Title: Hybrid Kinetic Monte Carlo Algorithm for Strongly Trapping Alloy Systems

Authors: Craig Daniels, Pascal Bellon

Affiliation: Department of Materials Science and Engineering, University of Illinois, Urbana-Champaign, Illinois 61801, USA

Abstract:

Selected solute atoms can strongly interact with and slow down the diffusion of point defects in alloy systems. While such additions can be beneficial, for instance to promote microstructural stability during thermal annealing or during irradiation by energetic particles, they create significant computational challenges when simulating these evolutions using atomistic techniques such as kinetic Monte Carlo (KMC) simulations. Point defect trapping in energy basins created by clusters of solute atoms leads to frequent re-visiting of states with short residence times, which dramatically reduces the efficiency of traditional KMC algorithms. We introduce here a hybrid algorithm that combines and expand on two prior KMC algorithms, the Chain KMC and the equilibrating basin KMC. This hybrid algorithm, referred to as the Equilibrating Chain algorithm, utilizes Chain KMC as previously reported, but leverages the data-handling framework to build an occupation distribution of the basin, allowing the equilibrating basin assumption to be statistically tested and applied. For a model A-B trapping alloy system on a face-centered cubic lattice, statistical comparisons of basin exit and cluster dissolution kinetics between traditional and accelerated KMC algorithms are presented to demonstrate the accuracy and the efficiency of the new algorithm. We also discuss our algorithm in the context of other accelerated KMC algorithms.

Word Count: 205.

Keywords: accelerated KMC; trapping; equilibration

1. Introduction

Atomistic simulations have become ubiquitous techniques for assessing the energetics and kinetics of alloys, leading to critical benefits for alloy design and for supporting the analysis of experimental results in complex alloy systems [1-5]. The temporal evolution of atomic configuration, which is the focus of the present study, is of particular importance for alloy systems that can undergo microstructure evolutions and phase transformations, for instance during isothermal annealing. While first principle methods and molecular dynamics simulations can provide very detailed and accurate information, they cannot reach timescale far in excess of that of atomic vibrations. For longer timescale, kinetic Monte Carlo (KMC) methods are widely used, provided that jump frequencies of all possible events are known or accessible. For thermally activated processes, for instance, these frequencies can be determined using transition state theory [6].

Algorithms used for equilibrium MC simulations, e.g., the Metropolis algorithm [7], are not well suited for KMC simulations owing to their high rejection rates. Bortz, Khalos and Lebowitz [8]

introduced the so-called n-fold way algorithm, also known as the BKL algorithm or as the residence time algorithm (RTA). Two important advantages of this algorithm are that each KMC step results in a transition from one state to another, and that it provides the physical time taken by the system to perform this transition. This algorithm is particularly effective for defect-driven evolutions, because the set of possible transitions from one configuration to the next is small. The efficiency of the BKL algorithm, however, can be strongly affected by the specifics of point-defect-atom interactions. In particular solute additions are sometimes used to slow down diffusion, and thus increase the stability of microstructures subjected to elevated temperature environments or to irradiation by energetic particles [9-13]. Strong vacancy-solute interactions, for instance, can lower the migration barriers for vacancy motion near solute atoms relative to the migration barriers for vacancies to move away from solute atoms, resulting in vacancy trapping. Pairs or clusters of vacancy-solute complexes may still migrate, and contribute significantly to the evolution of the system, but over a longer time scale. One issue observed in systems with bound vacancy-solute pairs is that of 'flickering', or immediate reversal to a previously explored state. The problem is exacerbated in the presence of clusters of solute, which can provide additional pathways for the vacancy to return to previous configurations. These repeatedly explored associated sites can be described as constituting energy basins where the system is trapped [14, 15], resulting in dramatic reduction in the BKL algorithm efficiency.

Four main strategies have been developed to increase the efficiency in exploring the configuration space for alloys with traps. The first one relies on considering transitions that include multiple defect jumps so that rates of jump sequences, rather than single jumps, can be calculated, such as in the 2nd Order Residence Time Algorithm [16, 17]. Similarly, Chain KMC [18], the Mean Rate Method [14, 19], and Kinetic Path Sampling (KPS) [20] accumulate state information during the simulation and apply it to higher order calculations . First Passage Time Analysis (FPTA) [14, 21, 22], which solves the master equation for a small set of states rather than tracking a specific path, has also been applied. These methods add complexity and cost to the KMC calculations, but are analytically equivalent to traditional 1st Order RTA. The second strategy is to modify the basin, for instance by increasing energy barriers within the basin [23]. In this last case, the algorithm deviates from the 1st Order RTA, but the introduced error can be estimated and controlled. The third strategy is to apply additional processing power to these traps to speed up the overall simulation, such as in Parallel Replica Dynamics [24, 25]. The fourth strategy is to apply a simplifying assumption, such as in the Quasi-Stationary [26] or Equilibrating Basin approaches [14, 15] both of which assume that the of long-term behavior of their respective forms will dominate exit modes. These methods can allow a vacancy to leave a trapping basin in a few KMC steps, but the validity of the assumptions vary from case to case, and conservative application of the assumptions will severely limit efficiency gains. In addition, the results are dependent on separate algorithms which define the energy basins, and therefore determine how the assumptions are applied.

Here we introduce a hybrid algorithm that combines the Chain KMC method with the equilibrating basin method. We refer to this novel algorithm as the Equilibrating Chain algorithm. We show that on-the-fly identification of basins and statistical testing of the equilibration assumption make it possible to both quickly explore energy basins and accurately identify when it is acceptable to exit the basin. We show that the Equilibrating Chain algorithm is particularly effective for overcoming point defect trapping on small solute clusters, thus making it an attractive

tool for studying kinetic evolutions in semi-dilute alloy systems undergoing clustering and precipitation during thermal annealing or during irradiation with energetic particles. In Section 2 we introduce the alloy system selected for this study, and briefly recall other standard KMC algorithms that will be used as reference points. The Equilibrating Chain algorithm is then introduced in Section 3, and results are presented and discussed in Section 4.

2. Method

2.1. Model Alloy System

We consider here the case of dilute A-B binary alloys, where B solute atoms provide effective trapping of vacancies. For simplicity, we restrict our work to alloys on a face centered cubic (FCC) lattice with atom interactions limited to first nearest neighbors (1st NN). This choice does not reduce the generality of the algorithm introduced in this work, see the discussion section, but greatly simplifies its implementation and its verification. Indeed, in the infinite dilution limit, there are only 5 independent jump frequencies, thus the so-called five-frequency model [27, 28]. These five frequencies, traditionally referred to ω_i (i = 0 to 4), correspond to a vacancy jump in the A matrix when the vacancy is not 1st NN of any B atom before or after the jump for i=0, a vacancy exchange with A atoms when the vacancy is 1st NN of one given B atom before and after the jump for i=1, a direct exchange between a vacancy and an isolated B atom for i=2, and vacancy exchange with A atoms that either destroy or create a solute- vacancy pair, respectively for i=3 and i=4. An additional benefit of this simplified model is that extensive data is available for metallic alloy systems. In particular D. Morgan and collaborators have determined from first principles calculations the corresponding activation energies (i = 1 to 5) for over 40 solutes in each of 5 different FCC matrices (Al, Cu, Ni, Pd, and Pt) [29] . Note that the mapping of these data into the five-frequency model assumes that vacancies do not interact with solute atoms beyond 1st NN distances. More detailed studies including longer interaction range are available for a few alloy systems, e.g. refs. [30-32], resulting in an increase in the number of independent jump frequencies that need to be considered. In infinitely dilute alloys, an additional benefit of a small number of independent frequencies is that it is possible to determine defect and solute microscopic transport coefficients directly from analytical methods [32-34]. Our work focuses on solutes that are strong trap for vacancies. A necessary condition is thus that the solute-vacancy binding energy exceeds $\approx 3k_BT$, where T is the temperature of interest. In a Cu-matrix, for instance, a solute such as Sb has a binding energy of about 0.33 eV, resulting in effective vacancy trapping up to ≈ 1000 K [35]. In dilute alloys, the effectiveness of the trapping can be directly assessed by calculating and comparing the transport coefficient of the bound vacancy to that of the free vacancy, see for instance Fig. 2 in ref. [32] for Sb in Cu.

In the non-infinitely-dilute regime, which is the focus of the present work, the conditions required for solute trapping are more complex since trapping of vacancies may occur at clusters of various size and shape. Consequently, the number of independent jump frequencies to consider is much larger, making it difficult if not impossible to study this trapping by extending current analytical methods. In this work, we will instead use KMC simulations to study vacancy trapping by solute clusters, although trapping, by nature, creates challenges for reaching long

physical time scales, as will be detailed in Sections 2.2 to 2.4. KMC algorithms require the determination of all independent jump frequencies, or, using the transition rate theory [6], the determination of all independent activation energies. We assume here for simplicity that the changes in the so-called attempt frequencies are small and thus can be ignored, as reported for dilute systems in ref. [36] . Since our objective is to compare the effectiveness and the consistency of several KMC algorithms, we use a broken-bond model [37, 38] to estimate these activation energies. This approach may break down when solute-defect and solute-solute interactions are large compared to solvent-solvent interactions. In order to avoid these effects, in this work we focus on the Cu-Ag alloy system, as the Ag-V binding energy is only 0.11 eV, and consider temperatures in the 175 K to 325 K range. The corresponding jump activation energies are given in Table 1, and are drawn from ref. [29]. Note that in addition to the 5 activation energies already mentioned, additional energies are required for parametrizing the broken bond model, namely the cohesive energy of the pure A and B elements, and the vacancy formation energy in the A matrix [39, 40]. Lastly, alloy systems such as Cu-Ag display a tendency toward phase separation. This tendency is quantified by the critical temperature of the corresponding miscibility gap, here assumed to be 600 K for simplicity. The actual estimation for in Cu-Ag is higher, ≈ 1200 K, but this would considerably reduce the B solubility at the temperatures considered here, and thus limit the interest of the present study as will be discussed later. These 9 input data, see Table 1, make it possible to calculate the 9 parameters of the kinetic model used in this study, consisting of 5 pairwise interaction energies, and 4 saddle point energy contributions as detailed in Appendix A. These model parameters are listed in Table 2.

