Exploring Potential Energy Surfaces with Saddle Point Searches

33

Vilhjálmur Ásgeirsson and Hannes Jónsson

Contents

V. Ásgeirsson
Science Institute of the University of Iceland, Reykjavík, Iceland
e-mail: via9@hi.is

H. Jónsson (✉)
Faculty of Physical Sciences, University of Iceland, Reykjavík, Iceland

Department of Energy Conversion and Storage, Technical University of Denmark, Lyngby, Denmark
e-mail: hj@hi.is

W. Andreoni, S. Yip (eds.), *Handbook of Materials Modeling*,
https://doi.org/10.1007/978-3-319-44677-6_28

Abstract

The energy surface of an atomic scale representation of a material contains the essential information needed to determine the structure and time evolution of the system at a given temperature. Local minima on the surface represent (meta)stable states of the system, while first-order saddle points characterize the mechanisms of transitions between states. While many well-known methods make it relatively easy to find local minima, the identification of saddle points is more challenging. In this chapter, methods for finding saddle points are discussed as well as applications to materials simulations. Both doubly constrained search methods, where the final and the initial state minima are specified, and singly constrained search methods, where only the initial state is specified, are discussed. The focus is on a classical description of the atom coordinates, but saddle points corresponding to quantum mechanical tunneling are also mentioned. An extension to magnetic systems where the energy surface depends on the orientation of the magnetic vectors is sketched.

1 Introduction

Atomic scale simulations of materials often involve finding likely arrangements of the atoms as well as identifying the mechanism and estimating the rate of transitions between different arrangements such as diffusion events, migration of dislocations, chemical reactions, and phase transformations. These can be obtained by analyzing the energy surface characterizing the system, i.e., the electronic ground-state energy of the system as a function of the location of the atoms (and possibly also the orientation of the moments of magnetic atoms) (Peters 2017). Each state corresponds to a minimum on the energy surface and the potential well surrounding it. Excited electronic states can, in some circumstances, be important, but the discussion here will focus on the electronic ground state. Even so, the energy surface is complex and multidimensional for most material systems of interest, and it requires efficient tools to navigate on the energy surface and extract the desired information.

Given some initial configuration of the atoms and a method for evaluating the energy and its gradient with respect to atom coordinates, the nearest local minimum can readily be identified with any number of available minimization methods. A significant challenge, however, is to find all the relevant minima for given conditions, such as temperature. The lower the local energy minimum of the state, the more likely the system is to be in that state, according to Boltzmann statistics. The width of the energy well around the minimum and the number of equivalent minima relate to the vibrational and configurational entropy, respectively, and also contribute to the stability of the state. The identification of the most likely state of a system requires navigation on the energy surface, for example, advancing from one energy minimum to another.

One way to discover new minima is to move through regions of dips in the energy ridge, i.e., regions around first-order saddle points on the energy surface. At a first-

order saddle point, the gradient of the energy vanishes, and the Hessian, the matrix of second derivatives, has one and only one negative eigenvalue. The eigenvectors are referred to as modes, and the one corresponding to the negative eigenvalue is the unstable mode. A general feature of energy surfaces for atomic systems is that a lower energy saddle point tends to lead to a lower-energy minimum on the other side (the basis of so-called Brønsted-Evans-Polanyi relation). So, more often than not, a low-energy minimum can be discovered by moving over the energy ridge through a low-energy saddle point. Methods for finding first-order saddle points can, therefore, be used to identify local minima and likely states of the system.

It may, however, not be possible to reach some states, such as the one corresponding to the global energy minimum, on a given time scale because of a large energy barrier. The rate at which the system moves from one state to another is the key quantity that is needed to predict the long time scale evolution of the system. On a given time scale, only transitions that are frequent enough are relevant. Classical dynamics simulations based on Newton's equation of motion can be carried out, but the time scale of such simulations is limited by the vibrational frequency of the atoms and even for the simplest description of the energy surface – an empirical potential function – such a simulation can only span a small fraction of a second, while the challenge is often to predict the evolution of a material over years. As is explained in the following section, the first-order saddle points on the energy surface and their vicinity can be used to estimate the rate of transitions between states of the system. The central focus of this chapter is methods for finding first-order saddle points for the purpose of exploring energy surfaces of material systems.

2 Estimation of Transition Rates

The dynamics of atoms in materials that are relatively stable under conditions of interest can be characterized by a large number of vibrations back and forth, mostly in the close vicinity of the local energy minimum but occasionally involving large excursions caused by rare fluctuations through coupling to the heat bath. Such fluctuations can be large enough for the system to enter a new state. Since the system spends a long time in each state, a Boltzmann distribution of energy is typically established in all degrees of freedom, and the rate at which the system can transform from one state to another can be estimated to a first approximation by using statistical mechanics.

The basic tool for estimating transition rates in materials is transition state theory (TST) (Wigner 1938; Peters 2017). Assuming the atoms can be described as classical particles and that Boltzmann distribution of energy has been established in all degrees of freedom, the lifetime of a given initial state and possible final states of transitions from it can be estimated in a two-step procedure referred to as Wigner-Keck-Eyring (WKE) dynamics (Wigner 1938; Keck 1967; Eyring 1935). In the first step, a dividing surface in configuration space separating the initial state from other states is defined. The dividing surface should be placed in such a way that it lies through regions of the energy surface where the system is least

likely to be found. The dimensionality of the dividing surface is $D - 1$ where D is the total number of degrees of freedom in the system (in three dimensions, $D = 3N$ where N is the number of atoms). By adding a small width around the dividing surface, a subspace of dimensionality D is generated, and this is referred to as the transition state. Transition state theory assumes that if the system makes it to the transition state and has velocity pointing away from the initial state, a transition will occur and the system enter a new state. This basic TST assumption is an approximation. Dynamical trajectories can be reflected back after crossing the transition state and reenter the initial state. They can, in fact, recross multiple times, and each recrossing to the initial state leads to an overestimate of the transition rate in the TST approximation. As a results, TST is guaranteed to give an upper bound to the transition probability. This gives a variational principle for the placement of the dividing surface. The dividing surface that leads to the lowest TST transition rate estimate is the optimal one. The first step of the WKE procedure should, therefore, involve a variational optimization of the dividing surface in addition to giving an approximation to the transition rate, k^{TST}.

The second step of the WKE procedure involves the calculation of short time scale trajectories started from the dividing surface. Since the transition state represents a region of low probability, the system advances quickly from there to either the initial state or some final state. The trajectories in the second WKE step reveal possible final states of the transition. Note that TST does not specify the final state. The trajectories also serve as a means to correct the TST rate estimate and obtain (in principle) the exact estimate of the rate by evaluating a dynamical correction factor, κ (Keck 1967; Voter and Doll 1985). The less the dividing surface was optimized in the first step, the larger the number of trajectories needed in the second step to obtain a statistically converged correction factor. A recrossing of the transition state can occur either because of the shape of the energy surface, such as a curved valley on the energy surface in the final state, or because of a fluctuation from the heat bath resulting in a reversal of the relevant component of the velocity. Both effects are included in the short dynamical trajectories of the second WKE step. The WKE procedure, therefore, provides essentially an exact estimate of the transition rate, $k = \kappa k^{\mathrm{TST}}$. Instead of involving impossibly long dynamical trajectories started near the initial state energy minimum, WKE uses a statistical estimate of the probability of making it to the transition state and then short time scale trajectories for the evolution from the transition state to possible final states.

The full implementation of the WKE procedure is in general challenging because of the need to represent and optimize a high-dimensional dividing surface. Efficient tools for such calculations have not been developed yet. The simplest approach is to use a hyperplanar dividing surface where both the location and the orientation are variationally optimized (Jóhannesson and Jónsson 2001). The orientation specifies which atoms are displaced and by how much at the transition state. The orientational optimization identifies the optimal transition mechanism. A single hyperplane is, however, in general not sufficient to specify the full dividing surface, and either a curved surface (Ciccotti et al. 1995) or a mosaic of hyperplanar segments is needed (Bligaard and Jónsson 2005).

The most commonly used form of TST involves an approximation to the energy surface. By expanding the energy to second order in the vibrational normal modes at both the initial state minimum and at a first-order saddle points on the energy rim surrounding the initial state, an approximation to the rate constant is obtained and referred to as harmonic TST (HTST) (Vineyard 1957):

$$k^{\mathrm{HTST}} = \frac{\prod_j^D \nu_{i,\mathrm{m}}}{\prod_j^{D-1} \nu_{i,\mathrm{s}}} \exp\left(-(E_\mathrm{s} - E_\mathrm{m})/k_\mathrm{B}T\right) \tag{1}$$

where E_m and E_s are the values of the energy the $\nu_{i,\mathrm{m}}$ and $\nu_{i,\mathrm{s}}$ are the vibrational frequencies at the minimum and at the saddle point, k_B is the Boltzmann constant, and T is the temperature. This expression agrees with the empirical Arrhenius dependence of the rate constant on temperature:

$$k = A \exp\left(-E_\mathrm{a}/k_\mathrm{B}T\right) \tag{2}$$

showing that the activation energy, E_a, is the energy difference between the first-order saddle point and the initial state minimum and the prefactor is related to the width of the potential energy well at the minimum and the energy valley at the saddle point.

