the low-temperature unknown rate as

$$\langle \delta k_i^{\text{un}}(\beta_L) \rangle = P(\mathcal{E})\langle \delta k_i^{\text{un}}(\beta_L) | \mathcal{E} \rangle + P(!\mathcal{E})\langle \delta k_i^{\text{un}}(\beta_L) | !\mathcal{E} \rangle. \quad (28)$$

Combining the above manipulations we can write the final objective function as

$$\begin{aligned} -\frac{d\langle k_i^{\text{un}}(\beta_L) \rangle}{dc_i(\beta_H)} &= \left(\frac{dc_i(\beta_H)}{d\tau_i(\beta_H)} \right)^{-1} \left[\langle k_{\text{new}} \rangle \langle k_i^{\text{un}}(\beta_H) \rangle \right. \\ &+ \left(\frac{\tau_i(\beta_L)}{\tau_i(\beta_H)} - \frac{\langle k_i^{\text{un}}(\beta_H) \rangle}{\langle k_i^{\text{un}}(\beta_L) \rangle} \right) \{ \langle [k_i^{\text{un}}(\beta_L)]^2 \rangle \\ &\left. - \langle k_i^{\text{un}}(\beta_L) \rangle^2 \right]. \quad (29) \end{aligned}$$

While this expression appears complex, all relevant quantities can be readily calculated using our Bayesian estimator and the results derived above. In our numerical implementation, we find the maximum of (29) to determine a different optimal β_H for every state in the system. This determination is periodically refined to ensure optimal performance.

V. ABSORBING MARKOV CHAIN ANALYSIS

In the preceding sections, we have described a scheme to estimate transition rates $k_{ij}(\beta)$ between known states $i, j \in \mathcal{K}$ and the unknown rate for each state $k_i^{\text{un}}(\beta)$. We have also derived the expected change (29) in the low-temperature unknown rate

$k_i^{\text{un}}(\beta_L)$ with additional computational work at a temperature β_H in order to determine the optimum temperature at which to carry out the TAD procedure. In this section we use the estimated rates to build an absorbing Markov chain [20], giving both the expected residence time $\langle \tau^{\text{res}} \rangle$ spent in the known state space and the expected change in $\langle \tau^{\text{res}} \rangle$ as a result of additional computational effort. As discussed in the introduction, the expected residence time $\langle \tau^{\text{res}} \rangle$ is an important *global* measure of sampling completeness, providing an estimate of the length of trajectories that can safely be generated from the CTMC; trajectories longer than $\langle \tau^{\text{res}} \rangle$ on the complete CTMC would have a significant probability of containing transitions that are not part of the estimated CTMC. One should therefore avoid using the CTMC to make predictions on times that exceed $\langle \tau^{\text{res}} \rangle$.

We emphasize that $\langle \tau^{\text{res}} \rangle$ is a *global* metric that accounts for the wider energy landscape. This is quite distinct from a statewise approach to uncertainty; for example, if a particular state has a high unknown rate, a statewise approach would always demand more computational work in this state to reduce the uncertainty. However, in our global approach, work would only be done in this state if it is sufficiently frequently visited to have a significant influence on the global trajectory distribution.

In our setting, $\langle \tau^{\text{res}} \rangle$ can be estimated as follows. Consider an absorbing CTMC in a discrete state space $\mathcal{K} \cup \Delta$, namely the set of observed states and an absorbing state Δ , as illustrated in Fig. 1. Let $\mathbf{P}(t) = \mathbf{P}_{\mathcal{K}}(t) \oplus \mathbf{P}_{\Delta}(t)$ give the probability that the system is in a state $i \in \mathcal{K} \cup \Delta$ at time t ; the continuous-time limit yields

$$\dot{\mathbf{P}}(t) = \mathbf{P}(t) \cdot \mathbf{Q} \Rightarrow \mathbf{P}(t) = \mathbf{P}(0) \cdot \exp(\mathbf{Q}t). \quad (30)$$

The absorbing transition matrix \mathbf{Q} , illustrated in Fig. 1, has a structure

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{\mathcal{K}} & \mathbf{k}^u \\ \mathbf{0}^T & 0 \end{bmatrix}, \quad (31)$$

where $(\mathbf{Q}_{\mathcal{K}})_{ij} \equiv k_{ij} - k_i^{\text{tot}}\delta_{ij}$ for $i, j \in \mathcal{K}$, $\mathbf{0}$ is a vector of zeros and $(\mathbf{k}^u)_i \equiv k_i^{\text{un}}$. From the structure of \mathbf{Q} one finds that

$$\mathbf{P}_{\mathcal{K}}(t) = \mathbf{P}_{\mathcal{K}}(0) \cdot \exp(\mathbf{Q}_{\mathcal{K}}t). \quad (32)$$

As the probability of transition to Δ from a state i at a time t is given by $[\mathbf{P}_{\mathcal{K}}(t)]_i k_i^{\text{un}}$, the expected residence time is simply

$$\langle \tau^{\text{res}} \rangle = \int_0^{\infty} t \mathbf{P}_{\mathcal{K}}(t) \cdot \mathbf{k}^u dt = \frac{\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-2} \cdot \mathbf{k}^u}{\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-1} \cdot \mathbf{k}^u}. \quad (33)$$

Defining a vector of ones $\mathbf{1}_{\mathcal{K}}$, it is simple to show that $\mathbf{Q}_{\mathcal{K}} \cdot \mathbf{1}_{\mathcal{K}} = -\mathbf{k}^u$, giving the further simplification

$$\langle \tau^{\text{res}} \rangle = -\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-1} \cdot \mathbf{1}_{\mathcal{K}}. \quad (34)$$

This expression for the residence time can be evaluated by solving the linear equation $\mathbf{Q}_{\mathcal{K}}^T \cdot \mathbf{x} = \mathbf{P}_{\mathcal{K}}(0)$ to give $\langle \tau^{\text{res}} \rangle = -\mathbf{x} \cdot \mathbf{1}_{\mathcal{K}}$.

