

EFFICIENT SAMPLING OF SADDLE POINTS WITH THE MINIMUM-MODE FOLLOWING METHOD*

ANDREAS PEDERSEN[†], SIGURDUR F. HAFSTEIN[‡], AND HANNES JÓNSSON[†]

Abstract. The problem of sampling low lying, first-order saddle points on a high dimensional surface is discussed and a method presented for improving the sampling efficiency. The discussion is in the context of an energy surface for a system of atoms and thermally activated transitions in solids treated within the harmonic approximation to transition state theory. Given a local minimum as an initial state and a small, initial displacement, the minimum-mode following method is used to climb up to a saddle point. The goal is to sample as many of the low lying saddle points as possible when such climbs are repeated from different initial displacements. Various choices for the distribution of initial displacements are discussed and a comparison made between (1) displacements along eigenmodes at the minimum, (2) purely random displacements with a maximum cutoff, and (3) Gaussian distribution of displacements. The last choice is found to give best overall results in two test problems studied, a heptamer island on a surface and a grain boundary in a metal. A method referred to as “skipping-path method” is presented to reduce redundant calculations when a climb heads towards a saddle point that has already been identified. The method is found to reduce the computational effort of finding new saddle points to as little as a third, especially when a thorough sampling is performed.

Key words. adaptive kinetic Monte Carlo, minimum-mode following method, saddle point determination, atomistic dynamics

AMS subject classifications. 70-08, 70F10, 82C22

DOI. 10.1137/100792743

1. Introduction. The problem of finding first-order saddle points (SPs), i.e., SPs where the Hessian has one and only one negative eigenvalue, on high dimensional surfaces appears in many contexts [1]. One example is the study of mechanism and rate of thermal transitions in materials where the first-order SPs on the energy surface as a function of atomic coordinates are of primary importance. A stable configuration of atoms corresponds to a minimum in the potential energy as a function of the atomic coordinates and rearrangements of the atoms to other, stable configurations can be identified by locating first-order SPs on the potential energy rim surrounding the initial state minimum. Note that from a state of equilibrium we want to find the least demanding ways of escaping its basin of attraction in measures of energy. Mathematically, this is an inverse problem from a singularity and thus, in general, classifies as a hard problem and is not solvable by any standard or simple method. Especially, one cannot resort to solving the equations of motion backwards in time.

In our model of thermal transitions in materials the transition probability can be estimated by transition state theory [2] which for solids can be further simplified by a harmonic approximation [4] unless the temperature is too high. In solids, the atoms are mostly confined to small regions around the minima and the bottle necks

*Submitted to the journal’s Computational Methods in Science and Engineering section April 19, 2010; accepted for publication (in revised form) December 13, 2010; published electronically March 3, 2011. This work was supported by the Icelandic Research Foundation (RANNIS), the University of Iceland research fund, and EC Integrated Project NES-SHy.

<http://www.siam.org/journals/sisc/33-2/79274.html>

[†]Faculty of Science and Science Institute, University of Iceland, 107 Reykjavík, Iceland (andreas@theochem.org, hj@hi.is).

[‡]School of Science and Engineering, Reykjavik University, Menntavegi 1, 101 Reykjavík, Iceland (sigurdurh@hr.is).

for transitions—the regions of smallest statistical weight that have to be traversed during transitions—are small regions around first order saddle points. A harmonic approximation to the energy surface in the vicinity of minima and first order SPs, therefore, gives a good approximation when the transition rate is estimated and the mode with negative eigenvalue at the first-order SP gives a clear description of the transition mechanism in the bottle neck region. The transition rate due to trajectories that pass through the region of a given SP is [4]

$$(1.1) \quad k^{HTST} = \nu \exp \left[-\frac{E_{SP} - E_{min}}{k_B T} \right],$$

$$(1.2) \quad \nu = \frac{\prod_i^D \nu_{R,i}}{\prod_i^{D-1} \nu_{SP,i}},$$

where E_{min} and E_{SP} are the values of the potential energy at the initial state minimum and at the SP, k_B is the Boltzmann constant, and T is the temperature. The prefactor, ν , is given by the ratio of frequencies for vibrational modes in the initial state, $\nu_{R,i}$, and at the SP, $\nu_{SP,i}$. The vibrational modes are found by constructing the mass weighted Hessian matrix, $\tilde{\mathbf{H}}$, and solving an eigenvalue problem representing the equations of motion for the atoms:

$$(1.3) \quad \tilde{\mathbf{H}} = \mathbf{LHL},$$

$$(1.4) \quad L_{i,j} = \frac{1}{\sqrt{m_{[i/3]}}} \delta_{i,j},$$

where i runs over all degrees of freedom (three per atom), $m_{[i/3]}$ is the mass of atom that corresponds to the degree of freedom, $\delta_{i,j}$ is the Kronecker delta, and \mathbf{H} is the Hessian, which is the second order derivative of the energy with respect to the degrees of freedom:

$$(1.5) \quad \mathbf{H} = \begin{bmatrix} \frac{\partial^2 E}{\partial x_1 \partial x_1} & \frac{\partial^2 E}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 E}{\partial x_1 \partial x_{3N}} \\ \frac{\partial^2 E}{\partial x_2 \partial x_1} & \frac{\partial^2 E}{\partial x_2 \partial x_2} & \cdots & \frac{\partial^2 E}{\partial x_2 \partial x_{3N}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 E}{\partial x_{3N} \partial x_1} & \frac{\partial^2 E}{\partial x_{3N} \partial x_2} & \cdots & \frac{\partial^2 E}{\partial x_{3N} \partial x_{3N}} \end{bmatrix}.$$

If the final state(s) of possible transitions are known, then various techniques for determining minimum energy paths between the initial state minimum and each final state minimum can be used, such as the nudged elastic band method [5, 6], and the relevant SP can be identified as the highest maximum along the minimum energy path. However, the final states are often not known or are only poorly known so that the search for SPs should preferably only make use of information about the initial state and not be biased by any preconceived notion about transition mechanism or final state. There are many examples where seemingly simple transitions turned out to prefer a more complex mechanism than initially assumed (see, for example, [7, 8]). Then, there are systems that are so complex, either because of the interatomic interactions and/or because of the atomic configurations, that it is practically impossible to predict the transition mechanism and final state configurations. One example of the latter is presented here in a study of grain boundary annealing where as many as 10 atoms are sometimes displaced by more than 0.5 Å in a concerted way as the system is displaced from a minimum to a low lying saddle point (see below).