Table 1 Physical parameters used for the fitting of the kinetic model for the Cu-Ag alloy. See Appendix A for definitions.

E_{coh}^{A}(eV)	E_{Vf}^{A}(eV)	E_0(eV)	E_1(eV)	E_2(eV)	E_3(eV)	E_4(eV)	E_{coh}^{B}(eV)	T_c(K)
-3.49	1.28	0.7174	0.9040	0.5785	0.7047	0.5895	-3.174	600K

Table 2 KMC simulation parameters corresponding to Table 1. The notation εAB refers to the pairwise interaction energy between A and B atom neighbors, while S_i refers to the saddle-point of the i-type jump, with index i referring to the 5-frequency notation, see details in Appendix A.

εAA	εAV	εAB	εBB	εBV	S_0	S_1	S_2	$S_{3,4}$
-0.582	-0.184	-0.545	-0.529	-0.263	-7.891	-7.746	-7.702	-7.982

2.2. First Order Residence Time Algorithm

In alloy systems where atomic diffusion is mediated by point defects, the most common KMC approach for generating system trajectories is based on the algorithm originally introduced by Bortz, Khalos, and Lebowitz [8]. This algorithm is often referred to as a residence time algorithm (RTA), as each KMC step consists in determining the residence time of the system in its current configuration and in selecting stochastically one exit path, i.e., one transition to a distinct state [16, 38, 41]. The standard implementation of this residence time algorithm considers transitions

that correspond to a single defect jump, and this implementation is thus referred to a 1st Order RTA [16]. We recall briefly here some key definitions relevant to the present work, considering the case of a system with a single vacancy for simplicity. Readers are referred to refs. [41, 42] for additional details. A vacancy occupying a lattice site I defines a configuration state *i*. The transition state theory [6] is used to calculate the rate of the vacancy jumping from site I to neighboring site J, denoted as $r_{i\to j}$ and defined as $r_{i\to j} = \nu_0 \exp(-E_{i\to j}^{act}/k_B T)$, where $E_{i\to j}^{act}$ is the energy barrier for a transition from configuration *i* to configuration *j*. As noted earlier, without loss of generality, we use here a constant pre-factor ν_0. We set $\nu_0 = 10^{14}\ s^{-1}$, fitted to vacancy diffusion in Cu. A vacancy can jump to several 1st NN sites, twelve here for an FCC lattice, and the probability of the vacancy jumping from site I to a given neighboring site J is $p_{i\to j} = r_{i\to j}/R_i$ with $R_i = \sum_j r_{i\to j}$. The residence time of the system in configuration *i* follows the exponential distribution $\Delta t_i = ln(\zeta)/R_i$, where ζ is a random variable uniformly distributed over (0,1]. Since we are interested in sequences involving a large number of jumps, following the simplification used in refs. [43, 44] we use the ensemble average timestep, which is $\Delta\overline{t_i} = 1/R_i$.

2.3. 2nd Order RTA and Chain KMC

Alloy systems where point defects interact strongly with solute atoms pose a significant challenge for first order RTA. Indeed, strong interactions often result in a high jump rate of a vacancy to a neighboring trapping solute atom, $r_{i\to trap}$. After this first jump, there is a high probability for the next vacancy jump to bring the vacancy back to its previous site, resulting in what is referred to a reversal or a "flicker". In alloys with strong trapping solutes, i.e. $R_i \approx r_{i\to trap}$, the system performs a large number of flicker sequences associated with short residence times, thus dramatically reducing the efficiency of the algorithm, as will be illustrated in Section 3 for the alloy system of interest here.

Athenes et al. proposed a 2nd Order Residence Time Algorithm to eliminate these flickers [16, 17]. The jump rates to vacancy neighbors are calculated as in the 1st Order RTA, but the jump rates of the vacancy from each of those neighboring sites are also calculated. This expanded exploration of the surrounding environment allows for the probability of flickering events to be determined analytically. Modified transition rates that include flickers can then be derived. The next vacancy jump is determined from these modified rates, excluding a direct return to the previous site. The 2nd Order RTA algorithm thus bypasses the flickers while including their contribution to the residence time and generating trajectories that are statistically equivalent to those of the 1st Order RTA. As a systematic extended exploration of transition rates after each vacancy jump is computationally demanding, Mason et al. modified this algorithm to take advantage of the fact that part of these transition rates may have been calculated in previous KMC steps [18]. Specifically, jump rates required for the 1st Order RTA are stored on the fly as the vacancy diffuses, allowing for the immediate calculation of second order jump probabilities to these previously visited sites. These sites form a subset of configurations that are connected by single vacancy jumps, and are referred to as a chain. In this so called Chain KMC, a direct reversal of the last vacancy jump may either be included or excluded from the next possible transitions, depending on whether the information to account for this reversal had been previously calculated

or not. We recall the main steps involved in implementing this algorithm in Appendix B. We keep the notations introduced in ref. [18], and the reader is referred there for more detail.

The Chain algorithm directly ensures that stored jump probabilities between states remain correct when the vacancy is swapping with a single atom type, such as when exploring a solute cluster. More careful bookkeeping is required whenever the vacancy swaps with an atom type different from the previous swaps, such as when a vacancy enters or leaves a solute cluster, and for that reason here we restrict the algorithm to store chain states corresponding to the solute-vacancy exchanges only. Care must also be taken to account for configurational changes that occur due to other events in the simulation, for instance in systems containing multiple migrating defects.

For strongly trapping alloy systems, Chain KMC can dramatically increase the effectiveness of the simulations for a modest increase in computational cost. In our implementation, each Chain KMC step is approximately 20% more computationally expensive than those for the 1st Order KMC. However, each Chain KMC step may effectively correspond to tens or hundreds of 1st Order KMC jumps, depending on the interactions and configuration. For example, with the alloy system selected here at 200 K, a vacancy interacting with a single solute performs an average of 215.6 exchanges with the solute atom before exchanging with a matrix atom. Chain KMC identifies and bypasses these exchanges in a maximum of 3 steps, thus resulting in an average speed-up factor of at least 60.

The Chain KMC method works exceptionally well when the vacancy is interacting with one single solute atom because any flickering events between transient sites, i.e., the vacancy-solute exchanges, are quickly accounted for without having to actually perform these jumps. This leaves only jumps away from or around the solute, corresponding to ω_3-type and ω_1-type jumps, respectively. Chain KMC however fails to efficiently manage larger clusters of solutes. As more solute atoms are added to the cluster, additional pathways connecting solute sites, such as loops, may become available. Loops allow the vacancy to remain trapped within the cluster without performing many, or perhaps any, direct reversals. As the probability of flickering is greatly reduced, the Chain KMC algorithm efficiency is reduced as well. We note that the possible formation of such loops is dependent upon the crystallographic lattice of the alloy as well as the jumps that a defect can perform. For instance, triangular and tetrahedral loops exist for a vacancy performing 1^{st} NN jumps on an FCC lattice, but not on a BCC lattice, unless 1^{st} and 2^{nd} NN jumps are considered. Even if no loops are present, navigating the branches of an extended cluster may be inefficient if the flicker probability is low. This type of trapping can be generally defined as trapping of the defect in a basin of connected low-energy configurations, and is detailed in the next section.

2.4. Equilibrating Basins

For the sake of simplicity, consider again a vacancy interacting with a single solute atom. There are two associated sites composing the energy basin, and these sites are termed ‘transient’ sites (e.g. refs. [14, 15]). As long as the vacancy stays trapped within these sites, the overall basin remains unchanged. When the vacancy jumps to a site distinct from transient sites, termed ‘absorbing’ sites, the vacancy is said to have exited or ‘escaped’ the basin. These definitions are straightforwardly extended to clusters comprised of more than one solute atom.

For deep energy basins, an equilibration assumption has been proposed [14, 15, 45, 46]. Assuming that the vacancy cannot exit the basin (a good approximation of very deep traps), the vacancy will perform a large number of jumps between the transient states and will reach a local equilibrium between those states. If local equilibrium is reached, the probability of occupying any transient state "*i*" is then given by a Boltzmann distribution for this set of sites.

$$P_i(t) = P_i = \frac{e^{\left(\frac{-E_i}{k_B T}\right)}}{\sum_i e^{\left(-\frac{E_i}{k_B T}\right)}} \quad (1)$$

Given this occupation distribution, the basin exit time and exit site can then be determined from the list of absorbing jump frequencies. Specifically, jump frequencies to absorbing states are weighted by the probability of occupying the associated transient state. The sum of these weighted exits rates, Q, becomes the scale parameter of the exponential distribution of basin exit time. As noted in ref. [15], the basin becomes analogous to a single site exited by a single KMC move.

$$Q = \sum_{i,j} P_i * R_{i \to j} \quad (2)$$

$$\Delta t_{exit} = \frac{-ln(\zeta)}{Q} \quad (3)$$

The main challenge of the equilibrating basin approach is to decide for each basin, when the equilibration assumption has become acceptable. For instance, Puchala et al. [14] chose to apply the assumption only when transient jump rates are 10^7 times higher than any absorbing jump rate, which severely limits its use. In ref. [45] the authors do not consider when the assumption might fail, though they limit the application to relatively large clusters in a strongly-trapping system. In another study, the equilibrating basin assumption was shown to be valid for two selected one-dimensional examples, but implementing the assumption more generally was not discussed [15]. If one were to inappropriately accept the equilibration assumption, the trajectories generated will not provide an accurate modeling of the system and, as reported by Van Siclen [46], the defect diffusion coefficient can be overestimated. A separate challenge in implementing an algorithm using the equilibrating basin approach is the need to identify and update these basins.