Since $\langle \tau^{\text{res}} \rangle$ quantifies the quality of the current CTMC, it is natural to use it as an objective function to guide further improvement given a computational effort δc that can be invested. To best harness massively parallel computational resources, the optimal allocation will be expressed as an allocation distribution $\{s_i\}$ which gives the proportion of

workers assigned to each state $i \in \mathcal{K}$. The computational effort c_i allocated to state i is therefore

$$\delta c_i \equiv s_i \delta c, \quad \sum_{i \in \mathcal{K}} s_i \equiv 1. \quad (35)$$

In the expression (34) for the residence time, only the unknown rates are affected by the additional computational work, giving to leading order in δc a change in the residence time of (see Appendix B)

$$\frac{\delta \langle \tau^{\text{res}} \rangle}{\delta c} = - \sum_{i \in \mathcal{K}} s_i \frac{\delta k_i^{\text{un}}(\beta_L)}{\delta c_i} [\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-1}]_i [\mathbf{Q}_{\mathcal{K}}^{-1} \cdot \mathbf{1}_{\mathcal{K}}]_i, \quad (36)$$

where $-\delta k_i^{\text{un}}/\delta c_i$ is precisely the maximized statewise cost function (29) found in the previous section, evaluated at its maximum, i.e., at the high temperature β_H which maximizes $-\delta k_i^{\text{un}}/\delta c_i$. As Eq. (36) takes the form of an inner product the optimal choice of s_i is simply

$$s_i = \eta \frac{\delta k_i^{\text{un}}(\beta_L)}{\delta c_i} [\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-1}]_i [\mathbf{Q}_{\mathcal{K}}^{-1} \cdot \mathbf{1}_{\mathcal{K}}]_i, \quad (37)$$

where $\eta^{-1} = \sum_{j \in \mathcal{K}} s_j$ ensures normalization. Solving the linear equation $[\mathbf{Q}_{\mathcal{K}}] \cdot \mathbf{y} = \mathbf{1}_{\mathcal{K}}$, one gets $s_i = \eta x_i y_i (\delta k_i^{\text{un}}/\delta c_i)$. This simple procedure ensures that additional resources are optimally invested in order to maximize $\langle \tau^{\text{res}} \rangle$ at the smallest computational cost. In practice, the optimal allocation is periodically updated using the latest CTMC.

The optimal allocation has a clear interpretation using two expressions that follow from Eq. (34) for the residence time. As the inner product with $\mathbf{1}_{\mathcal{K}}$ is simply a sum over the known states, the second term in (37), $-\mathbf{P}_{\mathcal{K}}(0) \cdot \mathbf{Q}_{\mathcal{K}}^{-1}$, is simply the expected time spent in a state i conditional on the initial distribution $\mathbf{P}_{\mathcal{K}}(0)$, which when summed over all states yields $\langle \tau^{\text{res}} \rangle$. If we instead take the initial distribution to be a delta function on a state i , the third term in (37), $-\mathbf{Q}_{\mathcal{K}}^{-1} \cdot \mathbf{1}_{\mathcal{K}}$, can be interpreted as the expected residence time in the known network, conditional on starting from a state i . The allocation of computational work to a state is thus a product of three factors: the degree to which the unknown rate will change under additional sampling, the amount of time (on average) spent in the state before absorption under the desired initial conditions, and the characteristic residence time of trajectories starting in the state. If a state is very well sampled, the last two factors might be large, but the change in the unknown rate with additional sampling will be very small, suppressing the allocation weight. Conversely, a poorly sampled state might be rarely visited and have a small residence time, but the change in the unknown rate will be very large, increasing the allocation weight. In this manner, TAMMBER is able to allocate computational work to a state according to a *global* measure of the state's influence on the ensemble of trajectories in the known state space, dependent only on the prescribed initial condition $\mathbf{P}_{\mathcal{K}}(0)$.

VI. TAMMBER SIMULATION CODE

We have implemented the TAMMBER workflow, illustrated in Fig. 1, within the ParSplice [31] simulation code, which provides the underlying framework for generating state-to-state

trajectories, state identification, and asynchronous control over the requested work using massively parallel computational resources. MD trajectories themselves are generated by the LAMMPS molecular dynamics package [32]; after a 1 ps thermalization and dephasing stage (which is repeated if a transition occurs [31]), a snapshot of the system is recorded 2–4 times over each ps trajectory segment, with the final snapshot relaxed and analyzed [31] to check for transitions between metastable states. If a transition is detected, the intermediate snapshots are relaxed and analyzed to find a more precise transition time and to check for multiple transitions, which can occur if a low barrier is found at a high temperature. Transition times and pathways are sent back to the central task manager, with new transitions submitted for climbing-image NEB calculations [24] and, if desired, Hessian prefactor calculations using LAMMPS force calls and the FIRE minimization routine [33].

The central task manager of TAMMBER analyzes, at regular intervals, all of the state-to-state trajectory data using the multitemperature TAD formalism outlined in Sec. II to produce a list of time-ordered first-passage times and final states for each state. The dynamical data $\{\tau_{ij}\}$ and static data v_{ij}^0, E_{ij} for each transition are then used to produce an estimate of the rate prefactor using the Bayesian estimators derived in Sec. III. With knowledge of the individual transition rates $k_{ij}(\beta) = v_{ij} \exp(-\beta \Delta E_{ij})$ at the desired temperature, we can estimate $\langle k_u(\beta) \rangle$ and $\langle k_u^2(\beta) \rangle$ using the Bayesian posterior distribution for the total escape rate (17) and therefore fully populate the matrix \mathbf{Q} for the absorbing Markov chain (30) at the low temperature β_L . The quality of this CTMC is assessed by computing $\langle \tau^{\text{res}} \rangle$ for a given initial distribution and further allocation of resources carried out according to the distribution (37) that maximizes the rate of increase of $\langle \tau^{\text{res}} \rangle$. The cycle then repeats until $\langle \tau^{\text{res}} \rangle$ is deemed sufficiently small, or computational resources are exhausted. In the next section, we first test TAMMBER against an exactly known total rate matrix using kinetic Monte Carlo to generate trajectories, then use TAMMBER to explore the evolution of interstitial clusters in iron.

