Journal Pre-proof

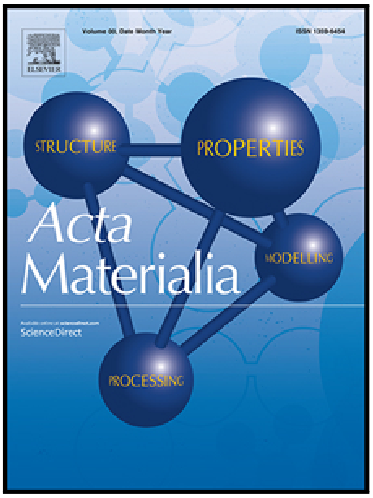

Accelerating off-lattice kinetic Monte Carlo simulations to predict hydrogen vacancy-cluster interactions in α–Fe

C.J. Williams, E.I. Galindo-Nava

PII: S1359-6454(22)00829-1
DOI: https://doi.org/10.1016/j.actamat.2022.118452
Reference: AM 118452

To appear in: *Acta Materialia*

Received date: 3 May 2022
Revised date: 13 October 2022
Accepted date: 18 October 2022

Please cite this article as: C.J. Williams, E.I. Galindo-Nava, Accelerating off-lattice kinetic Monte Carlo simulations to predict hydrogen vacancy-cluster interactions in α–Fe, *Acta Materialia* (2022), doi: https://doi.org/10.1016/j.actamat.2022.118452

Graphical Abstract

Accelerating off-lattice kinetic Monte Carlo simulations to predict hydrogen vacancy-cluster interactions in α–Fe

C.J. Williams, E.I. Galindo-Nava

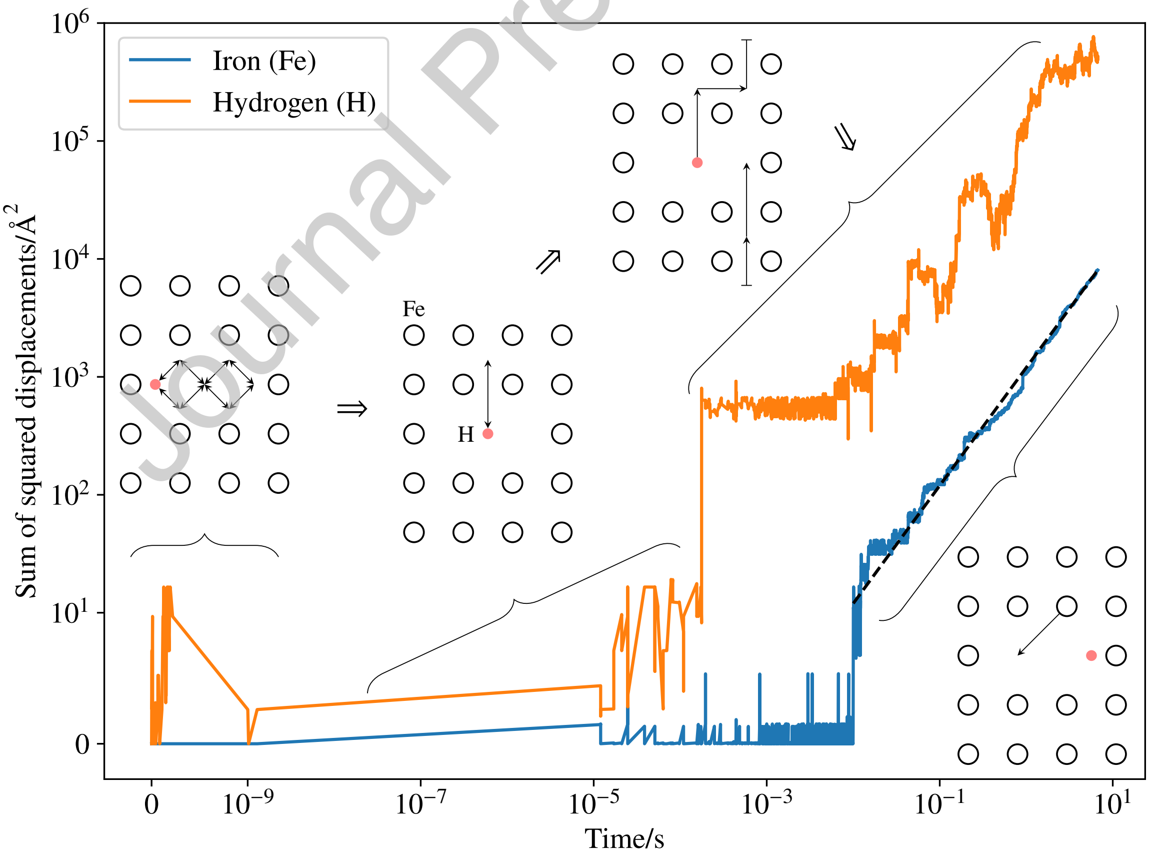

Accelerating off-lattice kinetic Monte Carlo simulations to predict hydrogen vacancy-cluster interactions in α–Fe

C.J. Williams[a,*], E.I. Galindo-Nava[a,b,*]

[a] Department of Materials Science and Metallurgy, University of Cambridge, Cambridge, CB3 0FS, UK
[b] Department of Mechanical Engineering, University College London, London, WC1E 6BT, UK

Abstract

We present an enhanced off-lattice kinetic Monte Carlo (OLKMC) model, based on a new method for tolerant classification of atomistic local-environments that is invariant under Euclidean-transformations and permutations of atoms. Our method ensures that environments within a norm-based tolerance are classified as equivalent. During OLKMC simulations, our method guarantees to elide the maximum number of redundant saddle-point searches in symmetrically equivalent local-environments. Hence, we are able to study the trapping/detrapping of hydrogen from up to five-vacancy clusters and simultaneously the effect hydrogen has on the diffusivity of these clusters. These processes occur at vastly different timescales at room temperature in body-centred cubic iron. We predict the diffusion pathways of clusters/complexes without *a priori* assumptions of their mechanisms, not only reproducing previously reported mechanisms but also discovering new ones for larger complexes. We detail the hydrogen-induced changes in the clusters' diffusion mechanisms and find evidence that, in contrast to mono-vacancies, the introduction of hydrogen to larger clusters can increase their diffusivity. We compare the effective hydrogen diffusivity to Oriani's classical theory of trapping, finding general agreement and some evidence that hydrogen may not always be in equilibrium with traps, when the traps are mobile. Finally, we compute the *trapping atmosphere* of meta-stable states surrounding non-point traps, opening new avenues to better understand and predict hydrogen embrittlement in complex alloys.

Keywords: Atomistic modelling, Off-lattice kinetic Monte Carlo, Point defects, Hydrogen embrittlement, Diffusion mechanism

1. Introduction

It has been known for over 100 years [1, 2] that the presence of hydrogen (H) in metals – particularly steels – can severely reduce ductility, leading to catastrophic failure below the yield-stress. The processes that cause these effects are collectively termed *hydrogen embrittlement* (HE). Despite a century of research, the core mechanisms of HE have yet to be fully understood and are still a topic of active research/debate [3]. The difficulty in understanding HE stems from its multi-scale nature; a full description of HE requires understanding of H-adsorption, H-diffusion/transport, and (most crucially) H interaction/influence with/on crystal defects. These processes span many orders

*Corresponding author
Email addresses: cw648@cam.ac.uk (C.J. Williams), e.galindo-nava@ucl.ac.uk (E.I. Galindo-Nava)

of length/time scales, which presents challenges when isolating/connecting the impact of H at the atomistic scale to the macroscopic results.

A breadth of mechanisms for HE have been proposed, most of them revolve around the interactions between H and crystal defects. For a more complete description see Ref. 4 and Ref. 3. A few of the most prominent mechanisms are: Hydrogen-induced decohesion (HID) [5–7], Adsorption-induced decohesion (AIDE) [8], Hydrogen-enhanced localised plasticity (HELP) [9–13] and Hydrogen-enhanced strain-induced vacancy (HESIV) formation [14–16]. Many of these mechanisms are supported by bodies of experimental work. As few are orthogonal to each other, it is likely that a full description of HE contains a combination of two/three of these mechanisms (alongside some yet undiscovered).

Theoretical and computational modelling play a crucial role in the study of the Fe-H system due to the inherent difficulty in experimental observations of atomic H [17]. The low solubility and high diffusivity [18] of H in body-centred cubic (BCC) iron (Fe), combined with the small 'nucleus' and low electron density, make direct experimental observations extremely challenging. Instead, techniques such as thermal desorption analysis (TDA) [19], electro-permeation (EP) experiments [20] and atom probe tomography (APT) [21] are employed. Many of these methods (with the notable exception of APT) are unable to directly investigate H diffusion and trapping behaviour within metals on the atomic scale thus, we must fall back to computation/theory to unravel the atomic mechanisms that cause HE.

Different computational modelling techniques have been used to investigate HE over varied assumptions and time/length scales. On the smallest length-scales, density functional theory (DFT) is used to study H binding sites [22–24] and occasionally combined with molecular dynamic (MD) in *ab initio* MD to study H diffusion at the highest accuracy [25]. Additionally, work has been done using path-integral MD [26, 27] to explore H diffusion in iron while incorporating quantum effects, which are known to be important at low temperatures [26]. Nevertheless, progress has been made modelling much-larger systems using classical approaches, the most popular of these is MD and its accelerated-variants using semi-empirical potentials. This has enabled the study of H-defect kinetics, such as grain-boundaries [28, 29] and dislocations [30, 31]. Molecular dynamics simulations must resolve atomic vibrations in order to accurately track the dynamics of atom-scale systems. This imposes a significant computational effort as the integration time-step must be of-the-order of these vibrations. Hence, even using today's computers, MD simulation timescales rarely exceed $\mathcal{O}(100\mu s)$. Continuing to the longest/largest scales, Monte-Carlo (MC) [32] and kinetic Monte-Carlo (KMC) [33] methods overcome this barrier by ignoring the explicit phase-space trajectory, instead focusing on state$\rightarrow$state transitions. This can significantly accelerate simulations. However, these methods are confined to discrete representation and require knowledge of mechanisms *a priori*.

Off-lattice[1] kinetic Monte Carlo (OLKMC) [34], an extended KMC method, is a general and unbiased tool (discovering mechanisms without any *a priori* input) being successfully applied to study the kinetics of various systems, e.g. Fe/Cu/Al, BBC/FCC, disordered systems, extended defects, point defects, etc. [35–42]. Off-lattice KMC automatically discovers the mechanisms available – using saddle-point (SP) searches to locate the transition states – and then applies the KM algorithm [33, 43] to advance the system state/time. Off-lattice KMC allows for the exploration of continuous systems, at previously inaccessible timescales, at atomic fidelity. As such, it is the

[1]Also known as: *"adaptive", "on-the-fly", "self-learning"* and *"self-evolving"*.

perfect tool to explore the uncertain and complex mechanisms controlling HE.

In this paper we develop enhancements to the OLKMC method. The motivation for our research is building a general simulation framework capable of modelling the complex interactions between crystal defects and H in iron, into the timescales required to study the mechanisms of HE. Our main contribution is an error-tolerant atomic local-environment (LE) identification/matching process to elide saddle-point searches.

We apply our OLKMC implementation to study the diffusion of vacancy clusters in the presence of H. We demonstrate OLKMC is capable of reaching embrittlement timescales, of-the-order-of seconds, while simultaneously resolving the atomic motion of H-atoms. With OLKMC we can study the atomic mechanisms through which H affects the diffusion of vacancy clusters. These are important first steps towards modelling the more complex H-defect interactions required to gain a full understanding of HE.

2. Background: off-lattice kinetic Monte Carlo

2.1. Saddle-point searches

The process of finding SPs, called the saddle-point search (SPS) procedure, is critical to the efficiency of OLKMC simulations. Several minimum-mode following algorithms were unified under one mathematical framework in Ref. 44 and compared. All investigated methods are bounded in efficiency by the Lanczos method [45]. We choose to use the superlinear dimer-method [46], owing to it requiring fewer force evaluations but converging almost as fast as the Lanczos method. The superlinear dimer-method contains several optimisations over the original dimer-method [47] – notably the improvements of Ref. 48. We discuss a minor modification in Section 3.3.

2.2. Saddle-point recycling

Each SPS requires many hundreds of calls to the force-field and many SPS must be carried out to ensure the completeness of the KMC catalogue. Due to the local nature of mechanisms, most of these SPS are unnecessary. For example, in a section of perfect lattice the LE around each atom is identical hence, the mechanisms that can occur at each atom are identical. Secondly, consider two atoms sufficiently far apart; a local mechanism centred on one will likely not change the LE around the second therefore, its accessible mechanisms remain the same. Finally, many atoms are in LEs differing only by an Euclidean transformation (of the form $\boldsymbol{r} \mapsto \boldsymbol{R}\boldsymbol{r} + \boldsymbol{c}$, with $\boldsymbol{R}$ an orthogonal matrix) hence, their mechanisms are related by the same transformation.

Multiple methods have been developed to reduce the cost of building the KMC catalogue by exploiting this locality. The simplest of these are system-wide methods, which attempt to reuse SPs discovered at the previous step [49] however, these do not exploit any relevant symmetries. Due to this inefficiency, they will not be discussed further. Alternatively, local methods seek to classify the LE around each atom in the system. It is then possible to associate mechanisms entirely within a LE. Mechanisms can then be cached and, when an equivalent LE is discovered, instead of launching new SPS, the mechanisms can be reconstructed from the cached information. If the LE classification is suitably invariant, these methods can account for all relevant symmetries hence, the focus shifts to LE classification, of which a number of methods have been proposed.

Space discretisation. Discrete pattern recognition methods for LE classification have been explored [50, 51] however, these often fail to account for (continuous) symmetries and, because of the discretisation of space, are sensitive to small changes in atomic positions (due to inexact energy minimisations).

Norm-based. Moving toward a tolerant classification, Ref. 52 presents a system that stores the LE of an atom at $\boldsymbol{r}_i$ as $\{\boldsymbol{r}_{ij} | r_{ij} < r_{\mathrm{env}}\}$. Local environments are considered equivalent when each atom in two superimposed LEs have a corresponding atom in the second environment within tolerance Δr_{tol}. This method gracefully allows for error on the positions of atoms in a LE. However, no method is presented for determining this equivalence between arbitrarily permuted/transformed LEs.

Topological. Graph-based topological methods, introduced in Ref. 37, fully exploit the symmetry of LEs. Atoms in the LE are used to draw a graph; atoms become nodes and atoms considered bonded (closer than some distance) are connected with an edge. LEs are equivalent if their graph representations are isomorphic. This is, in general, a problem in its own complexity class GI $\in$ NP which is not known to be in either P or NP-complete [53]. Fortunately, there exists implementations such as the `nauty`[2] software [54] which can solve this problem in polynomial time for many graphs. Although powerful, topological methods rely on a one-to-one correspondence between topology and geometry that may breakdown. Furthermore, they lose the tolerance of norm-based methods.

In Section 3.2, we introduce our own norm-based LE classification that combine the desirable properties of many of the previous methods.