3. Novel KMC Algorithm

We demonstrate below that by combining Chain KMC with an equilibrating basin approach one can detect these basins on the fly, at very little computational cost. We then demonstrate that by using suitable statistical tests, once a vacancy starts being trapped in a solute cluster, it is possible to define quantitative criteria to decide when the equilibration approach should be used, in contrast to previous implementations. We are able to dynamically test when the running occupation distribution matches the equilibrium distribution, and only apply the assumption if the distributions statistically match. This prevents the introduction of uncontrolled errors in generating the system trajectory.

3.1. Equilibrating Chain KMC

This novel algorithm relies on the Chain KMC algorithm but monitors potential trapping. When trapping occurs the algorithm occasionally tests whether the basin has equilibrated, in which case we exit the basin by applying equations 1-3. . We thus refer to this algorithm as Equilibrating Chain KMC algorithm. It leverages the framework that Chain KMC provides for tracking site-specific information as the simulation progresses. In addition to tracking jump rates and probabilities, we also track how often a site is visited, either explicitly or as part of a skipped flicker event. This is used to build a site occupation distribution. We define N_{chain} as the number of KMC steps made since the stored KMC chain was modified. The chain can be modified either by adding to it by visiting a new site, or deleting states that were changed as the simulation evolved. We recognize that when N_{chain} becomes large (see Appendix C), an energy basin has been found and is defined by the stored states. Because of how we choose to manage stored chain data (see Section 2.3), the basin corresponds to a vacancy trapped by a solute cluster. We choose to test for equilibration when $N_{chain} \geq 100 * N_{sites}$, where N_{sites} is the number of sites stored. The equilibrium occupation distribution is calculated from stored data by considering stored sites as transient, and all non-transient neighbors as absorbing sites. This implementation also simply handles sub-basins within solute clusters, as will be illustrated later. We then compare the expected equilibrium distribution with the observed occupation distribution using a G-test.

The G-test is a log-likelihood ratio statistical test, which is similar to the chi-squared test. The G-test statistic is calculated as $G = 2\sum_i O_i * ln(O_i/E_i)$, where O_i is the observed count, and E_i is the expected count in bin I (i.e. number of visits to the site I). We note that the observed count includes both explicit visits, and visits that were part of a skipped flicker event. The G-test statistic can be mapped to a p-value by the chi-squared distribution. As usually, the p-value refers to the probability of observing a test statistic that is equal or greater than the one calculated, assuming that the null hypothesis is true. A low observed p-value indicates high variance between the two distributions, meaning that the observed distribution is unlikely to have been sampled from the reference distribution. However a high observed p-value cannot be used to indicate a high probability of the observed distribution being sampled from the reference.

In practice, the G-statistic is compared to a critical statistic value. This value corresponds to a critical p-value, and indicates a specific confidence level. We choose our critical p-value to be 0.1 (see Appendix C). If the statistic is greater than the critical value, the null hypothesis (both samples drawn from the same distribution) is rejected. The simulation continues to evolve with Chain KMC only until equilibration is tested again. If the G-statistic is less than the critical value, the test fails to reject the null hypothesis and the equilibration assumption is applied. The G-test, like the chi-squared test, is known to be less reliable when some bins have low frequency. An often quoted rule of thumb is to require each bin to have an expected count of at least 5 [47]. To be conservative we require each basin site to have 10 visits both observed and expected before the G-test result is used (see Appendix C).

The time required to exit an equilibrated basin follows an exponential distribution, as given in Eq. (3). The time spent equilibrating in the basin does not affect the future exit time (see Appendix D). The full exit time distribution is composed of exits made using only Chain KMC moves and exits made after equilibration is reached and the equilibrating basin assumption is applied. Figure 1 illustrates the contributions of Chain KMC only and equilibrated exits to the overall exit time

distribution in an energy basin tested at two different temperatures. This figure illustrates that, as expected, the Equilibrating Chain algorithm applies the equilibration assumption more often at a lower temperature, where equilibration occurs more frequently before exiting. The details of this basin, which we refer to as the tetra-diamond basin, are described in section 3.2, and the quantitative assessment of the exit time will be discussed in section 4.

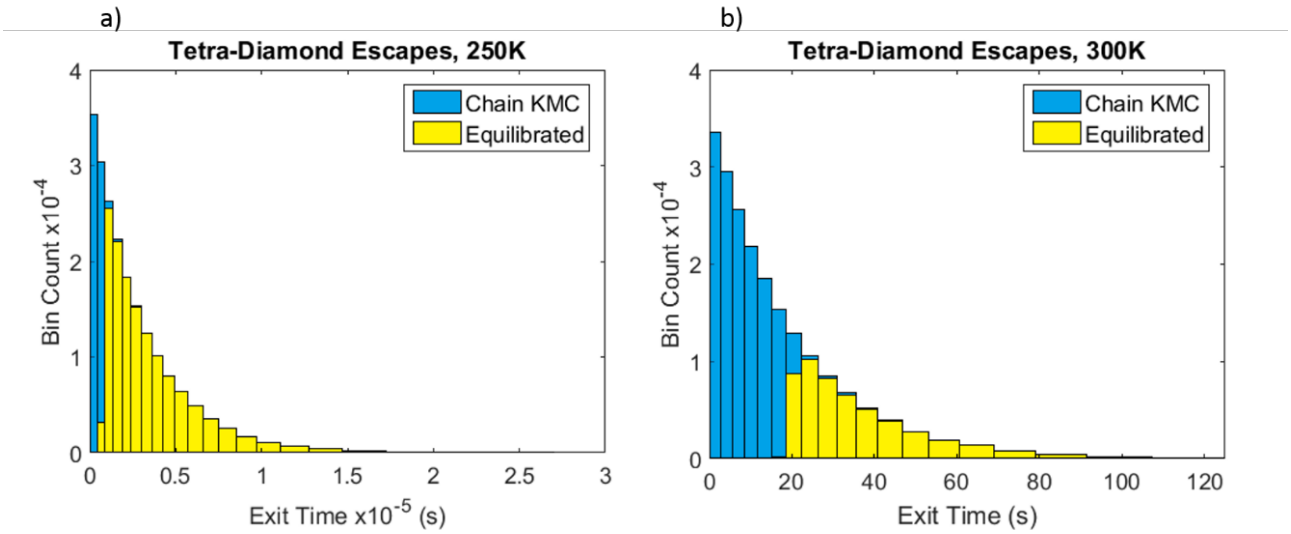

Figure 1 Relative contributions of equilibrated and Chain KMC only exits to the Equilibrating Chain algorithm exit time distributions at a) 250 K and b) 300 K.

It is possible that, due to large differences in occupation probability, a sub-basin is explored and reaches equilibration before all states that make up the larger basin are explored. The equilibrium assumption can still be applied because exit rates from the sub-basin to all absorbing sites are accurately calculated. If a new energy level of a larger basin is found through equilibration then the occupation distribution is reset. This is because an unknown number of visits in the sub-basin sites were skipped by applying the equilibration assumption. Before the assumption is applied to the full basin the occupation distribution must come to equilibrium again. Later equilibration testing is applied only to the larger basin.

Care should be taken when the basin is modified by another KMC event before the vacancy exits. This might occur when multiple vacancies are considered, or an independent event occurs such as ballistic mixing or defect production for alloys subjected to irradiation [38]. Before the equilibration assumption is accepted, the trajectory of the vacancy is controlled by Chain KMC so there is no introduced error. If the equilibration assumption for a given basin has been applied but the defect has not exited that basin before it is modified by an independent event, then a new vacancy site is selected according to the Boltzmann energy distribution of the basin, given in Eq. (1). Any modified jump parameters are then deleted, and the vacancy begins to explore the modified basin again using Chain KMC.

3.2. Evaluation Methods

In order to evaluate the accuracy and efficiency of the Equilibrating Chain KMC algorithm, we first compare it to 1st Order KMC and Chain KMC algorithms for the case of a vacancy escaping from static energy basins a priori generated by considering three clusters of solute atoms. The

simulation is initialized with the vacancy and the preset cluster embedded in a large box of matrix atoms. The simulation runs with the specified algorithm until the vacancy swaps with a matrix atom, when it is said to have escaped from the basin. The exit time, absorbing site, and number of KMC steps before exiting are recorded and the simulation is reset to the initial configuration. This is repeated 200,000 times for each set of parameters, which was found adequate here to sample the overall exit distributions.

The smallest basin we consider is comprised of 3 solute atoms and one vacancy forming a regular tetrahedron of 1st NN sites, see Fig. 2(a). This configuration is symmetric as the four possible vacancy sites correspond to degenerate sites, but it has low flicker probability since the vacancy has an equal probability to jump to any of the other three basin sites. Because of symmetry, the initial vacancy site is irrelevant. The next larger configuration, which we term 'tetra-diamond', consists of five sites, illustrated in see Fig. 2(b). It can be thought of as a four close-packed sites forming a diamond shape, with an additional solute at a tetrahedral site. This configuration is asymmetric and results in the basin having three distinct energy levels. We chose to initialize the configuration with the vacancy at one of the deep energy-level sites. These clusters have been selected since the tetrahedral and tetra-diamond configurations are the most compact, low-energy, and therefore the most probable configurations with four and five sites respectively. In contrast to these we examine what we refer to as the 'tetra-line' configuration, which consists of as a tetrahedron, with three additional sites extending along one of its edges. The configuration has a deep sub-basin, but also an extended shallow segment. This results in significantly more absorbing sites and a reduced probability of equilibrating the full basin. The vacancy is introduced either at the tetrahedral end, see Fig. 2(c), which is the deep sub-basin, or at the tip of the extended line in the shallow section of the energy basin, see Fig. 2(d). The energy basins corresponding to these clusters are shown in Figure 3. These four configurations are tested at various temperatures to investigate the effects of the overall trapping strength of the basin, the expected transient site occupation, and absorbing site distributions.