The vibrational frequencies are obtained from the positive eigenvalues of the Hessian. The number of vibrational modes at the saddle point is one less than at the minimum because the dividing surface does not contain the unstable mode, the vibrational mode corresponding to negative eigenvalue. The ratio of the products of vibrational frequencies represents the relative vibrational entropy of the transition state and the initial state. Effectively, the dividing surface is approximated in HTST by a hyperplane going through each one of the first-order saddle points on the energy rim surrounding the initial state minimum with normal vectors pointing in the direction of unstable modes. Strictly speaking these should be finite hyperplanar segments, but the full hyperplane is included for each saddle point in order to obtain a simple analytical expression for the rate constant.

The HTST approximation can give an accurate estimate of the rate constant for materials at not too high or too low temperature. For low enough temperature, nuclear quantum effects need to be accounted for (see below). The energy surface must be smooth enough and the first-order saddle points separated by regions of high enough energy for the harmonic approximation to accurately represent the energy surface in the neighborhood of the extrema where the Boltzmann statistical weight is appreciable. This means that the first-order saddle points on the energy ridge must be separated by second-order saddle points that are significantly higher in energy, by several $k_\mathrm{B}T$, compared to the first-order saddle points. Below, we will describe a method that can be used to check whether the energy landscape satisfies this criterion.

The main challenge is to find the relevant saddle points. The lower the energy of a saddle point, the more probable the corresponding transition is. While the number

of saddle points surrounding a given minimum on the energy surface can easily be enormous, it is enough to find only the lower lying saddle points in order to identify the relevant transitions. The task of finding saddle points can be divided into two categories. In the more general category, only the initial state is known, and the task is to find all relevant saddle points on the energy ridge surrounding the energy minimum. In some cases, the final state of interest is also known, and the task can then be cast in terms of finding the minimum energy path (MEP) between the initial and final state minima. The MEP is a path for which the energy is at a minimum in directions orthogonal to the path. A maximum along the MEP corresponds to a first-order saddle point. The activation energy for the overall transition is then given by the highest rise in energy along the path. We will first discuss methods for finding the MEP when both initial and final states are given. Then, we discuss methods for the more challenging problem of finding relevant saddle points when only the initial state is specified.

3 Initial and Final States Specified, CI-NEB

A path between a given initial and final state can be discretized by creating a number of replicas of the system and arranging them in such a way as to trace out a curve between the states. We will refer to the discretization points as "images" of the system. An image labeled i is specified by the coordinates of the atoms, R_i. The task is to first generate a reasonable initial path and then apply an optimization algorithm to iteratively move the images to the MEP. There can exist more than one MEP between the given minima, an issue that is addressed below. At this point, the task is just to find the MEP closest to the initial path.

The simplest way to generate an initial path is to make a linear interpolation in Cartesian coordinates between the two minima. This method is frequently used in materials simulations. It can, however, lead to unphysical configurations of the atoms. For example, two atoms may end up being very close to each other leading to a strong repulsive interaction. When the calculation of the energy and atomic forces is carried out using an electronic structure method, such strong overlap may slow down or even prevent the self-consistency calculation from converging. Furthermore, the initial path generated by a linear interpolation may be far from any MEP. A better approach is to generate the initial path by taking the pairwise distances between atoms into account. This can be done with the image dependent pair potential (IDPP) method (Smidstrup et al. 2014). There, the pairwise distances between neighboring atoms are interpolated linearly between the two minima and an initial path generated to match those distances as closely as possible. Since there are many more pairwise distances than atom coordinates, the matching can only be approximate and the initial path is found by minimizing the sum of squared deviations (Smidstrup et al. 2014). Another approach that avoids the problem of overlapping atoms is to generate the initial path by linear interpolation in internal coordinates (Goumans et al. 2009).

In order to find the MEP closest to the initial path, the images are moved iteratively in a direction obtained from the atomic forces, i.e., negative gradient of the energy with respect to atomic coordinates, $-\nabla E(R_i)$. But only the component perpendicular to the path should be used to modify the shape of the path. A force projection is, therefore, required based on an estimate of the local tangent to the path. While it seems natural to estimate the tangent at a given image, i, from the coordinates of atoms at the two adjacent images, $i-1$ and $i+1$, it turns out to be numerically more stable to use only the coordinates of the neighboring image that has the higher energy (Henkelman and Jónsson 2000). Letting the normalized tangent be denoted $\hat{\tau}_i$, the force acting on the shape of the path at image i is given by

$$F_i^{\mathrm{p}}|_{\perp} = -\nabla E(R_i) + \nabla E(R_i) \cdot \hat{\tau}_i \, \hat{\tau}_i . \tag{3}$$

By displacing the images in the direction of $F_i^{\mathrm{p}}|_{\perp}$, the perpendicular component of the atomic forces will vanish, and the images are then on the MEP. At each iteration of the optimization, the atomic forces of all images need to be calculated. The evaluation of the energy and atomic forces is typically the most computationally intensive part of the calculation. However, the calculations can readily be performed simultaneously, using parallel or distributed computing.

It is also necessary to specify how the images are distributed along the path. If the distribution is not controlled, the images tend to slide down to the local energy minima. The distribution is commonly controlled using a restraint method where a harmonic spring acts between adjacent images in the direction parallel to the path:

$$F_i^{\mathrm{s}}|_{\parallel} = \left(k_i^{\mathrm{s}} \left(R_{i+1} - R_i\right) - k_{i-1}^{s} \left(R_i - R_{i-1}\right)\right) \cdot \hat{\tau}_i \, \hat{\tau}_i . \tag{4}$$

The spring constants, k_i^{s}, can be chosen to produce a desired distribution of the images, for example, with higher density in regions of higher energy (Henkelman et al. 2000a). Most often, though, the spring constants are chosen to have the same constant value resulting in an even distribution of images along the path. For optimal convergence rate, the magnitude of the spring constant should be chosen such that the magnitude of $F_i^{\mathrm{p}}|_{\perp}$ and $F_i^{\mathrm{s}}|_{\parallel}$ is comparable. The effective force acting on image i in the path is then given by the sum of the two force components:

$$F_i = F_i^{\mathrm{p}}|_{\perp} + F_i^{\mathrm{s}}|_{\parallel} . \tag{5}$$

The force projections, referred to as "nudging," separate the distribution of the images along the path from the displacements of images affecting the shape of the path. The distribution of the images along the path can also be controlled using a constraint method based on an estimate of the total length of the path (Weiqing and Vanden-Eijnden 2002).

The most important part of the MEP is the highest energy point along the path, i.e., the highest first-order saddle point along the path connecting the initial and final

states. The activation energy of transitions between the two states can be estimated from the first-order saddle point using HTST. In order to obtain an accurate estimate of the saddle point, one of the images can be forced to "climb" upward along the path and converge to the highest saddle point. This image is referred to as the "climbing image." Generally, the highest energy image after a few normal iterations (or a certain user-defined threshold) is chosen to become the climbing image, $i = c$. The effective force acting on this image is (Henkelman et al. 2000a)

$$F_c = -\nabla E(R_c) + 2\left(\nabla E(R_c) \cdot \hat{\tau}_c\right)\hat{\tau}_c \tag{6}$$

The assumption here is that the tangent estimate at the climbing image gives an accurate estimate of the direction of the unstable mode at the saddle point. It is used to transform the gradient of the energy around the first-order saddle point so as to mimic the gradient around a minimum. An ordinary minimization algorithm can then be used with the transformed gradient to converge the climbing image on the first-order saddle point. This method is referred to as the climbing image nudged elastic band (CI-NEB) method (Henkelman et al. 2000a).

3.1 Optimization of the Path

Various iterative algorithms can be used to zero the effective force and hence move the images to the MEP.

Note, however, that the objective function, the function that gives the effective force by differentiation, is not known explicitly because of the force projections, i.e., the nudging. The optimization algorithm should, therefore, be based only on the forces, not on the objective function itself. This, for example, makes the usual implementation of line search approaches not applicable and hence complicates the implementation of some of the optimization methods.

In early implementations of the method (Mills et al. 1995; Jónsson et al. 1998), a velocity projection method based on the velocity Verlet dynamics algorithm (Andersen 1980) was used. There, the velocity, v, is zeroed in directions orthogonal to the force if $v \cdot F > 0$ and in all directions if $v \cdot F < 0$. A parametrized extension of the algorithm has been devised (Bitzek et al. 2006). The velocity projection algorithm gives fast and reliable performance in the initial stages where the images may be located far from the MEP and the effective forces large (especially when linear interpolation is used to generate the initial path). However, the convergence close to the MEP can be faster by using more advanced approaches (Chu et al. 2003; Sheppard et al. 2008). The limited memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) (Nocedal 1980) algorithm has proved to be efficient (Sheppard et al. 2008). If the second derivatives of the energy are available and can be computed in each iteration without large effort (which however is generally not the case), a highly efficient optimization algorithm can be used (Bohner et al. 2013).