2.3. Superbasins and the low-barrier problem

A common issue encountered during OLKMC simulations is the *low-barrier problem* (LBP) [37, 55]. This occurs when a collection of basins – often called a *superbasin* – are connected by a series of fast mechanisms. It requires many MC steps to escape from a superbasin. As the rate-sum is very large during this period, the simulated time advances very slowly.

The simplest methods to overcome the LBP effectively combine states connected by fast mechanisms into a single state and ignore all internal superbasin kinetics [56] – this is clearly not exact. Alternatively, TABU-like [57] methods that ban recent-transitions have been employed [37, 58]. These have been shown to be thermodynamically sound providing the total number of KMC steps is much greater than the oldest banned transition. Two exact solutions to the LBP are presented in Ref. 59; the key insight is the partitioning of states into *transient* and *absorbing* sets, followed by analytically solving the motion inside the transient states. A similar exact solution, the *mean-rate method* (MRM) [60], has been extended to OLKMC to form the basin auto-constructing MRM (bac-MRM) [61], which constructs superbasins on-the-fly. We discuss minor extensions in Section 3.5.

3. Methodology

3.1. Interatomic potentials

In order to reach HE timescales – of-the-order-of seconds – we employ embedded atom method (EAM) potentials [62]. These are short-range, fast, well tested, semi-empirical models of the

[2] https://pallini.di.uniroma1.it/

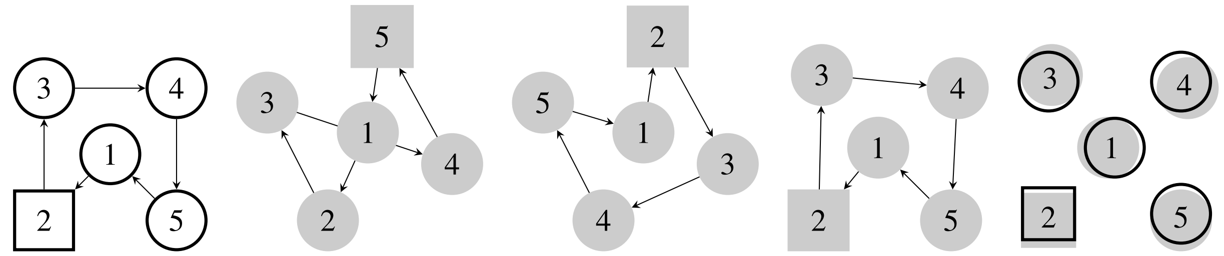

(a) Reference point-cloud. **(b)** Unclassified point-cloud. **(c)** Permutation of labels in Fig. 1b **(d)** Rotation of Fig. 1c **(e)** Fig. 1a over Fig. 1d.

Figure 1 Demonstration of the norm-based equivalence between reference point-cloud Fig. 1a (opaque shapes) and unclassified point-cloud Fig. 1b (grey shapes) via a permutation of the labels and a rigid-body rotation (transformation). Shape (square/circle) denotes the colour of each point and arrows act as a guide to the eyes for the permutation. In Fig. 1e we see all the points/atoms are close enough that the LEs can be considered equivalent.

potential energy of a collection of atoms. Although they are not without criticism [63], EAM potentials have become well established in the literature, particularly for metallic systems. We use the variation presented and fitted in Ref. 64, which generalise the EAM embedding function and are fit to first-principles (DFT) measurements and experimental data. Also, fit to a wide variety of targets, the potentials provide good reproduction of several crystal defect structures. We also include the modifications of Ref. 30, the introduction of additional H-H repulsion, to reduce the H clustering observed in the original potentials.

3.2. *New invariant and tolerant local-environment classification*

In previous work [65], we adopted a topological classification methodology however, this relied on the aforementioned one-to-one correspondence between topology and geometry. We found this correspondence to break-down in the Fe-H system due to the small size of the H-atom and small displacements during mechanisms. We tried to overcome this problem by allowing the bonding distance to vary with the species and colouring each atom as the pair formed from the atoms' atomic number and the local sum:

$$\left\lfloor c \sum_{j} G_{ji} r_{ij} \right\rfloor \tag{1}$$

with c a problem-dependant scaling constant and G_{ij} elements of the adjacency matrix. This encodes much more of the information contained within $\{r_{ij}\}$ into the coloured graph. Unfortunately, with the above modifications, infinitesimal perturbations in position are more likely to result in many keys representing the same geometry. Finally, there is no quantitative/qualitative link between topological keys and the similarity of LEs.

3.2.1. *Norm-based definition of equivalence*

We require a notion of equivalence between LEs that is invariant under:

1. Infinitesimal-perturbations of atomic positions.
2. Permutation of identical atoms.
3. Euclidean transformations of the group of atoms.

In order to overcome the aforementioned difficulties when dealing with the small errors in the atomic positions we move away from the graph-based representation of the atoms. Instead, we represent atoms as coloured points (2-tuples):

$$\left(\boldsymbol{p} \in \mathbb{R}^3, p \in \mathbb{Z}\right) \tag{2}$$

and a LE centred on the point $(\boldsymbol{p}_1, p_1)$ as the point cloud:

$$P = \{(\boldsymbol{p}_1, p_1), \ldots, (\boldsymbol{p}_n, p_n)\} \tag{3}$$

where (without loss of generality) we set the centroid of the points in P to the origin and $\|\boldsymbol{p}_1 - \boldsymbol{p}_i\| < r_{\text{env}}$ for all i with r_{env} the radius of the LE. The question of determining if two LEs, P and Q (of the same size), are *equivalent* is the same as asking if there exists a transformation matrix $\boldsymbol{O}$ and permutation π such that:

$$\sum_{i=1}^{n} \left\|\boldsymbol{p}_i - \boldsymbol{O}\boldsymbol{q}_{\pi(i)}\right\|^2 \leq \delta^2 \quad \text{and} \quad p_i = q_{\pi(i)} \tag{4}$$

subject to the constraints:

$$\boldsymbol{O}\boldsymbol{O}^{\mathsf{T}} = \boldsymbol{O}^{\mathsf{T}}\boldsymbol{O} = I \quad \text{and} \quad \pi(1) = 1 \tag{5}$$

where δ is the maximum ℓ_2 norm or *distance* between the point-cloud as well as the maximum inter-point separation:

$$\Delta_i = \left\|\boldsymbol{p}_i - \boldsymbol{O}\boldsymbol{q}_{\pi(i)}\right\| \tag{6}$$

The choice of δ controls how similar two LEs must be before they are considered equivalent. This equivalence is represented graphically in Fig. 1.

By design, relabelling a pair of identical points will always result in an equivalent environment. A desirable property, is to ensuring this relabelling results in a corresponding change in π (instead of being absorbed into δ^2) hence, for a useful definition of equivalence, we require:

$$\delta < r_{\text{min}} \tag{7}$$

where r_{min} is the minimum intra-point separation: $r_{ij} = \left\|\boldsymbol{r}_i - \boldsymbol{r}_j\right\|$. This ensures a consistent correspondence between points in equivalent LEs.

Construction of a point-cloud centred on a point (by selecting points within r_{env} from some larger set), is not fully tolerant of point perturbations near the edge of the LE, which could move atoms into/out of the LE. This is acceptable as the LEs will be used for mechanism reconstruction which requires a 1-to-1 correspondence between points. Otherwise, unbalanced equivalence is a natural extension, by simply adding the square of the distance between unmatched points and the boundary of the LE to Eq. (4).

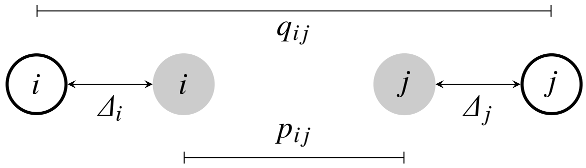

Figure 2 Diagram showing the orientation of two pairs of points i and j in point-clouds P (dark grey) and Q (white) that maximises $\varDelta_i^2 + \varDelta_j^2$, the sum of the square inter-point separations, under the constraint $\varDelta_i^2 + \varDelta_j^2 \leq \delta^2$ from Eq. (4).

3.2.2. *Connection to the potential energy*

An inequality on δ can be established by Taylor-expanding the potential energy, U, about a converged extrema:

$$\begin{aligned}\varDelta U &\approx \varDelta \boldsymbol{x}^{\mathsf{T}} \cancelto{\mathbf{0}}{\boldsymbol{\nabla} U} + \frac{1}{2}\varDelta \boldsymbol{x}^{\mathsf{T}} \boldsymbol{H} \varDelta \boldsymbol{x} \\ &\approx \frac{1}{2}\varDelta \boldsymbol{x}^{\mathsf{T}} \boldsymbol{Q}\boldsymbol{\varLambda}\boldsymbol{Q}^{\mathsf{T}} \varDelta \boldsymbol{x}\end{aligned} \tag{8}$$

where we have applied the eigendecomposition [66, p. 80] to the real symmetric Hessian, $\boldsymbol{H}$, forming $\boldsymbol{\varLambda}$ the diagonal matrix of eigenvalues and $\boldsymbol{Q}$, the orthogonal matrix of eigenvectors. Noting an orthogonal transformation does not change the magnitude of a vector. A (weak) upper-bound on $\varDelta U$ near a minima can be constructed from Eq. (8):

$$\begin{aligned}\varDelta U &\leq \frac{\lambda_{\max}}{2} \left\| \boldsymbol{Q}^{\mathsf{T}} \varDelta \boldsymbol{x} \right\|^2 \\ &\leq \frac{\lambda_{\max}}{2} \left\| \varDelta \boldsymbol{x} \right\|^2 \\ &\leq \frac{\lambda_{\max}}{2} \delta^2\end{aligned} \tag{9}$$

where $\lambda_{\max}$ is the maximum eigenvalue of $\boldsymbol{H}$ and the third line follows from Eq. (4). If two LEs are equivalent and δ is small enough (such that the mechanisms are transferable), then the energy barrier(s) of the reconstructed mechanism(s) should be of-the-order-of $\varDelta U$ off the true energy barrier(s). Ultimately, choosing a smaller value of δ increases the accuracy of the simulation, at the expense of increasing the number of SPS required. Hence, the largest value of δ that ensures Eq. (7) and that $\varDelta U$ is much less than the minimum relevant energy-barrier should be chosen.

3.2.3. *Point-cloud registration*

Simultaneously determining the orthogonal transformation and permutation that minimises the ℓ_2 norm between LEs is a variation of the well-studied rigid point-cloud registration problem [67, 68]. In general, this is not possible at the speeds required by OLKMC. However, we are only interested in finding a specific permutation/transformation that satisfies Eq. (4) and Eq. (5). Here we present a greedy method, that leverages the problem-specific distribution of points.

To find the minimising permutation, we first realise a relationship between the inter-point separations and intra-point separations $p_{ij} = \left\| \boldsymbol{p}_i - \boldsymbol{p}_j \right\|$ and $q_{ij} = \left\| \boldsymbol{q}_i - \boldsymbol{q}_j \right\|$ of the respective point-clouds P and Q. Studying Fig. 2 we see:

$$\left|p_{ij} - q_{ij}\right| \leq \Delta_i + \Delta_j \tag{10}$$

Then maximising $\Delta_i + \Delta_j$, subject to the constraint from Eq. (4), we find our intra-point tolerance criterion:

$$\left|p_{ij} - q_{ij}\right| \leq \sqrt{2}\delta \tag{11}$$

which can be used to match pairs of points in Q to P by recursively ordering Q. At each recursion we search for a point that ensures Eq. (11) holds for all previously ordered points in Q. As this method only requires the intra-point separations in P and Q (which are invariant under Euclidean transformations of P and Q), this method can match the order of the points in LEs that are related by arbitrary rotations/reflections *before* solving for the rotation/reflection.

Once the order of the points in P and Q match we must solve:

$$\min_{\boldsymbol{O} \in M_{3,3}(\mathbb{R})} \sum_{i=1}^{n} \|\boldsymbol{p}_i - \boldsymbol{O}\boldsymbol{q}_i\|^2 \quad \text{s.t.} \quad \boldsymbol{O}\boldsymbol{O}^\top = \boldsymbol{O}^\top\boldsymbol{O} = \boldsymbol{I} \tag{12}$$

this is equivalent to the *orthogonal Procrustes problem* [69] which can be efficiently solved using the singular value decomposition (SVD) see Appendix A.

The full permutation/ordering and equivalence-testing method is detailed in Algorithm 1; once the algorithm has matched the orders and colours it checks if the permutation satisfies Eq. (4). This is required as there may be degenerate permutations satisfying Eq. (11) $\forall i, j$ but not satisfying Eq. (4). A key consideration for the usefulness of Algorithm 1 is its complexity; for two randomly permuted point-clouds, we show in Appendix B (under moderate assumptions) provided:

$$\delta \lesssim \frac{2}{5} r_{\min} \tag{13}$$

the average-case time complexity is $\mathcal{O}\left(n^2\right)$.

3.2.4. Choosing δ values

Typically in the Fe-H system (using the perfect lattice as an order of magnitude estimation) $\lambda_{\max} \approx 10\text{eV Å}^{-2}$. Therefore, according to Eq. (9), we choose $\delta = 0.01$Å resulting in an energy tolerance of approximately $\Delta U \leq 5 \times 10^{-4}$eV. In practice we expect $\Delta U \ll 5 \times 10^{-4}$eV as Eq. (9) assumes $\Delta \boldsymbol{x}$ is parallel to the largest eigenvector of $\boldsymbol{H}$ which is unlikely. We see ΔU is much less than the energy barrier for H diffusion, around 5×10^{-2}eV, typically the fastest mechanism in the Fe-H system.

The choice of δ is continuously validated during a simulation. If δ is too large then, following a mechanism reconstruction, a relaxation of the lattice will result in a large energy change. If/when this is detected δ can be adjusted. Conversely, if no such energy changes are detected δ can be increased to try and increase the performance of the simulation. Furthermore, if we encounter a local environment that breaks Eq. (7) or Eq. (13) then δ can be reduced.

Algorithm 1 Function GREEDY_PERM and its subroutines attempt to permute elements of Q such that it is equivalant to P (Eq. (4) holds); returns `True` if P and Q are equivalent otherwise `False`.

Require: P and Q contain the same number of points, $n > 1$.

```
function GREEDY_PERM(P, Q)
    return _RECUR(P, Q, 2)

function _RECUR(P, Q, i)
    if i > n then
        O ← ROTOR_ONTO(P, Q)                          ▷ See Appendix A
        Δ² ← Σ_{i=1}^{n} ||p_i − O q_i||²
        return Δ² ≤ δ²
    for j ← i, . . . , n do
        if p_i = q_j then
            Swap points i and j in Q
            if _MATCH(P, Q, i) then
                if _RECUR(P, Q, i + 1) then
                    return True
            Swap points i and j in Q
    return False

function _MATCH(P, Q, i)
    for j ← 1, . . . , i − 1 do
        if |p_ij − q_ij| > √2δ then
            return False
    return True
```

3.2.5. Efficiency with δ

The number of SP searches required during an OLKMC simulation is proportional to the number of local environments, N_{env}, this is controlled by δ. Figure 3a shows how N_{env} varies smoothly with δ in a single frame of a simulation while Fig. 3b shows how N_{env} evolves over the course of a simulation for a selection of δ values.