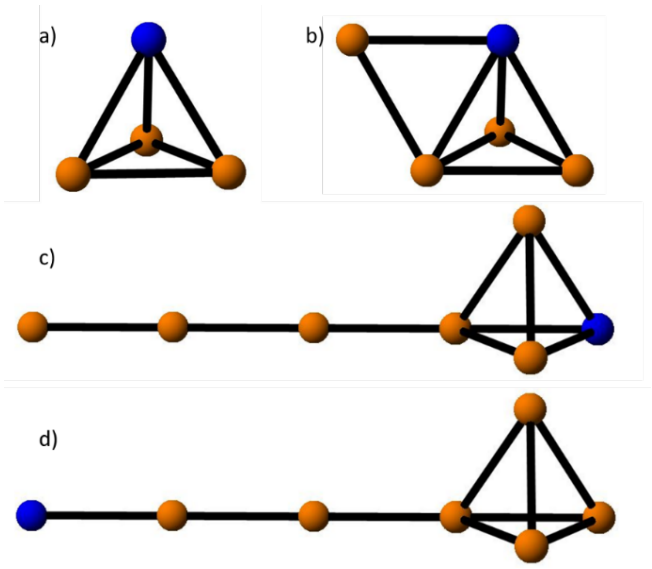

Figure 2 Images for the a) tetrahedral, b) tetra-diamond, and tetra-line c) deep and d) shallow configurations. Blue indicates a vacancy. Matrix atoms are not shown.

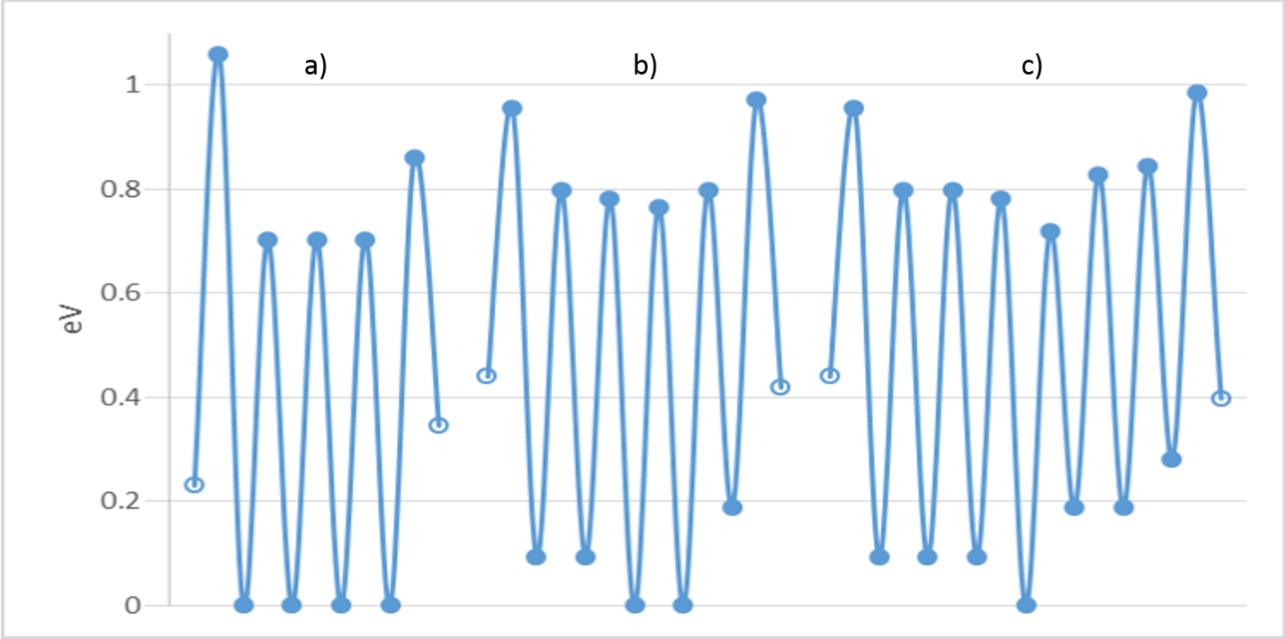

Figure 3 Relative saddle point and state energies defining the energy basins of the a) tetrahedral, b) tetra-diamond, and c) tetra-line clusters. The x-coordinate axis corresponds to various vacancy positions in the transient (•) and absorbing (o) sites. The saddle points to the absorbing sites shown are the lowest from the corresponding transient site. The tetrahedral basin only has two non-degenerate exits, so both are shown on opposite sides of the basin. The lowest energy level of each basin is arbitrarily set to zero in the figure for clarity.

These simulations on static clusters are used to determine whether the Equilibrating Chain algorithm introduced here reproduces both the exit time and exit site distributions of the 1st Order and Chain KMC algorithms. We compare the resultant distributions using the two-sample

Kolmogorov-Smirnov (KS) test and the two-sample chi-squared goodness-of-fit (2S-CS GoF) test respectively [48], always referencing Chain KMC results for the sake of consistency. Each of these tests produces a p-value, which we compare to our pre-selected critical confidence level of 0.05. We will also refer to a percent difference between the observed average exit time, and the equilibrated average exit time as calculated by Eq. 3. This is calculated as $\%Diff = abs(100 * (\Delta\bar{t}_{equil.} - \Delta\bar{t}_{obs.})/\Delta\bar{t}_{equil.})$.

4. Results and Discussion

4.1. Algorithm Equivalence and Speed-Up

In order to compare the three algorithms explained above (1st Order, Chain, and Equilibrating Chain), we present the statistical results of a vacancy escaping from a single basin, calculated at two representative temperatures, 250 K (Table 3) and 300 K (Table 4). The average exit time and average number of KMC steps to exit are listed for direct comparison between algorithms. Also listed where applicable are the fraction of exits that the basin was found to have come to equilibration. The uncertainty listed in these measurements is one standard deviation, and is estimated through the bootstrapping method [49]. The last columns are the p-value results of the KS-test and 2S-CS GoF tests as described above.

Table 3 Tetra-Diamond escape statistics at 200K. See text for column definitions.

Algorithm	Average Exit Time (s)	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
1st Order	1.344(3) x10^9	20220(50)	NA	0.128	0.156
Chain KMC	1.337(3) x10^9	9176(2)	NA	NA	NA
Equil. Chain	1.344(3) x10^9	512.2(4)	0.9439(5)	0.188	0.287

Table 4 Tetra-Diamond escape statistics at 250K. See text for column definitions.

Algorithm	Average Exit Time (s)	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
1st Order	22220(50)	2831(6)	NA	0.606	0.037
Chain	22300(50)	1364(3)	NA	NA	NA
Equil. Chain	22190(50)	434.8(4)	0.682(1)	0.567	0.314

The results presented in Tables 3 and 4 confirm that, as expected, the Chain algorithm represents a significant improvement over the 1st Order RTA algorithm, especially at low temperature, i.e., when reversals are more frequent. These results also demonstrate that the Equilibrating Chain algorithm is able to provide significant additional improvements, achieving comparable results with dramatically fewer KMC steps on average, especially at low temperature

where the equilibration assumption is often employed. Importantly, these results are obtained without introducing statistically significant errors since the post-run statistical analysis of the results yield confidence values in excess of the minimum desired value of 0.05. As a side comment, we note that the 2S-CS GoF test for the exit sites between the 1st Order and Chain KMC results fails for T=250 K marked in red in Table 4, despite the fact that the exit time KS-test is passed. At this confidence level, however, there is still a 5% percent chance of falsely rejecting a sample from the same distribution. Given the number of unique tests performed, 40 in this study (See Tables 3-8), and the confidence level of 0.05, the probability of false rejections can be calculated using the binomial distribution. Even assuming all the algorithms are equivalent there would be an 87.1% probability of at least one false rejection, and a 60.1% probability of at least two. We conclude that this case is an example of a false rejection. This conclusion is also supported by the fact that when comparing the 1st Order results to the Equilibrating Chain results, both the exit time and exit site tests are passed with p-values of 0.757 and 0.404 respectively.

4.2. Effect of Temperature

In order to more fully explore the effect of temperature on efficiency, we next present results for a single configuration over a wide temperature range, see Table 5. While results for the 1st Order algorithm are included for two temperatures in Tables 3-4, we did not generate results for all basins and temperatures with the 1st Order algorithm due to prohibitively long computation time, and therefore they are not included in Table 5. As in the previous section, the statistical tests are performed using the Chain KMC algorithm results as reference, and the columns are the same as in Tables 3 and 4.

Table 5 Tetra-Diamond escape statistics. See text for definitions.

Algorithm	Temperature (K)	Average Exit Time (s)	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
Chain	175	3.456(8) x10^{12}	36270(80)	NA	NA	NA
Equil. Chain	175	3.458(7) x10^{12}	525.1(4)	0.9857(3)	0.300	0.517
Chain	200	1.337(3) x10^{9}	9176(2)	NA	NA	NA
Equil. Chain	200	1.344(3) x10^{9}	512.2(4)	0.9439(5)	0.188	0.287
Chain	250	22300(50)	1364(3)	NA	NA	NA
Equil. Chain	250	22190(50)	434.8(4)	0.682(1)	0.567	0.314
Chain	300	14.42(3)	378.6(8)	NA	NA	NA
Equil. Chain	300	14.38(3)	283.0(4)	0.253(1)	0.329	0.283
Chain	325	0.859(2)	232.3(5)	NA	NA	NA
Equil. Chain	325	0.857(2)	207.4(4)	0.1064(7)	0.830	0.100

For each temperature, the average exit time has less than 1% difference between the Chain KMC and Equilibrating Chain algorithms. The Equilibrating Chain algorithm results also pass all

the statistical tests. Because the basin contains five sites, at least 500 KMC steps are performed before equilibration is tested. If the basin exits in less than 500 steps, as it often does above 300K, the equilibration assumption is never tested as it was not necessary to efficiently exit the basin. Because of this, the equilibration fraction decreases as the temperature increases. The average exit time at each temperature and for both algorithms is less than 0.5% different from the average equilibrated exit time, indicating that this basin equilibrates well within this entire temperature range. Because the equilibration fraction decreases the efficiency gain is also reduced, although the Equilibrating Chain algorithm always has fewer average KMC steps. The average number of KMC steps for escaping the tetra-diamond basin increases approximately exponentially as temperature decreases for both 1st Order and Chain KMC algorithms, as seen in Fig. 3. In contrast, the Equilibrating Chain algorithm quickly levels off and increases only marginally as temperature decreases. This demonstrates that as the equilibration assumption becomes more clearly acceptable, the efficiency of the new algorithm increases.