A. Validation using a known rate matrix

A key component of the TAMMBER code is to estimate $\langle k_i^{\text{un}} \rangle$, the unknown (or remaining) rates from each explored state, in order to construct an absorbing CTMC which both allocates resources and provides a metric for the degree of exploration. To validate our estimator for $\langle k_i^{\text{un}} \rangle$, we replaced the molecular dynamics engine with a simple kinetic Monte Carlo (kMC) routine [34] using a prescribed matrix rate matrix $k_{ij} = v_{ij} \exp(-\beta \Delta E_{ij})$ constructed at any temperature from a pre-specified list of energy barriers ΔE_{ij} and prefactors v_{ij} . To ensure the rate matrix satisfies detailed balance, we assign a free energy $F_i = E_i - \beta^{-1} \log \omega_i$ to each state and a symmetric-saddle-point free energy $F_{ij} = F_{ji} = E_{ij} - \beta^{-1} \log \omega_{ij}$, then build barriers and prefactors through $\Delta E_{ij} = E_{ij} - E_i$ and $v_{ij} = \omega_{ij}/\omega_i$. The energies were drawn from a uniform distribution and prefactors from a log uniform distribution between 0.01 THz and 100 THz.

When using the kMC back end, we have access to the exact remaining rate at any point in the simulation, which can be




FIG. 3. Comparison of TAMMBER and typical TAD sampling a single state at a target temperature of 300 K, using kMC to generate escape times, with an estimated computational cost in units of ps of MD. The optimal TAD scheme implemented in TAMMBER is able to find the optimum instantaneous TAD temperature to reduce the unknown rate for minimal computational cost, and thus is able to autonomously outperform constant-temperature TAD. Left inset: Optimal temperature during simulation. Right inset: The objective function (29) at the end of the simulation.

compared to our estimates $\langle k_i^{\text{un}} \rangle$. Figure 3 demonstrates the estimate of the unknown rate for a single state against the simulated computational cost (performing MD, identifying states, and NEB calculations) at a range of fixed TAD temperatures β_H^{-1} , and the TAMMBER process, which uses a variable TAD temperature determined by maximizing the benefit function $-\delta \langle k_i^{\text{un}} \rangle / \delta c_i$, Eq. (29). It can be seen that TAMMBER successfully adjusts the TAD temperature to decrease the unknown rate as fast as possible with computational effort, while the estimate $\langle k_i^{\text{un}} \rangle$ decreases with increasing sampling time. Importantly, the estimated unknown rate is greater than the actual remaining rate, meaning that we can have high confidence that the predicted residence times are conservative. This behavior emerges naturally from our Bayesian estimator; given only the knowledge that rare events are Poisson random variables (through the likelihood function) our estimate for the remaining rate cannot be significantly lower than the inverse time spent in the state; i.e., one cannot exclude the possibility of a given k_i^{un} remaining without running dynamics for a time of order $1/k_i^{\text{un}}$. While it is in principle possible to improve the estimator by encoding knowledge of the rate distribution into a Bayesian prior, such information is typically not available in atomistic simulation, so the estimator (17) is a good choice.

We have also used the kMC back end to test self-optimizing capability of TAMMBER beyond a single state (see Fig. 4). Two rate matrices were generated, each with 100 states and on average 40 connections per state, but with a different distribution of energy barriers. We chose a high connectivity to ensure each state has a similar spectrum of escape rates, while as before the target temperature was 300 K and TAD temperatures between 300 K and 1500 K were considered. To investigate the response of our control protocol 29, the first rate matrix (system 1) had barriers drawn between 0.25 eV and 1 eV, while the second rate matrix (system 2) had barriers drawn between 0.5 eV and 1.25 eV, suggesting a higher optimal




FIG. 4. Self-optimization of the TAMMBER code for two test systems. Left: Histogram of energy barriers. Right: Mean and standard deviation of the benefit function across the range of temperatures. System 2 has a systematically larger barrier spectrum than system 1, leading to an increase in the optimal TAD temperature.

temperature. As can be seen in 4, TAMMBER is able to self-optimize for these two systems; the mean optimal TAD temperature for system 1 is around 600 K, while for system 2 this rises to 1200 K.

For the example cases considered here, where each state has a similar spectrum of escape rates, the spread of optimal temperatures across the states is relatively narrow, but in a general case this can vary significantly as a function of rate spectrum and time spent in the state. In general, the optimal temperature will start at the lowest value, quickly rise as state time is accumulated before the first event is observed, then fall to a degree dependent on the discovered transition rates.

B. Interstitial capture by C15 clusters in iron

As a preliminary application of TAMMBER, we have investigated the capture of mono-interstitial dumbbell defects [35] by C15 tetra-interstitial clusters [36] using an embedded atom potential model of iron [37]. C15 clusters have been observed in irradiation damage simulations [38] and are known to be the most stable interstitial arrangement for small defect sizes [39], but their connection to the wider energy landscape of an irradiated material is still largely unexplored. In particular, C15 defects have been observed to act as sinks for mono-interstitials, resulting in C15 growth which is assumed to play an important role in the evolution of the defect population [40]. However, due to the vast energy landscape of a defective material quantitative statements on the nature of this capture process, beyond observation of individual trajectories, is very challenging to calculate by traditional methods.