We see in Fig. 3a there are three regimes in N_{env} for each complex. Below $\delta \approx 10^{-3}$Å, N_{env} levels off and converges to approximately $\frac{N}{2}$ for V_1H/V_2H and $\frac{N}{4}$ for the other complexes. This reflects the C_2/D_4 symmetries of the supercell but otherwise classifies every atom as in a different environment. At the upper limit of δ we approach Eq. (13) and environments further from the complex begin to be classified as equivalent. At much smaller values $\left(\delta \ll 10^{-4}\text{Å}\right)$, not included in Fig. 3a, every atom is classified into its own environment. This is because minimiser-convergence then machine precision introduce errors into the atomic positions, splitting the atoms that should be

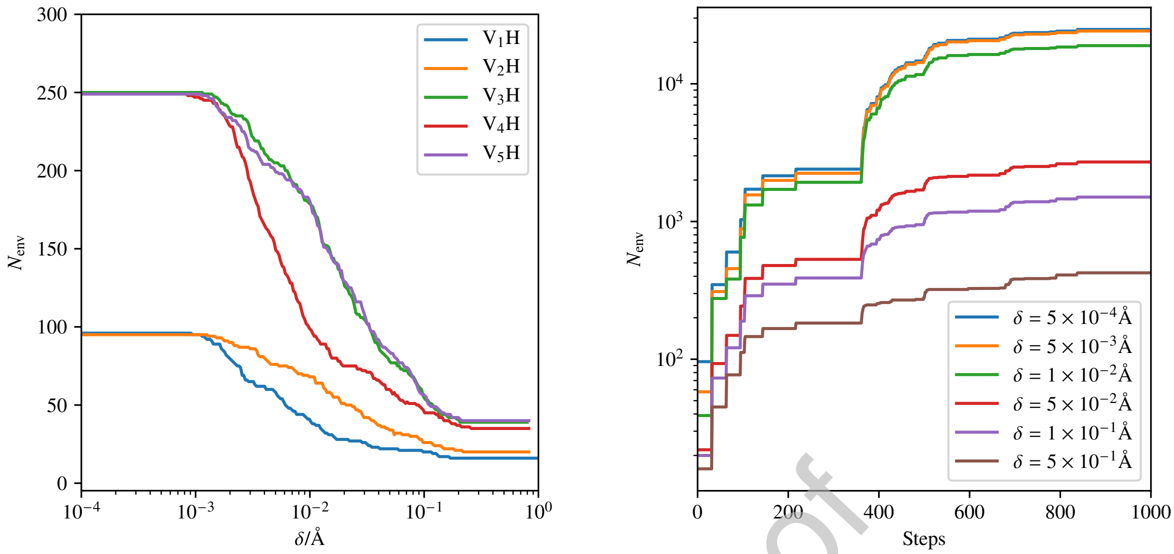

(a) Initial frame (H atom inside cluster). **(b)** Multiple frames/steps of V_1H diffusion.

Figure 3 Plots of the number of local environments, N_{env}, discovered during OLKMC simulations of V_nH diffusing. The 6^3 unit-cell supercell contained $N = 432 - n + 1$ atoms and $r_{\text{env}} = 5.2$Å.

related by symmetry into separate local environments.

Figure 3b show us, for V_1H, even for very small δ, the number of environments eventually goes to a constant. This means no more SP searches will be required after this point. The number of steps to reach this point is also independent of δ. This is a consequence of the tolerance of our method ensuring no equivalent LEs are ever incorrectly classified as distinct. In the $\approx$ 1000 steps it requires to reach this point, an OLKMC simulation with no SP recycling would have effectively encountered approximately half-a-million environments hence, our system which encounters between 400–25 000 (depending on δ) offers a 1–3 order-of-magnitude reduction in the number of SP searches in this region. The sharp increase in LE's near the 400$^{\text{th}}$ step corresponds to the hydrogen escaping the complex. Interestingly, we see a the final value of N_{env} in Fig. 3b is almost identical for $\delta = 5\times10^{-4}$Å and $\delta = 5 \times 10^{-3}$Å despite, N_{env} in Fig. 3a levelling off near $\delta = 10^{-3}$Å.

These results show the advantage of defining a norm-based tolerance method as, the tolerance (δ) can be optimised for each complex systems to minimise the number of LE's.

3.2.6. Heuristics

With the algorithms detailed thus far, a catalogue could be built that satisfies our requirements for LE classification however, although we ensure the condition of Eq. (13), a call to Algorithm 1 still takes $\approx$ 10µs with a LE containing 65 atoms. Hence, as we may call Algorithm 1 for every LE in the catalogue when encountering a new LE, this becomes prohibitively expensive. To reduce the search-space we partition the catalogue into sub-catalogues each indexed by a key, k:

$$k : P \longrightarrow (p_1, \{n^\alpha\}) \tag{14}$$

with n^α the number of points in P of the α^{th} colour. Due to its discrete nature, k can be used as the key to a hash-table (or other suitable key–value store) enabling $\mathcal{O}(1)$ look-up of the sub-catalogues.

The sub-catalogues may still become very large, especially in systems where all points are the same colour. As a simulation progresses, we can sort the order of the LEs in each sub-catalogue by its occurrence count. This significantly decreases the look-up time for a typical LE as many systems have most points in the same LE and only a small number of "active" points (e.g near defects) in rare LEs.

To further accelerate searches of the sub-catalogues we introduce a second discriminator, the *fingerprint* f, a collection of sorted sets/lists:

$$f : P \to \left\{ \{p_{1i} \mid i > 1\}^{\alpha}_{\leq}, \left\{p_{ij} \mid i > 1, j > i\right\}^{\alpha, \beta \leq \alpha}_{\leq} \right\} \tag{15}$$

where p_{ij} denotes the intra-point distances between points i and j and the superscripts α, β indicate the colour of the points in the point pair. For example, in the H-Fe system there are two possible point colours hence, f contains five ordered lists. Two each containing the intra-point distances between points of a particular colour and the central point; a further three lists containing the intra-point distances between pairs of atoms coloured H-H, Fe-H/H-Fe, and Fe-Fe. By construction, f is invariant under Euclidean transformations and permutations of the points in P. Two fingerprints can be compared for equivalence as follows:

Algorithm 2 Compare two fingerprints for equivalence under Eq. (4) subject to the constraints of Eq. (5).

Require: P and Q have matching keys.
  function EQUIV(f_P, f_Q)
    for each pair of ordered lists `p[]`, `q[]` in f_P, f_Q **do**
      for each pair of elements p, q in `p[]`, `q[]` **do**
        if $|p - q| > \sqrt{2}\delta$ **then**
          return `False`
    return `True`

By construction, it is necessary, but not sufficient, for two LEs fingerprints to be equivalent for Algorithm 1 to return `True`. In practice, equivalence of fingerprints is a very strong pre-conditioner for Algorithm 1. As comparison of fingerprints is orders of magnitude faster (typically taking tens of nanoseconds), this substantially accelerates searching the sub-catalogues.

3.2.7. Searching a catalogue of LEs

The full method for classifying a LE, represented by the point-cloud Q, and reconstructing the mechanisms discovered by previous SPS at an equivalent LE proceeds as follows:

Algorithm 3 Search a catalogue of reference LEs for an equivalent LE.

function SEARCH_CATALOGUE(Q)
  $k \leftarrow$ the key of Q
  $s \leftarrow$ the sub-catalogue corresponding to k in the catalogue
  for each P in s **do**
    if EQUIV(P, Q) **and** GREEDY_PERM(P, Q) **then**
      return P ▷ an equivalent LE, whose mechanisms can be reconstructed onto Q by multiplying their atomic displacement-vectors by $\boldsymbol{O}^{\mathsf{T}}$ (computed during GREEDY_PERM)
  return `NULL`

If no match is found then Q represents a new LE; append Q to the sub-catalogue and launch SPS centred on the LE in order to discover any mechanisms associated with it.

3.3. Saddle-point searches

In our implementation, we diverge slightly from the original formulation of the superlinear dimer method [46] during the dimer translation step. We still use the L-BFGS algorithm [70, 71] for determining the translation direction. Ideally a Wolfe condition [72, 73] satisfying line-search could be performed however, no explicit form for the potential that generates the effective force is available. Hence, we introduce a classical trust-radius based approach [74] to limit the step-size.

The maximum step size, s_{trust}, is scaled according to the success of the previous steps; the projection of the effective gradient on the search direction is calculated after a step:

$$P = -\boldsymbol{F}_{\text{eff}}^{\mathsf{T}}\boldsymbol{p} \tag{16}$$

where $\boldsymbol{p}$ is the approximate Newton step, computed using the L-BFGS method, and $\boldsymbol{F}_{\text{eff}}$ is the effective force acting on the dimer. An ideal step length would have $P = 0$. Hence, we increase s_{trust} when $P < -P_{\text{tol}}$ and decrease s_{trust} when $P > P_{\text{tol}}$. Additionally, we bound s_{trust} such that $s_{\text{min}} < s_{\text{trust}} < s_{\text{max}}$.

3.4. Kinetic prefactors

Most OLKMC implementations apply the constant pre-factor approximation, $\tilde{\nu}_{ij} = \nu$, to harmonic transition state theory (HTST) [37, 55] however, Ref. 75 identifies large variations in $\tilde{\nu}_{ij}$ during Al adatom diffusion events. This is to be expected when dealing with heterogeneous systems and evidence to start calculating $\tilde{\nu}_{ij}$, rather than relying on constant approximation. Applying HTST and assuming the PES is quadratic near the SP, we can calculate $\tilde{\nu}_{ij}$ for each mechanism [76, 77]:

$$\tilde{\nu}_{ij} = \frac{\prod_{k=1}^{N} \nu_k^i}{\prod_{k=1}^{N-1} \nu_k^\ddagger} \tag{17}$$

where $\nu_k^\ddagger$, ν_k^i are the real normal-mode frequencies at the SP and state i respectively. This requires computing the full mass-weighted Hessian for each LE. This can be done most efficiently analytically, the procedure is described in the supplementary material.

3.5. Superbasin caching

We elect to further extend the bac-MRM to incorporate superbasin caching. We follow Ref. 61 however, when a superbasin is exited, instead of discarding the superbasin, it is stored into a buffer. The implementation then checks if this new "exit" state is in any of the cached superbasins, if so the corresponding superbasin is loaded from the cache. This is particularly critical as it avoids re-exploring a superbasin that is re-entered immediately after it has been exited. We also dynamically set the tolerance (which the forward and reverse barriers of a mechanism must be less than) for exit state classification. This is achieved by lowering the tolerance when a superbasin gets too large and increasing it when the number of superbasins in the cache exceeds some user-defined threshold.

3.6. Measuring diffusivity

The diffusion coefficient of a collection of atoms over a time t, can be extracted from the mean-squared displacement (MSD):

$$\langle R^2 \rangle = \frac{1}{N} \sum_{\alpha}^{N} \left\| \boldsymbol{r}_{t=0}^{\alpha} - \boldsymbol{r}_{t}^{\alpha} \right\|^2 \tag{18}$$

of the atoms and the Einstein equation [78]:

$$D = \frac{\langle R^2 \rangle}{6t} \tag{19}$$

As we use a single H atom throughout our simulations, we compute the diffusivity of H by dividing the trajectory into intervals and averaging the diffusivity computed in each interval [79, 80]. Due to the complexity of explicitly tracking the position of vacancies during a simulation, we use the MSD of the Fe atoms to calculate the simulated diffusivity, D_{sim}, and effective diffusivity, D_{eff} of individual defects. These are related to the MSD via [81]:

$$D_{\text{sim}} = \frac{\langle R^2 \rangle_{\text{Fe}}}{6t} \quad \text{and} \quad D_{\text{eff}} = \frac{D_{\text{sim}}}{x_d} \tag{20}$$

with x_d the defect concentration. This has the additional benefit of averaging over many Fe atoms.

4. Results and discussion

4.1. Vacancy cluster diffusion

In order to test our implementation and validate our invariant and tolerant local-environment classification, we begin by studying small V_n clusters; we aim to thoroughly classify their diffusion pathways to allow comparisons upon the introduction of H. We construct a vacancy-cluster, V_n, in a (otherwise perfect) 6^3 unit-cell BCC supercell. After a series of mechanisms, the cluster is moved to an equivalent state (just rotated/translated). In combination with the periodic boundary conditions, our norm-based caching recognises this symmetry, reconstructs saddle-points and elides new searches. Hence, SP searches were only required during the initial *learning* phase of the simulation.

The V_n diffusion results are presented in Fig. 4 and summarised in Table 1. The energy profiles for the identified mechanisms are presented in Fig. 5 alongside the mechanisms themselves in Fig. 6.

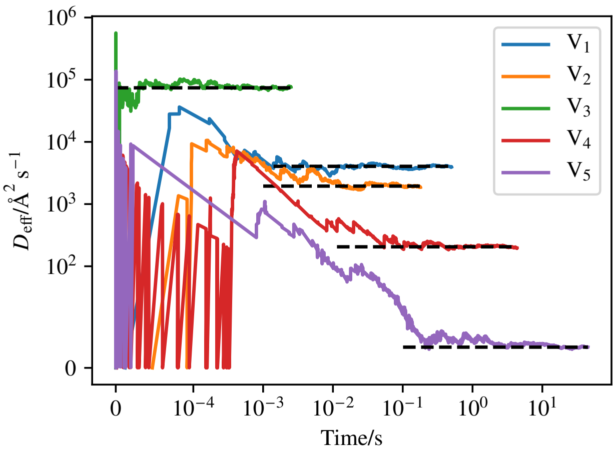

Figure 4 Vacancy clusters diffusing in a perfect 6^3 unit-cell supercell at 300K, dashed lines are fit to a constant.

In Fig. 4, we see convergence to diffusive behaviour for all clusters, verifying that we are reaching diffusive timescales. These timescales are a property of each system and range between 10^{-5}s and 1s. The expected behaviour for cluster diffusivity is for larger clusters to become less mobile. This trend is visible in Table 1 but, the diffusivity for V_3 seems to reverse this trend. We shall now discuss each cluster in detail to fully understand this behaviour.

V_1. The single vacancy diffuses by $\frac{1}{2}\langle 111\rangle$ vacancy hops (a feature common to all the clusters and complexes) with an activation energy of 0.65eV and kinetic pre-factor 7.44×10^{13}Hz. The energy barrier and mechanism are in good agreement with the literature [80, 81].