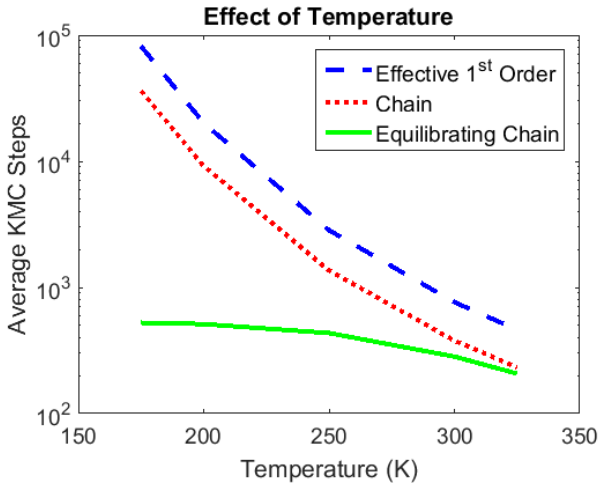

Figure 4 The average number of KMC algorithm steps to exit the tetra-diamond basin. 1st Order steps are calculated from the number of skipped flickers in the Chain KMC results.

4.3. Configurations

In order to test whether the Equilibrating Chain algorithm works well for a variety of basin shapes, we compare the results of the Chain and Equilibrating Chain algorithms for all of the basins described in Section 3.2, at the representative temperatures of 200 and 250 K. The same statistical tests are applied as described above.

Table 6 Escape statistics for various configurations at 200K. See text for definitions.

Configuration	Algorithm	Average Exit Time (s)	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
Tetrahedral	Chain	1.704(4) x10^7	4730(10)	NA	NA	NA
Tetrahedral	Equil. Chain	1.693(4) x10^7	391.0(2)	0.9172(6)	0.140	0.311

Tetra-Diamond	Chain	1.337(3) x10^9	9176(2)	NA	NA	NA
Tetra-Diamond	Equil. Chain	1.344(3) x10^9	512.2(4)	0.9439(5)	0.188	0.287
Tetra-Line Deep	Chain	1.039(2) x10^9	7900(20)	NA	NA	NA
Tetra-Line Deep	Equil. Chain	1.040(2) x10^9	2084(3)	0.746(1)	0.944	0.588
Tetra-Line Shallow	Chain	1.034(2) x10^9	7860(20)	NA	NA	NA
Tetra-Line Shallow	Equil. Chain	1.035(2) x10^9	2011(3)	0.741(1)	0.464	0.359

Table 7 Escape statistics for various configurations at 250K. See text for definitions.

Configuration	Algorithm	Average Exit Time (s)	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
Tetrahedral	Chain	775(2)	762(2)	NA	NA	NA
Tetrahedral	Equil. Chain	775(2)	316.3(3)	0.587(1)	0.521	0.630
Tetra-Diamond	Chain	22300(50)	1364(3)	NA	NA	NA
Tetra-Diamond	Equil. Chain	22190(50)	434.8(4)	0.682(1)	0.567	0.314
Tetra-Line Deep	Chain	15270(30)	1130(2)	NA	NA	NA
Tetra-Line Deep	Equil. Chain	15240(30)	793(1)	0.300(1)	0.867	0.576
Tetra-Line Shallow	Chain	14920(30)	1101(3)	NA	NA	NA
Tetra-Line Shallow	Equil. Chain	14890(30)	757(1)	0.306(1)	0.433	0.852

The tetra-line configuration was selected to investigate a basin that traps but does not equilibrate well. We note that an even less-well equilibrating example is considered in Appendix C. At 250 K the observed average exit times when the vacancy begins in the deep part of the basin are 0.53% different or less from the equilibrated average exit time (15190.9 s). However, when the vacancy begins at the shallow part of the basin the average escape time for the Chain and Equilibrating Chain algorithms are 1.78% and 1.97% different from the equilibrated average,

respectively. Despite the fact that this basin does not always equilibrate well, it equilibrates often enough that the Equilibrating Chain algorithm is still able to have a significant efficiency gain compared to the Chain algorithm. It is important to note that the Equilibrating Chain algorithm reproduces the slight difference in average exit time produced when the vacancy is initially located at different sites (shallow versus deep trapping sites) in the basin, with the Equilibrating Chain results differing from the Chain algorithm results by less than 0.5% in each case. The exit time and exit site distributions compared with the Chain KMC results all pass the respective statistical tests.

As noted above, the tetra-line basin has a sub-basin that may equilibrate before the full basin is explored. When this happens, the vacancy may exit the sub-basin either to the shallow portion of the full basin, or directly exit to a matrix absorbing site. Both scenarios were observed when the vacancy began at the deep part of the basin. When the vacancy begins in the shallow portion of the basin, the shallow part is explored first, so the full basin must equilibrate before the assumption is applied.

4.4. Cluster Dissolution Kinetics

The previous sections considered escape from basins of fixed size and shape. In order to investigate the effectiveness of the Equilibrating Chain algorithm in more realistic situations, we next apply the algorithms to the dissolution of a cluster of solute, i.e., in a microstructure containing evolving traps. The simulation is initialized in a cluster of a single vacancy, with all twelve nearest neighbors as B atoms, in a box of 32^3 A atoms with periodic boundary conditions. At the temperatures of interest, here 275, 300, and 325 K, the resulting B concentration (0.37 at.%) is below the thermal solubility (increasing from 0.5 at.% to 1.1 at.% in this temperature range), and thus the cluster dissolves during the simulation. The criterion for dissolution is that each of the twelve B atoms and the vacancy are all at least three nearest neighbor distances from each other. The simulation time, KMC steps, and final solute distribution are recorded for analysis. Statistics are collected over 10,000 dissolutions. The dissolution time distributions are compared using the KS-test, as above. The final solute distributions are projected along the normal of the (yz), (zx), and (xy) planes of the rhombohedral simulation cell. The results are compared between algorithms using the 2S-CS GoF test. Note that in this more generalized example, we report the average number of times the equilibration assumption was applied, rather than the fraction of times it was applied.

Table 8 Dissolution statistics for a 12-solute atom precipitate. See text for column definitions.

Algorithm	Temp. (K)	Average Time (s)	Average KMC Steps	Average Equilibrated Basins	Dissolution Time KS-Test P-value	B Distribution P-Value		
						(YZ)	(ZX)	(XY)
Chain	275	2.48(2) x10^{17}	1.281(6) x10^{7}	NA	NA	NA	NA	NA
Eq. Chain	275	2.52(2) x10^{17}	4.40(2) x10^{6}	628(3)	0.050	0.958	0.656	0.317

Chain	300	4.23(4) x10^{14}	3.59(2) x10^{6}	NA	NA	NA	NA	NA
Eq. Chain	300	4.30(4) x10^{14}	1.605(8) x10^{6}	255(1)	0.329	0.680	0.838	0.728
Chain	325	1.93(2) x10^{12}	1.263(6) x10^{6}	NA	NA	NA	NA	NA
Eq. Chain	325	1.95(2) x10^{12}	7.17(4) x10^{5}	113.4(5)	0.709	0.434	0.617	0.728

At lower temperatures, as expected, the gain from the Equilibrating Chain algorithm is larger, .e.g., requiring only ≈ 34% of the Chain algorithm steps at 275 K, owing to 628 equilibrated basins on average per dissolution. The distributions of dissolution times obtained from both algorithms are very similar, as illustrated in Fig. 5. The average dissolution times found with the Equilibrating Chain algorithm are less than 2% different than the Chain KMC results. We note that the exit time comparison at 275 K failed the KS-test at our confidence level. As explained in section 4.1, there was a significant probability of one or two false rejections when assuming that the algorithms are equivalent. The solute distributions all match and an example is shown in Fig. 6.

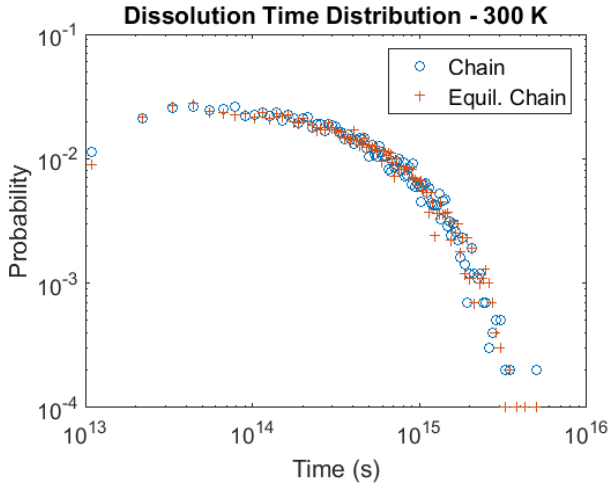

Figure 5 Histogram of the 12B precipitate dissolution time for both the Chain and the Equilibrating Chain algorithms at 300 K.