In principle, trajectories generated by classical dynamics (numerical solution to Newton equations), can provide information about mechanism and rate of thermally activated transitions. The problem, however, is that the time between typical activated events is many orders of magnitude longer than the fastest time scale, that of atomic vibrations, so the simulation would take an impossibly long time. Different algorithms are, therefore, needed to span the long time scales that are relevant for time evolution due to thermally activated transitions. The adaptive kinetic Monte Carlo algorithm (AKMC) [9] is one such approach. There, a table of possible transitions from a given initial state is generated by repeated SP searches. A random number is then used to pick a transition according to the relative transition probabilities and the system advanced to the final state of the transition. An important aspect of the method is that the table of possible events that can take place in the system is not constructed before hand, but is created on the fly as the simulation is carried out. The AKMC method has, for example, been used to simulate metal crystal growth by vapor deposition [10], the annealing of grain boundaries in metals [11] and the diffusion of hydrogen at grain boundaries in metals [12]. As long as the low energy SPs, corresponding to the more likely transitions, are identified and no systematic omission of SPs occurs, the method gives a good approximation to the long time scale evolution. The sampling efficiency discussed in this article is, therefore, particularly important for AKMC simulations. The set of SP-searches that are carried out for a given initial state are independent and can be farmed out to different computers that are simply connected by ethernet. Software for long time scale simulations of solids using such distributed computing has been developed and made available [13].

Several methods have been proposed to climb up from a given initial state minimum to converge onto a first-order SP. It is important to realize that a simple climb up the direction of slowest ascent will not necessarily lead to a SP (see a simple two-dimensional example in [14]), and even if it does, the question still remains how to sample other SPs on the potential energy rim. Cerjan and Miller [15] suggested using the eigenvectors of the Hessian to guide the climb, first using the vector for the lowest eigenvalue—the minimum-mode—and then successively the eigenvectors of the higher eigenvalues. Several variants of this mode following method have been presented (see, for example, [16, 17]). The problem with the mode following method is that the full Hessian needs to be constructed and an eigenvalue problem solved. Not only does the computational effort scale up rapidly as more degrees of freedom are included in the system, but also the construction of the Hessian can represent a large computational effort when analytical second derivatives are not available. A more economic approach is to follow only the minimum-mode, as this mode can be estimated without constructing and diagonalizing the Hessian [18, 19]. Different saddle points can be found by starting from different initial points, slightly displaced from the minimum. This minimum-mode following method is described in more detail below. It can be applied to large systems with many degrees of freedom (several thousands) and it can be used in combination with plane wave based density functional theory [20] evaluation of the energy and force where second derivatives of the energy are typically not available because of the large computational effort (see, for example, [21]). The choice of the initial displacements can strongly influence the efficiency of the SP searches. This issue is addressed below. Other methods for generating starting points for SP-searches, involving larger displacements, have also been suggested [22].

A general requirement for such SP-searches is to get a complete, or at least sufficiently thorough sampling of SPs that are below a given cutoff value. The number of




FIG. 2.1. A schematic illustration of three SPs connected to the same initial state minimum. The positive region, where all eigenvalues of the Hessian are positive, is shaded with gray. Outside of that, the lowest eigenvalue λ is negative, and the state-space included is the basins of attraction of the different first order SPs, as indicated by thin dotted lines. To the left, an example of a search path that converges on one of the SPs is shown. Typically, there is an abrupt change in the path as it exits the positive region and relaxation perpendicular to the minimum-mode direction sets in. To the right, the spherical regions around intermediate configurations along the path outside the positive region are indicated with a gray shade. These are used in the skipping-path method to terminate searches that later enter the same region to reduce computational effort wasted in redundant calculations and improve the sampling of SPs obtained for a given computational effort.

SPs is typically large and some may be hard to find so a large number of searches may need to be conducted. Each search can also take significant amount of computer time, especially when the energy and force is evaluated by solving some approximation to the wave equation for the electrons (density functional theory or quantum chemistry methods). It is, therefore, important to optimize the searches and reduce redundancy as much as possible. Two aspects of this issue are addressed in the present study. First, a comparison of different methods for generating displacements from the initial state is presented. Then, a new method is presented for reducing the amount of computational effort wasted when a search converges onto an already known SP. Before describing those, the minimum-mode following method is first reviewed and the two systems used in this study described.

2. Minimum-mode following. Several iterative algorithms can be used to reach a local minimum of a function by using only information about the gradient of the function and not its value. This includes, for example, the conjugate gradient method if the line searches are skipped. Such an algorithm can also be used to converge to a first-order saddle point if the component of the gradient in the direction of the negative curvature mode, i.e., the direction corresponding to the negative eigenvalue, is reversed. This component inversion maps a first order saddle point onto a local minimum. In the context of atomic systems, the force on each atom i is given by a three-dimensional vector \mathbf{F}_i . Because we are interested in equilibrium states of the whole system, i.e., in states where $\mathbf{F}_i = \mathbf{0}$ for $i = 1, 2, \dots, N$, where N is the number of the atoms in the system, it is profitable to change to a state-space (generalized coordinates) representation. In this representation we write down the components of the individual atomic forces \mathbf{F}_i for $i = 1, 2, \dots, N$ in the state-space force vector \mathbf{F} . Then \mathbf{F} is a $3N$ -dimensional vector, or more exactly a $3N$ -dimensional vector field of the $3N$ state variables $x = (x_1, x_2, \dots, x_{3N})$. The force acting on the i th atom is contained in the components at positions $3i - 2$, $3i - 1$, and $3i$ in \mathbf{F} and the position of the i th atom is similarly contained in the components at the same positions in x . The Hessian matrix \mathbf{H} of the energy E is the total differential of $-\mathbf{F}$, because, by definition $\mathbf{F} = -\nabla E$. The component inversion can now be given by the formula

$$(2.1) \quad \mathbf{F}^{eff} = \mathbf{F} - 2(\mathbf{F} \cdot \mathbf{v}_{min})\mathbf{v}_{min},$$

where \mathbf{v}_{min} is a normalized eigenvector corresponding to the lowest eigenvalue of the Hessian at the point. By making the atoms follow the transformed force, \mathbf{F}^{eff} , a minimization algorithm converges to a first order saddle point on the energy surface. The negative curvature mode can be found by constructing the Hessian matrix, the matrix of second derivatives of the potential energy, and finding a normalized eigenvector, \mathbf{v}_{min} , corresponding to the smallest eigenvalue. However, since only the mode with the smallest eigenvalue is needed, it is possible to get the required information more directly, thereby avoiding the construction of the Hessian matrix and the eigenvalue problem. In the dimer method [18], two replicas of the system—the dimer—are constructed and their separation kept fixed while the orientation of the dimer is optimized to minimize the total energy. This aligns the direction of the dimer to the minimum-mode. This is the method used here. Another possibility is to use the Lanczos method which is often used to converge to the lowest (and highest) eigenvalue of large matrices [24].