V_2. The minimum-energy configuration (MEC) for V_2 is the second nearest neighbour (NN) pair followed by the 1st NN then 4th NN orientations, this matches the literature [82]. The predominant V_2 diffusion mechanism was oscillations between the 2nd NN and 4th NN states, with an energy barrier of 0.65eV. The 1st NN pathway may be expected to be the dominant mechanisms, as one may predict the transition to the lower-energy 1st NN state to have a lower energy-barrier. However, the 2nd NN to 1st NN transition has an energy barrier of 0.72eV, making it kinetically less-favourable. The V_2

Table 1 Summary of vacancy-cluster diffusion results in the α–Fe lattice at 300K. All diffusivities have a fractional error less than one part in one hundred. Quoted kinetic pre-factors for multi-step mechanisms is that of the highest barrier step.

Cluster	ΔE/eV	$\nu/10^{13}$Hz	D_{eff}/m^2 s^{-1}
V_1	0.64(9)	7.44	4.05×10^{-17}
V_2	0.65(1)	10.4	2.07×10^{-17}
V_3	0.48(2)	5.22	7.40×10^{-16}
V_4	0.73(4)	3.41	2.05×10^{-18}
V_5	0.77(3)	3.12	9.01×10^{-20}

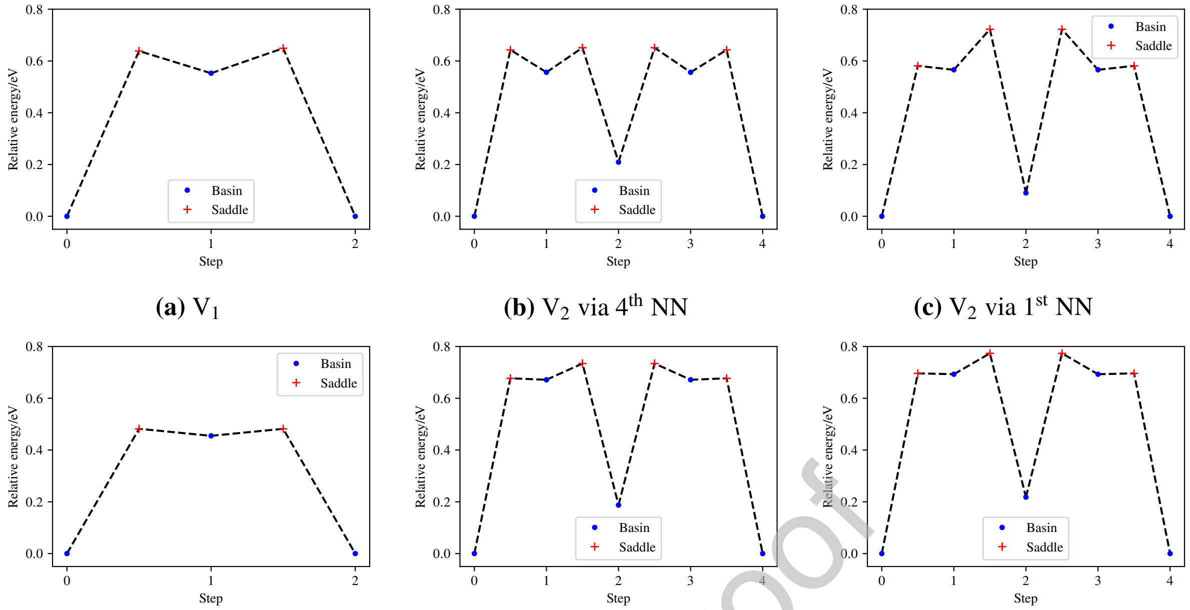

(a) V_1 **(b)** V_2 via 4^{th} NN **(c)** V_2 via 1^{st} NN

(d) V_3 **(e)** V_4 **(f)** V_5

Figure 5 Energy profiles for the vacancy cluster diffusion mechanisms sketched in Fig. 6 – extracted from OLKMC simulations at 300K. Figure 5a has energy barriers: 0.64 and 0.096eV and corresponding kinetic pre-factors: 7.44×10^{13} and 1.18×10^{13}Hz. Figure 5b has energy barriers: 0.64, 0.095, 0.44 and 0.087eV and corresponding kinetic pre-factors: 1.04×10^{14}, 1.06×10^{13}, 7.52×10^{13} and 1.11×10^{13}Hz. Figure 5c has energy barriers: 0.58, 0.16, 0.63 and 0.015eV and corresponding kinetic pre-factors: 2.97×10^{14}, 9.33×10^{12}, 1.16×10^{14} and 5.16×10^{12}Hz. Figure 5d has energy barriers: 0.48 and 0.027eV and corresponding kinetic pre-factors: 5.22×10^{13} and 5.24×10^{12}Hz. Figure 5e has energy barriers: 0.68, 0.063, 0.55 and 0.0056eV and corresponding kinetic pre-factors: 3.41×10^{13}, 7.62×10^{12}, 3.94×10^{13} and 3.01×10^{12}Hz. Figure 5f has energy barriers: 0.70, 0.081, 0.54, 0.37, 0.56 and 0.0034eV and corresponding kinetic pre-factors: 3.12×10^{13}, 7.17×10^{12}, 6.00×10^{13}, 6.96×10^{12}, 4.31×10^{13} and 2.58×10^{12}Hz.

diffusion barrier is very close to the diffusion barrier for V_1, this may be an artefact of the EAM potential used, as *ab initio* studies typically predict an energy barrier 0.05–0.11eV lower [83, 84]. Nevertheless, this agrees with other works that use similar semi-empirical potentials [80] hence, this discrepancy is an artefact of the potential and would be resolved with improved potentials. V_2 has a diffusivity approximately half V_1, this is due to the combination of a near-identical energy barrier but requiring two vacancy-hops to diffuse.

V_3. As previously hinted, V_3 defies the expectation and is the most mobile cluster with a diffusivity more than an order of magnitude higher than V_1 at 300K. This is due to the MEC permitting a vacancy hop with an energy barrier of 0.48eV, that almost immediately reforms the MEC just displaced/rotated. This means, similarly to V_1, V_3 can diffuse without changing its shape. The simulated MEC matches theoretical predictions, as does the mechanism [84] and the elevated diffusivity matches KMC simulations in Fe [85] and similar trends seen in FCC Ni [86].

V_4. The mobility of V_4 resumes the decreasing trend. This is predominantly due to the high energy barrier, 0.73eV, required to break apart the MEC. Reference 84 used DFT to compute the V_4 diffusion mechanism, they obtained a lower energy barrier of 0.48eV however, their mechanism matches ours. This could be due to image interactions introduced by the small 4^3 unit-cell supercell

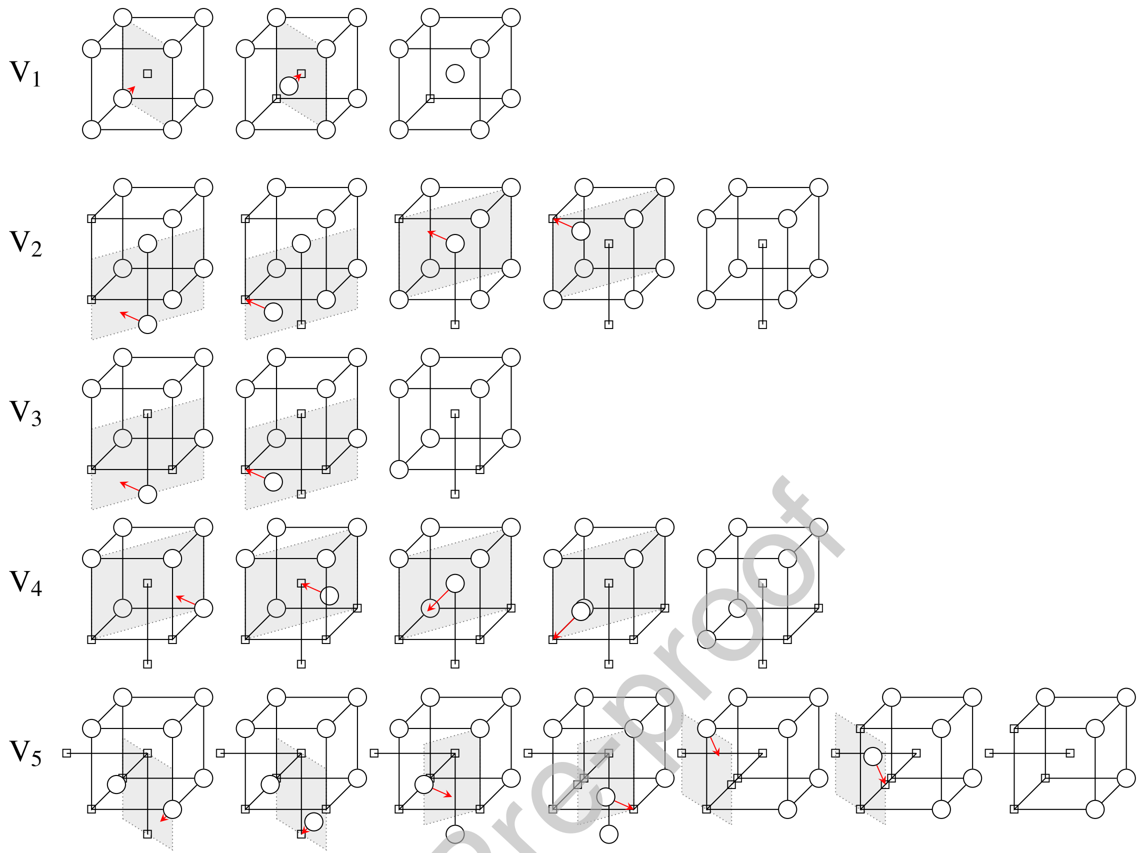

Figure 6 Diffusion mechanisms for vacancy-cluster diffusion in the α–Fe lattice. White circles represent an occupied lattice site; small □ symbols indicate an unoccupied BCC lattice site; arrows mark the path of an atom during a mechanism and transparent grey planes act as a guide to the eye containing the atomic path. Small perturbations away from lattice sites have been omitted for clarity. See Fig. 5 for the corresponding energy profiles. Note for V_2 only the lower energy mechanism (Fig. 5b) has been sketched.

used or indicate work is required on the EAM potential. Nevertheless, the match in mechanism pathway is reassuring and indicates the key physics is being captured by the potential.

V_5. Similarly to V_4, V_5 continues to become less mobile as its size increases. This is again predominantly due to the high energy barrier, now 0.77eV, required to break apart the MEC. Furthermore, additional steps are required to diffuse, increasing the probability of backtracking.

4.1.1. Discussion

Off-lattice KMC has successfully been applied to study the diffusion of vacancy clusters. The mechanisms predicted and diffusivity-trends match those seen in the literature [80, 82, 84, 87]. However, they have all been predicted, without *a priori* assumptions, by the highly general OLKMC framework. This is exemplified by the counter-intuitive diffusion mechanisms of V_2 that could easily be misidentified if using simpler models e.g. final-to-initial-state-energy (FISE) [80]. Although, this could have been captured with traditional KMC and careful DFT analysis, here it arises naturally without any special consideration or bias. The range in kinetic pre-factors, spanning almost two orders of magnitude, emphasises the need to compute pre-factors even in single element systems.

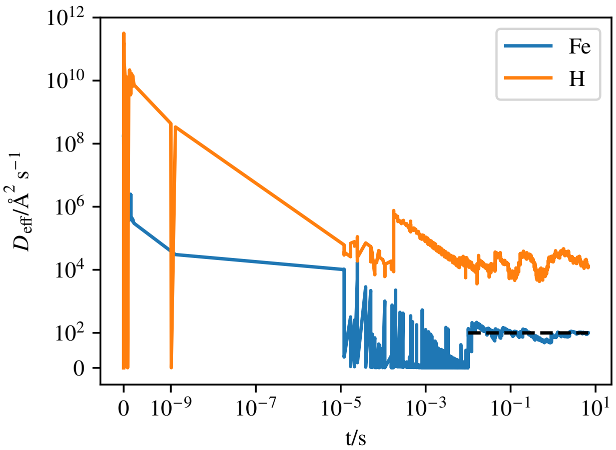

Figure 7 V_2H complex diffusing in a perfect 6^3 unit-cell supercell at 300K, dashed line is a fit to a constant.

A particularly relevant comparisons is with Ref. 87, they also use an OLKMC method to obtain the diffusivity of V_n clusters but using a different potential [88]. They report the diffusivities at 573K hence, a direct comparison is non-trivial. However, we do see good correlation in the trends observed: they see a drop in diffusivity from $V_1 \rightarrow V_2$ followed by an increase $V_2 \rightarrow V_3$ and then decreasing diffusivity for larger clusters. This matches our results and gives us confidence that this trend is not a quirk of the potentials used.

During our experiments at 300K, no dissociation of the vacancy clusters occurred. This confirms the dissociation barrier is higher than the diffusion barrier, which is in line with the literature [80]. Exploring diffusion across a range of (higher) temperatures with OLKMC (as we do for V_1H in Section 4.4) would offer an opportunity to confirm the energy barriers with a fit to the Arrhenius equation and study the dissociation behaviour. If larger clusters become increasingly immobile, dissociation may become the predominant form of diffusion.

4.2. *V_nH complex diffusion*

To build upon the cluster diffusivity results, we add a single H atom into each of the clusters described in Section 4.1, forming V_nH complexes. The cluster acts as a trap for the H atom, which can then detrap, diffuse rapidly through the lattice and re-trap at another (possibly the same) cluster. Alternatively, the complex itself can diffuse. We refer to the H trapping sites within the vacancy clusters as *deep* trapping-sites.

The parallel of Fig. 4 is presented in Fig. 7, for the representative V_2H complex. We see timescales range from 10^{-13}s to 10s, this is required to resolve the distinct regimes visible in Fig. 7. Firstly, below 10^{-9}s, the H atom explores the deep-trapping states within vacancy clusters. Analytical superbasin acceleration kicks in once all the states are explored, solves the flickering problem between the deep trapping-states and triggers the discontinuity to 10^{-5}s as the H atom escapes then rebinds to the cluster. Next, near 10^{-4}s, the H atom detraps and diffuses to another cluster. Finally, near 1×10^{-2}s, the complex displaces and beyond converges to diffuse behaviour. Like the experiments in Section 4.1, SP searches are only required during the learning phase of the simulations.

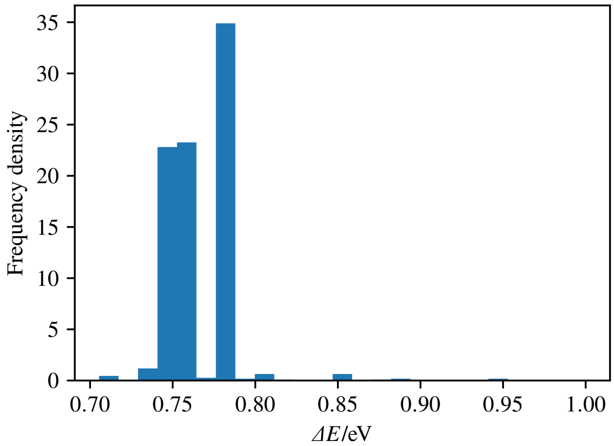

Figure 8 Histogram of energy barriers for V_2H complex diffusion in a perfect 6^3 unit-cell supercell at 300K.