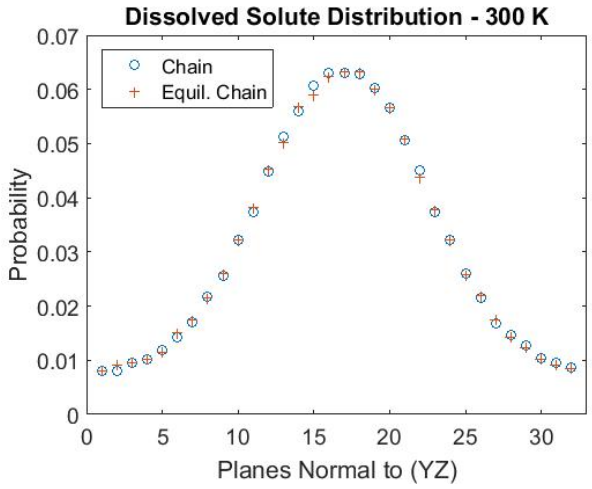

Figure 6 Histogram of the average dissolved solute distribution on the planes normal to (YZ) at 300 K.

4.5. Discussion

The results presented in sections 4.1-4 establish that, in the presence of traps, the Equilibrating Chain algorithm is always more efficient for than 1st Order or Chain RTA while maintaining statistical similar results. Specific efficiency gains will of course depend strongly on the system modeled, including atomic interactions, temperature, concentration, lattice structure and other possible events, e.g., in the case of alloys under irradiation. Therefore it is not possible to precisely quantify efficiency gains a priori. Nevertheless, it is clear that our algorithm provides larger gains in the case of microstructures containing strong traps, resulting in a frequent use of the equilibration approximation. An example system where we expect this algorithm to be highly effective is an irradiated alloy with strongly trapping but dilute nano-clusters, such as irradiated iron-copper [43].

While the equilibrium occupation distribution used here is calculated by assuming that the vacancy cannot leave the basin, a more accurate approach is to calculate the quasi-stationary distribution, or QSD. The QSD is the asymptotic limit of the no-passage distribution, which is the occupation distribution assuming that the vacancy does not leave the basin but can. It is calculated as the normalized eigenvector corresponding to the smallest eigenvalue of the transition matrix of the basin [50]. In the case of weaker traps, the QSD may deviate significantly from the equilibrium distribution, but the distributions are very similar for deep traps. We also note that while the equilibrium distribution only takes into account the site energy, the QSD predicts occupation differences between sites with the same energy if the basin exit rates from each site are different. We compared the differences between the QSD and the equilibrium site visit distributions for the basins and temperatures examined in this work. We find that the largest probability differences for any given site are of the order 10^{-3} and are often much less. The QSD predicted exit time is slightly longer than the equilibrated exit times, but still half a percent difference or less. This close agreement between the distributions is in agreement with our findings that the equilibrium distribution is often an acceptable approximation of long-term behavior. Compact (more inter-connected) basins and lower temperatures lead to smaller differences between distributions. See the supplementary material for more detailed comparisons. The core benefit of using the equilibrium distribution rather than the QSD is that it is less computationally demanding to calculate site energies than the eigenpairs of a transition matrix, especially for larger basins. If the observed distribution does deviate significantly from the equilibrium distribution, e.g. when the QSD and the equilibrium distribution are different, then the G-test will reject the equilibrium assumption and the algorithm will continue to explore and exit the basin using only error-free Chain KMC.

It is interesting to compare and contrast our algorithm with other accelerated KMC approaches that have been proposed, other than those used in the present study. Kinetic path sampling (KPS), proposed by Athenes and Bulatov [20], works in conjunction with standard 1st Order KMC by identifying a basin, sampling an exit from that basin, and then stochastically sampling a path to that exit in order to calculate an exit time. The KPS basin exit is statistically equivalent to 1st Order KMC and may be less computationally intensive than FPTA. However it cannot appropriately assign a position to a vacancy at a constrained time, which is required if the basin is modified by another event before exiting. This is important, for instance, for modeling the evolution of alloys under irradiation or plastic deformation. The authors are currently developing solutions to this

issue [51]. At this time, the Equilibrating Chain algorithm therefore appears to be advantageous for systems where basins can be modified by an independent event. In addition, the present algorithm does not require any significant overhead for its implementation. In 2010, Chatterjee and Voter proposed an accelerated superbasin (AS-KMC) algorithm that has similarities to the algorithm presented here [23]. As in our approach, the AS-KMC algorithm tracks events within basins, but instead of tracking site occupation it tracks the occurrence of specific jumps. Instead of using equilibration, AS-KMC progressively raises the activation barriers within the basin to reduce trapping. That algorithm tests if events within the basin have occurred often enough to limit the introduced errors. By using this approach, these authors are able to provide an analytical bound for the introduced errors, and show that in practical cases errors are indeed bounded by, and often much smaller than this upper value. Chatterjee and Voter considered a variety of examples of application of their algorithm, including extended bulk basins with gradual changes in the migration barriers and nanowire shape transitions. Our Equilibrating Chain algorithm performs best in the presence of deep traps. In contrast, in these cases, AS-KMC may be less effective since the transient energy barriers will have to be raised multiple times, rather than directly exiting. We note also that it would be detrimental to use the AS-KMC in conjunction with Chain KMC, since new jump rates would need to be calculated and accumulated every time barriers were raised. It will be informative to perform quantitative comparisons of the respective merits of the Equilibrating Chain, the AS-KMC and the FPTA algorithms. This will require dedicated and extensive tests, which are kept for future work.

Lastly we note that while our work has focused on the movement of vacancies near solute, the principles of the Equilibrating Chain KMC algorithm can be applied to any mobile object that experiences trapping. These include, but are not limited to, vacancies and vacancy clusters, direct interstitials, dumbbell interstitials and their clusters.

5. Conclusion

Standard KMC algorithms often become inefficient when applied to strongly trapping systems due to the long exploration of deep energy basins. We leverage the pre-existing data-tracking framework of Chain KMC to identify basins on the fly, to track the occupation of basin sites, and to evaluate if the mobile objects have reached equilibrium within the basin. This produces a flexible and robust algorithm that can safely apply the simplifying equilibrating basin assumption without a separate basin-finding algorithm and without a priori definition of basin characteristics. For selected basins, the Equilibrating Chain algorithm, tested here for a semi-dilute A-B FCC alloy, is found to generate escape time and site distributions in excellent agreement with the Chain KMC, used as a reference. In addition, excellent agreement is also obtained for the distribution of dissolution times of a small solute cluster. At the same time, the Equilibrating Chain algorithm can provide significant efficiency gains, particularly in simulation conditions where strong traps composed of several sites equilibrate. While this new algorithm is greatly accelerating, it should be less computationally intensive than other more exact algorithms, which makes it more suitable for systems where basins may be modified by independent events such as defect production in alloys under irradiation.

Acknowledgements

This research is supported by the NSF under Grant Nos. DMR-1306475 and DMR-1709857. We gratefully acknowledge stimulating discussions with Professor R.S. Averback, as well as to Manuel Athènes for an introduction to the QSD approach.

Appendix A

We present a summary of the relationships between the setting parameters and those used in the KMC calculations. A broken bonds model is used, where the energy barrier for a mobile object, such as a vacancy, jumping to a new site is calculated as the saddle point, minus the energy of all bonds 'broken' by the jump.

The S_0 saddle point is used when a vacancy is moving in A, away from any B atoms. S_1 is used when the vacancy swapping with an A atom but rotating around an associated B atom. S_2 is used for swapping with any B atom. $S_{3,4}$ is used when either associating or dissociating from any B atom. The jumps are reversed from each other, so they have the same saddle point energy. When the 5 frequency model scenarios are considered in the broken bonds framework, these relationships follow:

$$E_0 = S_0 - 11\varepsilon AA - 12\varepsilon AV$$
$$E_1 = S_1 - 10\varepsilon AA - 11\varepsilon AV - \varepsilon AB - \varepsilon BV$$
$$E_2 = S_2 - 11\varepsilon AV - 11\varepsilon AB - \varepsilon BV$$
$$E_3 = S_{3,4} - 11\varepsilon AA - 11\varepsilon AV - \varepsilon BV$$
$$E_4 = S_{3,4} - 10\varepsilon AA - 12\varepsilon AV - \varepsilon AB$$

Additional inputs include the cohesive energy of the A matrix, the cohesive energy of the B matrix, the formation energy of a vacancy in the matrix, and the critical temperature. The relationships between these parameters and the interaction energies are given below.

$$E_{coh}^{A} = \frac{Z}{2} * \varepsilon AA$$
$$E_{coh}^{B} = \frac{Z}{2} * \varepsilon BB = E_{coh}^{A} * \frac{T_m^B}{T_m^A}$$
$$E_{Vf}^{A} = Z * \varepsilon AV - \frac{Z}{2} * \varepsilon AA$$
$$w_{AB} = 2 * \varepsilon AB - \varepsilon AA - \varepsilon BB$$
$$T_c * k_B = \frac{3w_{AB}}{12K_c}$$
$$K_c = 0.102069$$

K_c includes a correction term between the T_c of the mean field regular solution model and the FCC lattice [52]. This system of equations is sufficient to define the all parameters based on the fitting parameters described in the text.

Appendix B

The probability that a flicker event will occur between two neighboring states *j* and *k* is defined as $\gamma_{j\leftrightarrow k} \equiv p_{j\to k} * p_{k\to j}$. A null state 0 is defined, such that $r_{j\to 0} = 0$ for all *j*. The probability of a flicker event occurring between any neighbors of state j is $\gamma_{j(0)} \equiv \sum_k \gamma_{j\leftrightarrow k}$, with *j(0)* indicating that a flicker to the null site is excluded, i.e. no neighbors are excluded. If a reversal to the previous state *i* is excluded, the probability becomes $\gamma_{j(i)} \equiv (\gamma_{j(0)} - \gamma_{j\leftrightarrow i})/(1 - p_{j\to i})$.