In our simulations, a C15 structure [36] was formed from 4 interstitial atoms in a $10 \times 10 \times 10$ cubic supercell before adding a further interstitial atom nearby, forming a dumbbell under further relaxation. Minimizing the hydrostatic pressure changed the final energy by less than 0.01 eV, consistent with the known small formation volume of these defects [36]. The final system, illustrated in Fig. 5(a), contained 2005 atoms. TAMMBER performed constant-volume TAD MD simulations using an underdamped Langevin thermostat [32], with a target temperature of 300 K and possible TAD

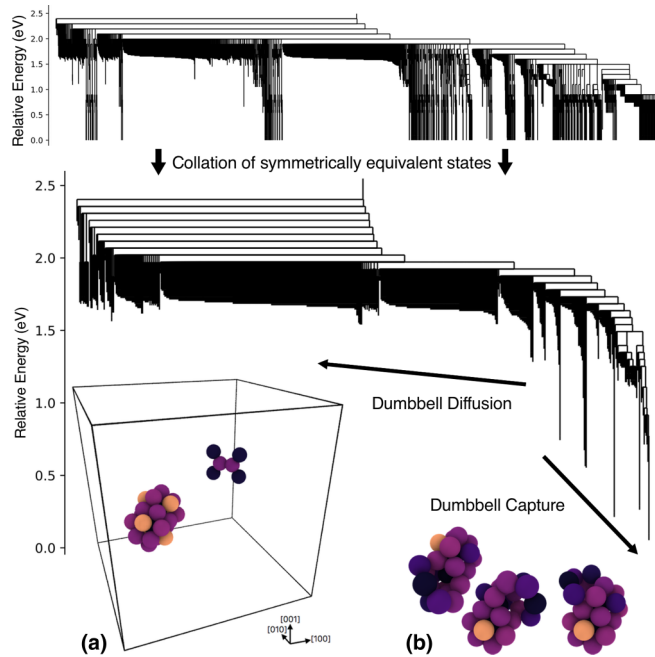


FIG. 5. Disconnectivity graph [41] for states found by TAMMBER with the tetra-C15 and dumbbell system studied shown in inset (a). As discussed in the main text, grouping symmetrically equivalent states leads to a significant reduction in the number of states and simplifies the graph structure. While all the dumbbell capture states (inset (b)) reside in the superbasin, distinct dumbbell states also exist inside this superbasin at relatively low energies, meaning this illustration is not a perfect representation of the coarse-grained landscape. Nevertheless, the Markov chain analysis shows the system remains in a penta-C15 state for multiple seconds at 300 K.

temperatures between 400 K and 900 K. Resource allocation was determined using the scheme detailed above, with the initial distribution being a delta function $[\mathbf{P}_{\mathcal{K}}(0)]_i = \delta_{ij}$ on the starting state of a separated dumbbell and C15 tetra-interstitial. The upper temperature threshold is limited by the presence of significant anharmonic effects on the transition rate which violate the harmonic approximation used in TAD; efficient anharmonic rate theory implementations [23] would therefore be extremely beneficial to further extend the range of TAD temperatures that can be used. As anharmonic vibrational effects typically act to increase transition rates, it can be shown that the inclusion of anharmonic effects would act to increase the expected residence time of the observed network and thus our present results can be considered a lower bound.

After 12 hours of operation on 2160 processors, TAMMBER had identified 2664 metastable states with 7676 connecting barriers from around $2 \mu\text{s}$ of high-temperature MD. The expected residence time conditional on $\mathbf{P}_{\mathcal{K}}(0)$ in the set of known states was found to be 43.4 seconds at 300 K, a testament to the timescales that can be accessed by the massively parallel temperature-accelerated dynamics controlled by TAMMBER.

The energy landscape, illustrated through a disconnectivity graph [41] in Fig. 5, consists of a large number of states corresponding to dumbbell diffusion [shown in Fig. 5(a)], with a smaller number of low-energy states corresponding to dumbbell capture [Fig. 5(b)]. As discussed in the next section,

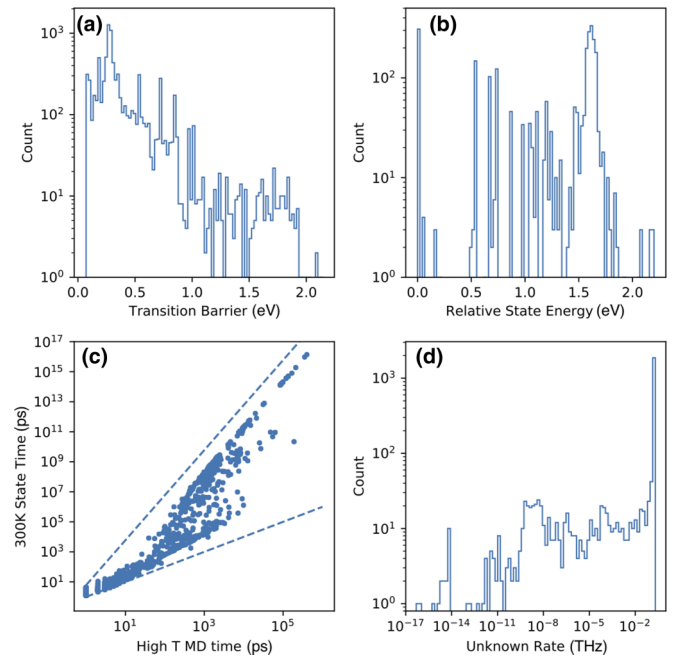


FIG. 6. Summary of the TAMMBER simulation for the C15-dumbbell system, discussed in the main text. (a) Histogram of energy barriers. (b) Histogram of state energies. (c) Effective low-temperature time versus high-temperature MD time. (d) Histogram of unknown rates.

while the clear superbasin shown in Fig. 5 contains all the dumbbell capture states, distinct dumbbell diffusion states also exist at relatively low energies, meaning the superbasin structure is an illustrative but imperfect representation of the coarse-grained landscape. Due to the high stability of the C15 tetra-interstitial, the number of states is expected to scale only linearly with the size of the system, resulting in the relatively low number of states found in this example.