An example of a CI-NEB calculation on a modified two-dimensional Müller-Brown energy surface using the L-BFGS algorithm is shown in Fig. 1. The algorithm

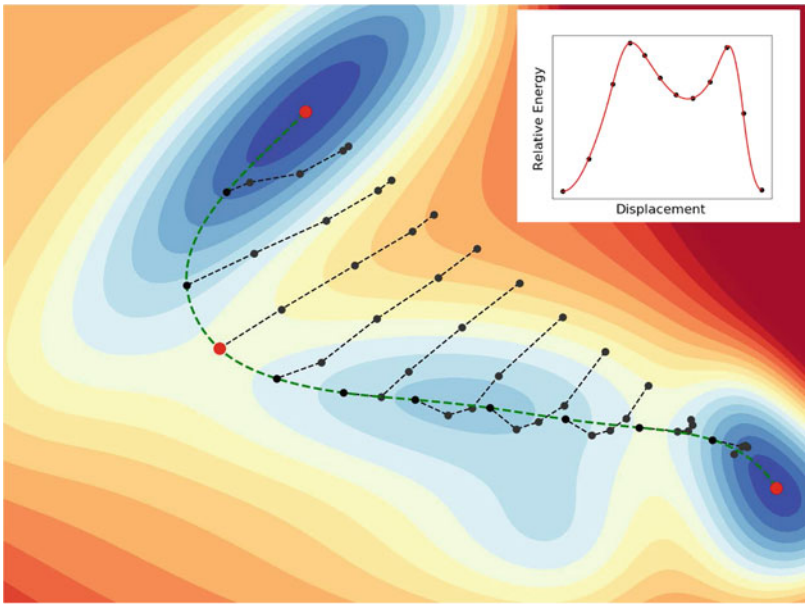

Fig. 1 A climbing image nudged elastic band calculation of the minimum energy path of a model two-dimensional potential surface (a modified Müller-Brown surface (Ásgeirsson and Jónsson 2018)). In addition to the initial and final states, which are fixed (red dots), the path is represented by ten discretization points, referred to as images (black dots). The initial path is generated from a straight line interpolation (in more realistic systems, it is better to use the IDPP method). The position of images is shown after 0, 5, 10, 15, and 21 L-BFGS iterations. The green dashed line shows the minimum energy path. A red dot marks the first-order saddle point to which the climbing image converges. Inset: Energy along the minimum energy path and energy of the images after convergence has been reached

is used here right from the start. A more reliable strategy is to start out with the more stable and conservative velocity projection algorithm and then switch to the faster L-BFGS method once the effective force acting on each atom has dropped below some user-defined threshold, for example, $||F_i|| < 0.5\,\text{eV/Å}$.

A calculation may be considered tightly converged to the MEP when the norm of the force acting on each atom perpendicular to the path $||F_i^{\text{p}}|_\perp||$ is below 0.01 eV/Å for all images. However, the MEP is mainly needed to identify where the highest energy point is and to estimate the tangent to the path at that point. The tight convergence is really only needed for the climbing image, and a looser threshold can be used for all other images, such as $||F_i^{\text{p}}|_\perp|| < 0.1\,\text{eV/Å}$.

When electronic structure methods, such as density functional theory (DFT) or ab initio quantum chemistry, are used in combination with a CI-NEB calculation, it is particularly important to reduce as much as possible the number of energy and atomic force evaluations. Recently, it has been shown that machine learning approaches can reduce the number of energy and force evaluations by an order of magnitude. Both neural networks (Peterson 2016) and Gaussian process regression

(GPR) (Koistinen et al. 2016, 2017) have been applied. There, an approximate energy surface is constructed with a machine learning model using all previous energy and force evaluations carried out by the electronic structure method. A CI-NEB calculation is carried out on the approximate energy surface and the electronic structure calculation then done for each image or only the image where the estimated uncertainty of the approximate energy surface is largest. The approximate energy surface is then refined using the new information and the CI-NEB calculation carried out again, etc. The GPR has the advantage of having a built-in error estimate which can be used to choose in an optimal way which image should be evaluated by the electronic structure method. The machine learning approach helps make optimal use of each electronic structure calculation, while traditional optimization methods, such as the velocity projection algorithm, only use the forces obtained at the current step and disregard all previous force evaluations. The L-BFGS keeps memory of the last M steps (typically $M = 20$) to construct an approximate Hessian matrix, but the machine learning algorithm constructs an approximate surface of a more general shape. As a result, the machine learning approach can use all previous energy and force calculations and reach convergence with fewer electronic structure calculations than L-BFGS.

3.2 Interpolation and Strategy

In order to analyze and visualize the results of a CI-NEB calculation as well as to monitor the progress of the calculation, it is important to not just interpolate the energy of the images but also the derivative of the energy along the path, i.e., the component of the atomic forces parallel to the path. It is convenient to use a cubic polynomial for each interval between adjacent images (Henkelman and Jónsson 2000). An interpolation using the derivative may reveal the presence of an intermediate minimum along the path, while a simple interpolation of only the energy does not. A good strategy in such a case is to locate the energy minimum with a separate minimization calculation and then proceed with CI-NEB calculations of the two segments of the path separately.

In an analogous manner, the atom coordinates can be interpolated to add new images and hence improve the resolution in certain regions along the path (Henkelman and Jónsson 2000). This especially applies to regions of high curvature and rapid changes in the tangent from one image to another which can cause inaccuracies and even non-convergence of the calculation. As any other numerical method that relies on discretization, the CI-NEB method requires the number of images to be large enough. However, instead of doubling the number of images in such a situation, it may be sufficient to simply add an image (or a few images) in the problematic region along the path. Note that the spring constants associated with the new image need to be twice as strong such that the location of the adjacent images does not get disrupted.

When the CI-NEB method is applied to materials where periodic boundary conditions are applied to the simulation cell, the size and shape of the cell can be

made part of the MEP calculation, in addition to the atom coordinates (Sheppard et al. 2012). This is needed if the crystal structure is changing during the transition. When the method is applied to finite systems such as clusters or molecules, it is important to remove the overall translation and rotation of the system from the available degrees of freedom. The optimization of the path could otherwise lead to an artificial lengthening of the path involving translation and/or rotation in order to enable images to slide down from high-energy regions. This can reduce the resolution of the path in regions of high energy and increase the computational effort or even prevent the calculation from converging. A method based on quaternions has been formulated for the purpose of constraining the translation and rotation (Melander et al. 2015).

3.3 Application of CI-NEB and WKE Dynamics

An application of the CI-NEB method is shown in Fig. 2. It illustrates how the method can reveal a mechanism that is entirely different from the one implicitly assumed in the initial path. Also, after the MEP and relevant saddle point had been

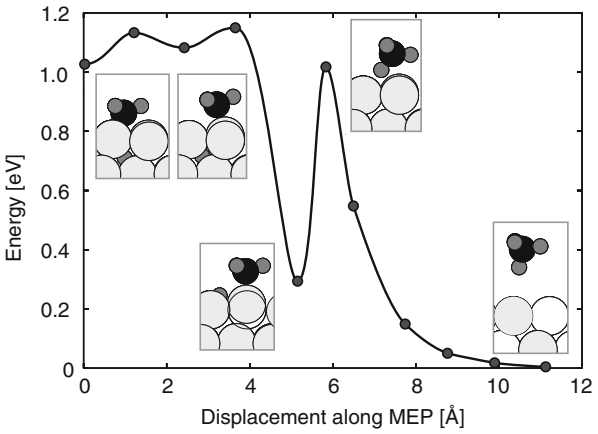

Fig. 2 An application of the climbing image nudged elastic band method to a surface reaction where a CH_4 molecule is formed on a Ni(111) surface from a subsurface H atom and a CH_3 admolecule. While such a transition mechanism had been postulated to be efficient, the calculation shows that a corresponding minimum energy path does not exist on the DFT/PW91 energy surface. While the initial path is consistent with direct addition of the subsurface H atom with the CH_3 admolecule, the optimization of the path reveals that the CH_3 first hops to the side, the H atom emerges to the surface to a state that corresponds to a deep intermediate minimum, and, finally, the H adatom and the CH_3 combine in a regular surface reaction mechanism to form the CH_4, which leaves the surface. The calculation illustrates how the method can find at a complex minimum energy path involving several elementary reaction steps even when the initial path is generated with an incorrect mechanism in mind. (From Henkelman et al. 2006)

located, short time scale dynamics were initiated from the transition state to obtain various information about the atomic dynamics.