An indirect move is defined as any move from state *j* to state *k* where the first two jumps after arriving in state *j* comprise a flicker, including moves with multiple flickers or flickers to multiple neighboring sites. A direct move is defined as any non-indirect move, and never includes any flickers. Excluding the reversal to state *i*, the probability of a direct move occurring from state *j* to state *k* is

$$p_{direct} = \sum_k \frac{(p_{j\to k} - p_{j\to k(i)} * p_{k\to j}) * (1-\delta_{ki})}{1-p_{j\to i}} = 1 - \gamma_{j(i)} \text{ .} \quad \text{(B.1)}$$

The Kronecker delta, δ_{ki}, evaluates as zero unless $k = i$ when it evaluates as 1. It is used in B.1 to exclude the probability associated with the excluded reverse move. Again, excluding the immediate reversal to state *i*, the probability of an indirect move occurring becomes

$$p_{indirect} = \sum_k \frac{\gamma_{j(i)} * (p_{j\to k} - p_{j\to k(i)} * p_{k\to j})}{1-\gamma_{j(0)}} = \gamma_{j(i)} \text{ .} \quad \text{(B.2)}$$

Instead of calculating all of these probabilities, the algorithm first stochastically selects the number of flickers as $f = int(ln[\zeta/\gamma_{j(i)}]/ln\,\gamma_{j(0)}) + 1$ if $\zeta \le \gamma_{j(i)}$ or $f = 0$ if $\zeta > \gamma_{j(i)}$. The next move is then selected, with the jump probabilities of that move based on whether the move is direct or indirect, i.e. $f = 0$ or $f > 0$. The sites involved in those flickers are selected stochastically based on the binomial distribution. The total timestep accounts for each jump made, including flickers, without explicitly performing these flickers in the KMC simulation.

It is important to stress that, as demonstrated by Mason et al. [18], the Chain KMC algorithm is statistically equivalent to the 1st Order RTA. In particular, we note that in the Chain KMC algorithm, for any neighbor *k* that does not have the return probability $p_{k\to j}$ stored, this probability is set to zero. Thus, if no neighbors have stored return jump probabilities, the first order jump probabilities are recovered.

Appendix C

The Equilibrating Chain algorithm has at least three input parameters that directly impact potential efficiency and accuracy as currently implemented. They are 1) the frequency of testing for equilibration, determined here by the number of in-basin KMC steps (N_{chain}), 2) the minimum number of visits and expected visits to each basin site before trusting the G-test, and 3) the critical p-value of the G-test to determine equilibration.

In our implementation, we chose to test for equilibration after $(100 * N_{sites})$ steps without visiting a new site. We chose this value as it is a clear indication that the Chain KMC algorithm is becoming inefficient. The scalar could be reduced from 100, as is done below, or a different trigger could be used. Testing for equilibration more frequently has a complex effect on efficiency, as it

will allow equilibrated basins to exit with fewer KMC steps, but may also result in more frequent rejected equilibrations, wasting the computation time spent testing. The ideal scenario is to test immediately after achieving equilibration. Unfortunately this is system specific and is unknown a priori.

The minimum number of visits to each site, observed and expected, affects the reliability of the G-test. As noted in the text, a generally accepted guideline is 5 counts-per-bin (visits-per-site). We chose to require 10 visits-per-site in order to increase reliability. However, this reduces how often the basin is considered equilibrated, especially for larger basins that may have sites with very low occupation probability, such as in the tetra-line case.

The last significant parameter is the critical p-value, or confidence level, of the G-test. As the critical p-value is increased, the probability of false passes, or Type II errors, decreases. However, the probability of Type I errors, or incorrectly rejecting equilibrated basins, increases. In other words, higher critical p-value results in greater confidence that the basin has indeed equilibrated, but at the expense of not accepting some basins that truly have equilibrated but with high variance. As with other parameter choices, the choice of critical p-value, p=0.1, was rooted in a desire to be conservative.

Further analysis is complicated by the fact that these parameters are not independent. For example, by requiring a large number of basin KMC steps before testing for equilibration, the observed variance will often be far lower than the critical value, making the algorithm relatively insensitive to decreases in the p-value.

To illustrate the trade-offs between accuracy and efficiency, we introduce a new basin. This basin is composed of two distantly connected trapping sub-basins, and does not equilibrate well at 300 K. Indeed the average escape time, when the vacancy begins as shown in the figure, has a 12% difference compared with the equilibrated average exit time of 5477.8 seconds. The results produced when one or two of the tuning parameters are relaxed are presented below. The first three columns refer to the tunable parameters described above. The exit time differences use the Chain KMC results as the reference. The other columns correlate to those shown above.

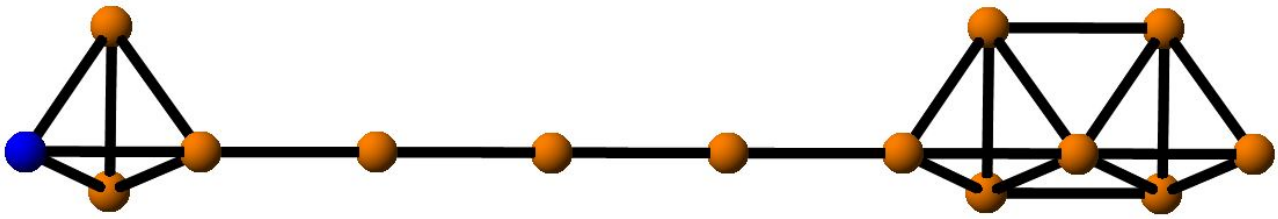

Figure C1 A selected non-equilibrating basin. As above, the orange atoms are solute, and the blue is the initial vacancy position. Matrix atoms are not pictured.

Table C1 Example non-equilibrating basin escape statistics at 300 K using the Equilibrating Chain algorithm with various tuning parameters. See text for column definitions.

G-test P-value	N_{chain}/N_{sites}	Min. Visits	Average Exit Time (s)	Exit Time % Diff	Average KMC Steps to Exit	Equilibration Fraction	Exit Time KS-Test P-value	Exit Site 2S-CS GoF P-value
NA	NA	NA	4820(10)	NA	1299(3)	NA	NA	NA

0.1	100	10	4790(10)	0.6	1217(3)	0.0477(5)	0.353	0.689
0.05	100	10	4800(10)	0.4	1215(3)	0.0522(5)	0.373	0.400
0.01	100	10	4790(10)	0.6	1203(3)	0.0597(5)	0.471	0.447
0.1	30	10	4770(10)	1.0	1174(3)	0.0834(6)	0.503	0.651
0.1	5	10	4800(10)	0.4	1152(3)	0.1625(9)	0.848	0.768
0.1	100	5	4790(10)	0.6	1165(3)	0.0849(6)	0.513	0.806
0.1	100	1	4780(10)	0.8	1159(3)	0.0890(6)	0.670	0.458
0.1	5	1	4870(10)	1.0	732(2)	0.929(2)	8.75E-25	0.038
0.01	100	1	4750(10)	1.4	1121(2)	0.1127(7)	0.209	0.270
0.01	5	10	4820(10)	<0.1	1136(3)	0.1798(9)	0.478	0.371

While these results are not a comprehensive study on the effect of the tuning parameters, they are instructive. For this example, the G-test p-value seems to be the weakest tuning parameter. Reducing the p-value alone only affects the average KMC steps slightly. In addition, the most inaccurate results were produced when the G-test p-value was the only strict parameter. While a minimum of 5 visits per site is recommended to provide proper sampling for the G-test [47], this had less effect than expected. This may be because the sites least often visited tend to have the least significant effect on the characteristics of the basin both before and after equilibration.

In addition to these parameters, other details of the algorithm could be modified that could have an effect on efficiency. For example, saved chain data that is far away from the current position could be periodically deleted. Limiting the number of saved sites in this way would reduce data management requirements, and allow strongly equilibrating sub-basins, such as in the example above, to equilibrate more frequently. However, this would prevent larger but still equilibrating basins from being recognized.

Another possible variation would be to apply the G-test to the distribution of occupation time, rather than the distribution of visits. The two distributions have an explicit relationship in our implementation, as each vacancy jump timestep is set to the average exit from that site. However, many KMC models draw the timestep from the exponential distribution rather than the average, which may lead to subtle differences for infrequently visited states.

Appendix D

Whenever the equilibration assumption is applied, the vacancy has already spent time exploring the basin and coming to equilibrium. Assume that this time, t_{basin}, is tracked and stored

with the other Chain KMC data. The correct exit time distribution could be recovered by selecting an exit from the portion of the exponential distribution that has not yet been explored by Chain KMC. The total basin exit time would be calculated simply by re-scaling the limits of the random number ζ from (0,1] to (0, ϕ] in Eq. 3, where $\phi = e^{(-t_{basin}*Q)}$. This prevents an exit time that is less than the time already spent in the basin, while still following the exponential distribution of equilibrated exits. The exit timestep would then be the difference between the selected total basin exit time and the time spent equilibrating.

The modified range random variable can be defined as $\zeta' = \zeta * \phi$. If we substitute ζ' into Eq. 3, we find:

$$\Delta t_{exit} = \frac{-ln(\zeta * \phi)}{Q} = \frac{-ln(\zeta * e^{(-t_{basin}*Q)})}{Q} = t_{basin} - \frac{\ln(\zeta)}{Q} \quad (D.1)$$

Thus, the total basin exit time is just the sum of the time spent in the basin and an equilibrated exit step. Therefore, the total exit time can simply be selected through an equilibrated exit, and the time spent in the basin does not need to be taken into account.

Data Availability

The raw data required to reproduce these findings are available to download from http://dx.doi.org/10.17632/dbff6r2p88.1 .