Figure 6 gives a more detailed presentation of the final state of the TAMMBER simulation. Figure 6(a) shows the wide distribution of energy barriers found, demonstrating the need for an adaptive parametrization to optimally sample the highly heterogeneous energy landscape. The peak around 0.26 eV corresponds to dumbbell migration, with higher energy barriers typically corresponding to escape pathways from the superbasin of captured states. Figure 6(b) shows the relative state energies, with a peak at the minimum energy for the superbasin of capture states, while the large higher-energy peak is for dumbbell migration states. Figure 6(c) shows the effective low-temperature TAD time versus high-temperature MD time for each state. The slope on the double-log plot is equal to the temperature ratio β_L/β_H , as can be seen from Eqs. (7) and (8). The deepest states have the highest optimum TAD temperature, and can clearly be seen as the upper envelope to the scatter low-temperature TAD times, while the lowest envelope is simply an equality ($\beta_L = \beta_H$). The scatter in slope is a demonstration of the range of optimal temperatures throughout the run; deep capture states were typically sampled at 900 K, while the dumbbell diffusion states were typically sampled at around 550 K. Finally, Fig. 6(d) shows a histogram of the low-temperature unknown rates $k_i^{\text{un}}(\beta_L)$. States which are not

deemed influential to the overall behavior by the Markov chain analysis receive little to no sampling and thus possess a high unknown rate, leading to the significant upper peak.

At the end of the initial TAMMBER simulation, it was observed that the resultant Markov chain predicts a long residence time in the superbasin of low-lying ‘‘dumbbell capture’’ states [shown in Fig. 5(b)]. To further explore this superbasin, TAMMBER was restarted using the previously generated trajectory and transition barrier information and ran for a further 4 hours on 2160 cores with a new initial distribution, namely a delta function on the best sampled dumbbell capture state, giving an expected residence time of $\langle \tau_{\text{res}} \rangle = 57.6$ s at 300 K, with 21 states having an expected visit time of more than 0.1 seconds, from a total effective low-temperature time of $\sum_i \tau_i(\beta_L) = 2.98 \times 10^4$ s. This scale separation between the total low-temperature time and the residence time is a consequence of the structure of the energy landscape; as superbasin states are frequently revisited, the unknown rate must be significantly lower than the total known escape rate to ensure long trajectories before absorption. AMD techniques are thus essential to provide efficient sampling of the energy landscape, as otherwise the raw sampling MD time greatly exceeds the typical residence time of the found transition network [17].

Upon a detailed investigation of the observed system configurations, it was found that a significant number of states were identical to each other up to a reindexing of atoms or an operation of the crystal’s symmetry group. Exploitation of these symmetries is clearly highly desirable, as well as the high-temperature MD trajectories and found escape times across all identical states to be collated, resulting in more efficient sampling, smaller unknown rates, and a more compact description of the transition network. As the effective state time is increased to the power of the temperature ratio used in TAD, consolidation of MD sampling can produce very large decreases in the unknown rates. Identification of symmetrically equivalent states is possible using graph isomorphism algorithms [42] on the connectivity graphs used to identify states in TAMMBER. Using the graph isomorphism algorithm to construct a map to the reduced set of symmetrically inequivalent states, we reprocessed the TAMMBER simulation output to construct new effective state times, transition rates, and unknown rates to build a new Markov chain in the symmetrically reduced state space. We find a new transition network of 626 states, illustrated in Fig. 6. The new residence time with a delta function on the same lowest energy superbasin state is now $\sum_i \tau_i(\beta_L) = 7.38 \times 10^6$ s with a residence time of 80.9 seconds. This very large difference between the total state and validity is due to the high degree of degeneracy (318 states) of the lowest-lying dumbbell capture state, resulting in an excessively long effective state time of 4.2×10^6 s, which would not be allocated in a symmetry-aware resource management scheme. The development of such a scheme in TAMMBER is clearly highly desirable but raises a number of subtle issues which are beyond the scope of the present paper. In our final section, we use the symmetrically reduced Markov Chain developed above to investigate superbasin escape times and explore the consequences of possessing, through the unknown rates, uncertainty quantification on the completeness of the discovered network.

VII. DISCUSSION: UNCERTAINTY QUANTIFICATION OF TRANSITION NETWORK OBSERVABLES

The central goal of the present paper was to construct, with rigorous uncertainty quantification, a transition network from atomistic simulations with a maximally long residence time in the found state space. The previous section demonstrated that extremely long residence times are readily accessible using our method. In this final section we provide a preliminary exploitation of the discovered transition network, in particular accounting for of uncertainty quantification provided by the unknown absorbing rates. A full exploration of these ideas, and a detailed examination of their use when transitioning to a higher-scale simulation scheme such as kinetic Monte Carlo, will be the subject of future work. A natural observable to extract from the transition network is the expected escape time from the dumbbell capture superbasin. This is clearly an important input for coarse-grained models of interstitial cluster evolution, informing the degree to which C15 clusters can be considered as pure sinks for mono-interstitial defects, which can otherwise collate into highly mobile prismatic dislocation loops. To calculate a superbasin escape time, we ask for the first escape time from a collection of states \mathcal{A} , here the lowest energy dumbbell capture states [Fig. 6(b)]. This can simply be achieved by artificially making all the remaining states $\mathcal{B} = \mathcal{K} \setminus \mathcal{A}$ an absorbing set. Similar ideas are regularly employed in the biochemical community [7], though the inclusion of an unknown rate to account for sampling incompleteness is novel to the best of our knowledge. Defining the known rate matrix on \mathcal{A} as $\mathbf{Q}_{\mathcal{A}}$ one can then define *two* sets of absorbing rates, namely the previously estimated unknown rates from \mathcal{A} to Δ and the sum of all rates from \mathcal{A} to \mathcal{B} :

$$[\mathbf{k}_{\Delta}^u]_i = k_i^u, \quad [\mathbf{k}_{\mathcal{B}}^u]_i = \sum_{j \in \mathcal{B}} k_{ij}, \quad i \in \mathcal{A}. \quad (38)$$

Restricting the initial distribution of states $\mathbf{P}_{\mathcal{A}}(0)$ to some distribution over \mathcal{A} , one can define a very useful convergence measure for averages over trajectories from \mathcal{A} to \mathcal{B} , namely the probability of absorbing to \mathcal{B} instead of Δ , given by

$$P_{\mathcal{B} < \Delta} = \mathbf{P}_{\mathcal{A}}(0) \cdot \mathbf{Q}_{\mathcal{A}}^{-1} \cdot \mathbf{k}_{\mathcal{B}}^u, \quad \lim_{k_{\Delta}^u \rightarrow 0} P_{\mathcal{B} < \Delta} = 1, \quad (39)$$

where the final limit corresponds to convergence to the complete model. The expected first-passage time from \mathcal{A} to \mathcal{B} , conditional on not absorbing to Δ , reads

$$\langle \tau_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{abs}} \rangle = \mathbf{P}_{\mathcal{A}}(0) \cdot \mathbf{Q}_{\mathcal{A}}^{-2} \cdot \mathbf{k}_{\mathcal{B}}^u / P_{\mathcal{B} < \Delta}. \quad (40)$$

However, when $P_{\mathcal{B} < \Delta}$ is small, absorption to Δ is much more likely and thus the true first-passage time from \mathcal{A} to \mathcal{B} is expected to be much greater than the current residence time, meaning (40) is likely to be a significant underestimate. One possible strategy to investigate the dependence of $\langle \tau_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{abs}} \rangle$ on sampling incompleteness is to ‘‘artificially’’ take the perfect sampling limit

$$\langle \tau_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{abs}} | k_{\Delta}^u = 0 \rangle \equiv \lim_{k_{\Delta}^u \rightarrow 0} \langle \tau_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{abs}} \rangle, \quad (41)$$

corresponding to the prediction of approaches without any uncertainty quantification. Another approach is to recognize that the conditional expectation in (40) is biased by sampling only from the subset of trajectories that absorb to \mathcal{B} before

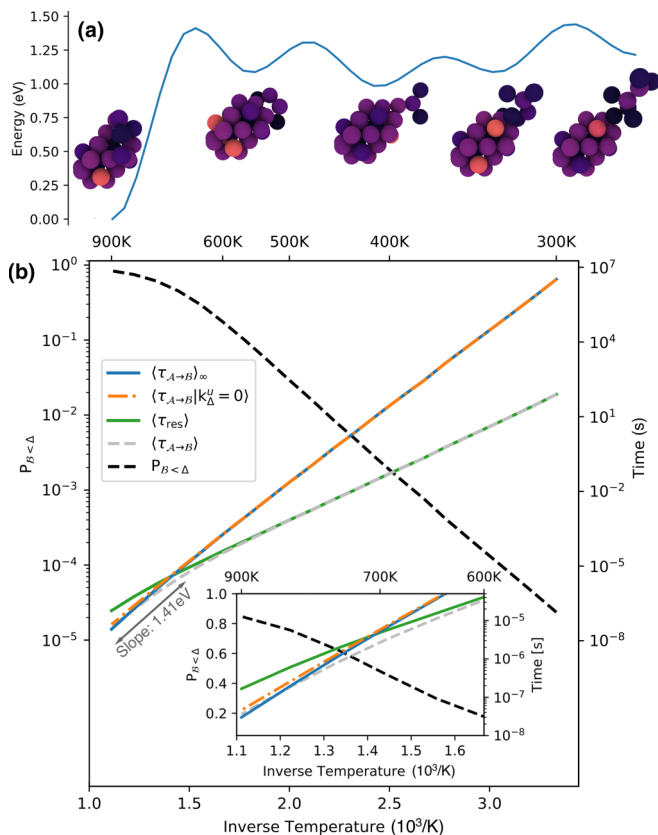


FIG. 7. Analysis escape from the dumbbell capture superbasis tetra-C15 and shown in Fig. 6(b). Lower panel: Calculated residence times and unbiased first-passage times across a range of temperatures. As discussed in the main text, the first-passage-time estimates all converge when the probability of escape before absorption is high. Upper panel: The minimum-energy path for superbasis escape. The highest saddle-point energy of 1.44 eV agrees well with the found Arrhenius slope of 1.41 eV.

Δ . An approximate form for the unbiased first-passage time $\langle \tau_{A \rightarrow B}^{\text{abs}} \rangle_{\infty}$ can be obtained by assuming absorption from \mathcal{A} to \mathcal{B} or Δ are two first-order Poisson processes with mean times $\langle \tau_{A \rightarrow B}^{\text{abs}} \rangle_{\infty}$ and $\langle \tau_{\text{res}} \rangle$; it is simple to show that this gives the approximate expression for the unbiased first-passage time of

$$\langle \tau_{A \rightarrow B}^{\text{abs}} \rangle_{\infty} = \left(\frac{1}{P_{B < \Delta}} - 1 \right) \langle \tau_{\text{res}} \rangle. \quad (42)$$

Coming back to the case of the C15 defects, while the disconnectivity graph has a clear superbasis structure, a detailed inspection shows that a number of lower-energy states have a distinct, separate dumbbell structure, meaning that the “true” capture superbasis has a more complex structure than that implied by the illustration in Fig. 5. The set \mathcal{A} of capture states was thus chosen to be the minimal set of connecting states to the found global minimum where a distinct dumbbell structure could not be found, consisting of 63 symmetrically inequivalent states, or around 500 states of the original network.

In Fig. 7(b), we plot $P_{B < \Delta}$ along with the residence time $\langle \tau_{\text{res}} \rangle$, the conditional first-passage time $\langle \tau_{A \rightarrow B}^{\text{abs}} \rangle$, and corrected first-passage times $\langle \tau_{A \rightarrow B}^{\text{abs}} | k_{\Delta}^u = 0 \rangle$ and $\langle \tau_{A \rightarrow B}^{\text{abs}} \rangle_{\infty}$

across temperatures from the low target temperature 300 K to 900 K, the highest temperature considered in our simulations. It can be seen that at low temperatures $P_{B < \Delta}$ is very small, leading the conditional first-passage time to converge to the network residence time, indicating that the current quality of the network is insufficient to “certify” that the predicted times are correct, even if the corrected times are essentially in perfect agreement with each other. In other words, the model cannot be used to exclude the possibility that other, yet undiscovered mechanisms could affect the predicted times at low temperatures. In contrast, at higher temperatures $P_{B < \Delta} \rightarrow 1$, the residence time exceeds the first-passage time, with all estimates for the first-passage time converging. From the Arrhenius gradient an effective energy barrier for the superbasis escape of 1.41 eV is found, which closely corresponds to the escape process illustrated in Fig. 7(a).

While the network produces a reliable superbasis escape time at high temperature, the large differences between the residence time and the corrected first-passage times at low temperature, or equivalently the small values of $P_{B < \Delta}$, demonstrate that care must be taken when constructing transition networks from atomistic simulations. The objective function (29) used in the current work was focused on optimizing a particular measure of transition network quality, the expected residence time. In future work, we will further develop the approach presented here to specifically address the issue of converging more targeted quantities such as superbasis escape times.

VIII. CONCLUSIONS

In this paper we have introduced a method to generate large networks of transition rates from atomistic simulations, sampling the energy landscape with a form of self-optimizing temperature-accelerated dynamics. Bayesian estimators were developed that quantify sampling incompleteness in the form of an absorbing unknown rate for each system state. Due to sampling incompleteness, trajectories in the observed rate network have a finite lifetime before absorption. The presented method, TAMMBER, determines the optimal allocation of computational resources “on the fly” in order to find new states and transition pathways, with the goal of maximize the expected time in the known transition network before absorption, conditional on a user-specified initial condition. After validation on exactly known transition networks TAMMBER was applied to the capture of interstitials by C15 clusters in an EAM model of iron, reaching expected absorption times of more than 80 s at 300 K. It was found that sampling completeness could be considerably improved by consolidating symmetrically equivalent states; incorporation of symmetry considerations into the TAMMBER allocation scheme is an immediate topic for future work.

The transition network was then used to explore superbasis escape times, with expressions derived for the average escape rate in terms of the absorbing rates. The uncertainty quantification indicated that while converged results can be produced when the predicted escape time is less than the network residence time, building statistical confidence on long-time, low-temperature behavior proves extremely challenging, as results can be strongly affected by the degree of sampling completeness, an observation which is likely to be widely

applicable across many coarse-grained modeling approaches. The further development of optimal strategies such as this one to reduce the often surprisingly large uncertainty sensitivity is therefore urgently needed.

ACKNOWLEDGMENTS

This material is based on work supported by the US Department of Energy, Office of Nuclear Energy, and Office of Science, Office of Advanced Scientific Computing Research, through the Scientific Discovery through Advanced Computing (SciDAC) project on Fission Gas Behavior, and used computing resources provided by the Los Alamos National Laboratory Institutional Computing Program. Los Alamos National Laboratory is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the US DOE under Contract No. DE-AC52-06NA25396.

APPENDIX A: MOMENTS OF THE POSTERIOR DISTRIBUTION

For N seen transitions, we define $a_j \equiv k_i^{\text{obs}} - k_{i,j}^{\text{obs}}$, $j \in [0, N_i - 1]$. As $k_{i,0}^{\text{obs}} = 0$ (no observed rate before the first event), Eq. (17) for the posterior distribution can then be written

$$\pi(k^{\text{un}}) = e^{-k^{\text{un}}\tau_i} \prod_{j=1}^{N-1} (k^{\text{un}} + a_j) \quad (\text{A1})$$

$$= e^{-k^{\text{un}}\tau_i} \sum_{r=0}^{N-2} (k^{\text{un}})^r A_r, \quad (\text{A2})$$

where A_r is the sum of all $N-2$ possible combinations of r elements from $\{a_j\}_1^{N-1}$. By considering the change in A_r when expanding the number of terms in the product, the A_r can be evaluated by a simple recursion. Using the integral relation $\int_0^\infty k^n e^{-kt} dk = n!t^{-(n+1)}$ we can thus write

$$\langle (k_i^{\text{un}})^n \rangle = \frac{\sum_{r=0}^{N-2} (r+n)! A_r \tau_i^{-r}}{\tau_i^n \sum_{r=0}^{N-2} r! A_r \tau_i^{-r}}. \quad (\text{A3})$$

APPENDIX B: DERIVATIVE OF AN INVERSE MATRIX ELEMENT

Consider the known derivative $\partial_l A_{ij}$ of an element of a matrix \mathbf{A} . To calculate the derivative of the inverse matrix element $\partial_l (A^{-1})_{ij}$, we apply the chain rule to the trivial result $\partial_l (\mathbf{A} \cdot \mathbf{A}^{-1}) \equiv 0$ then premultiply by \mathbf{A}^{-1} to obtain

$$\partial_l \mathbf{A}^{-1} = -\mathbf{A}^{-1} \cdot \partial_l \mathbf{A} \cdot \mathbf{A}^{-1}. \quad (\text{B1})$$

Note the nontrivial ordering of the matrix product. From the structure of the known rate matrix $\mathbf{Q}_{\mathcal{K}}$ we have

$$\frac{\partial}{\partial k_l^{\text{un}}} [\mathbf{Q}_{\mathcal{K}}]_{ij} = -\delta_{ij} \delta_{il}. \quad (\text{B2})$$