There had been some suggestions that subsurface H atoms could play an important role in the formation of CH_4 molecules from CH_3 molecules adsorbed on a catalyst surface. As can be seen from the inset showing the initial state, it seems reasonable that the H atom could add directly to the CH_3 as it moves up from the subsurface site. This mechanism appears to be more facile than the typical surface reaction mechanism since an H adatom on the surface cannot easily approach to the C atom in the CH_3 admolecule. A CI-NEB calculation for this reaction on a Ni(111) surface was carried out using DFT/PW91 to evaluate the energy and atomic forces (Henkelman et al. 2006). The calculation started from a linear interpolation assuming that the direct addition of the subsurface H atom to the C atom in CH_3 could take place. The initial path is simple and mainly involves the displacement of the H atom up from the subsurface site. Optimization of the path, however, revealed that an MEP corresponding to this direct reaction mechanism does not exist on the DFT/PW91 energy surface. Instead, the path converged on a complex multistep mechanism. First, the CH_3 hops to the side to a nearby bridge site, freeing up the surface site for the H atom. Then, the H atom hops to the surface. A deep minimum corresponds to this intermediate configuration. Finally, the H adatom and the CH_3 admolecule combine in the usual surface reaction step to form a CH_4 molecule that leaves the surface.

This is a good example of how a CI-NEB calculation can undo an incorrect preconceived notion of a reaction mechanism. While it can seem plausible that a subsurface H atom attaches directly to a CH_3 admolecule, the results obtained for the DFT/PW91 energy surface do not support this notion. The reason appears to be that the Ni atoms do not have catalytic activity where they are bonded to each other. The catalytic activity is confined to the side where the Ni atoms are exposed, i.e., to the surface.

Assuming the system is thermalized in the deep intermediate minimum corresponding to the H adatom and the CH_3 admolecule, the critical energy barrier for forming CH_4 is the one corresponding to the last saddle point. An interesting observation made from the CI-NEB calculation was a large displacement of the underlying Ni atom, by 0.25 Å. While a similar observation had been made in calculations of CH_4 formation on Ir(111) (Henkelman and Jónsson 2001), most surface reaction calculations were at that time carried out with frozen surface atoms.

This raises interesting questions about the atomic dynamics of the transition. For example, how is the 0.85 eV excess energy gained by the system as it advances past the transition state to a gas phase CH_4 molecule and clean Ni surface partitioned between the Ni atoms and the CH_4 molecule? Also, how is the energy of the CH_4 molecule partitioned among its vibrational modes? For the reverse reaction, dissociative sticking of CH_4, this reveals to what extent translation and the various vibrational modes of the molecule can enhance the sticking probability. To answer these questions, the second step of the WKE approximation was carried out, using the HTST approximation of the transition state. Only a minimal sampling of trajectories was carried out, starting with small displacements from the saddle

point. The calculations showed that 15% of the excess energy go into Ni atom vibrations, while 79% go into translation, 22% into vibrations, and 4% into rotation of the CH_4 molecule. Surprisingly, most of the vibrational energy goes into the stretching modes not the lower frequency deformation modes. This result indicates that excitation of the stretching modes is more likely to enhance sticking than excitation of the deformation modes. The translation is, however, by far most effective in enhancing sticking.

Another interesting dynamics issue that could be addressed in the second step of the WKE procedure is whether the high kinetic energy of the H adatom as it emerges at the surface could enhance the reaction with the CH_3. To address this question, dynamical trajectories started at the transition state for the subsurface to surface hop were calculated, but none of the trajectories led to a combination of the H atom with the CH_3. The conclusion is that subsurface H atoms do not play an important role in hydrogenation of CH_3 admolecules under the conditions simulated.

3.4 Variants of the Method

There are several different variants of the method for finding an MEP when both initial and final states are specified, and a few of them will be mentioned here briefly.

When the shape of the MEP has become clear enough for the approximate location of the highest energy saddle point to be located approximately, computational time can be saved by confining the calculations to the region in the vicinity of the saddle point and halting calculations of images in other regions. There have been several different implementations of this approach. In its simplest form, two images on opposite sides of the highest energy image of the partly converged path are chosen to be the fixed endpoints instead of the local minima, but a better approach is to make the new endpoints follow selected equipotential contours to the MEP (Zhu et al. 2007; Zhang et al. 2016). Another approach is to use two climbing images located on opposite sides of the highest energy image so as to bracket the saddle point (Zarkevich and Johnson 2015).

Alternatively, a single-ended saddle point search, discussed below, can be launched from the highest energy image on the partially converged path (Henkelman et al. 2000b).

When dealing with complex energy surfaces with multiple local minima and curved MEPs, the stability of the calculation can be improved by including part of the component of the spring force perpendicular to the path (Jónsson et al. 1998) or by using the so-called double nudging (Trygubenko and Wales 2004; Sheppard et al. 2008). Also, long paths and large number of images can cause the distance between images to become uneven. In such cases, a combination of the restraint and constraint approaches for distributing the images along the path can be used (Maras et al. 2016, 2017). Alternatively, it is good to divide the path calculation and focus on each segment separately whenever there is a sign of an intermediate minimum, as discussed above. An automatic algorithm for adding new images to the discrete representation of the path has been developed (Kolsbjerg et al. 2016).

4 Only Initial State Specified, MMF

It can be hard to predict the mechanism and even the possible final states that can be formed in a thermally activated transition. Ideally, a computer calculation should be able to reveal such information given only the initial state of the system. This is a harder problem than the one discussed above where the final state is specified as well as the initial state. One option is to raise the temperature of the system and identify new states that get visited, followed by an NEB calculation to find the minimum energy path from the initial state (Sørensen et al. 1996). But the increase in temperature will place larger emphasis on states with high entropy and the transitions relevant at the low temperature of interest may be hard to find. What is needed is a method for climbing up the energy surface from the initial state to converge onto a first-order saddle point. Such searches need to be carried out several times in order to identify all relevant, i.e., low-lying saddle points on the energy ridge surrounding the initial state minimum. Armed with such a method, the long time scale evolution of a materials system can be simulated, as discussed in chapter "The Chapter on AKMC Methods".

A first-order saddle point differs from a local minimum in that the Hessian matrix has one negative eigenvalue, corresponding to the unstable mode. The atomic forces near a first-order saddle point can be made to mimic atomic forces near a local minimum by reversing the component of the force in the direction of the unstable mode. More generally, let $\hat{\lambda}$ be a unit vector along the eigenvector corresponding to the lowest eigenvalue of the Hessian, the minimum mode (irrespective of the sign of the lowest eigenvalue). If the system is made to follow an effective force given by

$$F^{\text{eff}}(R) = -\nabla E(R) + 2\left(\nabla E(R) \cdot \hat{\lambda}(R)\right)\hat{\lambda}(R) \tag{7}$$

then convergence will be reached at a first-order saddle point. This is analogous to the force on the climbing image in the CI-NEB algorithm except that $\hat{\lambda}$ is used here instead of the path tangent, which is not known in this case. If the second derivatives of the energy can be evaluated easily, then the Hessian matrix can be constructed and its eigenvalues and eigenvectors computed quite readily, although the computational effort can be significant for large systems, as it scales with the third power of the number of degrees of freedom.

In most cases, however, the second derivatives of the energy cannot be obtained easily, and the challenge is to evaluate F^{eff} using only the first derivatives, i.e., the atomic forces, and without even evaluating the Hessian matrix. There are several ways to do this. One is to construct a dimer, i.e., two replicas of the system held at a small, fixed distance apart. The lowest energy orientation of the dimer is the direction of the minimum mode, $\hat{\lambda}$, and it can be found by iterative minimization of the energy using only the first derivatives (Henkelman and Jónsson 1999). An analogous method was formulated at a similar time (Munro and Wales 1999). Another way to find the minimum mode is to use the iterative Lanczos or,

even better, the Davidson method (Malek and Mousseau 2000; Olsen et al. 2004; Gutiérrez et al. 2016).

Since the system is displaced uphill in energy along the minimum mode, the method is often referred to as minimum mode following (MMF). The first step is to displace the atoms slightly from the initial state minimum and drive the system in some way into the region where the lowest eigenvalue of the Hessian is negative. From then on, the system is displaced in the direction of F^{eff} until the magnitude of the force drops below a given tolerance. The minimum mode will often change only slightly between iterations. Therefore, it is possible to save computational effort by only evaluating the minimum mode every few iterations (Gutiérrez et al. 2016). The MMF method is discussed in more detail in chapter "The Chapter on AKMC Methods".