The complex's diffusivities and energy barriers are summarised in Table 2; the energy profiles for the identified mechanisms are presented in Fig. 9 and the mechanisms themselves in Fig. 10. We expect higher energy barriers for small complexes, especially V_1H, due to the larger relative steric hindrance provided by the H atom in the smaller clusters. As the clusters get larger, one would predict that the 'regular' H-free mechanisms could occur whilst the H atom occupies the opposite side of the cluster, interacting minimally. Hence, larger complexes are expected to diffuse similarly to their corresponding clusters. Interestingly, we see for the larger clusters investigated that the introduction of H lowers the energy barrier, indicating a different mechanism may be occurring. We shall now discuss each complex in detail to fully understand this behaviour.

V_1H. As predicted by the steric argument, V_1H experiences the largest increase in diffusion barrier. The H atom can occupy 6 deep trapping-sites in the vacancy, close to the adjacent octahedral sites. The H atom must be 'pushed' out of the vacancy by an Fe atom during a $\frac{1}{2}\langle 111\rangle$ hop which increases the energy barrier, demonstrating co-dependence. The partially-escaped H atom then rebinds with the newly formed vacancy. This mechanism is asymmetric; from the perspective of the reverse direction, the H atom first escapes and then pushes a Fe atom into the vacancy. The energy barrier, energy profile and mechanism closely match the EAM results of Ref. 81, validating that

Table 2 Summary of cluster-H diffusion results in the α–Fe lattice at 300K. All diffusivities have a fractional error less than one part in one hundred. Quoted kinetic pre-factors for multi-step mechanisms is that of the highest barrier step.

Complex	ΔE/eV	$\nu/10^{13}$Hz	$D_{\mathrm{eff}}/\mathrm{m}^2\,\mathrm{s}^{-1}$
V_1H	0.75(8)	23.3	3.72×10^{-19}
V_2H	0.70(5)	8.70	1.00×10^{-18}
V_3H	0.49(1)	9.22	7.50×10^{-16}
V_4H	0.71(2)	7.86	3.95×10^{-18}
V_5H	0.74(8)	7.33	5.16×10^{-20}

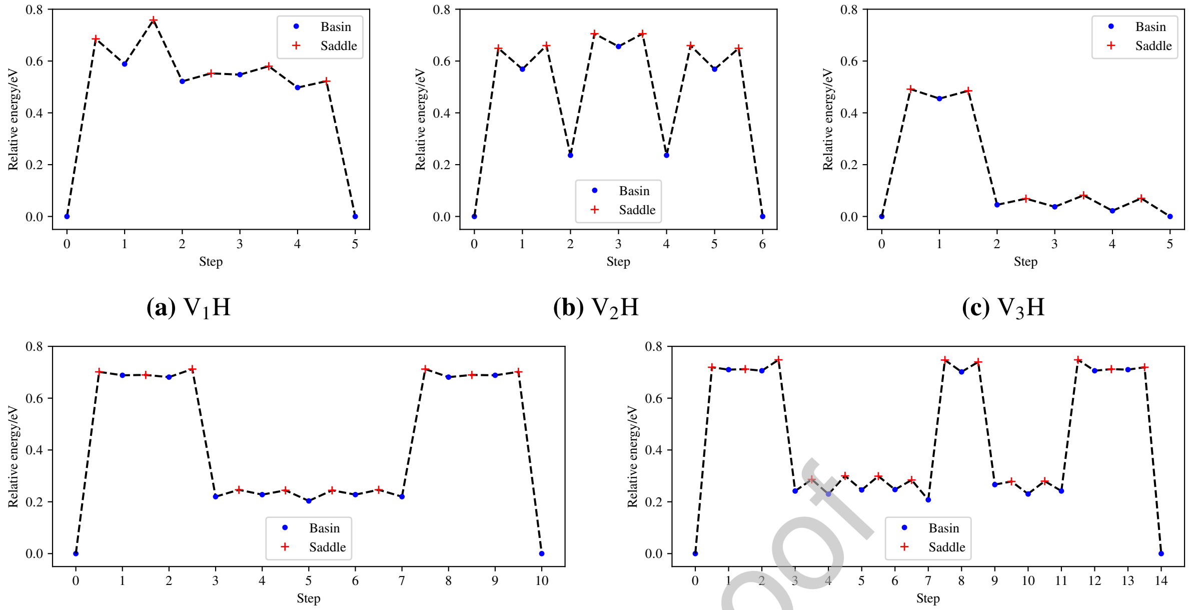

(a) V_1H (b) V_2H (c) V_3H

(d) V_4H (e) V_5H

Figure 9 Energy profiles for the cluster-H diffusion mechanisms sketched in Fig. 10 – extracted from OLKMC simulations at 300K. The individual barriers and kinetic pre-factors are omitted for brevity.

our norm-based classification generalises to small interstitials and heterogeneous systems. In the literature it is suggested that the reverse direction is the preferred mechanism [23, 81] however, in our experiments this only occurs ≈ 30% of the time. This could be due to a kinetic bias towards the forward direction at 300K.

V_2H. The H atom can occupy 14 deep trapping-sites in the V_2 2nd NN cluster, again close to the surrounding octahedral sites. The larger V_2H complex has a diffusion mechanism more similar to the H free case than V_1H; the vacancies first move into a 4th NN coordination leaving the H atom in the original vacancy, the H atom then jumps out and into the moved vacancy. Finally, the original vacancy hops and reforms a displaced 2nd NN complex. Less interaction between the H atom and moving Fe atom(s) results in a lower total energy barrier of 0.71eV, compared to V_1H. This explains the corresponding ten-fold increase in diffusivity compared to V_1H. While for V_1H only one diffusive mechanism occurred, for V_2H a group of diffusive mechanisms, with energy barriers 0.71–0.95eV were found. Figure 8 depicts the distribution of diffusion-mechanism energy-barriers identified. The most prolific mechanism at 300K had an energy barrier of 0.78eV. The minimum barrier mechanism was probably suppressed by the high likelihood of the vacancy hopping back to the initial configuration before the H atom could hop to the new vacancy. In contrast the 0.78eV mechanism progressed through a 1st NN intermediate, which is equally likely to move to a new state or backtrack as the H atom can freely move between the two vacancies. This highlights the significance of kinetics vs energetics at room temperature.

V_3H. The introduction of H to V_3 has little effect on the diffusion mechanisms, only fractionally increasing the exceptionally low-barrier V_3 mechanism to 0.49eV. This results in V_3H not disassociating for the entirety of the simulation. The V_3 cluster provides 18 deep trapping-sites for H.

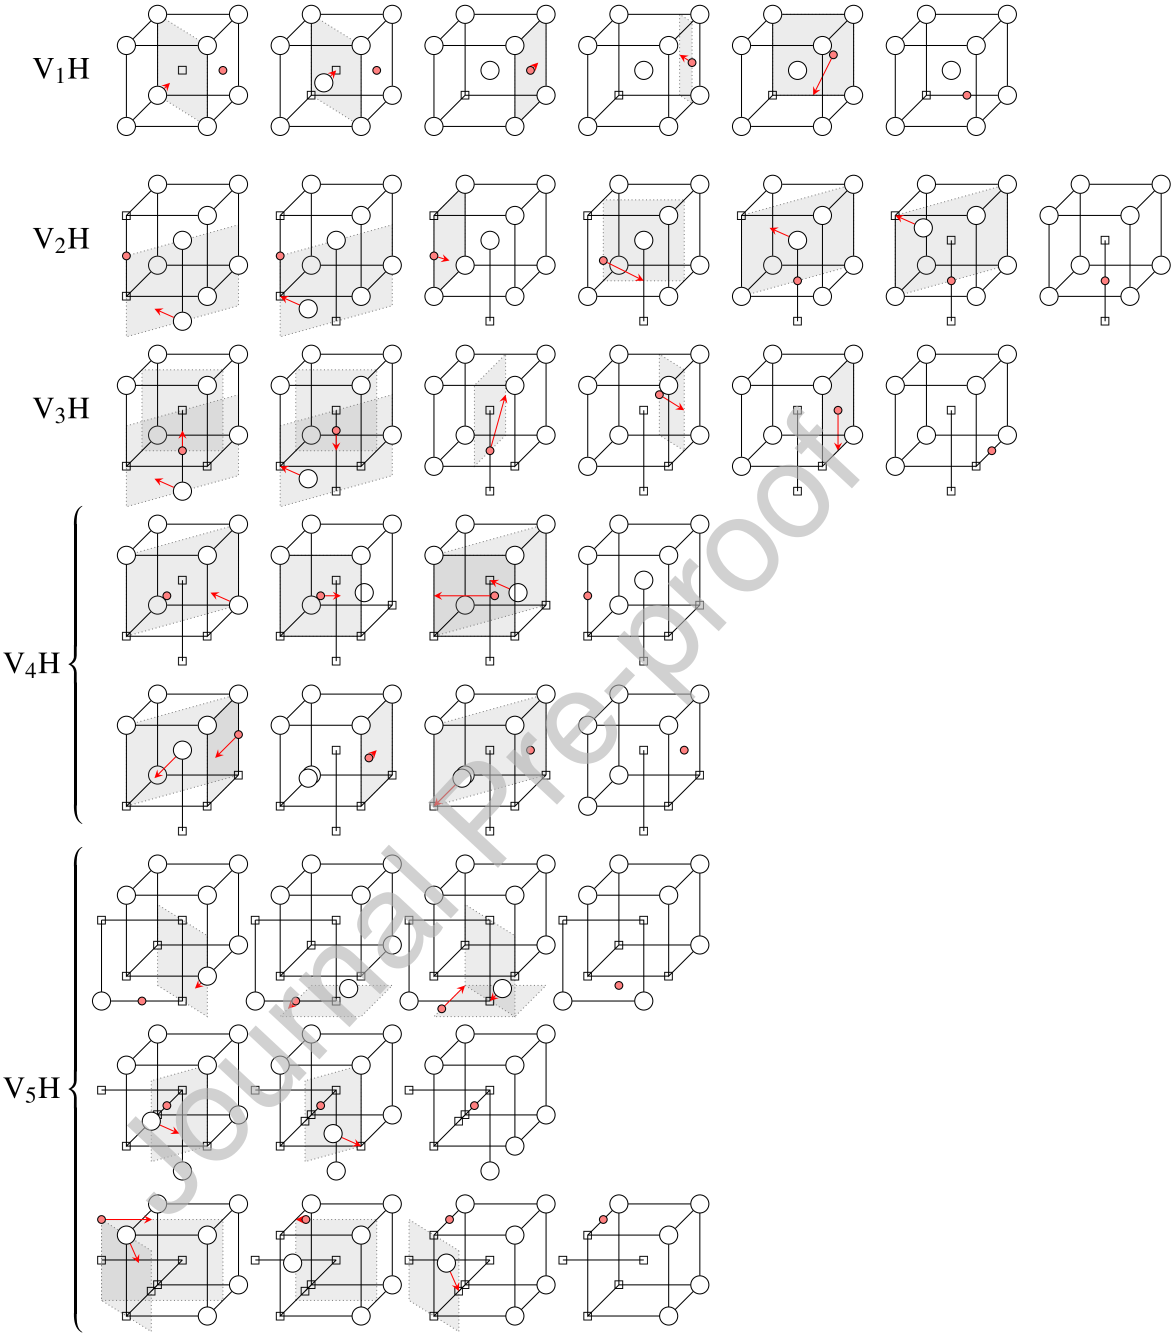

Figure 10 Diffusion mechanisms for V_nH_1 complexes in the α–Fe lattice. Small red circles mark a H atom; white circles represent an occupied lattice site; small □ symbols indicate an unoccupied BCC lattice site; arrows mark the path of an atom during a mechanism and transparent grey planes act as a guide to the eye containing the atomic path. Small perturbations away from lattice/octahedral/tetrahedral sites have been omitted for clarity. See Fig. 9 for the corresponding energy profiles. Note for V_4H only the frames corresponding to steps 0–3 and 7–10 in Fig. 9d are shown and similarly for V_5H with steps 0–2, 7–9 and 18–20 from Fig. 9e.

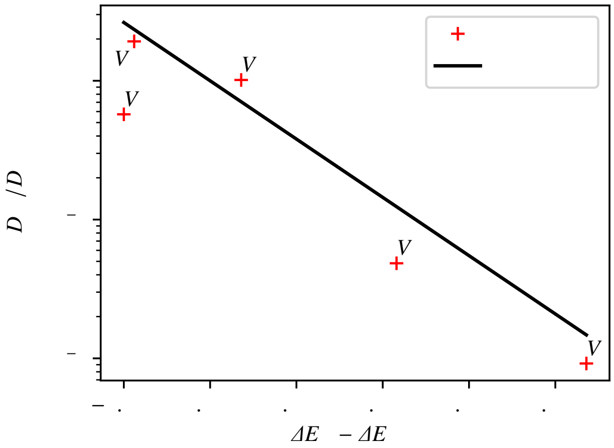

Figure 11 Comparison between V_n and V_nH diffusion barriers vs their corresponding effective diffusivities – expectation is Eq. (21) where we assume equal Arrhenius pre-factors.

Similarly to V_2H, a range of diffusion mechanisms were identified with energy barriers 0.49–0.57eV. This time the minimum energy mechanism accounted for a substantial fraction of the observed mechanisms but, was still not the most frequent.

V_4H. The V_4 cluster provides 20–32 deep trapping-sites for H (depending on the cluster configuration). For the first time, introduction of an H atom decreases the energy barrier and increases the diffusivity of the complex, compared to the cluster alone. Studying Fig. 6 and Fig. 10, we see the motion of Fe atoms remains the same as V_4. Although the H atom has extra space in the large cluster to 'avoid' the hopping Fe atom, as predicted by the steric argument, instead it remains close and lowers the energy of the saddle-point(s). This effect can be likened to the pathway provided by the partially-escaped H atom pushing a Fe atom in V_1H. However, due to the increased size of the cluster and its shape, the H atom does not need to escape before providing this push. Furthermore, the intermediate Fe configuration (steps 3–7 in Fig. 9d) is connected, meaning the H atom can easily reach all the deep traps in the cluster (unlike V_2H) hence, the forwards and reverse directions are equally likely.

V_5H. The V_5 cluster provides 28–36 deep trapping-sites for H and continues the V_4H trend of decreasing the energy barrier compared to its corresponding cluster. In Fig. 10 we see (again contrary to the steric hypothesis) the H atom remains close to the hopping Fe atom, with the first and last hops very similar to the corresponding steps of the V_4H mechanism. Despite the lower energy barrier, V_5H diffusivity is reduced compared to V_5. This can be attributed to the additional complexity of the mechanism (visible in Fig. 9e) and additional backtracking, all of which reinforces the importance of kinetic effects at room temperature. To the authors knowledge, this is the first-time mechanisms for V_5H have been reported.