References

1. Balbuena, J.P., M.J. Caturla, and E. Martinez, *Kinetic Monte Carlo Algorithms for Nuclear Materials Applications*, in *Handbook of Materials Modeling: Applications: Current and Emerging Materials*, W. Andreoni and S. Yip, Editors. 2018, Springer International Publishing: Cham. p. 1-22.
2. Martin-Bragado, I., et al., *Kinetic Monte Carlo simulation for semiconductor processing: A review.* Progress in Materials Science, 2018. **92**: p. 1-32.
3. Andersen, M., C. Panosetti, and K. Reuter, *A Practical Guide to Surface Kinetic Monte Carlo Simulations.* Frontiers in Chemistry, 2019. **7**.
4. Nordlund, K., *Historical review of computer simulation of radiation effects in materials.* Journal of Nuclear Materials, 2019.
5. Bellon, P., *Kinetic Monte Carlo simulations in crystalline alloys: Principles and selected applications.* Thermodynamics, Microstructures and Plasticity, 2003. **108**: p. 395-409.
6. Vineyard, G.H., *Frequency Factors and Isotope Effects in Solid State Rate Processes.* Journal of Physics and Chemistry of Solids, 1957. **3**(1-2): p. 121-127.
7. Metropolis, N., et al., *Equation of State Calculations by Fast Computing Machines.* Journal of Chemical Physics, 1953. **21**(6): p. 1087-1092.
8. Bortz, A.B., M.H. Kalos, and J.L. Lebowitz, *New Algorithm for Monte-Carlo Simulation of Ising Spin Systems.* Journal of Computational Physics, 1975. **17**(1): p. 10-18.
9. Leitnaker, J.M., E.E. Bloom, and J.O. Stiegler, *Effect of Minor Constituents on Swelling in Stainless-Steel.* Journal of Nuclear Materials, 1973. **49**(1): p. 57-66.
10. Koehler, J.S., *Decrease in Void Growth-Rate by Interstitial Trapping.* Journal of Applied Physics, 1975. **46**(6): p. 2423-2428.

11. Mansur, L.K. and M.H. Yoo, *Effects of Impurity Trapping on Irradiation-Induced Swelling and Creep.* Journal of Nuclear Materials, 1978. **74**(2): p. 228-241.
12. Kato, T., H. Takahashi, and M. Izumiya, *Grain-Boundary Segregation under Electron-Irradiation in Austenitic Stainless-Steels Modified with Oversized Elements.* Journal of Nuclear Materials, 1992. **189**(2): p. 167-174.
13. Gan, J., et al., *The effect of oversized solute additions on the microstructure of 316SS irradiated with 5 MeV Ni++ ions or 3.2 MeV protons.* Journal of Nuclear Materials, 2004. **325**(2-3): p. 94-106.
14. Puchala, B., M.L. Falk, and K. Garikipati, *An energy basin finding algorithm for kinetic Monte Carlo acceleration.* Journal of Chemical Physics, 2010. **132**(13).
15. Fichthorn, K.A. and Y.Z. Lin, *A local superbasin kinetic Monte Carlo method.* Journal of Chemical Physics, 2013. **138**(16).
16. Athenes, M., P. Bellon, and G. Martin, *Identification of novel diffusion cycles in B2 ordered phases by Monte Carlo simulation.* Philosophical Magazine a-Physics of Condensed Matter Structure Defects and Mechanical Properties, 1997. **76**(3): p. 565-585.
17. Athenes, M., P. Bellon, and G. Martin, *Effects of atomic mobilities on phase separation kinetics: A Monte-Carlo study.* Acta Materialia, 2000. **48**(10): p. 2675-2688.
18. Mason, D.R., R.E. Rudd, and A.P. Sutton, *Stochastic kinetic Monte Carlo algorithms for long-range Hamiltonians.* Computer Physics Communications, 2004. **160**(2): p. 140-157.
19. Dokukin, S.A., S.V. Kolesnikov, and A.M. Saletsky, *Efficient energy basin finding method for atomistic kinetic Monte Carlo models.* Computational Materials Science, 2018. **155**: p. 209-215.
20. Athenes, M. and V.V. Bulatov, *Path Factorization Approach to Stochastic Simulations.* Physical Review Letters, 2014. **113**(23).
21. Tokar, V.I. and H. Dreysse, *Accelerated kinetic Monte Carlo algorithm for diffusion-limited kinetics.* Physical Review E, 2008. **77**(6).
22. Oppelstrup, T., et al., *First-passage kinetic Monte Carlo method.* Physical Review E, 2009. **80**(6).
23. Chatterjee, A. and A.F. Voter, *Accurate acceleration of kinetic Monte Carlo simulations through the modification of rate constants.* Journal of Chemical Physics, 2010. **132**(19).
24. Perez, D., B.P. Uberuaga, and A.F. Voter, *The parallel replica dynamics method - Coming of age.* Computational Materials Science, 2015. **100**: p. 90-103.
25. Voter, A.F., *Parallel replica method for dynamics of infrequent events.* Physical Review B, 1998. **57**(22): p. 13985-13988.
26. Agarwal, A., et al., *Computing long time scale biomolecular dynamics using quasi-stationary distribution kinetic Monte Carlo (QSD-KMC).* Journal of Chemical Physics, 2019. **151**(7).
27. Leclaire, A.D., *Theory of Impurity Diffusion in Metals.* Philosophical Magazine, 1962. **7**(73): p. 141-+.
28. Philibert, J., *Atom Movements: Diffusion and Mass Transport in Solids*. Monographies de physique. 1991, Les Ulis, France: Editions de Physique. 577.
29. Morgan, D., T. Mayeshiba, and D. Morgan, *DFT dilute solute diffusion in Al, Cu, Ni, Pd, Pt, Mg, and W*. 2016.
30. Wolverton, C., *Solute-vacancy binding in aluminum.* Acta Materialia, 2007. **55**(17): p. 5867-5872.
31. Lee, E., F.B. Prinz, and W. Cai, *Enhancing ionic conductivity of bulk single-crystal yttria-stabilized zirconia by tailoring dopant distribution.* Physical Review B, 2011. **83**(5).
32. Schuler, T., et al., *Design principles for radiation-resistant solid solutions.* Physical Review B, 2017. **95**(17).

33. Allnatt, A.R., A.B. Lidiard, and Cambridge University Press., *Atomic transport in solids*. 1993, Cambridge University Press,: Cambridge England ; New York. p. 1 online resource (xxiv, 572 pages).
34. Garnier, T., et al., *Quantitative modeling of solute drag by vacancies in face-centered-cubic alloys.* Physical Review B, 2014. **89**(14).
35. Schuler, T., et al., *Modeling the long-term evolution of dilute solid solutions in the presence of vacancy fluxes.* Physical Review Materials, 2018. **2**(7).
36. Wu, H., T. Mayeshiba, and D. Morgan, *High-throughput ab-initio dilute solute diffusion database.* Scientific Data, 2016. **3**.
37. Doyama, M. and J.S. Koehler, *Relation between Formation Energy of a Vacancy and Nearest Neighbor Interactions in Pure Metals and Liquid-Metals.* Acta Metallurgica, 1976. **24**(9): p. 871-879.
38. Soisson, F., et al., *Atomistic Kinetic Monte Carlo studies of microchemical evolutions driven by diffusion processes under irradiation.* Journal of Nuclear Materials, 2010. **406**(1): p. 55-67.
39. Kittel, C., *Introduction to Solid State Physics*. 8th ed. 2005: John Wiley & Sons, Inc.
40. Erhart, P., *Atomic Defects in Metals*. Landholt-Bornstein New Series III/25, ed. H. Ullmaier. 1991, Berlin: Springer.
41. Abinandanan, T.A., F. Haider, and G. Martin, *Computer simulations of diffusional phase transformations: Monte Carlo algorithm and application to precipitation of ordered phases.* Acta Materialia, 1998. **46**(12): p. 4243-4255.
42. Voter, A.F. *INTRODUCTION TO THE KINETIC MONTE CARLO METHOD*. 2007. Dordrecht: Springer Netherlands.
43. Soisson, F. and C.C. Fu, *Cu-precipitation kinetics in alpha-Fe from atomistic simulations: Vacancy-trapping effects and Cu-cluster mobility.* Physical Review B, 2007. **76**(21).
44. Shu, S.P., P. Bellon, and R.S. Averback, *Complex nanoprecipitate structures induced by irradiation in immiscible alloy systems.* Physical Review B, 2013. **87**(14).
45. La Magna, A. and S. Coffa, *Accelerated Monte Carlo algorithms for defect diffusion and clustering.* Computational Materials Science, 2000. **17**(1): p. 21-33.
46. Van Siclen, C.D., *Stochastic method for accommodation of equilibrating basins in kinetic Monte Carlo simulations.* Journal of Physics-Condensed Matter, 2007. **19**(7).
47. McDonald, J.H., *G–test of goodness-of-fit*, in *Handbook of Biological Statistics*. 2014, Sparky House Publishing: Baltimore, Maryland. p. 53-58.
48. DeGroot, M.H. and M.J. Schervish, *Probability and statistics*. 3rd ed. 2002, Boston: Addison-Wesley. xv, 816 p.
49. Wang, H.T., W. , *Bootstrap Methods: The Classical Theory and Recent Developments*, in *Wiley StatsRey: Statistics Reference Online*. 2018.
50. Athenes, M., et al., *Elastodiffusion and cluster mobilities using kinetic Monte Carlo simulations: Fast first-passage algorithms for reversible diffusion processes.* Physical Review Materials (in press), 2019.
51. Athenes, M., C. Daniels, Editor. 2019.
52. Jindal, V., B.N. Sarma, and S. Lele, *An improvement of cluster variation method entropy functional for bcc alloys.* Calphad-Computer Coupling of Phase Diagrams and Thermochemistry, 2013. **43**: p. 48-51.