Figure 3 shows MMF calculations for two-dimensional test problems. There, a random displacement from the initial state minimum is first applied, and then it is advanced in the direction of the minimum mode until the lowest eigenvalue becomes

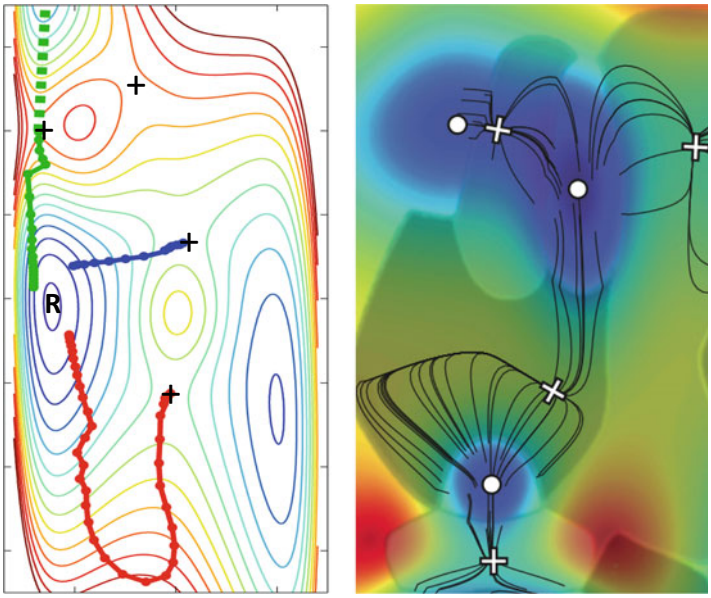

Fig. 3 An illustration of the minimum mode following method for finding first-order saddle points given only the initial state minimum (blue). Saddle points are marked with + signs. Left: Three saddle point searches started from different, random displacements from the minimum. In one of the searches (green), a kink on the path is evident when the lowest eigenvalue becomes negative and the perpendicular relaxation is turned on. After locating the saddle point, the system is displaced in the direction of the unstable mode, followed by minimization to locate the final state minimum. Right: A larger number of saddle point searches started from three different initial states. The region around each saddle point where at least one eigenvalue is negative is illustrated with a lighter color. In some cases, several different search paths trace the same final approach to the saddle point

negative. From that point on, the system is advanced in the direction of F^{eff}, and relaxation perpendicular to the minimum mode is included. This can result in a kink in the path. The figure also shows that several paths can approach a saddle point in a similar way even when they start out being quite different.

An application to a materials problem is shown in Fig. 4 where various different migration mechanisms of a kink on a dislocation in silicon crystal are identified (Pedersen et al. 2009a). The calculation made use of the EDIP empirical potential function (Justo et al. 1998). Three different mechanisms were found from multiple saddle point searches. The optimal mechanism, the one corresponding to lowest activation energy, involves an intermediate state where three Si atoms have fivefold coordination. Two of the mechanisms, including the optimal one, involve an intermediate state.

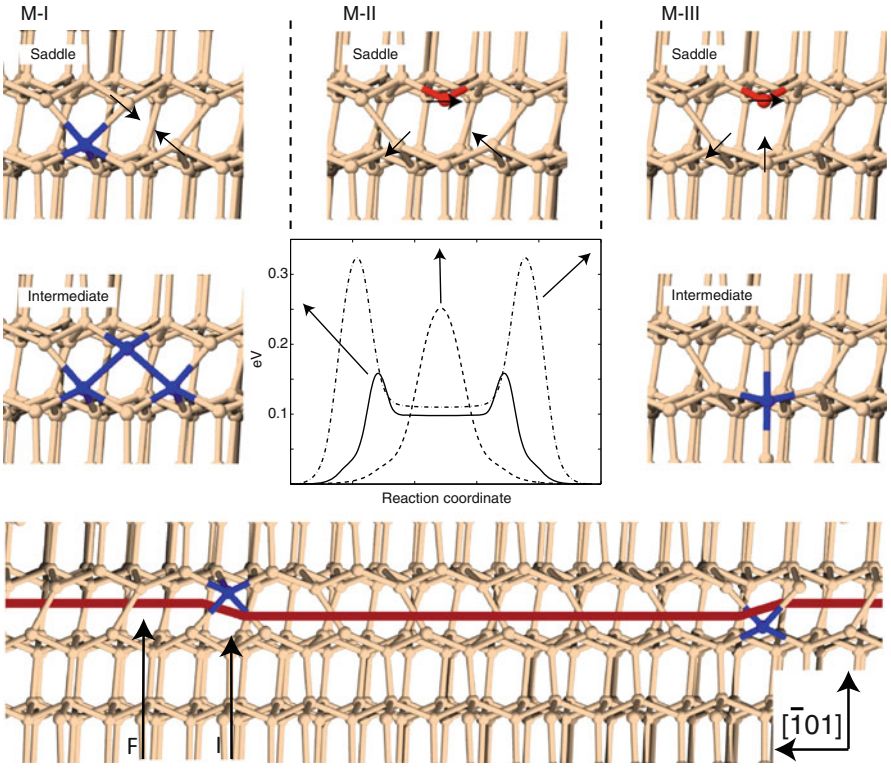

Fig. 4 Application of the minimum mode following method to find the mechanism of dislocation kink migration in crystalline silicon, based on the EDIP empirical potential function. Fivefold coordinated atoms are shown in blue and threefold coordinated atoms in red. The three lowest energy saddle points are shown (top row), two of them lead to an intermediate state (middle row). The initial state with two kinks separated by a large distance is shown (bottom), with a red line indicating the dislocation. (From Pedersen et al. 2009a)

Minimum mode saddle point searches have been used in various studies, for example, of the binding and diffusion of H_2O molecules on the surface of proton-disordered ice (Pedersen et al. 2014, 2015) and structure and diffusion in metal grain boundaries (Pedersen et al. 2009b; Pedersen and Jónsson 2009). Several other applications are discussed in chapter "AKMC Methods".

5 Additional Characterization of the Energy Surface

In this section, a few additional saddle point search methods that can be used to further characterize the energy surface of a system are discussed.

5.1 Search for Optimal MEPs

So far, the discussion of the MEP calculations has focused on finding only the MEP closest to the initial path. Sometimes the energy surface is simple enough, and the initial and final states are similar enough that only a single MEP connects the two. But, more generally, there can be two or more MEPs for a given pair of endpoints. One example is the CH_4 formation shown in Fig. 3. The CH_3 admolecule could jump to the left instead of jumping to the right in the first step to make room for the H atom on the surface. Also, the subsurface H atom could jump to an adjacent subsurface site before going to the surface, analogous to an MEP found for H_2 formation (Henkelman et al. 2006). These various paths would all have the same critical step, namely, the combination of an H adatom and a CH_3 admolecule on the surface, so they would, for practical purposes, be equivalent.

For other systems, it may be important to sample the MEPs to find the one that corresponds to lowest overall activation energy. One example of such a system is the nucleation of dislocation on a strained Ge overlayer (Maras et al. 2017). Since the atomic model of the dislocation involves a large number of atoms, there are many possible mechanisms, and each one involves a large number of intermediate minima. A global optimization strategy based on an evolutionary algorithm has been applied to this problem (Maras et al. 2017). There, segments of different paths are mixed and matched and a catalog maintained of all the intermediate minima found so far. More work needs to be done to develop optimal strategies for such complex systems.

5.2 Energy Ridge Tracking

The accuracy of the HTST approximation rests on the assumption that second-order saddle points are significantly higher in energy, on the scale of $k_B T$, than the first-order saddle points. In order to check that, a calculation of a path lying along the energy ridge between two first-order saddle points can be carried out. The method, ridge tracking nudged elastic band (RT-NEB), is an extension of the

CI-NEB method where each discretization point is a dimer of images (Maronsson et al. 2012). It combines, in a sense, the technique discussed above for finding an MEP and the dimer method for finding a saddle point given only the initial state. Just as the CI-NEB method can reveal the presence of intermediate minima that were not known beforehand, the RT-NEB method can reveal the presence of first-order saddle points in between known saddle points. It can, therefore, help complete the table of possible transitions a system can undergo (Maronsson et al. 2012).

5.3 Saddle Points for Quantum Tunneling

At low enough temperature, quantum mechanical tunneling becomes the dominant transition mechanism rather than hops over the energy barrier. An estimate of the crossover temperature can be obtained from the MEP (Gillan 1987). The larger the curvature of the MEP and the lower the effective mass along the unstable mode at the saddle point, the higher the crossover temperature.

The quantum statistical mechanics of the system can be obtained conveniently from thermal path integrals (Feynman and Hibbs 1965). Within that formalism, the classical energy surface gets replaced by an energy surface for a ring polymer where the beats are replicas of the system connected with temperature-dependent springs. This gives an effective, quantum mechanical energy surface that depends on temperature. The higher the temperature and the larger the mass of the atoms, the stiffer the springs and more classical the system becomes (Feynman and Hibbs 1965). A quantum mechanical extension of transition state theory can be formulated in terms of the rate of transitions of the ring polymer from the initial state to the final state (Mills et al. 1997, 1998; Richardson and Althorpe 2009; Hele and Althorpe 2013). Within a harmonic approximation, the optimal tunneling path is a first-order saddle point on the quantum mechanical energy surface (Benderskii et al. 1994; Mills et al. 1998), and the rate of thermally activated tunneling can be estimated from its properties (Benderskii et al. 1994; Richardson 2016). The optimal tunneling path is often referred to as an "instanton."