4.2.1. Discussion

As we move from V_n clusters to V_nH complexes, we may expect the diffusivities to scale with the energy barriers according to the classical Arrhenius behaviour:

$$\frac{D^{H}_{eff}}{D_{eff}} \simeq e^{-\beta\left(\Delta E^{H}-\Delta E\right)} \tag{21}$$

with $\beta = \frac{1}{k_B T}$ and H superscripts denoting the complexes. This assumes both the cluster and the complex have the same Arrhenius pre-factor, this assumption can be motivated by the equivalent pathways of Fe atoms during diffusion in Fig. 6 and Fig. 10. The comparison is drawn in Fig. 11, where we see a general conformance to expectation. We see the largest deviations for V_2H and V_5H, these complexes require three high-barrier steps during their diffusion mechanism hence, these deviations are probably due to the increased likelihood of backtracking.

For all the complexes beyond V_1H, a number of alternative mechanisms were accessible at 300K. Off-lattice KMC was able to discover these on-the-fly; traditional KMC simulation that use (typically small) pre-determined lists of mechanisms could easily omit mechanisms that contribute to interesting behaviour. Capturing only the lowest energy barrier mechanisms may not be sufficient. Figure 8 exemplifies this for V_2H, which is the simplest example of this increasing complexity. This effect will become worse at higher temperatures where alternate, higher-energy mechanisms become ever more common.

Similarly to Section 4.1, the vacancies in the complexes did not dissociate during the simulation. This suggests the dissociation barrier remains higher than the complex's diffusion barriers but, sheds little light on the effect of H on the dissociation barrier. Extending the simulations across a range of temperatures (as we do for V_1H in Section 4.4) would offer an opportunity to confirm the energy barriers with a fit to the Arrhenius equation and (through comparisons with the H free case) study the effect of H on the dissociation barrier.

It has also been found that the interaction between H and vacancy-clusters has not followed the pattern if extrapolating from V_1H; perhaps the most interesting behaviour is the reduction in diffusion barrier, upon introduction of H into the larger clusters V_4H and V_5H. A similar phenomenon, dubbed the *hydrogen lubrication effect*, has been explored computationally in FCC metals [89]. If this trend continues and the gap continues to increase, it could help to explain the experimentally observed H-induced nano-void migration [90] and could have implications for HE by contributing towards vacancy agglomeration during the HESIV mechanism [14, 91, 92]. A similar atomic mechanism that lowers the diffusion barrier for Vacancy-H complexes could be responsible for the predicted increase in dislocation velocity/mobility [93, 94], which could directly support the HELP mechanism [9] of HE. As ever, we should be wary of extrapolating; nano-sized cluster and dislocations could be simulated with OLKMC and would need to be studied before drawing conclusions about macroscopic phenomena.

4.3. Effective hydrogen diffusivity vs classical approaches

Alongside the diffusion of the complexes in Section 4.2, we also investigate the motion of the H atom between the network of defects. To explore these results we make comparisons with Oriani's theory of equilibrium-trapping [95], which predicts the effective diffusivity of H, D_{Or}, is related to

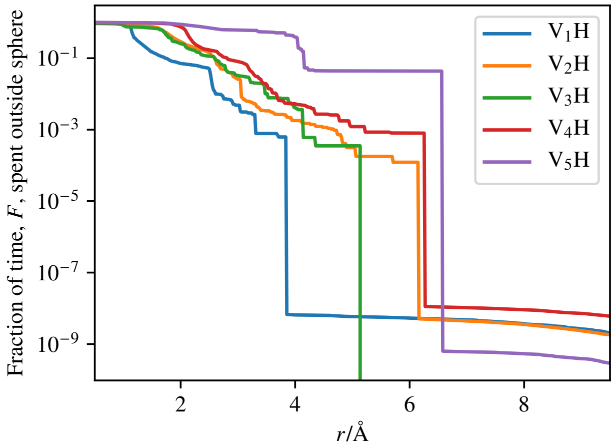

Figure 12 Fraction of the total-time the H atom spent outside a sphere of radius r centred on the vacancy cluster at 300K.

the diffusivity in the perfect lattice, D, via [43]:

$$D_{\text{Or}} = D\frac{n_L\theta_L}{n_L\theta_L + n_x\theta_x(1-\theta_x)} \tag{22}$$

where θ_x, θ_L are the fractional occupancy of available point-trap and regular lattice sites respectively, and n_x, n_L the corresponding number of sites. Oriani's theory assumes the traps are stationary hence, D_{Or} only takes account for H displacement while the H atom moves between traps.

To obtain the average of θ_x and θ_L over the course of the entire simulation we identify the locations of the deep trapping-sites in each frame and determine their centroid. Then, defining r as the distance from the H atom to the centroid, we partition the frames into two sets: trapped frames with the $r < r_x$, a *cut off* trapping radius, and lattice frames with $r \geq r_x$. We now discuss how r_x can be chosen and give us information about the size of the defect. Associating trapping sites to the location of the H atom at trapped frames and similarly for lattice sites/frames, the total time spent in each partition, t_x and t_L, are related directly to θ and n via:

$$n_x\theta_x = \frac{t_x}{t_x + t_L} \quad \text{and} \quad n_L\theta_L = \frac{t_L}{t_x + t_L} \tag{23}$$

and we assume n_x and n_L are approximately constant. As we expect the H atom to spend most of its time bound to the vacancy cluster, we apply $t_L \ll t_x$ to obtain:

$$D_{\text{Or}} = \frac{t_L}{t_x}\frac{D}{1-\frac{1}{n_x}} \tag{24}$$

In order to determine the appropriate value of r_x for each complex, we plot the fraction of the time, F, the H atom spent outside a sphere of radius r centred on the vacancy cluster, as a function of r. The results are presented in Fig. 12: for small r, as we expect, $F \to 1$ as no trap/lattice sites are inside the sphere. Conversely, at large r, F levels-out as all the states outside the sphere are linked by approximately-equal low-barrier mechanisms and the H atom must diffuse an approximately

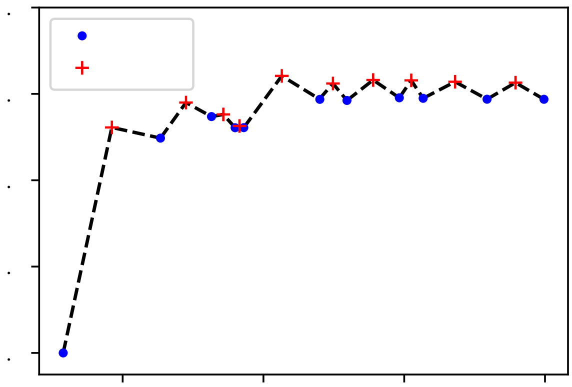

Figure 13 Energy profile of the V_1H dissociation pathway as a function of the distance of the H atom from the centre of the vacancy.

constant (but slowly decreasing hence the trail off in Fig. 12) distance before rebinding. Between the two regions is a sharp discontinuity which marks the bounding sphere that contains all the trapping sites hence, this discontinuity defines r_x, an effective size for each defect or *trapping atmosphere*. This is recorded in Table 3 alongside the diffusivity of the H atom and the relevant energy barriers.

Oriani's theory was postulated assuming traps as single points in the lattice interacting with H. In the present work, we are able not only to predict D_{eff} but also calculate the effective trapping-distance or *atmosphere* of traps with arbitrary structure i.e. non-point traps.

4.3.1. Discussion

As expected, in Table 3 we see increasing the number of vacancies leads to a higher r_x; V_2H has a larger jump in r_x than the rest of the clusters, this is due to extended size of the 4$^{\text{th}}$ NN intermediate state. All of the clusters have a larger effective size than the cluster radius would suggest. This is well demonstrated by V_1H, where we see r_x = 3.9Å is between the 2$^{\text{nd}}$ and 3$^{\text{rd}}$ NN distances, indicating the tetrahedral sites close to the vacancy also act as a weak trap for H. If we examine the energy profile of the dissociation mechanism in Fig. 13, we see a collection of

Table 3 Summary of V_nH binding energies and detrapping barriers in the α–Fe lattice extracted from OLKMC simulations at 300K. Additionally, the diffusivity of the H atom is included alongside the trapping radii and Oriani diffusivities. Note: V_3H did not detrap during the simulation hence, the corresponding values are omitted.

Complex	r_x/Å	ΔE_{detrap}/eV	$\Delta E - \Delta E_{\text{detrap}}$/eV	E_B/eV	D_{H}/m^2 s^{-1}	D_{Or}/m^2 s^{-1}
V_1H	3.9	0.63(3)	0.13	0.59	$1.57 \pm 0.21 \times 10^{-16}$	8.47×10^{-17}
V_2H	6.2	0.66(7)	0.04	0.62	$1.25 \pm 0.06 \times 10^{-16}$	5.69×10^{-17}
V_3H	5.2	-	-	-	$1.24 \pm 0.11 \times 10^{-15}$	-
V_4H	6.3	0.70(3)	0.01	0.65	$4.88 \pm 0.80 \times 10^{-16}$	1.22×10^{-16}
V_5H	6.6	0.70(8)	0.04	0.66	$4.01 \pm 0.27 \times 10^{-17}$	6.84×10^{-18}

metastable states just below $r = 3.9$Å. These metastable states join the vacancy's superbasin; as the state-to-state dynamics are not preserved inside superbasins the time fractions at $r < r_x$ in Fig. 12 do not correspond to the Boltzmann distribution. Understanding the depth and density of hydrogen's weak-trapping/metastable states introduced by defects is of key importance to designing HE resistant steels that utilise H traps. Our results show that even the simplest defects have a complex secondary-structure of surrounding metastable states. Furthermore, the large effective size of the defects means they will alter the diffusive cross section to a greater extent than would be expected for point defects and such variations are not monotonic with the defect complexity (i.e. the number of vacancies).

The detrapping barriers and binding energies in Table 3 increase with the number of vacancies and the two are separated by ≈ 0.047eV for each complex. Therefore, larger clusters act as deeper traps, this trend should level-out as the inside of the cluster approaches a surface.

The diffusivities of the H atom follow a more complex trend. In general, D_{H} and D_{Or} fall in accordance with $\Delta E_{\mathrm{detrap}}$ and ΔE_B rising but, D_{Or} under-predicts the diffusivity. This is because D_{Or} assumes the traps are immobile while in reality the energy barrier for the complex diffusing can be very close to the detrapping barrier. This would present a minor contribution to the diffusivity of the H atom if the mean free path (MFP) in the lattice was large, as the diffusion in the lattice is very fast. However due to the high defect-concentration limiting the MFP, the two effects are competing. Furthermore, D_{Or} becomes a worse estimator of D_{H} as the defect size increases, this is because the approximations made by Oriani (point trapping, single trapping barrier and no change in diffusive cross-section) become less true as the defect's size increases. The trend is broken by V_4H, despite having a higher detrapping barrier, the H atom diffuses faster than in V_1H and V_2H. This could be partially explained by fact that V_4H has the smallest gap between ΔE of the complex and $\Delta E_{\mathrm{detrap}}$ thus, complex diffusion is contributing more to D_{H}. The corresponding rise in D_{Or} (which should not take into account complex diffusion) could then be attributed to the H atom not having enough time to reach equilibrium between the lattice/defect before the defect diffuses again. Nevertheless, the detrapping barriers for V_1H and V_2H are still lower than the complex-diffusion barrier of V_4H; perhaps the difference is made up by the kinetic pre-factors of the rate limiting step(s).

Interpolating Table 3 we would expect the detrapping barrier for V_3H to be 0.67–0.70eV, this is much higher than the diffusion barrier, which explains the lack of detrapping. If we compare the H diffusivities between the complexes, we see that V_3H complex-diffusion offers a pathway for H transport that is about an order of magnitude faster at 300K compared to trapping and detrapping from the other complexes. This is partly due to the aforementioned high defect-concentration – due to the periodic nature of the system – limiting the MFP of lattice-diffusing H. Nevertheless, this unexpected result could explain elevated H transport in some scenarios. Furthermore, with the energy barrier for H transport via V_3H complex-diffusion being so much lower, this could be the predominant H transport mechanism at lower temperatures.

4.4. Variation with temperature

In order to make comparison of our predictions with molecular dynamics, which require higher temperatures to reach diffusive timescales, we extend the V_1H simulations from Section 4.2 across a range of temperatures. Focusing on V_1H avoids entanglement of diffusion with dissociation effects for multi-vacancy complexes, to better illustrate the agreement with other results. The results are

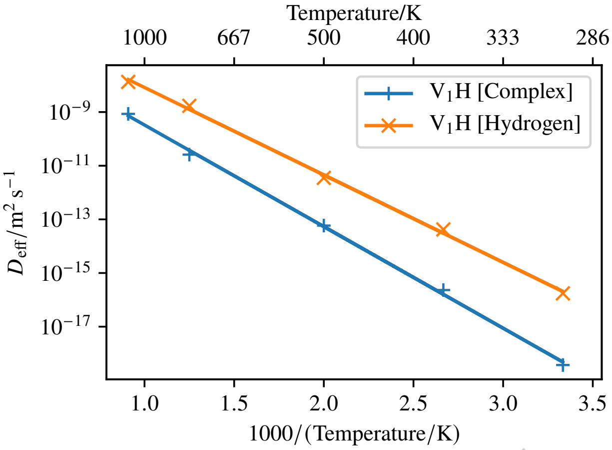

Figure 14 Arrhenius plot of the effective diffusivity of the V_1H complex diffusing in an otherwise perfect 6^3 unit-cell supercell. The blue data corresponds to the effective diffusivity, D_{eff}, of the complex computed from the motion of the Fe atoms while the orange data corresponds to the diffusivity of the H atom, D_{H}.

presented in Fig. 14, where we have fit to an Arrhenius curve:

$$D_{\mathrm{eff}} = D_0 e^{-\beta \Delta E} \tag{25}$$

For the complex we obtain $D_0 = 2.0 \pm 0.8 \times 10^{-6}\mathrm{m}^2\,\mathrm{s}^{-1}$ and $\Delta E = 0.75 \pm 0.02\mathrm{eV}$ and for H we obtain $D_0 = 1.5 \pm 0.6 \times 10^{-5}\mathrm{m}^2\,\mathrm{s}^{-1}$ and $\Delta E = 0.64 \pm 0.01\mathrm{eV}$. During the simulation, even at the highest temperatures, the high H-mobility and effective concentration ensure that the H always rebinds to a vacancy before the vacancy has an opportunity to diffuse independently.

The Arrhenius-fitted energy-barrier for V_1H diffusion is within-error of our direct measurement in Table 2. Furthermore, we see the effective energy-barrier for hydrogen diffusion from the Arrhenius plot is within-error of our measurement of the detrapping barrier in Table 3. This confirms, under these conditions, the diffusivity of the hydrogen is entirely mediated by detrapping events.

Hydrogen diffusion in Fe is separated into two regimes: H has a very low mass and large de Broglie extension hence, at low temperatures, H diffuses via a quantum tunnelling mediated mechanism [26, 27, 96]; at higher temperatures, classical diffusion dominates. Therefore, to ensure we are in the classical region we should choose a high temperature. Unfortunately, the barrier for H diffusion in the lattice is relativity low, as $k_B T$ approaches this barrier the dynamic pressure effects make HTST a poor description of the system – the HTST approximation may begin to fail. This is a fundamental limitation to all KMC simulations which effectively quench the system to 0K at each step. Therefore, we should choose a low temperature to ensure the validity of the HTST approximation.