This gives the inverse derivative as

$$\frac{\partial}{\partial k_l^{\text{un}}} [\mathbf{Q}_{\mathcal{K}}^{-1}]_{ij} = [\mathbf{Q}_{\mathcal{K}}^{-1}]_{il} [\mathbf{Q}_{\mathcal{K}}^{-1}]_{lj}. \quad (\text{B3})$$

-
- [1] A. F. Voter, *Phys. Rev. Lett.* **78**, 3908 (1997).
 [2] A. F. Voter, *Phys. Rev. B* **57**, R13985(R) (1998).
 [3] M. So and A. Voter, *J. Chem. Phys.* **112**, 9599 (2000).
 [4] D. Perez, B. P. Uberuaga, Y. Shim, J. G. Amar, and A. F. Voter, *Annu. Rep. Comput. Chem.* **5**, 79 (2009).
 [5] G. Henkelman, *Annu. Rev. Mater. Res.* **47**, 199 (2017).
 [6] L. K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, and N. Mousseau, *Phys. Rev. E* **84**, 046704 (2011).
 [7] D. J. Wales, *Mol. Phys.* **100**, 3285 (2002).
 [8] V. S. Pande, K. Beauchamp, and G. R. Bowman, *Methods* **52**, 99 (2010).
 [9] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder *et al.*, *APL Mater.* **1**, 011002 (2013).
 [10] S. T. Chill, J. Stevenson, V. Ruehle, C. Shang, P. Xiao, J. D. Farrell, D. J. Wales, and G. Henkelman, *J. Chem. Theory Comput.* **10**, 5476 (2014).
 [11] L. M. Ghiringhelli, J. Vybiral, S. V. Levchenko, C. Draxl, and M. Scheffler, *Phys. Rev. Lett.* **114**, 105503 (2015).
 [12] R. J. Zamora, B. P. Uberuaga, D. Perez, and A. F. Voter, *Annu. Rev. Chem. Biomol. Eng.* **7**, 87 (2016).
 [13] M. A. Novotny, *Phys. Rev. Lett.* **74**, 1 (1995).
 [14] G. C. Boulougouris and D. Frenkel, *J. Chem. Theory Comput.* **1**, 389 (2005).
 [15] G. C. Boulougouris and D. N. Theodorou, *J. Chem. Phys.* **127**, 084903 (2007).
 [16] S. T. Chill and G. Henkelman, *J. Chem. Phys.* **140**, 214110 (2014).
 [17] A. Bhoutekar, S. Ghosh, S. Bhattacharya, and A. Chatterjee, *J. Chem. Phys.* **147**, 152702 (2017).
 [18] A. Chatterjee and S. Bhattacharya, *J. Chem. Phys.* **143**, 114109 (2015).
 [19] We note that in general all possible connections in \mathcal{K} will not have been observed, i.e., $\mathcal{K}_i \neq \mathcal{S}_i \cap \mathcal{K}$.
 [20] S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence* (Wiley, 2009).
 [21] G. Di Gesù, T. Lelièvre, D. Le Peutrec, and B. Nectoux, [arXiv:1706.08728](https://arxiv.org/abs/1706.08728).
 [22] P. Hänggi, P. Talkner, and M. Borkovec, *Rev. Mod. Phys.* **62**, 251 (1990).
 [23] T. D. Swinburne and M.-C. Marinica, *Phys. Rev. Lett.* **120**, 135503 (2018).
 [24] G. Henkelman, B. P. Uberuaga, and H. Jonsson, *J. Chem. Phys.* **113**, 9901 (2000).
 [25] C. Huang, A. F. Voter, and D. Perez, *Phys. Rev. B* **87**, 214106 (2013).
 [26] The generalization to multitemperature data is straightforward and allows an additional estimate ΔE_{ij} .
 [27] I. J. Myung, *J. Math. Psychol.* **47**, 90 (2003).
 [28] We note that this minimum variance estimator will systematically overestimate the rate, as $P(\frac{N}{\tau} \exp(\beta \Delta E) > \nu) = \Gamma(N, N) / \Gamma(N) > 1/2$.

- [29] H. Jeffreys, *Proc. R. Soc. London A* **186**, 453 (1946).
- [30] Y. Shim and J. G. Amar, *J. Chem. Phys.* **134**, 054127 (2011).
- [31] D. Perez, E. D. Cubuk, A. Waterland, E. Kaxiras, and A. F. Voter, *J. Chem. Theory Comput.* **12**, 18 (2015).
- [32] S. Plimpton, *J. Comput. Phys.* **117**, 1 (1995).
- [33] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, *Phys. Rev. Lett.* **97**, 170201 (2006).
- [34] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, *J. Comput. Phys.* **17**, 10 (1975).
- [35] C.-C. Fu, J. Dalla Torre, F. Willaime, J.-L. Bocquet, and A. Barbu, *Nat. Mater.* **4**, 68 (2005).
- [36] M.-C. Marinica, F. Willaime, and J.-P. Crocombette, *Phys. Rev. Lett.* **108**, 025501 (2012).
- [37] L. Malerba, M.-C. Marinica, N. Anento, C. Björkas, H. Nguyen, C. Domain, F. Djurabekova, P. Olsson, K. Nordlund, A. Serra *et al.*, *J. Nucl. Mater.* **406**, 19 (2010).
- [38] E. Zarkadoula, S. Daraszewicz, D. Duffy, M. Seaton, I. Todorov, K. Nordlund, M. Dove, and K. Trachenko, *J. Phys.: Condens. Matter* **25**, 125402 (2013).
- [39] L. Dézerald, M.-C. Marinica, L. Ventelon, D. Rodney, and F. Willaime, *J. Nucl. Mater.* **449**, 219 (2014).
- [40] Y. Zhang, X.-M. Bai, M. R. Tonks, and S. B. Biner, *Scr. Mater.* **98**, 5 (2015).
- [41] D. J. Wales, *Energy Landscapes* (Cambridge University Press, 2003).
- [42] B. D. McKay and A. Piperno, *J. Symb. Comput.* **60**, 94 (2014).