Saddle point search methods, such as the MMF method, can be used to find instantons (Andersson et al. 2009; Rommel and Kästner 2011). However, a more efficient approach is to use a path optimization method that minimizes a line integral along the path. This line integral nudged elastic band (LI-NEB) method can be used to find the optimal tunneling path at a given energy based on the semiclassical JWKB approximation. There, the images can be distributed evenly. Then, after the path has been found, a larger number of images need to be distributed unevenly along the path to represent an instanton, the optimal Feynman path for tunneling at a given temperature (Einarsdóttir et al. 2012; Ásgeirsson and Jónsson 2018). The beads of the ring polymer have a higher density in the lower-energy regions, making the direct search for the instanton using the MMF method less efficient than the LI-NEB.

Examples of tunneling paths at two different values of the temperature are shown in Fig. 5 for the same two-dimensional energy surface as in Fig. 1. The calculation

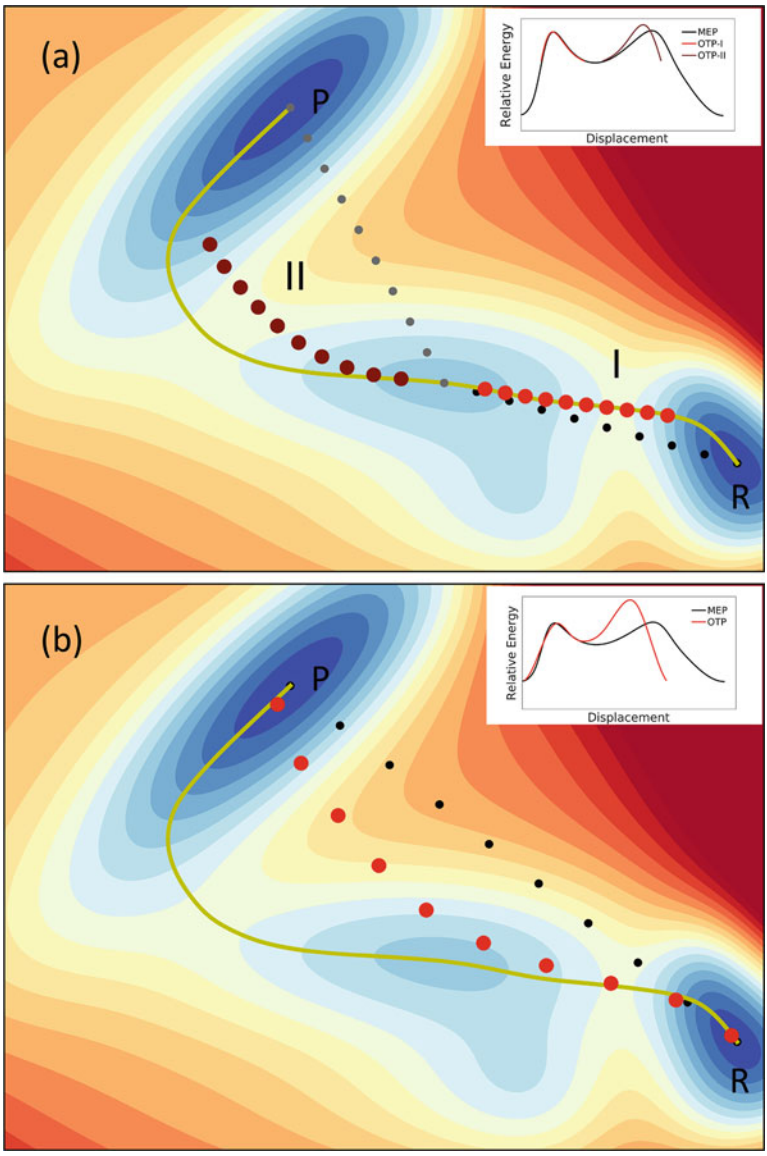

Fig. 5 Optimal tunneling paths found using the line integral nudged elastic band method. The energy surface is the same as in Fig. 1. (**a**) Tunneling occurs in two steps, first into the intermediate state and then the final state. The initial, straight line interpolation path is shown with black and gray dots, while the optimal tunneling paths are shown with red and brown dots. The corner cutting of the tunneling path is particularly clear at saddle point II where the MEP has a larger curvature. (**b**) Tunneling occurs in one step between states R and P, without visiting the intermediate state. At low enough temperature, this becomes the dominant transition mechanism. The initial, straight line interpolation path is shown with black dots, while red dots show the optimal tunneling path. Insets: Energy along the minimum energy path and along the optimal tunneling paths. (From Ásgeirsson and Jónsson 2018)

of the first-order saddle points on the quantum mechanical surface, the instantons, serves as further explorations of the energy surface. In regions where the MEP is curved, the tunneling path cuts corners. At the higher temperature, the tunneling takes place in two steps, first from the initial state, R, to the intermediate state and then from there to the final state, P. At the lower temperature, the tunneling takes the system directly between states R and P without entering the intermediate state. Results of the LI-NEB calculation are shown where the images are evenly distributed along the optimal tunneling path.

6 Saddle Points for Magnetic Transitions

Transitions from one magnetic state of a material to another can also be slow on the time scale of the vibrations of magnetic moments (Braun 2012). The same considerations about time scale differences and rare events discussed above in the context of atomic rearrangements can apply to such transitions. As a result, the characterization of the energy surface as a function of the magnetic degrees of freedom using saddle point searches is also useful when studying magnetic systems. It turns out that a semiclassical, adiabatic approximation can usually be made where the angles specifying the orientation of the magnetic momentum vectors are treated as slow variables while the length of the magnetic momentum vectors is a fast variable, determined by the electronic structure (Antropov et al. 1996). A full semiclassical specification of the configuration of an atomic scale system should, therefore, include the location of the atomic nuclei as well as two angles, for example, the polar and azimuthal angles (θ and ϕ) for each atom. A transition can involve simultaneous rearrangement of atoms and reorientation of magnetic moments. Within a harmonic approximation, the mechanism of such a transition can be characterized by a saddle point on this higher-dimensional energy surface.

A brief account will be given here of tools for navigating on magnetic energy surfaces, analogous to the tools presented above for atomic rearrangements. For simplicity, the atom coordinates will be assumed to be fixed and only the orientation of the magnetic moments treated as variables. The configuration space of N magnetic moments is a direct product of N two-dimensional spheres. The fact that the configuration space is curved brings up special considerations for the saddle point search methods. Again, the activation energy for thermally activated transitions is given by the energy difference between a first-order saddle point and the initial state minimum. But the equation of motion for magnetic moments, the Landau-Lifshitz equation, gives a different pre-exponential factor in the HTST approximation for the rate constant (Bessarab et al. 2012, 2013).

6.1 Geodesic CI-NEB Method

The distance between two points in the configuration space of magnetic moments is given by the length of the geodesic path connecting the two points. This should be

taken into account in the CI-NEB calculations of minimum energy paths between two local minima on the magnetic energy surface. Also, the nudging needs to include a projection onto the tangent space of the N two-dimensional spheres. This variant of the method for finding MEPs for magnetic transitions is referred to as the geodesic nudged elastic band (GNEB) method (Bessarab et al. 2015).

An example of a calculation of an MEP for a two-dimensional magnetic system is shown in Fig. 6. The Hamiltonian here is of an extended Heisenberg form including the Dzyaloshinsky-Moriya interaction (Uzdin et al. 2018). The system initially

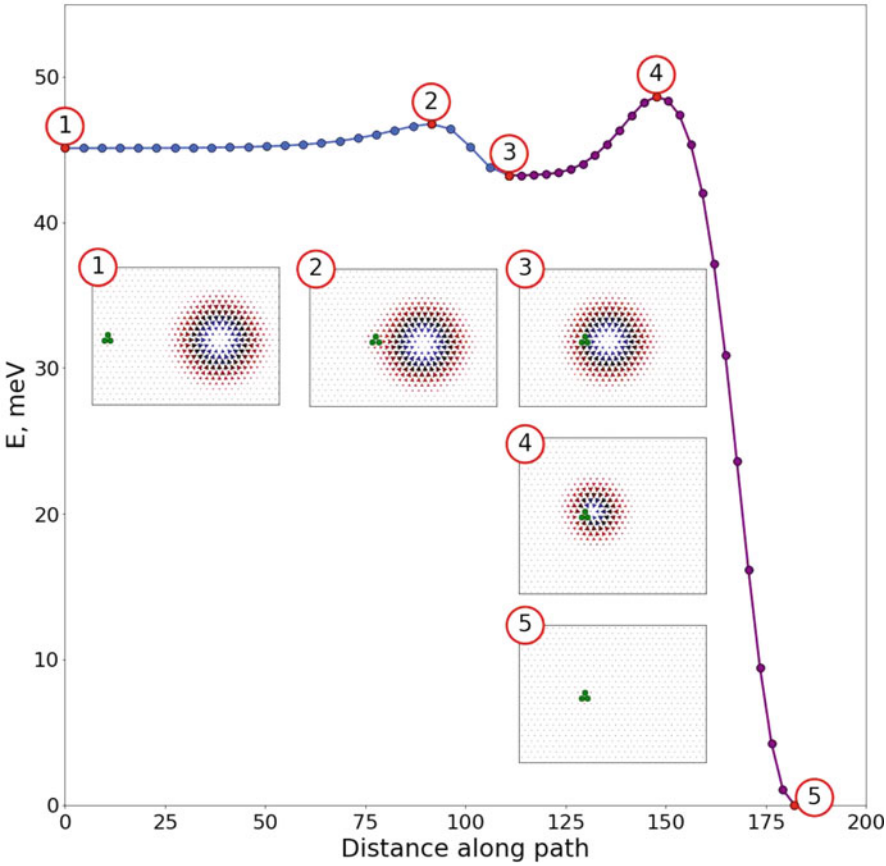

Fig. 6 Geodesic climbing image nudged elastic band method calculation of a minimum energy path for a magnetic transition where a skyrmion first attaches to a three-atom nonmagnetic impurity and then collapses. Red (blue) arrows show spins with z-component pointing up (down). Initially, the skyrmion and the nonmagnetic impurity are separated by a large distance, (1). By overcoming a small energy barrier, (2), a more stable state is reached where the skyrmion incorporates the impurity in a region where its magnetic moments lie in the plane (black arrows), (3). The skyrmion then collapses by overcoming a larger energy barrier (4) resulting in the ground-state, ferromagnetic phase (5). (From Uzdin et al. 2018)

contains a localized non-collinear structure called a skyrmion. There is great interest in such states as they represent a pseudo-particle that can be quite stable and could be used for future data storage and even data manipulation. The MEP shows the skyrmion approaching a triatomic, nonmagnetic impurity. By overcoming a small energy barrier, the impurity becomes incorporated in the skyrmion, in a region where the magnetic moments lie in the plane of the system. This lowers the energy of the system slightly. While a skyrmion in a continuous field is rigorously stable due to topological protection, its collapse can occur with a finite energy barrier on a discrete lattice (Bessarab et al. 2015). Here, the collapse occurs in the presence of the impurity. The mechanism involves shrinkage of the skyrmion and finally rotation of the core. For this set of parameters, the uniform ferromagnetic phase is lower in energy.

6.2 Mode Following in Curved Space

Since transitions involving reorientations of magnetic moments can be complex, many different types of transitions can be possible; a saddle point search method where only the initial state is specified is useful, just as for atomic rearrangements. The curvature of the configuration space needs to be taken into account when constructing the Hessian and calculating the eigenvalues and eigenvectors. An extension of mode following has recently been formulated for magnetic systems and applied to skyrmions (Müller et al. 2018). In addition to collapse and escape through a boundary, the method identified a transition where the skyrmion divides and forms two equivalent skyrmions. The activation energy for this duplication process was for some parameter values similar to that of collapse and escape. A dynamics simulation based on the Landau-Lifshitz equation showed that a time-dependent external field could induce the duplication process (Müller et al. 2018).

More work is needed to complete the implementation of the saddle point searches for systems of atoms with magnetic moments. First of all, an efficient method for constraining spin orientations in nonstationary configurations in electronic structure calculations needs to be developed. Then, a formulation of these various saddle point search methods for energy surfaces that are functions of both atomic coordinates and magnetic momentum orientations needs to be developed and implemented.

7 Conclusion

This chapter has summarized briefly various frequently used methods for finding saddle points. Such calculations can be used to reveal new states of the system by identifying local energy minima, not just saddle points. The height of the saddle point and the curvature of the energy surface around it as well as around the minimum can, within HTST, be used to estimate the rate of transitions between states. The simplest and most stable way to find saddle points is by finding MEPs

and identifying maxima along the path. This is now routinely done in the context of electronic structure calculations, DFT or ab initio. The MMF method where only the initial state is specified is more challenging and computationally demanding, and most calculations so far have been carried out with empirical potential functions. This will likely change in the future when more computational cores become available since the calculations involve several independent saddle point searches that can easily be carried out simultaneously via parallel or distributed computing (Chill et al. 2014a). Machine learning such as Gaussian process regression will also likely reduce the number of electronic structure calculations, as has already been demonstrated in CI-NEB calculations. Saddle point searches may even become an important component of global optimization algorithms (Pedersen et al. 2012; Plasencia et al. 2014). It is important to keep improving the algorithms to reduce the computational effort. A web page that collects algorithms and performance numbers for saddle point searches has been set up and is a valuable resource in this field (Chill et al. 2014b).

Acknowledgments This work was supported in part by the Icelandic Research Fund (grant 185405-051) and by the Academy of Finland (grant 278260). V.Á. acknowledges support from a Doctoral Grant of the University of Iceland Research Fund.

References

Andersen HC (1980) Molecular dynamics simulations at constant pressure and/or temperature. J Chem Phys 72(4):2384–2393

Andersson S, Nyman G, Arnaldsson A, Manthe U, Jónsson H (2009) Comparison of quantum dynamics and quantum transition state theory estimates of the H + CH_4 reaction rate. J Phys Chem A 113:4468

Antropov VP, Katsnelson MI, Harmon BN, van Schilfgaarde M, Kusnezov D (1996) Spin dynamics in magnets: equation of motion and finite temperature effects. Phys Rev B 54:1019

Ásgeirsson V, Jónsson H (2018) Efficient evaluation of atom tunneling combined with electronic structure calculations. J Chem Phys 148:102334

Benderskii VA, Makarov DE, Wight CA (1994) Chemical dynamics at low temperatures. Adv Chem Phys 88:1

Bessarab PF, Uzdin VM, Jónsson H (2012) Harmonic transition state theory of thermal spin transitions. Phys Rev B 85:184409

Bessarab PF, Uzdin VM, Jónsson H (2013) Potential energy surfaces and rates of spin transitions. Zeitschrift für Physikalische Chemie 227:1543

Bessarab PF, Uzdin VM, Jónsson H (2015) Method for finding mechanism and activation energy of magnetic transitions, applied to skyrmion and antivortex annihilation. Comput Phys Commun 196:335

Bitzek E, Koskinen P, G´ahler F, Moseler M, Gumbasch P (2006) Structural relaxation made simple. Phys Rev Lett 97(17):170201

Bligaard T, Jónsson H (2005) Optimization of hyperplanar transition states: application to 2D test problems. Comput Phys Commun 169:284

Bohner MU, Meisner J, Kastner J (2013) A quadratically-converging nudged elastic band optimizer. J Chem Theory Comput 9(8):3498–3504

Braun H-B (2012) Topological effects in nanomagnetism: from superparamagnetism to chiral quantum solitons. Adv Phys 61(1):1–116

Chill ST, Welborn M, Terrell R, Zhang L, Berthet J-C, Pedersen A, Jónsson H, Henkelman G (2014a) EON: software for long time simulations of atomic scale systems. Model Simul Mater Sci Eng 22:055002

Chill ST, Stevenson J, Ruehle V, Cheng S, Xiao P, Farrel JD, Wales DJ, Henkelman G (2014b) Benchmarks for characterization of minima, transition states, and pathways in atomic, molecular, and condensed matter systems. J Chem Theory Comput 10:5476–5482

Chu J-W, Trout BL, Brooks BR (2003) A super-linear minimization scheme for the nudged elastic band method. J Chem Phys 119(24):12708–12717

Ciccotti G, Ferrario M, Laria D, Kapral R (1995) Simulation of classical and quantum activated processes in the condensed phase. In: Reatto L, Manghi F (ed) Progress in computational physics of matter: methods, software and applications. World Scientific, Singapore, p 150

Einarsdóttir DM, Arnaldsson A, Óskarsson F, Jónsson H (2012) Path optimization with application to tunneling. Lect Notes Comput Sci 7134:45

Eyring H (1935) The activated complex in chemical reactions. J Chem Phys 3:107

Feynman RP, Hibbs AR (1965) Quantum mechanics and path integrals. McGraw-Hill, New York

Gillan MJ (1987) Quantum-classical crossover of the transition rate in the damped double well. Phys C Solid State Phys 20(24):3621–3641

Goumans TPM, Catlow CRA, Brown WA, Kästner J, Sherwood P (2009) An embedded cluster study of the formation of water on interstellar dust grains. Phys Chem Chem Phys 11(26):5431–5436

Gutiérrez MP, Argáez C, Jónsson H (2017) Improved minimum mode following method for finding first order saddle points. J Chem Theory Comput 13(1):125–134

Hele TJH, Althorpe SC (2013) Derivation of a true (t $\rightarrow$ 0+) quantum transition-state theory. I. Uniqueness and equivalence to ring-polymer molecular dynamics transition-state-theory. J Chem Phys 138:084108

Henkelman G, Jónsson H (1999) A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. J Chem Phys 111:7010

Henkelman G, Jónsson H (2000) Improved tangent estimate in the NEB method for finding minimum energy paths and saddle points. J Chem Phys 113(22):9978–9985 [Note: There is a typographical error in the Appendix, $2V_{i+1} - V_i$ should be $-2(V_{i+1} - V_i)$]

Henkelman G, Jónsson H (2001) Theoretical calculations of dissociative adsorption of methane on an Ir(111) surface. Phys Rev Lett 86:664

Henkelman G, Uberuaga BP, Jónsson H (2000a) A climbing image nudged elastic band method for finding saddle points and minimum energy paths. J Chem Phys 113(22):9901–9904

Henkelman G, Jóhannesson G, Jónsson H (2000b) Theoretical methods in condensed phase chemistry, methods for finding saddle points and minimum energy paths. In: Schwartz SD (ed) Progress in theoretical chemistry and physics. Kluwer Academic Publishers, Dordrecht, pp 269–302

Henkelman G, Arnaldsson A, Jónsson H (2006) Theoretical calculations of CH_4 and H_2 associative desorption from Ni(111): could subsurface hydrogen play an important role? J Chem Phys 124:044706

Jóhannesson GH, Jónsson H (2001) Optimization of hyperplanar transition states. J Chem Phys 115:9644

Jónsson H, Mills G, Jacobsen KW (1998) Nudged elastic band method for finding minimum energy paths of transitions. In: Berne BJ (ed) Classical and quantum dynamics in condensed phase simulations. World Scientific, Singapore, pp 385–404

Justo JF, Bazant MZ, Kaxiras E, Bulatov VV, Yip S (1998) Interatomic potential for silicon defects and disordered phases. Phys Rev B 58:2539

Keck JC (1967) Variational theory of reaction rates. J Chem Phys 13:85

Koistinen O-P, Maras E, Vehtari A, Jónsson H (2016) Minimum energy path calculations with Gaussian process regression. Nanosyst Phys Chem Math 7:925

Koistinen O-P, Dabgjartsdóttir F, Ásgeirsson V, Vehtari A, Jónsson H (2017) Nudged elastic band calculations accelerated with gaussian process regression. J Chem Phys 147(15): 152720

Kolsbjerg EL, Groves MN, Hammer B (2016) An automated nudged elastic band. J Chem Phys 145(9):094107
Malek R, Mousseau N (2000) Dynamics of Lennard-Jones clusters: a characterization of the activation-relaxation technique. Phys Rev E 62(6):7723–7728
Maras E, Trushin O, Stukowski A, Ala-Nissila T, Jónsson H (2016) Global transition path search for dislocation formation in Ge on Si(001). Comput Phys Commun 205:13
Maras E, Pizzagalli L, Ala-Nissila T, Jónsson H (2017) Atomic scale formation mechanism of edge dislocation relieving lattice strain in a GeSi overlayer on Si(001). Sci Rep 7:11966
Maronsson JB, Jónsson H, Vegge T (2012) A method for finding the ridge between saddle points applied to rare event rate estimates. Phys Chem Chem Phys 14:2884
Melander M, Laasonen K, Jónsson H (2015) Removing external degrees of freedom from transition-state search methods using quaternions. J Chem Theory Comput 11(3): 1055–1062
Mills G, Jónsson H, Schenter GK (1995) Reversible work based transition state theory: application to H_2 dissociative adsorption. Surf Sci 324:305–337
Mills G, Schenter GK, Makarov DE, Jónsson H (1997) Generalized path integral based quantum transition state theory. Chem Phys Lett 278:91
Mills G, Schenter GK, Makarov DE, Jónsson H (1998) RAW quantum transition state theory. In: Berne BJ, Ciccotti G, Coker DF (eds) Classical and quantum dynamics in condensed phase simulations. World Scientific, Singapore, pp 405–421
Müller GP, Bessarab PF, Vlasov FM, Lux F, Kiselev NS, Blügel S, Uzdin VM, Jónsson H (2018) Duplication, collapse and escape of magnetic skyrmions revealed using a systematic saddle point search method. Phys. Rev. Lett. 121:197202
Munro LJ, Wales DJ (1999) Defect migration in crystalline silicon. Phys Rev B 59:3969
Nocedal J (1980) Updating quasi-Newton matrices with limited storage. Math Comput 35(151):773–782
Olsen RA, Kroes G-J, Henkelman G, Arnaldsson A, Jónsson H (2004) Comparison of methods for finding saddle points without knowledge of the final states. J Chem Phys 121(20):9776–9792
Pedersen A, Jónsson H (2009) Simulations of hydrogen diffusion at grain boundaries in aluminum. Acta Materialia 57:4036
Pedersen A, Pizzagalli L, Jónsson H (2009a) Finding mechanism of transitions in complex systems: formation and migration of dislocation kinks in a silicon crystal. J Phys Condens Matter 21:084210
Pedersen A, Henkelman G, Schioetz J, Jónsson H (2009b) Long timescale simulation of a grain boundary in copper. New J Phys 11:073034
Pedersen A, Berthet J-C, Jónsson H (2012) Simulated annealing with coarse graining and distributed computing. Lect Notes Comput Sci 7134:34
Pedersen A, Wikfeldt KT, Karssemeijer LJ, Cuppen HM, Jónsson H (2014) Molecular reordering processes on ice (0001) surfaces from long timescale simulations. J Chem Phys 141:234706
Pedersen A, Karssemeijer LJ, Cuppen HM, Jónsson H (2015) Long-timescale simulations of H_2O admolecule diffusion on Ice Ih(0001) surfaces. J Phys Chem C 119:16528
Peters B (2017) Reaction rate theory and rare events. Elsevier Science & Technology, Amsterdam
Peterson AA (2016) Acceleration of saddle-point searches with machine learning. J Chem Phys 145(7):074106
Plasencia M, Pedersen A, Arnaldsson A, Berthet J-C, Jónsson H (2014) Geothermal model calibration using a global minimization algorithm based on finding saddle points as well as minima of the objective function. Comput Geosci 65:110
Richardson JO (2016) J Chem Phys 144:114106
Richardson JO, Althorpe SC (2009) Ring-polymer molecular dynamics rate-theory in the deep-tunneling regime: connection with semiclassical instanton theory. J Chem Phys 131:214106
Rommel JB, Kästner J (2011) Adaptive integration grids in instanton theory improve the numerical accuracy at low temperature. J Chem Phys 134:184107
Sheppard D, Terrell R, Henkelman G (2008) Optimization methods for finding minimum energy paths. J Chem Phys 128(13):134106

Sheppard D, Xiao P, Chemelewski W, Johnson DD, Henkelman G (2012) A generalized solid-state nudged elastic band method. J Chem Phys 136(7):074103
Smidstrup S, Pedersen A, Stokbro K, Jónsson H (2014) Improved initial guess for minimum energy path calculations. J Chem Phys 140(21):214106
Sørensen MR, Jacobsen KW, Jónsson H (1996) Thermal diffusion processes in metal tip-surface interactions: contact formation and adatom mobility. Phys Rev Letters 77(25):5067–5070
Trygubenko SA, Wales DJ (2004) A doubly nudged elastic band method for finding transition states. J Chem Theory Comput 120(5):2082–2094
Uzdin VM, Potkina MN, Lobanov IS, Bessarab PF, Jónsson H (2018) Energy surface and lifetime of magnetic skyrmions. J Magn Magn Mater 459:236–240
Vineyard GH (1957) Frequency factors and isotope effects in solid state rate processes. J Phys Chem Solids 3:121
Voter AF, Doll JD (1985) Dynamical corrections to transition state theory for multistate systems: surface self-diffusion in the rare-event regime. J Chem Phys 82(1):80–92
Wigner E (1938) The transition state method. Trans Faraday Soc 34:29
Weinan E, Weiqing R, Vanden-Eijnden E (2002) String method for the study of rare events. Phys Rev B 66(4):052301
Zarkevich NA, Johnson DD (2015) Nudged-elastic band method with two climbing images: finding transition states in complex energy landscapes. J Chem Phys 142(2):024106
Zhang J, Zhang H, Ye H, Zheng Y (2016) Free-end adaptive nudged elastic band method for locating transition states in minimum energy path calculation. J Chem Phys 145(9):094104
Zhu T, Li J, Samanta A, Kim HG, Suresh S (2007) Interfacial plasticity governs strain rate sensitivity and ductility in nanostructured metals. Proc Natl Acad Sci USA 104(9):3031–3036