These two temperature requirements are opposing so, in Section 4.2, we choose the compromising temperature of 300K. This is close to room temperature, relevant for H embrittlement and used in other KMC studies [43]. To verify our temperature is sufficiently low, such as that the HTST approximation is a good description of the system we must make comparisons to a model that does not make this assumption. The ideal comparison is provided by Ref. 81, they simulate the diffusivity of V_1H complexes using MD and parallel replica dynamics (PRD) and obtain $D_{\text{eff}}^{800\text{K}} = 1.7\text{–}2.0 \times 10^{-11}\,\text{m}^2\,\text{s}^{-1}$. We can make a direct comparison by interpolating the results in Section 4.4, we obtain: $D_{\text{eff}}^{800\text{K}} = 3.8 \pm 2.3 \times 10^{-11}\,\text{m}^2\,\text{s}^{-1}$. This is within-error of the MD/PRD results and gives us good confidence that the HTST approximation is still valid at 0K and the results are not very different from reality.

In summary, we demonstrate through direct quantitative-comparison with other methods that our results are not significantly distorted by the HTST approximation of low-barrier events.

5. Conclusions

We have developed and implemented a tolerant norm-based LE classification method for OLKMC, that is invariant under Euclidean-transformations and permutations of atoms. This has enabled the simulation of the Fe-H system into HE timescales at room temperature, predicting features not previously reported using a single modelling framework. The introduction of hydrogen has produced a system that presents many challenges to model with OLKMC: small interstitials, multi-stage mechanisms, frequent flickering-problems, varied harmonic pre-factors and sensitive energy-barriers. Nevertheless, these have been overcome and OLKMC has proved an invaluable and capable tool for the study of Fe-H with no *a priori* assumptions of the underlying mechanisms. This is extremely promising for the future of modelling H-defect interactions and improving our understanding of HE at the atomic scale. Specifically, we have:

- Investigated the diffusion of small (less than six) vacancy-clusters, with and without the addition of H and found evidence that H can increase the diffusivity of larger clusters.
- Thoroughly classified the diffusion pathways of these cluster/complexes (energetically and mechanistically) and understood how H changes their diffusion mechanisms.
- Simultaneously, obtained the trapping/detrapping barrier(s) of H from – and its effective diffusivity in – the presence of these clusters. We have made comparisons to Oriani's theory, testing the equilibrium hypothesis in the presence of mobile traps and expanded the conclusions to also include predicting trapping atmospheres in arbitrarily defined non-point traps.
- Quantified the trapping atmospheres surrounding vacancy clusters and begun to demonstrate the kinetic effects of shallow traps surrounding point-defects.
- Found examples of harmonic pre-factors varying by two orders-of-magnitude, reinforcing that the constant pre-factor approximation should always be carefully verified (particularly in multi-element systems).

Finally, OLKMC is a materials-agnostic method. Our norm-based classification should find applications in many materials systems requiring long-term atomistic predictions.

Code availability

The software used to perform the simulations in this work, including an implementation of the norm-based classification developed, is open-source[3] or available upon request.

[3] https://github.com/ConorWilliams/onthefly

Acknowledgements

We gratefully acknowledge the funding received from the EPSRC via the CDT in Computational Methods for Materials Science (grant number EP/L15 552/1) and grant EP/T8687/2. E. I. Galindo-Nava acknowledges the Royal Academy of Engineering for his research fellowship funding. We also acknowledge Rolls-Royce Plc for the provision of funding. All information and foreground intellectual property generated by this research work is the property of Rolls-Royce Plc.

References

[1] W. H. JOHNSON, On some remarkable changes produced in iron and steel by the action of hydrogen and acids, Nature 11 (1875) 393–393. doi:10.1038/011393a0.

[2] O. Reynolds, On the effect of acid on the interior of iron wire, Journal of the Franklin Institute 99 (1875) 70–72. doi:10.1016/0016-0032(75)90215-x.

[3] O. Barrera, D. Bombač, Y. Chen, T. D. Daff, E. Galindo-Nava, P. Gong, D. Haley, R. Horton, I. Katzarov, J. R. Kermode, C. Liverani, M. Stopher, F. Sweeney, Understanding and mitigating hydrogen embrittlement of steels: a review of experimental, modelling and design progress from atomistic to continuum, Journal of Materials Science 53 (2018) 6251–6290. doi:10.1007/s10853-017-1978-5.

[4] F. F. Dear, G. C. G. Skinner, Mechanisms of hydrogen embrittlement in steels: discussion, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 375 (2017) 20170032. doi:10.1098/rsta.2017.0032.

[5] L. B. Pfeil, The effect of occluded hydrogen on the tensile strength of iron, Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character 112 (1926) 182–195. doi:10.1098/rspa.1926.0103.

[6] H. Vehoff, P. Neumann, Crack propagation and cleavage initiation in fe-2.6%-si single crystals under controlled plastic crack tip opening rate in various gaseous environments, Acta Metallurgica 28 (1980) 265–272. doi:10.1016/0001-6160(80)90161-3.

[7] W. W. Gerberich, R. A. Oriani, M.-J. Lji, X. Chen, T. Foecke, The necessity of both plasticity and brittleness in the fracture thresholds of iron, Philosophical Magazine A 63 (1991) 363–376. doi:10.1080/01418619108204854.

[8] S. Lynch, Hydrogen embrittlement and liquid-metal embrittlement in nickel single crystals, Scripta Metallurgica 13 (1979) 1051–1056. doi:10.1016/0036-9748(79)90202-3.

[9] H. Birnbaum, P. Sofronis, Hydrogen-enhanced localized plasticity—a mechanism for hydrogen-related fracture, Materials Science and Engineering: A 176 (1994) 191–202. doi:10.1016/0921-5093(94)90975-x.

[10] P. Ferreira, I. Robertson, H. Birnbaum, Hydrogen effects on the interaction between dislocations, Acta Materialia 46 (1998) 1749–1757. doi:10.1016/s1359-6454(97)00349-2.

[11] G. Lu, Q. Zhang, N. Kioussis, E. Kaxiras, Hydrogen-enhanced local plasticity in aluminum: AnAb InitioStudy, Physical Review Letters 87 (2001). doi:10.1103/physrevlett.87.095501.

[12] I. Robertson, The effect of hydrogen on dislocation dynamics, Engineering Fracture Mechanics 68 (2001) 671–692. doi:10.1016/s0013-7944(01)00011-x.

[13] S. Bak, M. Abro, D. Lee, Effect of hydrogen and strain-induced martensite on mechanical properties of AISI 304 stainless steel, Metals 6 (2016) 169. doi:10.3390/met6070169.

[14] M. Nagumo, Hydrogen related failure of steels – a new aspect, Materials Science and Technology 20 (2004) 940–950. doi:10.1179/026708304225019687.

[15] R. McLellan, Z. Xu, Hydrogen-induced vacancies in the iron lattice, Scripta Materialia 36 (1997) 1201–1205. doi:10.1016/s1359-6462(97)00015-8.

[16] R. Nazarov, T. Hickel, J. Neugebauer, Ab initiostudy of h-vacancy interactions in fcc metals: Implications for the formation of superabundant vacancies, Physical Review B 89 (2014). doi:10.1103/physrevb.89.144108.

[17] S. J. Kim, J.-W. Lee, Metallurgical evidence of stress induced hydrogen diffusion and corresponding crack nucleation behaviors of high-strength ferritic steel used in sour environment, Journal of Materials Research and Technology 8 (2019) 5036–5040. doi:10.1016/j.jmrt.2019.07.030.

[18] Y. Fukai, The metal-hydrogen system: basic bulk properties, Springer, Berlin New York, 2005.

[19] A. Raina, V. Deshpande, N. Fleck, Analysis of thermal desorption of hydrogen in metallic alloys, Acta Materialia 144 (2018) 777–785. doi:10.1016/j.actamat.2017.11.011.

[20] A. Raina, V. S. Deshpande, N. A. Fleck, Analysis of electro-permeation of hydrogen in metallic alloys, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 375 (2017) 20160409. doi:10.1098/rsta.2016.0409.

[21] J. Takahashi, K. Kawakami, T. Tarui, Direct observation of hydrogen-trapping sites in vanadium carbide precipitation steel by atom probe tomography, Scripta Materialia 67 (2012) 213–216. doi:10.1016/j.scriptamat.2012.04.022.

[22] W. Counts, C. Wolverton, R. Gibala, First-principles energetics of hydrogen traps in α-fe: Point defects, Acta Materialia 58 (2010) 4730–4741. doi:10.1016/j.actamat.2010.05.010.

[23] E. Hayward, C.-C. Fu, Interplay between hydrogen and vacancies inα-fe, Physical Review B 87 (2013). doi:10.1103/physrevb.87.174103.

[24] R. Nazarov, T. Hickel, J. Neugebauer, First-principles study of the thermodynamics of hydrogen-vacancy interaction in fcc iron, Physical Review B 82 (2010). doi:10.1103/physrevb.82.224104.

[25] J. Sanchez, J. Fullea, M. C. Andrade, P. L. de Andres, Ab initiomolecular dynamics simulation of hydrogen diffusion inα-iron, Physical Review B 81 (2010). doi:10.1103/physrevb.81.132102.

[26] H. Kimizuka, H. Mori, S. Ogata, Effect of temperature on fast hydrogen diffusion in iron: A path-integral quantum dynamics approach, Physical Review B 83 (2011). doi:10.1103/physrevb.83.094110.

[27] I. H. Katzarov, D. L. Pashov, A. T. Paxton, Fully quantum mechanical calculation of the diffusivity of hydrogen in iron using the tight-binding approximation and path integral theory, Physical Review B 88 (2013). doi:10.1103/physrevb.88.054107.

[28] S. Wang, M. L. Martin, I. M. Robertson, P. Sofronis, Effect of hydrogen environment on the separation of fe grain boundaries, Acta Materialia 107 (2016) 279–288. doi:10.1016/j.actamat.2016.01.067.

[29] S. Teus, V. Gavriljuk, Grain-boundary diffusion of hydrogen atoms in the α-iron, METALLOFIZIKA I NOVEISHIE TEKHNOLOGII 36 (2016) 1399–1410. doi:10.15407/mfint.36.10.1399.

[30] J. Song, W. A. Curtin, Atomic mechanism and prediction of hydrogen embrittlement in iron, Nature Materials 12 (2012) 145–151. doi:10.1038/nmat3479.

[31] A. Tehranchi, X. Zhang, G. Lu, W. A. Curtin, Hydrogen–vacancy–dislocation interactions inα-fe, Modelling and Simulation in Materials Science and Engineering 25 (2016) 025001. doi:10.1088/1361-651x/aa52cb.

[32] W. K. Hastings, Monte carlo sampling methods using markov chains and their applications, Biometrika 57 (1970) 97–109. doi:10.1093/biomet/57.1.97.

[33] A. Bortz, M. Kalos, J. Lebowitz, A new algorithm for monte carlo simulation of ising spin systems, Journal of Computational Physics 17 (1975) 10–18. doi:10.1016/0021-9991(75)90060-1.

[34] G. Henkelman, H. Jónsson, Long time scale kinetic monte carlo simulations without lattice approximation and predefined event table, The Journal of Chemical Physics 115 (2001) 9657–9666. doi:10.1063/1.1415500.

[35] L. Xu, D. Mei, G. Henkelman, Adaptive kinetic monte carlo simulation of methanol decomposition on cu(100), The Journal of Chemical Physics 131 (2009) 244520. doi:10.1063/1.3281688.

[36] A. Pedersen, H. Jonsson, Simulations of hydrogen diffusion at grain boundaries in aluminum, Acta Materialia 57 (2009) 4036–4045. doi:10.1016/j.actamat.2009.04.057.

[37] F. El-Mellouhi, N. Mousseau, L. J. Lewis, Kinetic activation-relaxation technique: An off-lattice self-learning kinetic monte carlo algorithm, Physical Review B 78 (2008). doi:10.1103/physrevb.78.153202.

[38] H. Xu, Y. N. Osetsky, R. E. Stoller, Self-evolving atomistic kinetic monte carlo: fundamentals and applications, Journal of Physics: Condensed Matter 24 (2012) 375402. doi:10.1088/0953-8984/24/37/375402.

[39] H. Xu, R. E. Stoller, L. K. Béland, Y. N. Osetsky, Self-evolving atomistic kinetic monte carlo simulations of defects in materials, Computational Materials Science 100 (2015) 135–143. doi:10.1016/j.commatsci.2014.12.026.

[40] C. Lu, K. Jin, L. K. Béland, F. Zhang, T. Yang, L. Qiao, Y. Zhang, H. Bei, H. M. Christen, R. E. Stoller, L. Wang, Direct observation of defect range and evolution in ion-irradiated single crystalline ni and ni binary alloys, Scientific Reports 6 (2016). doi:10.1038/srep19994.

[41] O. A. Restrepo, N. Mousseau, F. El-Mellouhi, O. Bouhali, M. Trochet, C. S. Becquart, Diffusion properties of fe–c systems studied by using kinetic activation–relaxation technique, Computational Materials Science 112 (2016) 96–106. doi:10.1016/j.commatsci.2015.10.017.

[42] A. Ervin, H. Xu, Mesoscale simulations of radiation damage effects in materials: A SEAKMC perspective, Computational Materials Science 150 (2018) 180–189. doi:10.1016/j.commatsci.2018.03.054.

[43] D. Bombač, I. H. Katzarov, D. L. Pashov, A. T. Paxton, Theoretical evaluation of the role of crystal defects on local equilibrium and effective diffusivity of hydrogen in iron, Materials Science and Technology 33 (2017) 1505–1514. doi:10.1080/02670836.2017.1310417.