Let us explain the component inversion method more closely. A system is at a local equilibrium at z if there is no force acting on it. That is $\mathbf{F}(z) = \mathbf{0}$. Without loss of generality, we can assume $z = 0$. If we discard nonhyperbolic equilibrium points, i.e., equilibria whose potential cannot be approximated locally by a quadratic function, then, by the Hartman–Grobman theorem [25] a complete topological description of the flow close to the equilibrium is given by the linear differential equation

$$\frac{d}{dt} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} = \mathbf{M} \begin{bmatrix} x \\ \dot{x} \end{bmatrix}, \quad \mathbf{M} = \begin{bmatrix} \mathbf{O}_{3N} & \mathbf{I}_{3N} \\ -\mathbf{H} & \mathbf{O}_{3N} \end{bmatrix}.$$

Here $\dot{x} = \frac{dx}{dt}$, \mathbf{H} is the Hessian of the potential E , and \mathbf{I}_{3N} and \mathbf{O}_{3N} are the $3N \times 3N$ identity- and zero-matrix respectively. If all eigenvalues of \mathbf{M} are positive, then the equilibrium at the origin is a local minimum of the total energy, if all are negative, then it is a local maximum, and if some are positive and some are negative, then it is a saddle point. Clearly \mathbf{M} is fixed by \mathbf{H} , so the flow close to the equilibrium is determined uniquely by the Hessian \mathbf{H} . Because the determinant of a matrix is the product of its eigenvalues and $\det \mathbf{M} = \pm \det \mathbf{H}$, the equilibrium at the origin is hyperbolic, if and only if the Hessian \mathbf{H} does not have zero as an eigenvalue.

It is interesting to take a look at the solution of $\ddot{x}(t) = -\mathbf{H}x(t)$ in this case ($\ddot{x} = \frac{d^2x}{dt^2}$). Because \mathbf{H} is symmetric, we can write

$$x(t) = \sum_i a_i(t) p_i,$$

where the p_i , $i = 1, 2, \dots, 3N$, are an orthonormal set of eigenvectors of \mathbf{H} and the a_i are real-valued functions. Hence,

$$\sum_i \ddot{a}_i(t) p_i = \ddot{x}(t) = -\mathbf{H}x(t) = -\mathbf{H} \sum_i a_i(t) p_i = -\sum_i \lambda_i a_i(t) p_i,$$

where λ_i is the eigenvalue corresponding to the eigenvector p_i . Because the p_i form an orthonormal set, we are left with $3N$ decoupled one-dimensional differential equations:

$$\ddot{a}_i(t) = -\lambda_i a_i(t) \quad \text{for } i = 1, 2, \dots, 3N.$$

Consider what happens if the system is initially at rest in the eigenspace of λ_j , close to the equilibrium at the origin. We solve the initial-value problem

$$\ddot{a}_j(t) = -\lambda_j a_j(t), \quad \dot{a}_j(0) = 0, \quad a_j(0) = \varepsilon \neq 0 \text{ but small},$$

and by simple calculations we get

$$a_j(t) = \begin{cases} \varepsilon \cdot \cosh(t\sqrt{-\lambda_j}) & \text{if } \lambda_j < 0, \\ \varepsilon \cdot \cos(t\sqrt{\lambda_j}) & \text{if } \lambda_j > 0. \end{cases}$$

Hence, if $\lambda_j < 0$, the system gets repelled away from the equilibrium, and if $\lambda_j > 0$, it oscillates about the equilibrium. Thus, the system needs kinetic energy to pass through the equilibrium in an eigenspace corresponding to a negative eigenvalue, but in an eigenspace corresponding to a positive eigenvalue, it does not need kinetic energy.

Now, if a system is at a local minimum at x and we want to calculate the probability of the system to leave this minimum due to as small perturbations as possible to come to rest at another local minimum y , it is clear, e.g., from the fact that the potential E of the system at x is locally approximated by a positive definite quadratic form, that the system goes through an equilibrium point of the potential in nonpathological cases. By simple geometrical reasoning it seems logical to concentrate on first order saddle points, i.e., saddle points where all eigenvalues of the Hessian are positive except for one. First, as the probability of the system to climb over an energy-barrier decreases exponentially with the barrier height, one can concentrate on low barriers. Second, a saddle point of k th degree, i.e., the Hessian has k negative eigenvalues, implies that the quadratic approximation of the energy at the saddle point includes a k -dimensional paraboloid, of which the saddle point is a local maximum, laying in the subspace spanned by the eigenvectors corresponding to the negative eigenvalues. If $k > 1$, then a path crossing the saddle point can take a shortcut on the surface of the paraboloid avoiding the maximum at the saddle point and, thus, a smaller perturbation is able to cause this transition.

A consequence of the eigenvalues of the Hessian being continuous functions of the state is that a first order saddle point is contained in an open region of the state-space where the Hessian has one negative eigenvalue and all others are positive. This implies that if we use the transformed force \mathbf{F}^{eff} in the vicinity of the saddle point, then this saddle point is transformed into a local minimum and function minimizing methods will converge to this particular saddle point. A region containing the saddle point forms its basin of attraction (see Figure 2.2). This is illustrated schematically in Figure 2.1 showing a region of a potential energy ridge where three saddle points are present, corresponding to three different mechanisms for adatom diffusion on a solid surface. As soon as a search path has entered one of these basins of attraction, the final outcome of the search has been determined. This feature will be exploited in a method described below to reduce the redundancy in repeated SP-searches.

The initial state minimum is in a region where all eigenvalues of the Hessian are positive. This will be referred to as the positive region. A saddle point search path needs to get out of this region and into one of the basins of attractions of a first order saddle point. If too much minimization using the transformed force \mathbf{F}^{eff} is carried out in the positive region, the search path will tend to slide down to the slowest ascent path. This is not desired since different SP-searches will then all tend to give the same result. It is, therefore, important to move rather aggressively out of the positive region. This is accomplished in the minimum-mode following method by zeroing the gradient perpendicular to the minimum-mode. The effective force then becomes

$$(2.2) \quad \mathbf{F}^{eff} = \begin{cases} -(\mathbf{F} \cdot \mathbf{v}_{min})\mathbf{v}_{min} & \text{if } \lambda_{min} > 0, \\ \mathbf{F} - 2(\mathbf{F} \cdot \mathbf{v}_{min})\mathbf{v}_{min} & \text{if } \lambda_{min} < 0, \end{cases}$$




FIG. 2.2. Two-dimensional energy surface and multiple saddle point searches. Minima are denoted by a filled circle and saddle points by a white X where the arms point along the eigenvectors, gray-shaded regions mark the saddle points attraction basins. In each minima searches are conducted until the lowest saddle point has been determined ten times. The paths tend to merge together close to a saddle point. This is the motivation for the skipping-path algorithm where later searches get terminated if they come close to a previous search path.

where λ_{min} is the smallest eigenvalue of the Hessian at that point on the surface and \mathbf{v}_{min} is a corresponding normalized eigenvector.

The larger the displacements of the atoms in the positive region, the less tendency there is to slide onto the slowest ascent path, but the larger the probability is to land on a saddle point that does not connect with the initial state minimum (see discussion in [23]). After a saddle point has been identified, it is important to slide down the surface to check whether it connects with the initial state from which the search started. Such a slide can be performed with any of a number of minimization algorithms after a small displacement from the saddle point along the eigenvector corresponding to the negative eigenvalue, both forwards and backwards.

To locate a first order saddle point, a point where the effective force of (2.2) vanishes and one eigenvalue is negative, an iterative method is applied. Convergence is obtained when the RMS value of the effective force drops below a predefined value, typically 0.01 eV/\AA . Each iteration step involves the determination of a conjugated search direction and a single displacement towards the zero point of the effective force along this direction as originally proposed by Henkelman and Jónsson [18]. The search direction in iteration n is determined by a linear combination of the current gradient, \mathbf{F}_{eff} , and the previous direction, \mathbf{d}_{n-1} , as

$$(2.3) \quad \mathbf{d}_n = \mathbf{F}_{\text{eff}}(\mathbf{r}_n) + \beta_n \mathbf{d}_{n-1},$$

where the coefficient β_n is obtained using the Polak–Ribière formula:

$$(2.4) \quad \beta_n = \frac{\mathbf{F}_{\text{eff}}(\mathbf{r}_n) \cdot (\mathbf{F}_{\text{eff}}(\mathbf{r}_n) - \mathbf{F}_{\text{eff}}(\mathbf{r}_{n-1}))}{\mathbf{F}_{\text{eff}}(\mathbf{r}_{n-1}) \cdot \mathbf{F}_{\text{eff}}(\mathbf{r}_{n-1})}.$$

But, since an object function corresponding to the effective force is not known, the line search that is typically carried out in conjugate gradient minimization cannot be used. Instead, only one displacement along the direction, \mathbf{d}_n , is made based on a finite difference estimate of the derivative of the effective force along this line. The displacement is estimated using Newton’s method.

Algorithm 1. Force only conjugate gradient.

Require: The current configuration, $\mathbf{r}_n \in \mathbb{R}^{3N}$, the normalized search direction, $\mathbf{d}_n \in \mathbb{R}^{3N}$, an infinitesimal step size, δ , and a max displacement Δ_{max} .

- 1: $f^{\parallel} = \mathbf{F}_{\text{eff}}(\mathbf{r}_n) \cdot \mathbf{d}_n$
 - 2: $f_{\delta}^{\parallel} = \mathbf{F}_{\text{eff}}(\mathbf{r}_n + \delta \mathbf{d}_n) \cdot \mathbf{d}_n$
 - 3: $k^{\parallel} = \frac{f^{\parallel} - f_{\delta}^{\parallel}}{\delta}$
 - 4: $\Delta = \frac{f^{\parallel}}{k^{\parallel}}$, assuming k^{\parallel} to be constant
 - 5: **if** $k^{\parallel} < 0$ **or** $\Delta_{\text{max}} < |\Delta|$ **then**
 - 6: $\Delta = \text{sign}(f^{\parallel})\Delta_{\text{max}}$
 - 7: **end if**
 - 8: $\mathbf{r}_{n+1} = \mathbf{r}_n + \Delta \mathbf{d}_n$
-

A simple two-dimensional example energy surface and multiple SP-searches using the algorithm described above is shown in Figure 2.2. Seven distinct local minima and seven SPs are present on the surface. The figure illustrates how different initial displacements can lead to convergence on different SPs. In some cases, several different search paths tend to merge outside the positive region and have a long stretch where they lie very close until they reach the same SP. This is the motivation for the skipping-path method. There, later search paths are truncated after they have come within a certain distance of a previous search path. This can save a considerable amount of computations. Only points outside the positive region and within the basin of attraction of a SP are stored and used for possible termination of later search paths.

3. Simulated systems. Two systems which differ substantially in both the atomic arrangement and the atomic interactions were studied. One is a heptamer island (HI) on a close-packed crystal surface and the other a twisted and tilted grain boundary (GB) in a metal crystal. The atomic interactions are described by a Morse pair-potential in the HI system with parameters chosen to roughly describe platinum (this same system was studied in [23]) but for the GB system the effective medium theory [26, 27] potential function was used to reproduce properties of copper. The HI system is much simpler since the transitions involve mostly the seven island atoms, although the surface atoms are also allowed to move. The GB system is more complex since a large number of atoms can move and there are many low barrier transition paths on the energy surface due to the disorder in the grain boundary layer. In either case, the relevant transitions are local in that they mainly involve atoms in a limited region of the system, especially atoms with an anomalous coordination (i.e., whose neighbors are not arranged according to a FCC lattice).




FIG. 3.1. Heptamer island on the surface of a FCC solid, the HI test problem. (I) Green atoms show the seven atoms in the hexagonal island. (II) Blue atoms are the lowest coordinated atoms, with just six neighbors, and are chosen as possible centers for initial displacements. (III) Red atoms are within a given cutoff distance from a selected edge atom and have been displaced to move the system from the initial state minimum. This includes some of the surface atoms as well as some of the island atoms.

3.1. Heptamer island. The HI system is shown in Figure 3.1. A total of 343 atoms are included in the simulation. The crystal is represented by six layers of atoms arranged in an FCC lattice. The three bottom layers are kept frozen (a total of 168 atoms). The heptamer island on top of the surface has a hexagonal shape. The simple Morse potential is used to roughly (but not accurately) describe the atomic interactions because it is easy to implement and the calculations can, therefore, easily be reproduced. This system has been proposed as a general test problem for comparing various methods for finding mechanisms of transitions and estimating rates of transitions [28]. It is relatively simple since only a small number of low energy SPs are present and they mainly involve displacement of the seven island atoms. The initial state is the lowest energy, compact island with hexagonal shape. The lowest energy transitions to escape this minimum involve SPs with energy of 0.61 eV above the minimum. These correspond to either a clockwise or a counter-clockwise rotation of the whole island on the surface. Since the transition probability decreases exponentially as the SP energy increases (see (1.1)), a cutoff energy maximum is defined for what will be considered as relevant SPs. A SP with energy of 1.2 eV corresponds to a transition with a rate that is 10^{-10} smaller than the island rotation, assuming the prefactor is the same. Any SP with larger energy than 1.2 eV will be considered as irrelevant. There are 12 relevant SPs in this system.

3.2. Grain boundary. The GB system represents a junction between two FCC crystal grains which are both twisted and tilted with respect to each other. This is shown in Figure 3.2. Common neighbor analysis [29] is used to identify atoms that are not coordinated by their neighbors as in an FCC crystal. To minimize surface effects, the two upper most layers and the two lower most layers are kept frozen, while periodic boundary conditions are applied in the two directions parallel to the interface. The total number of atoms in the system is 1309 and of those 215 are kept frozen. The configuration has previously been annealed for nearly 0.1 ms using the AKMC method [11] and is, therefore, compact and quite stable. The lowest energy SP connected with this initial state has a height of 0.13 eV. It involves concerted displacement of four atoms by more than 0.25 Å. Using the same cutoff for minimum relative transition probability as in the HI system (10^{-10}), the largest relevant SP energy is 0.73 eV. A total of 459 SPs have been determined in this system from 2000 converged searches, of these are 92 within the above mentioned range and shown in Figure 3.3 together with curves for less strict cutoff energies. Compared with the HI system, this is a more challenging case, both because of the larger number of SPs and the close proximity of some of these SPs both in energy and in state-space.

3.3. Simulation parameters. The parameter values used here in the SP-searches are summarized below. The values were not tuned to optimize performance for ei-




FIG. 3.2. *The boundary between two grains of FCC metal, the GB test problem. Atoms that are not locally in an FCC environment as determined by common neighbor analysis are shown in blue. They are potential centers for initial displacements since they are likely in a less stable site.*

ther the HI or GB systems but are rather chosen to perform well for a general, atomic scale system. In the dimer method for determining the lowest eigenmode, the distance between the two images was fixed at 0.1 \AA (δx_D). Initially, 20 rotational iterations were carried out to find the lowest eigenmode, but after that four rotational iterations were used at each iteration before the system was translated towards the SP. To determine the change in rotational-force, a finite rotation of 0.005 rad ($\delta\theta_D$) was carried out. After the effective force had been determined, the system was translated using a conjugate gradient algorithm without line search and a maximum of 0.2 \AA (Δx^{max}) in displacement was allowed. Finally, a convergence criterion of 10^{-5} eV/\AA in the magnitude of the total force was used when SPs or minima were determined. When a weaker criterion of 10^{-4} eV/\AA was used, two points were mistakenly identified as SPs in the GB system.

In some cases, SP-searches require an excessive number of iterations to converge and when they finally do, have likely moved to a potential energy rim corresponding to a different initial state. In order to reduce the number of wasted force evaluations due to such lengthy searches, three types of abortion criteria have been implemented.




FIG. 3.3. The relevant saddle points for the GB system determined within 2000 converged saddle point searches. Different cutoff criteria are shown accounting for mechanisms with a barrier below 0.73 eV (92 relevant saddle points), 0.64 eV (65), 0.52 eV (39), 0.40 eV (27) and 0.28 eV (12), respectively, rightmost axis. A total of 459 distinct saddle points were located and 3236 searches were carried out, leftmost axis.

It is important to realize that a search can enter a high energy region and then come back down to converge on a low lying SP. It is, therefore, not advisable to abort a search when the energy reaches the cutoff energy for a relevant SP (1.2 eV for HI and 0.75 eV for GB). The choice of energy cutoff used here, 20 eV, is more than an order of magnitude higher. After a search has reached such a high energy, it will likely not converge to a SP that connects with the initial state. The same consideration applies to a search that has already included a large number of iterations without converging. Searches are, therefore, aborted after 500 iterations. Finally, there is also a maximum number of iterations allowed for a search path in the positive region. This condition is applied because searches sometimes get stuck in the positive region and circle around the minimum without the lowest eigenvalue becoming negative, i.e., into a basin of attraction for a first order SP. A maximum of 50 iterations were allowed in the positive region.

4. Initial displacements. Given the minimum-mode following algorithm for climbing up from a given initial point on the surface to a first order SP, the sampling of the variety of relevant SPs in a set of searches depends on the way the initial points are generated. We have tested three different distributions for the initial displacements: (1) along the eigenmodes of the Hessian evaluated at the initial state minimum, (2) a random distribution with a maximum cutoff, and (3) a Gaussian distribution.

When method (1) was applied, the Hessian matrix at the minimum was first evaluated and the eigenvalue problem solved. From the corresponding eigenvector $\hat{\mathbf{v}}_i$, a more simple displacement vector \mathbf{D}^{disp} was created by keeping only the n components largest in absolute value of the $3N$ components and setting others equal to zero:

$$(4.1) \quad d_j^{disp} = \begin{cases} v_j & \text{if } |v_j| \text{ is one of the } n \text{ largest values of } |v_k|, k = 1, 2, \dots, 3N, \\ 0 & \text{otherwise.} \end{cases}$$

The displacement vector with $n = 30$ was then scaled to a length of $\delta^{disp} = 0.65 \text{ \AA}$,

furthermore, it was assured that no atom was displaced by more than 0.2 Å:

$$(4.2) \quad \mathbf{D}^{disp} = \delta^{disp} \frac{\mathbf{d}^{disp}}{\|\mathbf{d}^{disp}\|}.$$

An advantage of this method is that no a priori knowledge of which degrees of freedom are most likely displaced in low lying SPs is needed.

In systems consisting of a large number of atoms, the displacement of coordinates in going from the minimum to a low energy SP typically involve only a small subset of the atoms, mostly atoms that are not optimally coordinated by their neighbors (due to the local nature of the bonding between atoms, a consequence of the “near sightedness” of the electrons [20]). It is, therefore, most efficient to focus the initial displacements on such atoms. This can be done explicitly in methods (2) and (3), while method (1) will implicitly do this since low frequency vibrational modes involve mostly displacements of poorly coordinated atoms. In the HI system, this means atoms with few neighbors, i.e., the island atoms. In the GB system, this means atoms that are locally not in a FCC environment, i.e., atoms that are in the GB layer. The initial displacement of the atoms in methods (2) and (3) is generated by selecting one such atom at random and then assigning a small displacement to all atoms within a spherical region centered at the atom. It should be pointed out, however, that it is quite possible to find SPs involving, for example, rearrangements of the island atoms in the HI system after displacing initially only atoms in a subsurface layer. But, the number of iterations needed to converge on the SP will then be very large. One example of a selection of atoms for initial displacement in the HI system is shown in Figure 3.1. The non-FCC atoms in the GB system, which are potential centers for spherical displacement regions, are shown and highlighted in Figure 3.2. In method (2), all atoms within a sphere of radius 4.0 Å were displaced by drawing random numbers from a continuous uniform distribution covering the interval -0.2 Å and 0.2 Å, while in method (3) the displacements were drawn from a Gaussian distribution:

$$(4.3) \quad \varphi(x)_{\sigma^2} = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right),$$

where the standard deviation, σ , was set to 0.2 Å.

After the initial displacement has been determined, a pair of replicas of the system—a dimer—is constructed separated by a small distance, here 0.1 Å. The direction of the separation vector is chosen to be along the eigenmode when the initial displacement was chosen according to method (1). In the other two cases, methods (2) and (3), a random initial direction of the dimer was used.

To test and compare the performance of the three methods, 500 converged SP-searches were carried out for both the HI and GB systems. The results for the HI system are shown in Figure 4.1. When method (3), the Gaussian distribution, was used, a total of 79 SPs had been found after 500 converged searches. This included all 12 relevant SPs. Method (2), the random distribution, led to the identification of 46 SPs including 10 relevant ones. The eigenmode scheme gave the largest number of SPs, 82, but this included only 7 of the 12 relevant ones. As seen from the figure, no new SPs were determined in the last 168 searches based on that method. The computational effort, which is best quantified by counting the number of force evaluations needed to do the calculation (since the force evaluation is by far the most demanding part of the calculation), was quite similar for the three types of initial displacements, about




FIG. 4.1. Results of saddle point searches for the HI system using three different methods for initial displacements. Black: Gaussian distribution. Blue: random distribution. Red: displacements along eigenmodes. The solid lines show how many different saddle points have been found after a given number of searches (left vertical axis). Dashed lines show the number of saddle points with energy less than 1.20 eV, i.e., relevant saddle points (right vertical axis). Only with the Gaussian distribution are all 12 relevant saddle points found. Also, for a rather small number of SP-searches, the largest number of different saddle points is found when the initial displacements are Gaussian distributed.




FIG. 4.2. Results of saddle point searches for the GB system using three different methods for initial displacements. Black: Gaussian distribution. Blue: random distribution. Red: displacements along eigenmodes. The solid lines show how many different saddle points have been found after a given number of searches (left vertical axis). Dashed lines show the number of saddle points with energy less than 0.73 eV, i.e., relevant saddle points (right vertical axis). Only with the Gaussian distribution are all 59 relevant saddle points found. However, for a rather small number of SP-searches, the largest number of different saddle points is found when the initial displacements are randomly distributed, but Gaussian displacements give almost as many.

2000 force calls for each converged search. From this test it is clear that method (3) using the Gaussian displacements is the best option out of the three.

The GB system is much richer than the HI system since many more SPs are present both in low and high energy region and this system is, therefore, more challenging for the search methods. The results are shown in Figure 4.2. Again, the Gaussian scheme outperforms the other two by determining 217 unique SPs in 500 searches, including all 59 relevant SPs. The eigenmode scheme appears to saturate as it did for the HI system. The last 133 searches did not result in any new finds. Another disadvantage of the eigenmode scheme is that many of the searches did not converge. A total of 4042 searches were needed to collect 500 searchers that converged on SPs that connected to the initial state minimum. For the Gaussian distribution, 825 searches were needed and 939 searches were needed for the random distribution. The most common reason for a failed search is convergence to a SP that does not connect with the initial state when a minimization is carried out. The random scheme gives the highest number of different low energy SPs if only a limited number of searches are performed, only after 400 searches does the Gaussian scheme give more results. The number of force calls needed for each successful SP-search is larger for the GB system than the HI system, about 5000 for the Gaussian and random distributions, but much larger, nearly 20000 for the eigenmode scheme because of the large number of failed searches.

5. The skipping-path method. An important aspect of the minimum-mode following method is that the low lying SPs are most likely to be found [18]. A necessary consequence of this is that the same SPs tend to be found several times when multiple searches are carried out. In order to reduce the computational effort that is wasted because of this redundancy, we have implemented a simple method to detect when a search has entered the basin of attraction of a SP that has already been found in a previous search or a region that appears to be far from such regions. We will refer to the method as a “skipping-path” method and it is illustrated in Figure 2.1. When a search exits the positive region, the value of a few coordinates that have changed the most from the initial minimum are stored. This is repeated each time the search path has advanced by a certain distance. The region on the surface specified by these points and their neighborhood become part of a chart that contains information about basins of attraction for known SPs and regions that do not correspond to such regions. A spherical region with a small, assigned radius around these points is used to estimate the neighborhood and define a marked region. In a subsequent SP-search, the values of the coordinates with largest displacement are compared with these stored values and the search is terminated as soon as it has entered a marked region. The intermediate points in this search path are then added to the chart. This is illustrated in Figure 5.1. While the basin of attraction of a SP is large when many degrees of freedom are involved, the searches tend to retrace the same area and this method turns out to be quite successful in the test problems studied here, as shown below. The increase in computational effort is negligible and the memory needed to store the information is small. Information about unsuccessful searches, terminated by one of the abortion conditions, is also stored and used to terminate subsequent searches that enter the same region on the surface. Usually these correspond to basins of attraction to SPs that do not connect to the given initial state.

When collecting information about intermediate points along a search, only a few coordinates, i.e., those that have changed most, are stored. We refer to this subset as finger print (FP) coordinates. For the test problems studied here, this subset is




FIG. 5.1. Illustration of three saddle points corresponding to three different mechanisms for adatom diffusion on a solid surface, and the skipping-path method for reducing wasted computational effort when a saddle point search converges to a known saddle point. (1) Saddle point corresponding to a diffusion hop of an adatom on a solid surface. The original search is shown (as in Figure 2.1) as well as three subsequent searches that were terminated when they came within a given tolerance of points stored from a previous search. (2) Saddle point corresponding to a concerted displacement of an adatom with a surface atom. An original search and a second terminated search are shown. While this basin of attraction involves the same finger print coordinates as (1), namely coordinates of atoms 1 and C, the two can be distinguished by the distance to stored points. (3) Saddle point corresponding to a concerted displacement process involving three atoms. This process can already be distinguished from (1) and (2) by the identity of finger print coordinates, here atoms 1 and D.

defined in terms of coordinates of atoms and it turns out that the number of atoms needed, N_{FP} , is around 7. A smaller set appears to be insufficient to distinguish between SPs that are close in state-space.

A more detailed description of the algorithm is as follows. A flowchart is shown in Figure 5.2. During a minimum-mode following climb outside the positive region, where at least one eigenvalue of the Hessian is negative, the FP-coordinates that have changed most from the initial state minimum are identified. The index of these coordinates is compared with those of the known SPs. In the case that a match is found, the current values of the FP-coordinates are compared with the stored ones, first to check whether the distance from one of the stored points, i , is less than D :

$$(5.1) \quad D^2 < \sum_{h=N_{FP}} (r_{h,i} - r_h)^2,$$

and then to check whether each one of the FP-coordinates, h , are within d of the stored point:

$$(5.2) \quad d^2 < (r_{h,i} - r_h)^2.$$

Once a search has been completed, either by being terminated by these tests, or by reaching a new SP or being aborted, the intermediate values of the FP-coordinates are stored and assigned to the final result of the search. However, all configurations in the positive region are discarded and, furthermore, if the positive region is revisited during a search, all configurations prior that point are discarded. During repeated searches, the list of configurations assigned to the basin of attraction for a SP gets extended and the probability that a subsequent search reaches a mapped out region increases.

5.1. Performance analysis. Two important criteria for the success of the method is (1) success in the reduction of computational effort needed to find new SPs and (2) possible failure as a search is terminated prematurely. To analyze the performance, a set of SP-searches were first carried out without the use of the method.




FIG. 5.2. Flowchart illustrating the algorithm for the skipping-path method for terminating a search when a region has been entered which is known to be part of the basin of attraction of a known saddle point. Left column: The current point is compared to the m known saddle points. The degrees of freedom that have been displaced most as compared with the initial minimum, k , are compared with the h “finger print” coordinates characterizing each saddle point. If there is a match, the distance from the current point from each of the n termination points is calculated. If it is less than D , (5.1), then the value of each finger print coordinate is compared and if the difference is less than d , (5.2), the saddle point search is terminated.

The initial displacements were Gaussian distributed and the random number seeds stored. The set contains a total of 100 successful SP-searches. Then, searches were restarted from the same random number seeds which would lead to retracing of these searches, but now the skipping-path method was applied and the results compared. A range of values for the parameters used to terminate a search were tested, values of D up to 5.0 \AA were tried and the number of FP-coordinates, N_{FP} , was varied from five to ten. For both test problems, $N_{FP} = 7$ turned out to give good results. For the HI system $D = 0.50 \text{ \AA}$ was chosen whereas for the GB system, where more atoms are involved in the transitions, a larger value of 0.75 \AA was found to be safe. The larger termination distance could be used since the different mechanism’s FP-atoms served as more distinct keys; see Figure 5.1 (m) and (1–2). For the inspected parameter sets d was given by D/N_{FP} .

The results for the HI system are shown in Figure 5.3. The first thing to notice is how rapidly the computational effort needed to find a new SP increases as the sampling of SPs progresses. The reason is that the searches converge more often to known SPs the more SPs have been found. It takes four times as much computational effort (measured by the number of gradient evaluations) to identify SP number 80 as it took to find SP number 20. The skipping-path method reduces this redundancy. Already for SP number 20, the skipping-path method reduces the computational effort to nearly a half, but for SP number 80 the reduction is to a third as compared with the original nonterminated searches. The use of seven FP-atoms is optimal. When




FIG. 5.3. Computational effort measured in units of gradient evaluations when searching for saddle points in the HI system with and without the skipping-path method. Solid line: Set of original, unterminated searches giving a total of 100 distinct saddle points. Dashed and dotted lines: Searches using skipping-path method with 6, 7, or 8 finger print atoms. A significant reduction in computational effort due to reduced redundancy is obtained with the skipping-path method, down to a third of the original effort for saddle point number 80. The graphs are smoothed using Gaussian averaging, so no values are given for the intervals $[0, 14]$ and $[86, 100]$. The inset shows the percentage of computational effort as compared with the original, nonterminated searches. The best performance is obtained when using 7 finger print-atoms (FP-atoms).

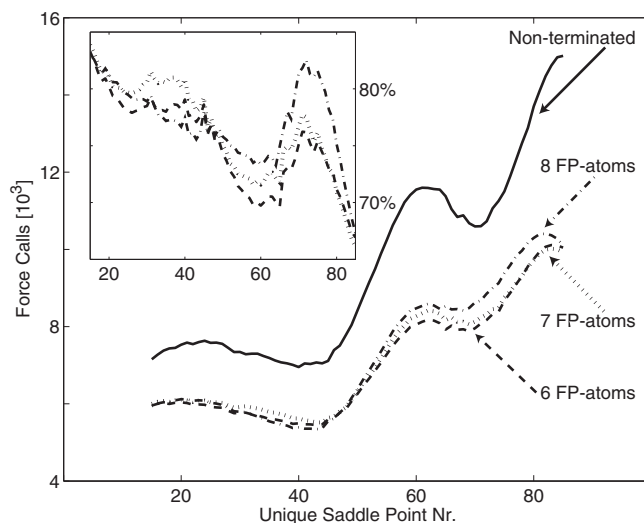


FIG. 5.4. Computational effort measured in units of gradient evaluations when searching for saddle points in the GB system. Solid line: Set of original, unterminated searches giving a total of 100 distinct saddle points. Dashed and dotted lines: Searches using skipping-path method with 6, 7, or 8 finger print atoms. As the GB system includes many saddle points, the skipping-path method does not give as large a reduction in computational effort as in the HI system when the first 100 saddle points are determined. A thorough sampling of the saddle points should include more searches, however, and then more savings are obtained.




FIG. 5.5. Computational effort measured in units of gradient evaluations needed to find saddle points in the GB system. Here, the sampling is continued until the lowest saddle point has been found ten times. The curves are smoothed using Gaussian averages. The thin vertical lines mark when a search converges on the lowest saddle point. The computational effort is reduced to 39% when the last saddle point is found as compared with the original, nonterminated searches.

six FP-atoms were used, one SP with energy 1.195 eV was missed as a search that originally converged on it got terminated and assigned to a different SP. With seven FP-atoms this did not occur. When more atoms were included in the FP-coordinates, the stopping criterion was stricter and the searches not terminated as early, so the savings in computational effort were slightly smaller.

Similar analysis is shown for the GB system in Figure 5.4. This is a more difficult system than the HI system. The atomic disorder at the GB leads to the existence of many different SPs and some are quite close in energy and in state-space. After 100 different SPs have been found, the searches are still not converging too often to known SPs and the reduction in computational effort by using the skipping-path method is only about 30%. No SPs were missed in these calculations by premature termination of a search. A more complete sampling of the SPs is, however, probably required in such a system. One possible criterion for stopping the sampling is to require that the lowest SP has been found ten times. The results of such a calculation is shown in Figure 5.5. A total of 1056 distinct SPs had been determined by the time the lowest one had been found ten times. The computational effort of finding new SPs was then down to a third when the skipping-path method was used as compared with the original, nonterminated searches. One SP with energy of 0.73 eV, right at the boundary of relevant SPs as defined above, was missed when the skipping-path was used with the parameter values listed above.

6. Summary. A study of the efficiency of repeated searches for finding low lying saddle points connected to a given initial state minimum is presented. Three different methods for choosing initial displacements from the minimum and a new method for reducing redundancy in the searches is presented. Two applications were studied, both involving thermally activated transitions in atomic scale models of solids: A heptamer island on a crystal surface and a boundary between two crystalline grains. The

results show that a Gaussian distribution of initial displacements outperforms random displacements and displacements along eigenvectors. The Gaussian distribution gives the widest diversity of saddle points, has highest success in finding low lying saddle points, leads to searches that converge most often, and has lowest computational cost per successful search.

A method is presented that reduces wasted computational effort when a search converges on a saddle point that has previously been found or does not converge within the given constraints. The method is based on the idea that a search should be terminated when it enters the basin of attraction of any of the previously found saddle points, or regions that are known not to be of interest (for example, too high in energy, or do not connect with the given initial state). Coordinates of points along search paths outside the positive region are stored and when a subsequent search comes close enough to one of those points, the search is terminated and the new, intermediate coordinates along the new path assigned to the appropriate region. Application of the method to the two test systems shows that it can reduce the computational cost of finding new saddle points down to a third. The method is particularly useful when a thorough search for all relevant saddle points is carried out. For a given amount of computational effort, this method leads to a more complete sampling of the relevant saddle points.

Acknowledgment. We thank Prof. G. Henkelman, Dr. J.-C. Berthet, and Dr. P. Klüpfel for helpful discussions.

REFERENCES

- [1] R. W. BUTLER, *Saddlepoint approximations with applications*, Camb. Ser. Stat. Probab. Math., Cambridge University Press, Cambridge, UK, 2007.
- [2] E. WIGNER, *Effects of the electron interaction on the energy levels of electrons in metals*, Trans. Faraday Soc., 34 (1938), pp. 678–685.
- [3] H. EYRING, *The activated complex in chemical reactions*, J. Chem. Phys., 3 (1935), pp. 107–116.
- [4] G. H. VINEYARD, *Frequency factors and isotope effects in solid state rate processes*, J. Phys. Chem. Solids, 3 (1957), pp. 121–127.
- [5] G. HENKELMAN, B. UBERUAGA, AND H. JÓNSSON, *A climbing image nudged elastic band method for finding saddle points and minimum energy paths*, J. Chem. Phys., 113 (2000), pp. 9901–9904.
- [6] G. HENKELMAN AND H. JÓNSSON, *Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points*, J. Chem. Phys., 113 (2000), pp. 9978–9985.
- [7] P. J. FEIBELMAN, *Diffusion path for an Al adatom on Al(001)*, Phys. Rev. Lett., 65 (1990), pp. 729–732.
- [8] M. VILLARBA AND H. JÓNSSON, *Low-temperature homoepitaxial growth of Pt(111) in simulated vapor deposition*, Phys. Rev. B, 49 (1994), pp. 2208–2211.
- [9] G. HENKELMAN AND H. JÓNSSON, *Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table*, J. Chem. Phys., 115 (2001), pp. 9657–9666.
- [10] G. HENKELMAN AND H. JÓNSSON, *Multiple time scale simulations of metal crystal growth reveal importance of multi-atom surface processes*, Phys. Rev. Lett., 90 (2003), article 116101.
- [11] A. PEDERSEN, G. HENKELMAN, J. SCHIÖTZ, AND H. JÓNSSON, *Long time scale simulation of a grain boundary in copper*, New J. Phys., 11 (2009), article 073034.
- [12] A. PEDERSEN AND H. JÓNSSON, *Simulations of hydrogen diffusion at grain boundaries in aluminum*, Acta Mater., 57 (2009), pp. 4036–4045.
- [13] A. PEDERSEN AND H. JÓNSSON, *Distributed implementation of the adaptive kinetic Monte Carlo method*, Math. Comput. Simul., 80 (2010), pp. 1487–1498.
- [14] H. JÓNSSON, G. MILLS, AND K. W. JACOBSEN, *Methods for finding saddle points and minimum energy paths*, in Classical and Quantum Dynamics in Condensed Phase Simulations, E. D.

- B. J. Berne, G. Ciccotti, and D. F. Coker, eds., World Scientific, River Edge, NJ, 1998, pp. 385–404.
- [15] C. J. CERJAN AND W. H. MILLER, *On finding transition states*, J. Chem. Phys., 75 (1981), pp. 2800–2807.
- [16] A. BANERJEE, N. ADAMS, J. SIMONS, AND R. SHEPARD, *Search for stationary points on surfaces*, J. Phys. Chem., 89 (1985), pp. 52–57.
- [17] J. P. K. DOYE AND D. J. WALES, *Surveying a potential energy surface by eigenvector-following: Applications to global optimisation and the structural transformations of clusters*, Z. Phys. D., 40 (1997), pp. 194–197.
- [18] G. HENKELMAN AND H. JÓNSSON, *A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives*, J. Chem. Phys., 111 (1999), pp. 7010–7022.
- [19] Y. KUMEDA, D. WALES, AND L. J. MUNRO, *Transition states and rearrangement mechanisms from hybrid eigenvector-following and density functional theory*, Chem. Phys. Lett., 341 (2001), pp. 185–194.
- [20] W. KOHN, *Nobel lecture: Electronic structure of matter-wave functions and density functionals*, Rev. Mod. Phys., 71 (1999), pp. 1253–1266.
- [21] D. MEI, L. XU, AND G. HENKELMAN, *Dimer saddle point searches to determine the reactivity of formate on Cu(111)*, J. Catal., 258 (2008), pp. 44–51.
- [22] A. NAKANO, *Pathfinder: A parallel search algorithm for concerted atomistic events*, Comput. Phys. Commun., 176 (2007), pp. 292–299.
- [23] R. A. OLSEN, G. J. KROES, G. HENKELMAN, A. ARNALDSSON, AND H. JÓNSSON, *Comparison of methods for finding saddle points without knowledge of the final states*, J. Chem. Phys., 121 (2004), pp. 9776–9792.
- [24] R. MALEK AND N. MOUSSEAU *Dynamics of Lennard-Jones clusters: A characterization of the activation-relaxation technique*, Phys. Rev. E, 62 (2000), pp. 7723–7728.
- [25] J. PALIS AND W. DE MELO, *Geometric Theory of Dynamical Systems*, Springer-Verlag, Heidelberg, 1982.
- [26] K. W. JACOBSEN, J. K. NØRSKOV, AND M. J. PUSKA, *Interatomic interactions in the effective-medium theory*, Phys. Rev. B, 35 (1987), pp. 7423–7442.
- [27] K. W. JACOBSEN, P. STOLTZE, AND J. K. NØRSKOV, *A semi-empirical effective medium theory for metals and alloys*, Surf. Sci., 366 (1996), pp. 394–402.
- [28] G. HENKELMAN, G. JÓHANNESSON, AND H. JÓNSSON, *Methods for finding saddle points and minimum energy paths*, in Progress on Theoretical Chemistry and Physics, E. D. S. D. Schwartz, ed., Kluwer Academic Publishers, Dordrecht, 2000, pp. 269–300.
- [29] A. S. CLARKE AND H. JÓNSSON, *Structural changes accompanying densification of random hard sphere packings*, Phys. Rev. E, 47 (1993), pp. 3975–3984.