[44] Y. Zeng, P. Xiao, G. Henkelman, Unification of algorithms for minimum mode optimization, The Journal of Chemical Physics 140 (2014) 044115. doi:10.1063/1.4862410.
[45] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, Journal of Research of the National Bureau of Standards 45 (1950) 255. doi:10.6028/jres.045.026.
[46] J. Kästner, P. Sherwood, Superlinearly converging dimer method for transition state search, The Journal of Chemical Physics 128 (2008) 014106. doi:10.1063/1.2815812.
[47] G. Henkelman, H. Jónsson, A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives, The Journal of Chemical Physics 111 (1999) 7010–7022. doi:10.1063/1.480097.
[48] A. Heyden, A. T. Bell, F. J. Keil, Efficient methods for finding transition states in chemical reactions: Comparison of improved dimer method and partitioned rational function optimization method, The Journal of Chemical Physics 123 (2005) 224101. doi:10.1063/1.2104507.
[49] L. Xu, G. Henkelman, Adaptive kinetic monte carlo for first-principles accelerated dynamics, The Journal of Chemical Physics 129 (2008) 114104. doi:10.1063/1.2976010.
[50] S. I. Shah, G. Nandipati, A. Kara, T. S. Rahman, Extended pattern recognition scheme for self-learning kinetic monte carlo simulations, Journal of Physics: Condensed Matter 24 (2012) 354004. doi:10.1088/0953-8984/24/35/354004.
[51] G. Nandipati, A. Kara, S. I. Shah, T. S. Rahman, Off-lattice pattern recognition scheme for kinetic monte carlo simulations, Journal of Computational Physics 231 (2012) 3548–3560. doi:10.1016/j.jcp.2011.12.029.
[52] D. Konwar, V. J. Bhute, A. Chatterjee, An off-lattice, self-learning kinetic monte carlo method using local environments, The Journal of Chemical Physics 135 (2011) 174103. doi:10.1063/1.3657834.
[53] S. Fortin, The Graph Isomorphism Problem, Technical Report, University of Alberta Libraries, 1996. doi:10.7939/R3SX64C5K.
[54] B. D. McKay, A. Piperno, Practical graph isomorphism, II, Journal of Symbolic Computation 60 (2014) 94–112. doi:10.1016/j.jsc.2013.09.003.
[55] M. Trochet, N. Mousseau, L. K. Béland, G. Henkelman, Off-lattice kinetic monte carlo methods, in: Handbook of Materials Modeling, Springer International Publishing, 2018, pp. 1–29. URL: https://doi.org/10.1007/978-3-319-42913-7_29-1. doi:10.1007/978-3-319-42913-7_29-1.
[56] A. Ramasubramaniam, M. Itakura, M. Ortiz, E. Carter, Effect of atomic scale plasticity on hydrogen diffusion in iron: Quantum mechanically informed and on-the-fly kinetic monte carlo simulations, Journal of Materials Research 23 (2008) 2757–2773. doi:10.1557/jmr.2008.0340.
[57] F. Glover, M. Laguna, Tabu search, in: Handbook of Combinatorial Optimization, Springer US, 1998, pp. 2093–2229. URL: https://doi.org/10.1007/978-1-4613-0303-9_33. doi:10.1007/978-1-4613-0303-9_33.
[58] M. V. Chubynsky, H. Vocks, G. Barkema, N. Mousseau, Exploiting memory in event-based simulations, Journal of Non-Crystalline Solids 352 (2006) 4424–4429. doi:10.1016/j.jnoncrysol.2006.01.117.
[59] K. A. Fichthorn, Y. Lin, A local superbasin kinetic monte carlo method, The Journal of Chemical Physics 138 (2013) 164104. doi:10.1063/1.4801869.
[60] B. Puchala, M. L. Falk, K. Garikipati, An energy basin finding algorithm for kinetic monte carlo acceleration, The Journal of Chemical Physics 132 (2010) 134104. doi:10.1063/1.3369627.
[61] L. K. Béland, P. Brommer, F. El-Mellouhi, J.-F. Joly, N. Mousseau, Kinetic activation-relaxation technique, Physical Review E 84 (2011). doi:10.1103/physreve.84.046704.
[62] M. W. Finnis, J. E. Sinclair, A simple empirical n-body potential for transition metals, Philosophical Magazine A 50 (1984) 45–55. doi:10.1080/01418618408244210.
[63] J. J. Möller, E. Bitzek, Comparative study of embedded atom potentials for atomistic simulations of fracture inα-iron, Modelling and Simulation in Materials Science and Engineering 22 (2014) 045002. doi:10.1088/0965-0393/22/4/045002.
[64] A. Ramasubramaniam, M. Itakura, E. A. Carter, Interatomic potentials for hydrogen inα–iron based on density functional theory, Physical Review B 79 (2009). doi:10.1103/physrevb.79.174101.
[65] C. J. Williams, Developing an off-lattice kinetic Monte Carlo model to simulate hydrogen-crystal defect interactions in α–iron, Master's thesis, University of Cambridge, Cambridgeshire UK, 2020.
[66] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics), Springer-Verlag, Berlin, Heidelberg, 2006.
[67] H. Zhu, B. Guo, K. Zou, Y. Li, K.-V. Yuen, L. Mihaylova, H. Leung, A review of point set registration: From pairwise registration to groupwise registration, Sensors 19 (2019) 1191. doi:10.3390/s19051191.

[68] B. Maiseli, Y. Gu, H. Gao, Recent developments and trends in point set registration methods, Journal of Visual Communication and Image Representation 46 (2017) 95–106. doi:10.1016/j.jvcir.2017.03.012.

[69] P. H. Schönemann, A generalized solution of the orthogonal procrustes problem, Psychometrika 31 (1966) 1–10. doi:10.1007/bf02289451.

[70] D. C. Liu, J. Nocedal, On the limited memory BFGS method for large scale optimization, Mathematical Programming 45 (1989) 503–528. doi:10.1007/bf01589116.

[71] J. Nocedal, Updating quasi-newton matrices with limited storage, Mathematics of Computation 35 (1980) 773–773. doi:10.1090/s0025-5718-1980-0572855-7.

[72] P. Wolfe, Convergence conditions for ascent methods, SIAM Review 11 (1969) 226–235. doi:10.1137/1011036.

[73] P. Wolfe, Convergence conditions for ascent methods. II: Some corrections, SIAM Review 13 (1971) 185–188. doi:10.1137/1013035.

[74] J. Nocedal, S. J. Wright, Numerical Optimization, 2 ed., Springer New York, 2006. URL: https://doi.org/10.1007/978-0-387-40065-5. doi:10.1007/978-0-387-40065-5.

[75] S. T. Chill, G. Henkelman, Molecular dynamics saddle search adaptive kinetic monte carlo, The Journal of Chemical Physics 140 (2014) 214110. doi:10.1063/1.4880721.

[76] G. H. Vineyard, Frequency factors and isotope effects in solid state rate processes, Journal of Physics and Chemistry of Solids 3 (1957) 121–127. doi:10.1016/0022-3697(57)90059-8.

[77] P. Hänggi, P. Talkner, M. Borkovec, Reaction-rate theory: fifty years after kramers, Reviews of Modern Physics 62 (1990) 251–341. doi:10.1103/revmodphys.62.251.

[78] G. Lv, M. Zhang, H. Zhang, Y. Su, Hydrogen diffusion and vacancy clusterization in iron, International Journal of Hydrogen Energy 43 (2018) 15378–15385. doi:10.1016/j.ijhydene.2018.06.075.

[79] M. Pascuet, N. Castin, C. Becquart, L. Malerba, Stability and mobility of cu–vacancy clusters in fe–cu alloys: A computational study based on the use of artificial neural networks for energy barrier calculations, Journal of Nuclear Materials 412 (2011) 106–115. doi:10.1016/j.jnucmat.2011.02.038.

[80] L. Messina, L. Malerba, P. Olsson, Stability and mobility of small vacancy–solute complexes in fe–MnNi and dilute fe–x alloys: A kinetic monte carlo study, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 352 (2015) 61–66. doi:10.1016/j.nimb.2014.12.032.

[81] S. E. Restrepo, H. Lambert, A. T. Paxton, Effect of hydrogen on vacancy diffusion, Physical Review Materials 4 (2020). doi:10.1103/physrevmaterials.4.113601.

[82] F. Djurabekova, L. Malerba, R. Pasianot, P. Olsson, K. Nordlund, Kinetics versus thermodynamics in materials modeling: The case of the di-vacancy in iron, Philosophical Magazine 90 (2010) 2585–2595. doi:10.1080/14786431003662515.

[83] F. Djurabekova, L. Malerba, C. Domain, C. Becquart, Stability and mobility of small vacancy and copper-vacancy clusters in bcc-fe: An atomistic kinetic monte carlo study, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms 255 (2007) 47–51. doi:10.1016/j.nimb.2006.11.009.

[84] C.-C. Fu, J. D. Torre, F. Willaime, J.-L. Bocquet, A. Barbu, Multiscale modelling of defect kinetics in irradiated iron, Nature Materials 4 (2004) 68–74. doi:10.1038/nmat1286.

[85] V. Borodin, P. Vladimirov, Diffusion coefficients and thermal stability of small helium–vacancy clusters in iron, Journal of Nuclear Materials 362 (2007) 161–166. doi:https://doi.org/10.1016/j.jnucmat.2007.01.019, e-MRS 2006: Symposium N.

[86] S. Mahmoud, M. Trochet, O. A. Restrepo, N. Mousseau, Study of point defects diffusion in nickel using kinetic activation-relaxation technique, Acta Materialia 144 (2018) 679–690. doi:https://doi.org/10.1016/j.actamat.2017.11.021.

[87] P. Brommer, L. K. Béland, J.-F. m. c. Joly, N. Mousseau, Understanding long-time vacancy aggregation in iron: A kinetic activation-relaxation technique study, Phys. Rev. B 90 (2014) 134109. doi:10.1103/PhysRevB.90.134109.

[88] L. Malerba, M. Marinica, N. Anento, C. Björkas, H. Nguyen, C. Domain, F. Djurabekova, P. Olsson, K. Nordlund, A. Serra, D. Terentyev, F. Willaime, C. Becquart, Comparison of empirical interatomic potentials for iron applied to radiation damage studies, Journal of Nuclear Materials 406 (2010) 19–38. doi:https://doi.org/10.1016/j.jnucmat.2010.05.017, fP6 IP PERFECT Project: Prediction of Irradiation Damage Effects in Reactor Components.

[89] J.-P. Du, W. T. Geng, K. Arakawa, J. Li, S. Ogata, Hydrogen-enhanced vacancy diffusion in metals, The Journal of Physical Chemistry Letters 11 (2020) 7015–7020. doi:10.1021/acs.jpclett.0c01798.

[90] K. Arakawa, A. Kageyama, H. Hiroshima, H. Yasuda, S. Ogata, Hydrogen effects on the migration of nanoscale cavities in iron, ISIJ International 61 (2021) 2305–2307. doi:10.2355/isijinternational.isijint-2021-140.

[91] S. Li, Y. Li, Y.-C. Lo, T. Neeraj, R. Srinivasan, X. Ding, J. Sun, L. Qi, P. Gumbsch, J. Li, The interaction of dislocations and hydrogen-vacancy complexes and its importance for deformation-induced proto nano-voids formation in α-fe, International Journal of Plasticity 74 (2015) 175–191. doi:10.1016/j.ijplas.2015.05.017.

[92] M. Nagumo, K. Takai, The predominant role of strain-induced vacancies in hydrogen embrittlement of steels: Overview, Acta Materialia 165 (2019) 722–733. doi:10.1016/j.actamat.2018.12.013.

[93] P. Gong, I. H. Katzarov, J. Nutter, A. T. Paxton, W. M. Rainforth, The influence of hydrogen on plasticity in pure iron—theory and experiment, Scientific Reports 10 (2020). doi:10.1038/s41598-020-66965-z.

[94] I. H. Katzarov, D. L. Pashov, A. T. Paxton, Hydrogen embrittlement i. analysis of hydrogen-enhanced localized plasticity: Effect of hydrogen on the velocity of screw dislocations in α–fe, Physical Review Materials 1 (2017). doi:10.1103/physrevmaterials.1.033602.

[95] R. Oriani, The diffusion and trapping of hydrogen in steel, Acta Metallurgica 18 (1970) 147–157. doi:10.1016/0001-6160(70)90078-7.

[96] A. T. Paxton, I. H. Katzarov, Quantum and isotope effects on hydrogen diffusion, trapping and escape in iron, Acta Materialia 103 (2016) 71–76. doi:10.1016/j.actamat.2015.09.054.

[97] Z. Zhang, A flexible new technique for camera calibration, IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (2000) 1330–1334. doi:10.1109/34.888718.

[98] T. C. Hales, Historical overview of the kepler conjecture, Discrete & Computational Geometry 36 (2006) 5–20. doi:10.1007/s00454-005-1210-2.

Appendix A. Orthogonal Procrustes problem

Algorithm 4 Solves the orthogonal Procrustes problem, find and return the optimal orthogonal transformation to map the point-cloud Q onto the point-cloud P [69, 97].

Require: P and Q contain the same number of points, n.

function ROTOR_ONTO(P, Q)

$\boldsymbol{H} \leftarrow \sum_{i=1}^{n} \boldsymbol{q}_i \boldsymbol{p}_i^\top$

Compute $\boldsymbol{U}, \boldsymbol{V}$ from the SVD of $\boldsymbol{H}$ such that $\boldsymbol{H} = \boldsymbol{U\Sigma V}^\top$

return $\boldsymbol{VU}^\top$

Appendix B. Complexity analysis of GREEDY_PERM

The time complexity of the _RECUR function fromAlgorithm 1, $\mathcal{O}(R_i)$, when called with integer i and point clouds P and Q of size n is:

$$\mathcal{O}(R_i) = \begin{cases} (n-i)\left(\mathcal{O}(M_i) + \mathcal{O}(R_{i+1})\right) & i \leq n \\ n & i > n \end{cases} \tag{B.1}$$

with $\mathcal{O}(M_i)$ the time complexity of _MATCH called with integer i and point clouds P and Q. This is clearly exponential in n for general point clouds. However, in 3D, each intra-point pair in P defines a (thin) shell of thickness $2\sqrt{2}\delta$ around the corresponding points in Q, which the next point must fall inside of for _MATCH to return `True`. As the distance from four non-coplanar points in 3D define a unique point, once we have matched a small, constant number of points that span the LE, we expect the intersections of these shells to converge to a sphere. Therefore, the volume of space that the tolerance defines in Q is:

$$V \approx \frac{8}{3}\sqrt{2}\pi\delta^3 \tag{B.2}$$

and point the density, ρ, is bound in the worst case by closely-packed spheres of radius $\frac{r_{\min}}{2}$ [98]:

$$\rho \leq \frac{\sqrt{2}}{r_{\min}^3} \tag{B.3}$$

hence, the expected number of points inside the volume of tolerant space, m, is approximately:

$$m \approx \frac{16\pi}{3}\left(\frac{\delta}{r_{\min}}\right)^3 \tag{B.4}$$

Each iteration will on average: call _MATCH m times for points that will match, each with $\mathcal{O}(n)$ runtime; trigger m recursions; call _MATCH $\mathcal{O}(n)$ times for points that will **not** match, each with constant runtime. Hence, we can rewrite Eq. (B.1):

$$\begin{aligned} \mathcal{O}(R) &= nm + m\,(nm + m\,(\ldots)) \\ \mathcal{O}(R) &= \begin{cases} n^2 & m \leq 1 \\ nm^n & m > 1 \end{cases} \end{aligned} \tag{B.5}$$

In order to avoid exponential complexity, we require $m \leq 1$. Rearranging Eq. (B.4) and using the approximate nature to simplify the constants, this requires:

$$\delta \lesssim \frac{2}{5} r_{\min} \tag{B.6}$$

For the case of α-Fe with $r_{\min} \approx 1$Å, this requires $\delta \lesssim 0.4$Å. This matches our empirical experiments with typical LEs in α-Fe where, we observe the runtime of Algorithm 1 blowing-up beyond $\delta \gtrsim 0.4$Å.

Declaration of interests

☒ The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

☐The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: